10.1098/rsta.2003.1356

THE ROYAL SOCIETY

Foundations of chaotic mixing

By Stephen Wiggins[1] and Julio M. Ottino[2]

[1]*School of Mathematics, University of Bristol, University Walk, Bristol BS8 1TW, UK* (s.wiggins@bristol.ac.uk)

[2]*Departments of Chemical and Biological Engineering and Mechanical Engineering, R. R. McCormick School of Engineering and Applied Sciences, Northwestern University, Evanston, IL 60208, USA* (jm-ottino@northwestern.edu)

Published online 11 March 2004

The simplest mixing problem corresponds to the mixing of a fluid with itself; this case provides a foundation on which the subject rests. The objective here is to study mixing independently of the mechanisms used to create the motion and review elements of theory focusing mostly on mathematical foundations and minimal models. The flows under consideration will be of two types: two-dimensional (2D) 'blinking flows', or three-dimensional (3D) duct flows. Given that mixing in continuous 3D duct flows depends critically on cross-sectional mixing, and that many microfluidic applications involve continuous flows, we focus on the essential aspects of mixing in 2D flows, as they provide a foundation from which to base our understanding of more complex cases. The baker's transformation is taken as the centrepiece for describing the dynamical systems framework. In particular, a hierarchy of characterizations of mixing exist, *Bernoulli* $\rightarrow$ *mixing* $\rightarrow$ *ergodic*, ordered according to the quality of mixing (the strongest first). Most importantly for the design process, we show how the so-called linked twist maps function as a minimal picture of mixing, provide a mathematical structure for understanding the type of 2D flows that arise in many micromixers already built, and give conditions guaranteeing the best quality mixing. Extensions of these concepts lead to first-principle-based designs without resorting to lengthy computations.

Keywords: microfluidics; mixing; chaos; diffusion; fluid mechanics; dynamical systems theory

1. Introduction

In the past few years microfluidics has come of age. The applications are many, and with the applications came the need to mix small volume of fluids in low-Reynolds-number flows. Many of the applications are described in this issue (Ottino & Wiggins 2004). To this one can add many applications in the area of micro-reactors (Jensen 1999; Losey *et al*. 2002), although the length-scale in these applications can be considerably larger and often referred to as 'mini' rather than 'micro'. It is nevertheless in this territory that the elements of chaotic-mixing theory find fertile ground. Numerous experiments and theory developed since the early 1980s suggest ways to produce efficient mixing in viscous-dominated flows.

One contribution of 11 to a Theme 'Transport and mixing at the microscale'.

The objective of this paper is to review in a brief manner a few elements of the necessary theory focusing mostly on mathematical foundations. Many recent applications of mixing in microfluidic applications can benefit by a closer linkage and use of basic theory. Relevant material exists (Aref 1984), including presentations by the present authors themselves (Ottino 1989; Wiggins 1992). Nevertheless, it may be argued that mathematical advances that have taken place in the last few years are not properly accounted for in these works and that often the presentations do not qualify as tutorial. At the same time it seems of some benefit to have most of the relevant terminology condensed in a single source. More importantly we want to show 'what can be done' and 'how it can be done' rather than showing precisely 'how it is done' to the point of transforming the theory into an engineering design procedure. It should nevertheless be clear that a detailed design procedure is feasible and that concepts presented here lead to first-principle-based designs without resorting to lengthy computations.

Everything we will say corresponds to the case of a single fluid and vanishingly small Reynolds numbers. Given the length-scales of typical microfluidic applications (*ca.* 100 μm) flows are inertialess and what happens for finite Reynolds number is, in general, of limited applicability. The restriction to low Reynolds number is crucial; time does not enter directly into the equations and may be considered a parameter.

The mechanisms used to create the motion are of no concern here. Suffice it to say that there are many: electro-osmotic effects, use of patterned walls, systems based on clever use of geometric re-orientations (as in static mixers), etc. The important aspect is that flows can be driven by various effects that are simply ineffective at larger scales. It may be argued that the science base for microfluidics is already largely developed and what remains is technological development. This is a technological area dominated by ingenuity, and exploitation of increasingly clever mechanisms and device development may be expected in the near future. The room for creativity is significant; several possible mixing mechanisms, such as using small particles as mixing devices and mixing within droplets, exploit phenomena that are manifestly inefficient at larger scales.

The goal is to study the consequences of the motion *independently of the mechanisms used to create the motion itself*, and even the character of the fluid. The point of view here is to focus on the very foundation of chaotic mixing (the simplest possible case of stretching and folding or stretching and cutting) and to establish clear definitions. The point of view is solely kinematical; diffusion of passive scalars, though important in many applications, can simply be added to the stretching and folding template (Ottino 1989). What remains is non-trivial though. We argue that a certain type of map studied in dynamical systems theory, the so-called 'linked twist map' (LTM), is the key to the understanding a large class of microfluidic systems. Most importantly, LTMs provide an analytical framework where it is possible to design mixers for a rigorously provable strong form of mixing: Bernoulli. The pay-off for simplicity is generality. As with anything that is based on mathematics and theorems, results are absolute; it is in the application of the results where engineering judgment should be exercised.

We address two questions.

(i) What is the best mixing (and how do we define and quantify 'best')?

(ii) What are the necessary and sufficient conditions that lead to the best mixing?

To answer the first question we review elements of theory including some elements of ergodic theory. Some of the more technical details are placed in Appendix A. As to the second question, we will show that there is theoretical guidance in assuring conditions for best mixing, when the term *mixing* is defined in a mathematically rigorous way. We will exemplify concepts in terms of brief sketches belonging to recent examples from the literature.

Mixing in many microfluidic applications takes place largely in continuous flows and the key to effective mixing in continuous flows resides in the effectiveness of the cross-sectional flow. In the limit of perfect mixing in the cross-sectional flow, say in the plane x, y, elements sample all regions of the axial flow and this has consequences such as (qualitatively) uniform residence times. In general, a very large class of mixers can be regarded as a *piecewise succession of duct flows* (Franjione & Ottino 1991; Mezić *et al*. 1999), where the cross-flow, v_x, v_y, and the axial flow, v_z, are independent; i.e. neither the cross-flow nor the axial flow depend on the axial coordinate (note that this cannot happen if *Re* is finite). The simplest case corresponds to $v_z = \text{const}$. We will therefore focus on the key aspect of mixing, which is mixing in the two-dimensional (2D) cross-flow (how to incorporate the case of v_z being a function of x, y is explained in Appendix B).

2. Foundations

(*a*) *Dynamical systems terminology and concepts for describing mixing*

The following section reviews definitions that are needed in the rest of the presentation. What we will say is mathematically rigorous; however, we will strive to not make the presentation needlessly formal, and occasionally we will offer physical interpretations of the definitions. Arnold & Avez (1968) provide an accessible introduction to the ergodic theory of dynamical systems. Newer and more up-to-date texts are hard to follow on a first exposure to the subject. The relevant background from dynamical systems theory can be found in Wiggins (1990), and many other places as well, whereas the kinematical aspects of mixing are covered in Ottino (1989).

(i) *Terminology for general fluid kinematics*

Mappings. The motion of fluid particles is described mathematically with a *map* or *mapping*. Let R denote the region occupied by the fluid. We refer to points in R as *fluid particles*. The flow of fluid particles is mathematically described by a smooth, invertible *transformation*, or *map*, of R into R, denoted by S, also having a smooth inverse. The particles are labelled by their initial condition at some arbitrary time, usually taken as $t = 0$. The application of S to the domain R, denoted $S(R)$, is referred to as one *advection cycle*. Similarly, n advection cycles are obtained by n repeated applications of S, denoted $S^n(R)$. Let A denote any subdomain of R. Then $\mu(A)$ denotes the volume of A (if we are in a 2D setting, read 'volume' as 'area'). Thus μ is a function that assigns to any (mathematically well-behaved) subdomain of R its volume. In mathematical terminology the function μ is known as a *measure*. Incompressibility of the fluid is expressed by stating that, as any subdomain of R is stirred, its volume remains unchanged, i.e. $\mu(A) = \mu(S(A))$. In the language of dynamical systems theory, S is an example of a *measure-preserving transformation*. We will assume that R has finite volume, which is natural for the applications we

have in mind. In this case we can normalize the function that assigns the volume to a subdomain of R, so that without loss of generality we can assume that $\mu(R) = 1$. This will make certain mathematical definitions simpler, and is a standard assumption in the mathematics literature.

Orbits. For a specific fluid particle p, the *trajectory*, or *orbit*, of p is the sequence of points $\{\ldots, S^{-n}(p), \ldots, S^{-1}(p), p, S(p), S^2(p), \ldots, S^n(p), \ldots\}$. So the orbit of a point is nothing more than the sequence of points corresponding to the point, where the point has been (under the past advection cycles), and where the point will go (under future advection cycles).

(ii) *Specific types of orbits*

The next three definitions refer to specific orbits that often have a special significance for transport and mixing.

Periodic orbit. This orbit consists of a finite number of points (where the number of points is the period of the orbit) and has the property that during each application of the advection cycle each point on the orbit shifts to another point on the orbit. Periodic orbits may be distinguished by their stability type. In the typical case, periodic orbits are either stable (i.e. nearby orbits remain near the periodic orbit), referred to as elliptic, or unstable of saddle type (i.e. meaning that typical nearby orbits either move away from the periodic orbit, or approach the periodic orbit for a time, but ultimately move away), referred to as hyperbolic. If one is interested in the design of a micromixer, it may not seem particularly relevant to focus on particles that undergo periodic motion during the advection cycle. However, they are often the 'template' of the global mixing properties. For example, elliptic periodic orbits are bad for mixing, as they give rise to regions that do not mix with the surrounding fluid ('islands'). Hyperbolic periodic orbits provide mechanisms for contraction and expansion of fluid elements, and they can also play a central role in the existence of Smale horseshoe maps, which may lead to efficient global mixing properties.

Homoclinic orbit. This is an orbit that, asymptotically in positive time (i.e. forward advection cycles), approaches a hyperbolic periodic orbit, and asymptotically in negative time (i.e. inverse advection cycles) approaches the same periodic orbit. These types of orbits are significant because in a neighbourhood of such an orbit a Smale horseshoe map can be constructed (the 'Smale–Birkhoff homoclinic' theorem).

Heteroclinic orbits and cycles. This is an orbit that, asymptotically in positive time (i.e. forward advection cycles), approaches a hyperbolic periodic orbit, and asymptotically in negative time (i.e. inverse advection cycles) approaches a *different* periodic orbit. If two or more heteroclinic orbits exist and are arranged in a *heteroclinic cycle* (i.e. roughly, this means that a closed curve can be constructed from the union of the periodic orbits and their intersecting stable and unstable manifolds), then it is generally possible to construct a Smale horseshoe map near the heteroclinic cycle in the same way that it is constructed near a homoclinic orbit.

(iii) *Behaviour near a specific orbit*

Lyapunov exponents are a ubiquitous diagnostic in the chaotic dynamics literature. It is important to understand that they are numbers associated with *one orbit*. Additional information (such as ergodicity, to be discussed below) may allow us to extend this knowledge to larger regions of the domain of the map.

Lyapunov exponent. This number is associated with an orbit and describes its stability in the linear approximation (i.e. the growth rate of 'infinitesimal' perturbations). Elliptic periodic orbits have zero Lyapunov exponents. Hyperbolic periodic orbits have some positive and some negative Lyapunov exponents. In incompressible flows, the sum of all Lyapunov exponents for an orbit must be zero. It is important to realize that a Lyapunov exponent is an infinite-time average. Consequently, it can only be approximated in general. So-called *finite-time Lyapunov exponents* have been considered by various people (see, for example, Lapeyre 2002); however, it is important to understand that they are not on the same rigorous mathematical footing as standard Lyapunov exponents (Oseledec 1968), and their applicability to mixing must often be assessed on a case-by-case basis.

The limitations associated with the fact that Lyapunov exponents characterize *infinitesimal* separations of orbits have been addressed with the development of *finite-size Lyapunov exponents* (see, for example, Boffeta *et al*. 2001). This is an interesting, and potentially important, development, but it should be realized that, at present, there are no rigourous mathematical foundations for this concept. It is mainly a computational tool whose effectiveness must be addressed on a case-by-case basis, although it should be noted that work of Yomdin (1987) and Newhouse (1988) on the growth of curves in 2D flows and curves and surfaces in three-dimensional (3D) flows is certainly relevant to our needs.

(iv) *Collections of fluid particles that give rise to 'flow structures'*

'Flow structures' are routinely and commonly seen in many flow visualization experiments. Dynamical systems theory provides a way of describing what is seen in these experiments, and also a way of predicting their evolution and their dependence on parameters.

Invariant set. Let A be a subdomain of R. A is then said to be *invariant* under the advection cycle (or an *invariant set*) if $S(A) = A$. That is, all points in A remain in A under repeated applications of the advection cycle. Clearly, invariant sets strictly smaller than R are bad for mixing, since they represent subdomains of the flow that do not mix with the rest of R, *except* via the mechanism of molecular diffusion (but we are restricting our discussion solely to kinematical mechanisms for mixing). The orbit of a point p and a homoclinic orbit are examples of invariant sets (but ones with zero volume).

Kolmogorov–Arnold–Moser (KAM) theorem. The KAM theorem is concerned with the existence of quasi-periodic orbits in perturbations of integrable Hamiltonian systems, or volume-preserving maps. These orbits densely fill out tori or 'tubes'. These

tubes are therefore material surfaces, and fluid particles cannot cross them. Consequently, these tubes trap regions of fluid that cannot mix with their surroundings (without molecular diffusion). The theorem may seem essentially useless for direct application in the sense that it is rare to be in an 'almost integrable' situation. Moreover, even if that were the case, rigorous verification of the hypotheses may be quite difficult, and even impossible. The theorem does imply the existence of *islands* in the neighbourhood of an elliptic periodic orbit (discussed below). In this sense, the KAM theorem is surprisingly 'effective' and describes a phenomenon that has been observed to occur very generally in Hamiltonian systems and is present in virtually every computed example of Poincaré sections resulting in area-preserving maps. It has become traditional to refer to all such material tubes or tori as 'KAM tori', even if they are observed in situations where the theorem does not rigorously apply, or cannot be applied. Such tubes have been observed experimentally (Kusch & Ottino 1992; Fountain *et al*. 1998).

Island. Elliptic periodic orbits are significant because, according to the KAM theorem, they are surrounded by 'tubes' which trap fluid. Moreover, these tubes exhibit a strong effect on particles outside the tube, but close to the tube. In a mathematically rigorous sense, these tubes are 'sticky' (Perry & Wiggins 1994). The tubes and the neighbouring region that they influence in this way are referred to as 'islands'. Clearly, islands inhibit good mixing.

Barriers to transport and mixing. In certain circumstances there can exist surfaces of one fewer dimension than the domain R that are made up entirely of trajectories of fluid particles, i.e. *material surfaces*. Consequently, fluid-particle trajectories cannot cross such surfaces and in this way they are *barriers to transport*. *KAM tori* are examples of *complete* barriers to transport: fluid-particle trajectories starting inside remain inside forever. *Partial* barriers to transport are associated with hyperbolic periodic orbits. The collection of fluid-particle trajectories that approach the hyperbolic periodic orbit asymptotically as time goes to positive infinity forms a material surface called the *stable manifold of the hyperbolic periodic orbit*. Similarly, the collection of fluid-particle trajectories that approach the hyperbolic periodic orbit asymptotically as time goes to negative infinity forms a material surface called the unstable *manifold of the hyperbolic periodic orbit*. *Homoclinic orbits* can therefore be characterized as orbits that are in the intersection of the stable and unstable manifolds of a hyperbolic periodic orbit. Similarly, *heteroclinic orbits* can be characterized as orbits that are in the intersection of the stable manifold of one hyperbolic periodic orbit with the unstable manifold of another hyperbolic periodic orbit.

Lobe dynamics. As mentioned above, the stable and unstable manifolds of hyperbolic period orbits are material curves and, therefore, fluid-particle trajectories cannot cross them. However, they can deform in very complicated ways and result in intricate flow structures. The resulting flow structure is the spatial, geometrical template on which the transport and mixing takes place in time. *Lobe dynamics* provides a way of describing the geometrical structure and quantifying the resulting transport. See Rom-Kedar & Wiggins (1990) and Wiggins (1992) for the general theory,

Camassa & Wiggins (1991) and Horner *et al.* (2002) for an application to a time-dependent 2D cellular flow and experiments and Beigie *et al.* (1994) for further applications and development.

(v) *Characterizing global aspects of fluid particle kinematics*

By 'global aspects' we mean a feature related to the collective motion of a 'large' (this will usually mean non-zero volume) set of orbits. All of the definitions above refer to specific fluid-particle trajectories, or collections of fluid-particle trajectories, of the advection cycle. An understanding of these different types of fluid-particle trajectories is important for quantifying the quality of mixing (up to now the term 'mixing' has been only loosely defined). As we shall see, mixing can be good in the region of a Smale horseshoe, and bad in the region of islands. However, it would be helpful if we could provide a characterization of the *global* mixing properties of the advection cycle. In particular, we would like to characterize *optimal mixing*. The following notions from ergodic theory can accomplish this task, in principle, but verifying that they occur in specific mixers, or how to design mixers so that they hold, requires additional considerations.

Ergodicity. Let x denote a point in R and let $G(x)$ denote a function defined on R. Then the time average of $G(x)$ along the orbit of x, denoted $G^*(x)$, is defined by

$$G^*(x) = \lim_{N\to+\infty} \frac{1}{N} \sum_{n=0}^{N-1} G(S^n(x)).$$

The space average of $G(x)$, denoted $\bar{G}$, is defined by

$$\bar{G} = \int_D G(x)\,\mathrm{d}\mu$$

(recall that we have normalized the measure of the volume of R so that $\mu(R) = 1$). The transformation S is said to *ergodic* if, for a sufficiently large class of functions, denoted by the set of functions $G(x)$, we have $G^*(x) = \bar{G}$ (except, possibly, on certain sets of zero volume). In other words, the time average of functions along an orbit (with the possible exception of a set of orbits of zero volume, or *measure zero*) is equal to the space average. This is significant because clearly the space average of a function is just a number. From its definition, the time average may vary from point to point, but *not* if the transformation S is ergodic. We remark that even if the transformation is not ergodic on all of R, it may be ergodic on invariant subsets of R. In fact, there is a mathematically rigorous way of partitioning R into invariant sets, where the transformation is ergodic on each invariant set. This *ergodic partition* (an idea originating in the work of von Neumann, Halmos and Rokhlin) is described in Mezić & Wiggins (1999), where it is further developed as a flow visualization method.

It is instructive to mention *Lyapunov exponents* in this context. As stated in § 2 *a* (iii), Lyapunov exponents are the time averages of quantities related to the stretching and contraction rates along *individual orbits*. Therefore, if the transformation is ergodic, almost every orbit has the same Lyapunov exponents. ('Almost every' means that the orbits that do not satisfy this property have zero volume.)

Of course, in mixing applications we are generally not interested in time averages and space averages of just *any* function, but functions of physical relevance to the mixing process.

Mixing. This is a critical concept. In fact, the astute reader will have noted that we have used the phrases 'good mixing', 'efficient mixing' and 'poor mixing' above, but we have never defined what we mean by the term 'mixing' (a situation that is all too common in practice). Within the domain R let B denote a region of, say, black fluid and let W denote any other region within R. Mathematically, we denote the amount of black fluid that is contained in W after n applications of the mixing process by $\mu(S^n(B) \cap W)$, that is the volume of $S^n(B)$ that ends up in W after n advection cycles. Then the fraction of black fluid contained in W is given by

$$\frac{\mu(S^n(B) \cap W)}{\mu(W)}.$$

Intuitively, the definition of mixing (as the number of applications of the advection cycles is increased) is that for *any* region W we would have the same fraction of black fluid as for the entire domain R, i.e.

$$\frac{\mu(S^n(B) \cap W)}{\mu(W)} - \frac{\mu(B)}{\mu(R)} \to 0 \quad \text{as } n \to \infty$$

or, since we have taken $\mu(R) = 1$, $\mu(S^n(B) \cap W) - \mu(B)\mu(W) \to 0$ as $n \to \infty$. In fact, this is the mathematical definition of a *mixing* transformation. For *any* subdomains B and W, $\mu(S^n(B) \cap W) - \mu(B)\mu(W) \to 0$ as $n \to \infty$. Thinking of this in a probabilistic manner, this means that, given any subdomain, upon iteration it becomes (asymptotically) independent of any other subdomain. Mixing transformations are *ergodic*, but *ergodic* transformations need not be mixing.

It is important to point out (although it was probably apparent to the practically minded reader) that the definitions of ergodic and mixing are 'infinite-time' quantities in the sense that certain limits as the number of advection cycles approach infinity must be considered. (Indeed, many mathematicians define the subject of dynamical systems theory as the study of the 'asymptotic in time' or 'long-time' behaviour of the system.) Clearly, this is not desirable for applications where one would like the number of advection cycles to be as small as possible in order to achieve the limit. The study of the rate of approach to the limit that defines mixing transformations is a very difficult problem and a subject of much current interest in the ergodic theory community. This topic is referred to as the 'decay of correlations'.

Decay of correlations. In many areas of applications the *decay of correlations of a scalar field* is used as a diagnostic for quantifying mixing. We now show how this concept is related to the definition of mixing given above. First, we need a few technical definitions and notation. For a region B in R the function χ_B, referred to as the *characteristic function associated with* B, assigns a '1' to a point that is in B, and '0' to all other points. Then the volume of B can be written as $\int \chi_B \, \mathrm{d}\mu = \mu(B)$ (think of $\mathrm{d}\mu$ as the 'infinitesimal volume element' of integration). Similarly $\chi_{S^n(B)\cap W}$ assigns 1 to a point that is at the intersection of the nth mapping of B with W, and 0

to all other points. So in terms of integrals over characteristic functions, the limit in the definition of mixing above can be written as

$$\int \chi_{S^n(B)\cap W}\,\mathrm{d}\mu - \int \chi_W\,\mathrm{d}\mu \int \chi_B\,\mathrm{d}\mu \to 0 \quad \text{as } n \to \infty,$$

and we can view this as an *integral formulation* of the mixing condition. This expression can be written in a more useful form. First, note that

$$\chi_{S^n(B)\cap W} = \chi_W \chi_{S^n(B)}.$$

Now $S^n(B)$ is the set of points that map to B under S^{-n}. Therefore,

$$\chi_{S^n(B)} = \chi_B \circ S^{-n}.$$

Putting all of this together, the integral expression derived above can be rewritten as

$$\int \chi_W(\chi_B \circ S^{-n})\,\mathrm{d}\mu - \int \chi_W\,\mathrm{d}\mu \int \chi_B\,\mathrm{d}\mu \to 0 \quad \text{as } n \to \infty.$$

The latter expression suggests a modification of the definition where we might replace the characteristic functions with arbitrary functions (from a certain class of functions of interest). So for functions f and g (from the class of interest) we define the *correlation function*

$$C_n(f,g) \equiv \left| \int g(f \circ S^{-n})\,\mathrm{d}\mu - \int g\,\mathrm{d}\mu \int f\,\mathrm{d}\mu \right|.$$

In the language of ergodic theory, f and g are referred to as *observables*, and the decay of the correlation function for general observables is considered. In applications, it is usually a specific observable that is of interest. In applications one typically takes $f = g =$ 'a scalar field', and considers the *decay of correlations of a scalar field*. Of course, if the transformation is not mixing then we should not expect the correlations to decay to zero. If the transformation is mixing, however, then the rate of decay of correlations is a quantifier of the *speed of mixing*.

A study of the decay of correlations of different classes of mappings is currently at the forefront of research in dynamical systems theory. An excellent review is given by Baladi (1999). Results have been obtained where the decay of correlations is exponential, and where the decay is polynomial. Unfortunately, there are no results that rigorously apply to the types of maps that arise in mixing applications. However, it has been proven that on the (zero volume) invariant set associated with horseshoe maps (technically, and importantly, the invariant set needs to be *uniformly hyperbolic*), the decay of correlations is exponentially fast (horseshoe maps are discussed in § 2 *a* (vi)). Evidently, this should influence nearby orbits, but the nature of this influence is unclear at the present time.

Entropy. There are many different types of entropy in the dynamical systems literature. The one we speak of here is 'topological entropy'. Quoting from Newhouse & Pignataro (1993), 'The topological entropy of a system is a quantitative measure of its orbit complexity. In a certain sense, it is the maximum amount of information lost per unit time by the system using measurements with finite precision'. Hence,

positive topological entropy indicates orbit complexity in this sense. Newhouse & Pignataro present a novel numerical method for estimating the topological entropy by estimating the logarithmic growth rates of curves. The theoretical basis for this numerical method is the work of Yomdin (1987) and Newhouse (1988) mentioned earlier. This is why the dynamical systems notion of entropy is relevant to mixing as it is related to the growth of material curves in the flow. A relation between positive topological entropy and Smale horseshoes was established by Burns & Weiss (1995). The famous Pesin formula (Pesin 1977) gives a relationship between metric entropy and Lyapunov exponents.

Chaos. The term 'chaos', associated with *chaotic mixing* and *chaotic advection*, has been missing from our list of dynamical systems concepts and terms up to this point. From its ubiquitous presence in the literature, one might think that it should have a position of prominence. However, with everything we have defined and discussed so far, it really is unnecessary. The reason for this is that chaos, in much of the literature, is more of a descriptive term than a quantitative term (such as 'ergodicity' or 'exponentially decaying correlations for mixing maps'). Even though one can find the term 'chaos' defined in textbooks (although not always consistently), it is probably fair to say that there is still no universal acceptance of a single definition (and there may never be). Nevertheless, practically speaking, a good working definition might be that a map is *chaotic on a bounded invariant set* if it is ergodic and the orbits have some positive Lyapunov exponents. Some people may want to weaken this definition. The Bernoulli property discussed below is probably a stronger property than what one might want to call 'chaos', even though the Bernoulli property certainly embodies how one might imagine a deterministic chaotic system to behave. In any case, this should give an idea of the difficulties of pinning down a universally accepted definition of the term 'chaos'. However, the good news is that this is of little practical consequence. The precise, quantitative definitions given above serve not only to define but also to quantify the concept of mixing.

(vi) *Specific classes of maps*

Bernoulli shift. After reading a few lines of this next definition, the reader interested in designing a micromixer may feel an overwhelming urge to skip the rest of the section. We strongly urge them to forge ahead because the Bernoulli shift embodies the essence of how the fluid particles move in an efficient mixing flow. The Bernoulli shift is a dynamical systems theory description of 'tossing a (fair) coin', a process that in the minds of most people is truly random. Still, it may appear to some as quite a stretch to relate (in some way) coin tossing to fluid mixing. However, it can be argued that under certain conditions (that can be realized and measured in the laboratory) a 'change of coordinates' can be constructed that transforms a mixing process to the Bernoulli shift. We will see that actually carrying out such a coordinate change is not critical. Rather, the fact that such a coordinate change *exists* means that our micromixer (or, some subdomain of the micromixer) immediately inherits the remarkable properties of the Bernoulli shift; properties that would be essentially impossible to verify by working solely with the mixer itself. In order to understand what these properties are we must first describe the Bernoulli shift as a dynamical system.

Consider the set of bi-infinite sequences, where each element in the sequence is either '0' or '1'. We call this set Σ^2, where the superscript two denotes the two 'symbols', 0 or 1. So an element of Σ^2 has the form

$$s = \{\cdots s_{-n}s_{-n+1}\cdots s_{-2}s_{-1}.s_0s_1s_2\cdots s_{n-1}s_n\cdots\},$$

where s_i is either 0 or 1, for all i. The two infinite sequences in our bi-infinite sequence are separated by the period. The reason for two infinite sequences will become apparent when we relate the Bernoulli shift to the baker's transformation in Appendix A.

Mathematicians typically put some additional 'structure' on Σ^2. For example, a distance function, or *metric*, may be defined so that closeness of symbol sequences can be considered. However, the mathematical arena of ergodic theory is probability theory and Σ^2 is then equipped with the necessary structure to make it a probability space. This is done by first assigning a probability to each symbol. For our example we assign 0 and 1 equal probabilities of $\frac{1}{2}$, and from this it is possible to define subsets of Σ^2 in such a way that probabilities can be assigned to the subsets.

We define a map from Σ^2 into itself, called the *shift map*, or *Bernoulli shift*. The shift map, denoted by σ, acts on a bi-infinite sequence by shifting the period one place to the right, i.e.

$$\sigma(s) = \{\cdots s_{-n}s_{-n+1}\cdots s_{-2}s_{-1}s_0.s_1s_2\cdots s_{n-1}s_n\cdots\}.$$

From the probability interpretation that can be put on Σ^2, it is clear that as we shift we are equally likely to have a '0' or '1' immediately to the right of the period. Hence, the analogy with tossing a fair coin should be clear. This was first studied by Jacques Bernoulli and is why his name has come to be associated with this formalism.

Many properties that are difficult, or seemingly impossible, to prove for the typical types of maps that arise in applications can be (relatively) easily proved for shift maps. Mixing and ergodicity are such properties. As an example, it is easy to see that the Bernoulli shift has an infinite number of periodic orbits of all periods. These are just the periodic sequences.

It is often possible to show that another map has the same properties as the Bernoulli shift by showing that it is *isomorphic* to a Bernoulli shift. An isomorphism is a one-to-one correspondence between points in the space of symbol sequences and points in the domain of the map that preserves the essential structures in both spaces. In more applied terminology, an isomorphism is a coordinate transformation (but one must take care here because, often, there are certain technical properties that this coordinate transformation must satisfy, e.g. volume preservation, measurability, etc.). For our purposes, it will transform our micromixer into the Bernoulli shift (and vice versa).

Bernoulli transformations. A measure-preserving transformation is called *Bernoulli* if it is isomorphic to a Bernoulli shift.

The Bernoulli shift is the paradigm for *deterministic chaos*. As we show in Appendix A, the baker's transformation is Bernoulli (see figure 1). In general, the following chain of implications holds:

$$\text{Bernoulli} \rightarrow \text{mixing} \rightarrow \text{ergodic},$$

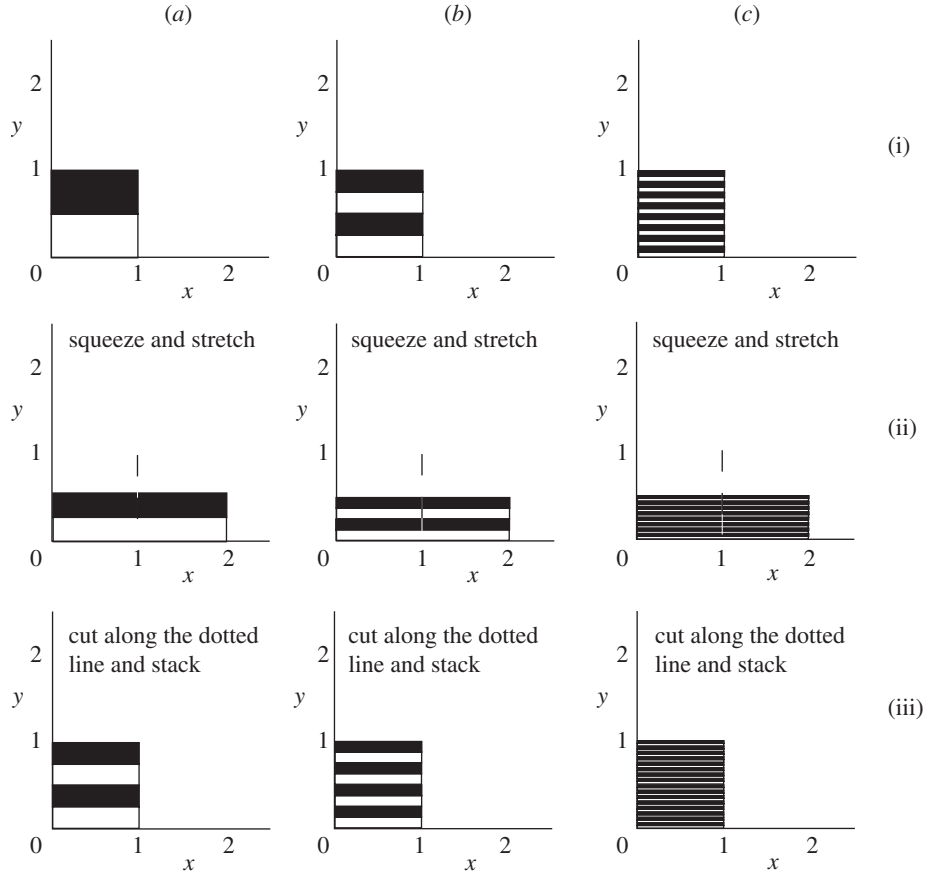

Figure 1. (*a*) One iteration of the baker's transformation on the unit square, (i) at the beginning the upper half of the square is black and the lower half is white. The square is squeezed, stretched, cut, and re-stacked. (*b*) The second iteration of the baker's transformation. (*c*) The fourth iteration of the baker's transformation. Note how quickly the black and white material is 'mixed'. Also note that there is no loss of material from the original domain (the unit square).

and the direction of the arrows *cannot* be reversed. Thus, Bernoulli is the most desired property for mixing and the baker's transformation is the best mixing transformation.

Smale horseshoe. The Smale horseshoe map has some similarities to the baker's transformation, but also some key differences from the point of view efficient mixing. We describe the map geometrically, following figure 2. We begin with the unit square exactly as in our description of the baker's transformation. The square is squeezed in the vertical direction and expanded in the horizontal direction. Now comes the key difference from the baker's transformation: rather than *cutting* off the part that has squeezed out of the original square and putting *all* of it back into the square (as for the baker's transformation), we *fold* back into the square part of what has been squeezed out, with a part remaining outside the square. If we consider the square to be the region in which we are interested in the mixing of fluid, it is clear that the

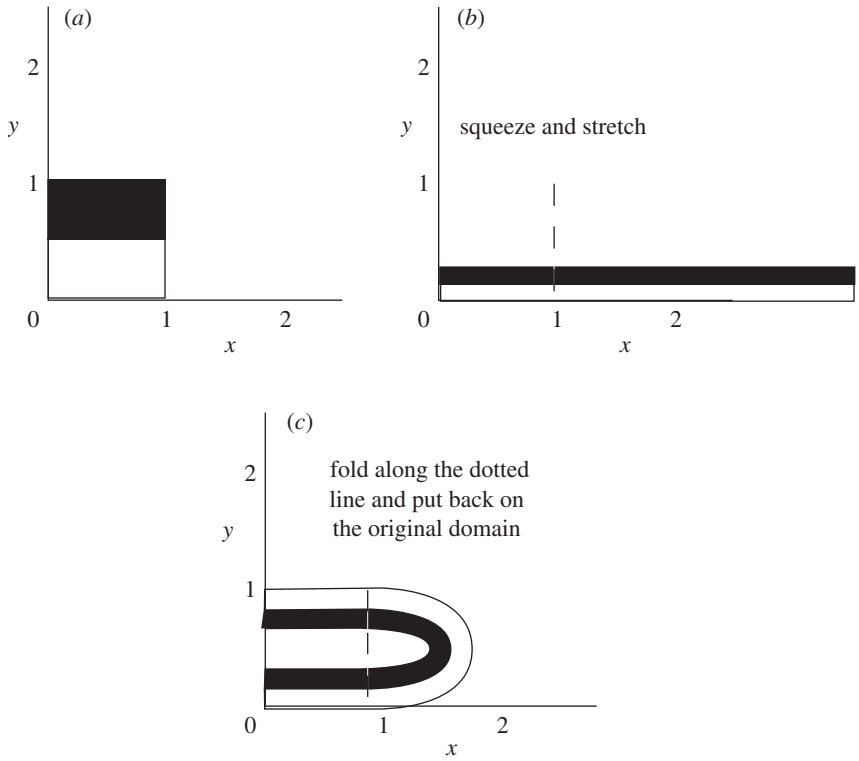

Figure 2. The Smale horseshoe map acting on the unit square.
Note the loss of material from the original domain (the unit square).

Smale horseshoe map has 'lost' some of the fluid that we were interested in mixing. With further mapping, the amount of original material that remains in the unit square becomes smaller and smaller (an experimental study of horseshoe detection in a fluid flow is presented in Chien *et al*. (1986)).

One might argue that we did not need to leave part of the fluid outside the square to begin with. We could have found some way to 'stuff it back inside'. Incompressibility is obviously a hindrance. However, even if this were possible, it would dramatically affect the mixing properties. The baker's map solves this issue by 'stuffing the fluid back inside' in a way that results in good mixing. It 'cuts', rather than 'folds'. In this way, the contracting and stretching directions are always the same at every point in the fluid. This is highly desired for good mixing. When one 'folds', these contracting and expanding directions can rotate, resulting in lack of control of the mixing process.

Still, the Smale horseshoe map can have a tremendously positive influence on mixing. It has an invariant set of zero volume and, if we restrict the map to this set of zero volume, it is isomorphic to a Bernoulli shift. This seems to be vastly inferior to the baker's transformation, where the invariant set was the entire square, and the baker's transformation is isomorphic to a Bernoulli shift *on the entire square*. Nevertheless, the *chaotic invariant set of zero volume* of the Smale horseshoe certainly has a 'randomizing influence' on fluid in its neighbourhood. However, quantifying that

influence can be difficult, and it can be quite hard to eliminate islands (sometimes very small) of unmixed fluid.

(vii) *An example: the baker's transformation*

Consider now a number of the ideas introduced above in the context of a specific mapping (or transformation).

The baker's transformation is defined on the unit square, with periodic boundary conditions. In figure 1*a*, we show the unit square, with the upper half of the square black, and the lower half white. The map contracts the square in the y-direction by a factor of $\frac{1}{2}$, and expands the square in the x-direction by a factor of two (hence, it is area preserving). More precisely, it has the form

$$S(x,y) = \begin{cases} (2x, \frac{1}{2}y) \mod 1, & \text{if } 0 \leqslant x < \frac{1}{2}, \\ (2x, \frac{1}{2}(y+1)) \mod 1, & \text{if } \frac{1}{2} \leqslant x < 1. \end{cases}$$

The term 'mod 1' refers to the enforcement of the periodic boundary conditions. Geometrically, it means we 'cut off' the part of the square protruding from the right-hand vertical boundary of the square ($x = 1$) and stack it on top of the part remaining in the square. After this, the square now has four alternating black and white strips, each of width $\frac{1}{4}$, as we illustrate in figure 1*b*, *c*; repeated applications produce $8, 16, 32, \ldots, 2^n$ strips.

What happens in the limit of an infinite number of iterations? In this limit the unit square becomes completely filled out with an infinite number of alternation black and white *lines*. In this way, the black and white material has become completely mixed (recall that this is purely kinematics; there is no molecular diffusion).

The baker's transformation is a highly efficient mixing transformation. We can get an idea of this by applying the mathematical definition of mixing given above to the black and white regions. Let B denote the region of black material and W the region of white material. The area of B, denoted $\mu(B)$, is equal to $\frac{1}{2}$. The area of W, denoted $\mu(W)$, is also equal to $\frac{1}{2}$. Consider the situation after n advection cycles. Mathematically, $S^n(B) \cap W$ denotes the black material that is in the region originally occupied entirely by white material. By construction, this consists of 2^{n-1} black strips, each of length 1 and width $1/2^{n+1}$. Hence, we have $\mu(S^n(B) \cap W) = \frac{1}{4}$, and therefore

$$\lim_{n\to\infty} \mu(S^n(B) \cap W) = \mu(B)\mu(W)$$

(of course, the definition of mixing requires this limit to hold *for every* choice of sets W and B in R).

The Lyapunov exponents of orbits of the baker's transformation can also easily be computed. For a given point (x, y) in R, the nth iterate of (x, y) under S (i.e. n successive applications of the advection cycle) is denoted by

$$S^n(x,y) \equiv \underbrace{S \circ S \circ \cdots \circ S(x,y)}_{n \text{ map compositions}}.$$

The Jacobian of the nth iterate of the point (x, y) under S is given by

$$DS^n(x,y) = DS(S^{n-1}(x,y))DS(S^{n-2}(x,y)) \cdots DS(S(x,y))DS(x,y),$$

where D denotes the derivative operator. Then the *Lyapunov exponents of the orbit of* (x, y) are given by the logarithms of the eigenvalues of the matrix:

$$\lim_{n\to\infty} (DS^n(x,y)^{\mathrm{T}} DS^n(x,y))^{1/(2n)} \equiv \Lambda.$$

This limit exists under fairly general conditions and is discussed in the fundamental paper of Oseledec (1968). The Lyapunov exponents have the following interpretation. Consider an *infinitesimal* circle centred at (x, y). As it evolves, the circle deforms to an ellipsoid; the Lyapunov exponents are the average logarithmic expansion rates of the principal axes of this ellipsoid. In the case of the baker's transformation—*except* where $x = p/2^k$, where p and k are non-negative integers (this condition arises form the discontinuity associated with the 'cutting')—we have

$$DS^n(x,y) = \begin{bmatrix} 2^n & 0 \\ 0 & \dfrac{1}{2^n} \end{bmatrix}$$

and therefore,

$$\Lambda = \lim_{n\to\infty} \begin{bmatrix} 2^{2n} & 0 \\ 0 & \dfrac{1}{2^{2n}} \end{bmatrix}^{1/(2n)} = \begin{bmatrix} 2 & 0 \\ 0 & \frac{1}{2} \end{bmatrix}.$$

So the Lyapunov exponents for *almost all* (i.e. except for the measure-zero set of countable points where the expression for the Jacobian does not hold) orbits are $\log 2$ and $\log \frac{1}{2}$.

Another interesting point is that the stretching and contraction directions are *constant* and *uniform*. That is, they do not 'rotate' from orbit to orbit. This 'control' over the stretching and contraction directions for all orbits is essential for understanding and analysing the mixing properties. We shall have more to say about this issue below.

It is instructive to ask what makes the baker's transformation such an efficient mixing mechanism. The key feature is the compression and expansion along perpendicular directions *at every point of the domain*, then the cutting and stacking. Two points should be emphasized.

(i) The directions for the saddle point behaviour, i.e. the compressing direction and the expanding direction, are the same for every point of the domain. In general maps this is not true. In general these directions 'rotate' with respect to each other so that at some points the expansion direction could be close to the compressing direction at other points.

(ii) The previous issue is related to the idea of folding. There is no folding (just cutting and stacking) in the baker's transformation. Folding can be bad for mixing, as regions of the fold may mix poorly with the rest of the domain. Repeated folding in different locations can aid in the elimination of this bad mixing behaviour, but it cannot eliminate it entirely (this is why, to date, for such generic area-preserving maps as the standard map it has not been proved that they have a region of non-zero area on which the map is mixing (see Sinai 1994)).

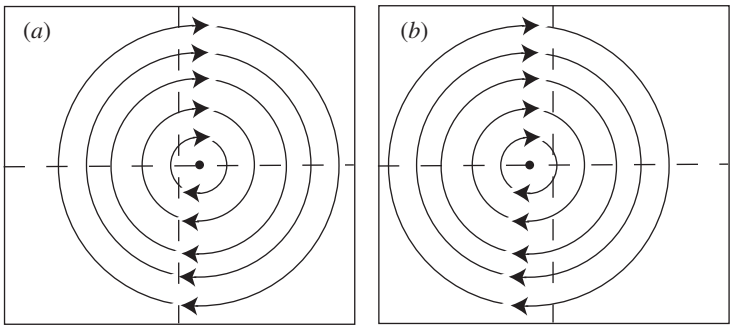

Figure 3. Geometry of the flow patterns for the two halves of the advection cycle. (*a*) Streamlines in the first half-cycle; (*b*) streamlines in the second half-cycle.

The message here is clear. Squeezing, stretching and folding fluid is good for mixing. However, squeezing, stretching and cutting can be *much* better. Consider now some ways of realizing this in the design of micromixers.

3. LTMs: a basic building block for chaotic mixers

The obvious question is, How do we design a mixer so that it has the Bernoulli property? To a mathematician, this might seem a naive question since rigorous proofs that maps possess the Bernoulli property on regions of non-zero volume are notoriously difficult. A few examples are known, the baker's transformation being one of them. To a physicist or engineer, the known examples have a rather artificial flavour, epitomized by the baker's transformation, and, consequently, tend not to provide inspiration for digging out 'useful' concepts from mathematics that are difficult and often less than user-friendly. The situation is not bleak, however; in this paper we describe a new (from the point of view of applications to fluid mixing) type of map, or *advection cycle*, that has been rigorously shown to possess the Bernoulli property. Most surprisingly (and very fortuitously for design purposes) many previously built micromixers, as well as a variety of potentially new micromixers, can be optimized for the Bernoulli property if they are designed so that the flow patterns give rise to an LTM. In order to show how this can be done we begin by describing the flow patterns and how we derive an LTM from them, and, in the process, explain precisely what an LTM is and, most importantly, its mixing properties. The material presented here is based on a series of papers by Burton & Easton (1980), Devaney (1980), Wojtkowski (1980) and Przytycki (1983) published in the pure mathematics literature. The (loose) connection of this material with physical systems is in the context of celestial mechanics. As we shall see, however, the theory is tailor-made for mixing applications (Wiggins 1999).

(*a*) *LTMs: definitions and relation to flow patterns*

Consider a region of fluid, possibly a 2D cavity with solid boundaries (e.g. Chien *et al*. 1986) or the cross-flow in a 3D flow in a channel with an axial flow in the direction normal to the page, containing a stagnation point surrounded by closed streamlines, as shown in the figure 3*a*. The fact that we are showing our region to be an (enclosed) square with circular streamlines is irrelevant for our conclusions. The

horizontal and vertical dashed lines are axes centred in the middle of the region and merely serve as an (essentially arbitrary) reference point. The stagnation point is on the horizontal axis, offset to the left of the centre.

Imagine that the flow pattern is altered at some later time and by some mechanism (the details of which are unimportant to the argument). The alteration involves moving the stagnation point to the right of the centre, as shown in figure 3*b*. We remark that the flow domain could contain multiple 'recirculation cells'. Our arguments can be applied to any number of them.

The flow cycles between the upper and lower patterns in a periodic fashion, which is the *advection cycle* of interest to us. The cycling could occur as a result of the imposition of periodic time dependence in a cavity flow, or as a result of the flow pattern changing periodically in the cross-section of a flow as in a discontinuous duct flow. The utility of this approach is that it is independent of the details of how the flow is created, and relies only on the geometry of the flow patterns. Clearly, the length of the period of cycling will also play an important role in the rate of mixing, and we will address this later.

The fluid particle motion from the beginning of a half-cycle to the end of the same half-cycle is described by a *twist map*. For the closed streamlines in each half-cycle let (r, θ) denote *streamline coordinates* (the actual trajectories in the (x, y)-plane need not be circular; they can be made circular by some nonlinear transformation). That is, on a streamline r is constant and θ is an angular variable that increases monotonically in time. The map of particles from the beginning to the end of a half-cycle is given by $S(r, \theta) = (r, \theta + g(r))$. We will see that the function $g(r)$ is the key here. It provides the angular displacement along the streamline during one half of the advection cycle. Typically it varies from streamline to streamline, and it is from this property that the phrase *twist map* arises. For example, $g(r)$ may increase as r moves from the elliptic stagnation point to a particular streamline (i.e. $g(r)$ achieves a unique maximum), and then it decreases monotonically to zero on the boundary.

We can now define an LTM over the entire advection cycle. Let A_1 denote an annulus whose inner (denoted r_{i}^1) and outer (denoted r_{o}^1) boundaries are streamlines centred at $(c, 0)$ in the first half of the advection cycle. Let A_2 denote an annulus constructed in the same way and centred at $(-c, 0)$ in the second half of the advection cycle. We show the two annuli in figure 4. The two annuli must be chosen so that they intersect each other *transversally*, in the sense that they intersect each other in two disjoint regions, as shown in the figure.

Obvious questions arise. How do we choose the annuli A_1 and A_2? Or which annuli A_1 and A_2 do we choose? There are clearly many such choices of pairs of annuli that will satisfy the transversal intersection condition. The theory of LTMs does not answer this question. It is only concerned with the behaviour on a pair of annuli satisfying the transversal intersection property, and some additional properties described below. For mixing purposes, in a given situation we will need to find the largest and/or largest number of annuli in the flow domain satisfying the hypotheses. We will next discuss these hypotheses on a given pair of annuli. Also, in each half-cycle depicted in figure 3 we show one 'recirculation region', i.e. one region of closed streamlines. In some applications a half-cycle may contain multiple recirculation regions (separated by heteroclinic and/or homoclinic orbits). This poses no difficulty in the application of the LTM approach, since we merely need to find transversally intersecting annuli in each half-cycle. However, it does raise questions about how

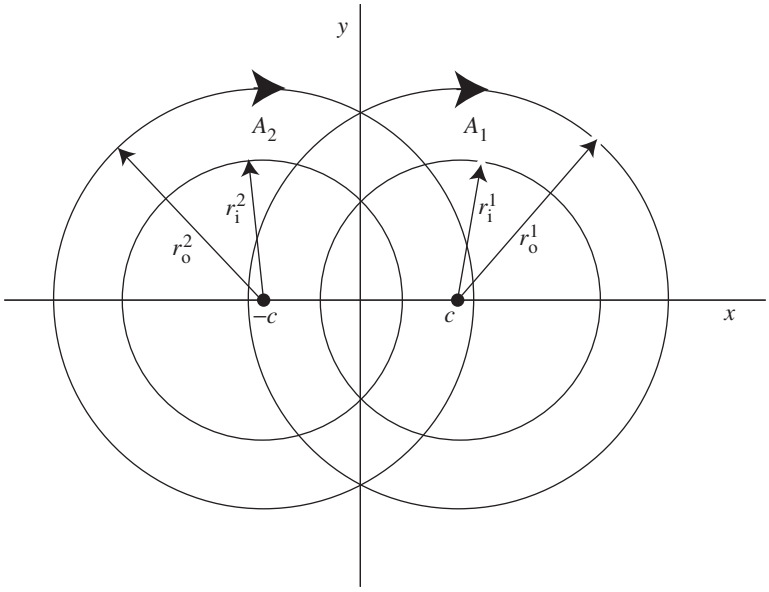

Figure 4. Geometry of the annuli that make up the LTM from each half-cycle of the advection cycle.

many annuli can be found and, as we shall see, the sense of rotation may also be important (i.e. co- versus counter-rotation).

Let $S_1(r,\theta) = (r, \theta + g_1(r))$ be a twist map defined on A_1 with $\mathrm{d}g_1/\mathrm{d}r \neq 0$ and $g_1(r_\mathrm{i}^1) = 2\pi n$, for some integer n. Then r_o^1 is chosen such that $g_1(r_\mathrm{o}^1) = 2\pi n + k_1$, where k_1 is an integer whose value will be discussed shortly.

Let $S_2(r,\theta) = (r, \theta + g_2(r))$ be a twist map defined on A_2 with $\mathrm{d}g_2/\mathrm{d}r \neq 0$ and $g_2(r_\mathrm{i}^2) = 2\pi m$, for some integer m, and $g_2(r_\mathrm{o}^2) = 2\pi m + k_2$. Furthermore, we suppose that the annuli intersect transversally in two disjoint components in the sense that $r_\mathrm{i}^1 \cap r_\mathrm{i}^2 \neq \emptyset$ and $r_\mathrm{o}^1 \cap r_\mathrm{o}^2 \neq \emptyset$ (see figure 4).

The map defined by $S_2 \circ S_1$ on $A_1 \cup A_2$ is referred to as an *LTM*.† Let

$$\alpha_i \equiv \sup_{r_\mathrm{i}^i \leqslant r \leqslant r_\mathrm{o}^i} \frac{\mathrm{d}g_i}{\mathrm{d}r}.$$

The value of α_i is a measure of the strength of the differential winding of the map. $\alpha_i \neq 0$ is just referred to as the 'twist condition'. Regarding the annuli as chosen and fixed, there are four important parameters for the LTM: k_1 and k_2, which describe the number of twists for each twist map, and α_1 and α_2, which describe the strength of each twist.

† Strictly speaking, the map $S_2 \circ S_1$ is not defined on all of $A_1 \cup A_2$ since S_1 is only defined on A_1 and S_2 is only defined on A_2. In this situation we can define S_1 (respectively, S_2) to be the identity map on the part of $A_1 \cup A_2$ that is *not* A_1 (respectively, the part of $A_1 \cup A_2$ that is *not* A_2). Since S_1 (respectively, S_2) is the identity map on the boundary circles defining A_1 (respectively, A_2) this *extension* of the maps to a larger domain can be done smoothly. Conceptually, from the physical point of view, this point can be ignored. It says nothing more than that we consider S_1 (respectively, S_2) only to be acting on A_1 (respectively, A_2), and it does nothing to points that are in the part of A_2 (respectively, A_1) that does not intersect A_1 (respectively, A_2).

We can now state the main results on mixing. There are two cases to consider. One is where the annuli rotate in the same sense. This corresponds to k_1 and k_2 having the same sign (the *co-rotating* case). In this case, if each map is at least a double twist ($|k_i| \geqslant 2$) and $\alpha_1\alpha_2 > 0$, then on $A_1 \cup A_2$ the LTM has the Bernoulli property.

Interestingly, if the annuli rotate in the opposite sense, i.e. k_1 and k_2 have opposite signs (the *counter-rotating* case), then the conditions that the LTM has the Bernoulli property on $A_1 \cup A_2$ are more restrictive. If the product of α_1 and α_2 is negative, then we must have $\alpha_1\alpha_2 < -C_0 \approx 17.244\,45$.

As mentioned earlier, these results were proved by Burton & Easton (1980), Devaney (1980), Wojtkowski (1980) and Przytycki (1983). They also indicate that, all things being equal, it is easier to achieve the Bernoulli property in the co-rotating case, as opposed to the counter-rotating case.

It is then apparent that, in designing a flow for optimal mixing, the key quantities we need to understand are $g_i(r)$, $i = 1, 2$, on the annuli of choice, since these functions determines the rotation properties of the annuli, the radii of the annuli, and the strength of the twist, α_i, which in the context of fluids is the shear rate. This is significant because the design of a mixer with the Bernoulli property boils down to the properties of one function describing closed streamlines in each half-cycle of the advection cycle.

Finally, there is a technical problem that must be addressed in applying the currently known LTM results for the purpose of concluding that the mixing has the Bernoulli property on the two chosen annuli. Let us describe the situation above a little more carefully. We are concerned with two flow patterns that alternate (blinking flow): call them pattern 1 and pattern 2. Now the LTM formalism and results as developed by mathematicians, they are not focusing on all of pattern 1 and pattern 2. Rather, they are focusing on one annulus (in isolation) in pattern 1 and one annulus (in isolation) in pattern 2 that intersect 'transversely'.

Hence, in the mathematicians' formalism, the LTM maps particles between the annuli as we alternate the application of the twist maps to each annulus, and the theorems describe mixing of particles in the two annuli. Now in order for this to make sense, the same particles must remain in the two annuli for all time. This is not true for two arbitrarily chosen flow patterns. However, it is true if pattern 2 is a *rigid rotation* of pattern 1 (this is precisely what is done in the analysis of the partitioned pipe mixer described in Khakhar *et al*. (1987)). Since the flow is bounded, and we know where all the particles go, we suspect that some strong mixing results can be proven in the case where this is not true. But this is a problem that awaits further mathematical analysis. Further, we remark that if one is only interested in constructing Smale horseshoes, this does not require this condition as the invariant set associated with the horseshoe is directly constructed in the overlap regions between two appropriately chosen annuli (see Devaney 1978; Wiggins 1999; Khakhar *et al*. 1986).

(*b*) *Optimization of mixing: qualitative arguments*

Consider now the issue of optimization of mixing in terms of the size of the two annuli in the discussion above. Here we give qualitative arguments for circular streamlines. Nevertheless, the approach is general. Here the maximal size will be determined by the transversal intersection condition and the size of the domain. This is com-

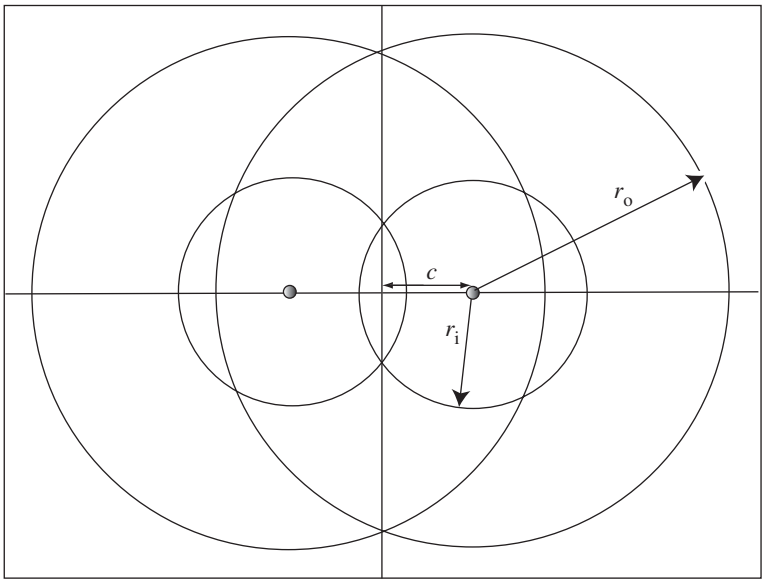

Figure 5. Geometry associated with choosing an annulus in each half-cycle so that their superposition has the largest area.

pletely independent of the properties of $g_i(r)$. Once this is chosen, the 'twist properties', i.e. double twists, sense of the twists, and strength of the twists, must still be verified.

As above, consider an LTM. The horizontal and vertical axes are at the centre of the cavity, and we will assume that the streamlines in the second half of the advection cycle are obtained by reflecting the streamlines in the first half of the advection cycle about the vertical axis (see figure 5).

The distance of the stagnation point (assumed to be on the horizontal axis) from the origin in the first half of the advection cycle is denoted by $c > 0$. The distance of the stagnation point from the origin in the second half of the advection cycle is then $-c$. Denote the radii of the two circular streamlines defining the annulus in the first half of the advection cycle by r_i and r_o. Let A denote the area of the domain. For the moment, assume it is a square with sides of length b, so that $A = b^2$. Then we must have

$$c + r_\mathrm{o} < \tfrac{1}{2}b \tag{3.1}$$

(this condition ensures that the outer circles of the annuli cannot protrude outside the domain).

For a given $c > 0$, the condition that the two annuli intersect transversely (as described earlier) is guaranteed if

$$r_\mathrm{i} - c > 0 \tag{3.2}$$

(this ensures that the two inner circles intersect transversely), and

$$r_\mathrm{i} + 2c > r_\mathrm{o} \tag{3.3}$$

(this ensures that the outer circle of one annulus intersects the inner circle of the other annulus transversely).

Conditions (3.2) and (3.3) together guarantee the transversal intersection property for the two annuli. If the various conditions are satisfied, then mixing is Bernoulli on the union of the two annuli. Hence the mixing area is given by

$$A_{\text{mix}} = 2\pi(r_{\text{o}}^2 - r_{\text{i}}^2) < A. \tag{3.4}$$

The idea now is to maximize this area, subject to the constraints (3.1)–(3.3). We want to simplify this a bit. We claim that (3.4) is maximized provided that

$$r_{\text{o}} - r_{\text{i}} \tag{3.5}$$

is maximized.

It follows from (3.2) and (3.3) that we can make (3.5) large by choosing some $\varepsilon > 0$ very small, and

$$r_{\text{i}} = c + \varepsilon, \qquad r_{\text{o}} = 3c - \varepsilon. \tag{3.6}$$

Then (3.1) will be satisfied provided

$$4c < \tfrac{1}{2}b. \tag{3.7}$$

The way to interpret this is as follows. For a given $c > 0$, we have optimized the region of mixing, subject to c satisfying the constraint (3.7). Using (3.4) and (3.6), we have

$$A_{\text{mix}} = 16\pi(c^2 - c\varepsilon). \tag{3.8}$$

Now $\varepsilon > 0$ can be taken as small as we like, so the region of mixing scales like c^2. From (3.7), if we choose the maximum value of c as $c = \frac{1}{8}b$, then substituting this into (3.8) gives

$$A_{\text{mix}}^{\max} = \tfrac{1}{4}\pi(b^2 - c\varepsilon). \tag{3.9}$$

The $c\varepsilon$ term can be taken as small as we like. The area of the cavity is b^2; therefore $\pi/4 \sim 78\%$ of the cavity is Bernoulli mixing, provided the twist conditions hold. Clearly, there is considerable scope for improving this.

4. Examples of mixers that can be analysed as LTMs

Many mixers fit within the LTM framework. This is significant because LTMs provide an analytical approach to the design of devices producing a mathematically optimal, and precisely defined, type of mixing.

Two comments first. The first example of a chaotic flow, the blinking vortex flow (Aref 1984), is also the most transparent and the most immediately analysable example. In this case the flow itself is already in the form of an LTM and the functions $g_i(r)$ can be controlled at will. This connection was described in Wiggins (1999). It is remarkable that, in some sense, this example encompasses a large number, if not all, of other examples. It is also important to stress that the most conceptually efficient way to think about mixing is in terms of maps and not in terms of deviations for integrability. The most useful heuristic is 'streamline crossing', i.e. streamlines in a bounded domain at two different times must intersect. This precisely the central message of the LTMs.

Let us consider a few examples, the first two from recent devices intended for microfluidic applications, and others from older systems that illustrate mechanisms that may be used in future microfluidic applications.

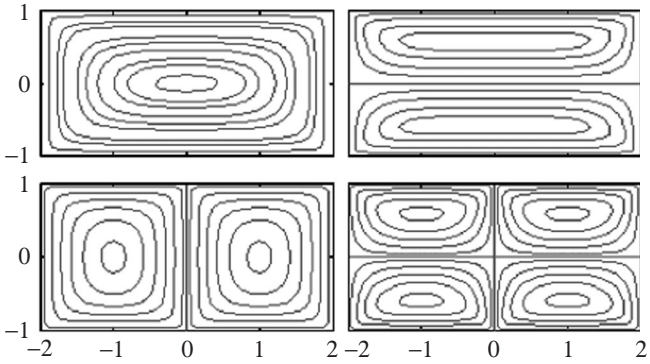

Figure 6. Flow patterns computed for different ζ-potential distributions on the walls of the cavity. (Reproduced with permission from Qian & Bau (2002). Copyright (2002) American Chemical Society.)

Qian & Bau (2002, fig. 4) considered flows generated by electro-osmotic flow (EOF) in cavities (until recently EOFs have been used primarily as an alternative to pressure-driven flow in microchannels, the simplest case corresponding to uniformly charged walls). However, several other scenarios are possible; an early study considering the effects on non-uniform charge was made by Anderson & Idol (1985). The authors computed flow patterns for specific (non-uniform) ζ-potential distributions on the walls of the cavity. Different ζ-potential distributions gave rise to different cellular flow fields in the cavity, as shown in figure 6. Qian & Bau also suggested that one could switch between different flow patterns through 'judicious control of embedded electrodes' in the walls of the cavity. In this way a blinking flow can be realized. Not surprisingly, they demonstrated numerically that such flows can give rise to chaotic fluid particle orbits.

Clearly, such a scheme also fits squarely within the LTM formalism. If we superimpose two chosen flow patterns that are rigid rotations of each other, the structure of the LTM is clear. One way of applying the results on LTMs is to choose annuli in one flow pattern and other annuli in the other flow pattern, such that the annuli intersect pairwise 'transversely' in two disjoint components, as described earlier. Then the switching time between patterns, T, is chosen such that for each annulus the outer circle rotates *twice* with respect to the inner circle during the time $T/2$. We then need the twists to be 'sufficiently strong', which will also depend on whether or not the chosen annuli pair are co- or counter-rotating. If this can be done, then, appealing to the dynamical systems results described earlier, the flow will be Bernoulli in the region defined by the chosen annuli. Of course, there are numerous open problems. For example, which ζ-potential distributions lead to the maximal region on which the LTM results hold? However, such an analysis is possible using formulae for the flow patterns given by Qian & Bau (2002).

Another possibility for flow manipulation is the use of patterned walls. In fact, a patterned wall, at the time of writing, is a more robust mechanism than the EOF for flow manipulation. Stroock *et al.* (2002) built and conducted experiments in a micromixer consisting of a straight channel with ridges placed on one of the walls of the channel at an oblique angle with respect to the axis of the channel. When the fluid is driven axially by a pressure gradient, the ridges on the floor of the channel give

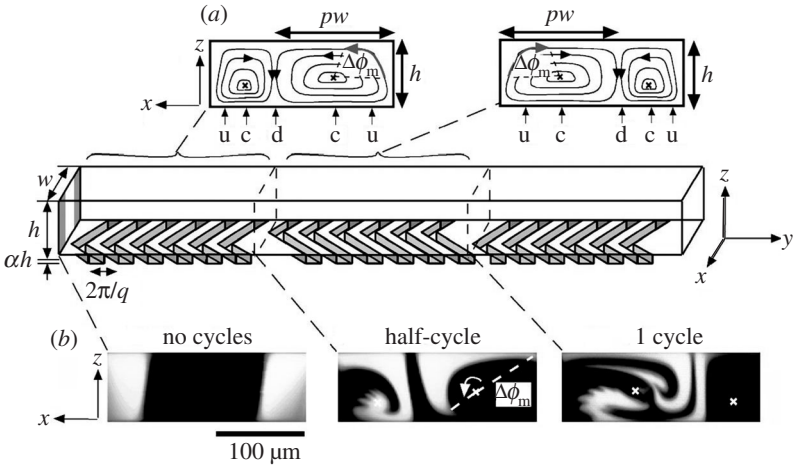

Figure 7. Closed streamlines in the cross-section of each half-cycle. (*a*) The mixer; (*b*) flow visualizations in the cross-section. (Reproduced with permission from Stroock *et al*. (2002). Copyright (2002) AAAS.)

rise to a transverse flow. In the (x, y)-plane or cross-flow the streamlines are closed and helical in three dimensions. If the ridges are arranged in a periodic pattern down the axis of the channel, a herringbone pattern zigzagging to the right and to the left, each period consists of two half-cycles, producing two cells. If the pattern is such that the two cells are asymmetric with respect to the y-axis, then, if one looks at the mixer along the axial path, the elliptic points corresponding to the centre of the cells switch positions after one cycle. The overall map then consists of the composition of two maps: the maps between each half-cycle and the map between cycles is an LTM. This is shown in figure 7 (Stroock *et al*. 2002, fig. 2). The LTM results described here provide a basis for design and analysis of rigorously defined mixing properties in such flows.

Several extensions become apparent. For example, if the cross-section is not mirror symmetric as in a rectangle, but is, for example, trapezoidal, then a herringbone pattern is unnecessary. In this case a patterned wall is all that is needed. The key idea is to shift the location of the elliptic point (see figure 8).

There are many other examples that can be fitted within the LTM framework. In fact, some of these examples are older, going back to static mixing concepts. The partitioned pipe mixer (PPM) of Khakhar *et al*. (1987), analysed experimentally by Kusch & Ottino (1992), is representative of a large class of spatially periodic flows, and the first continuous flow that was shown to be chaotic. The PPM consists of a pipe partitioned into a sequence of semicircular ducts by means of orthogonally placed rectangular plates (figure 9). A cross-sectional motion is induced through rotation of the pipe wall. At every length L along the pipe axis, the orientation of the dividing plate shifts by 90°. Thus a series of two co-rotational flows is followed by two co-rotational flows but shifted by 90°. This particular geometry of the PPM is just one of three possible spatially periodic configurations which may be realized. Franjione & Ottino (1992) considered two variants of this flow. One qualitatively captures the motion in a sequence of pipe bends (the 'twisted-pipe' flow of Jones *et al*. (1989)): two counter-rotational vortices followed by two counter-rotational

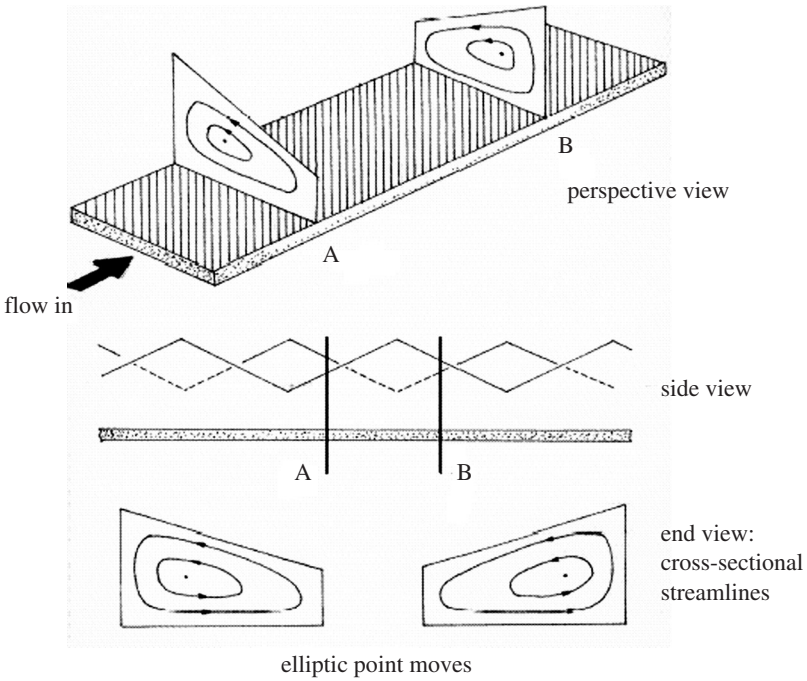

Figure 8. Micromixer with a periodical varying trapezoidal cross-section driven by a grooved bottom wall. 'A' and 'B' denote the beginning of each half cell where the cross-sectional flow is shown.

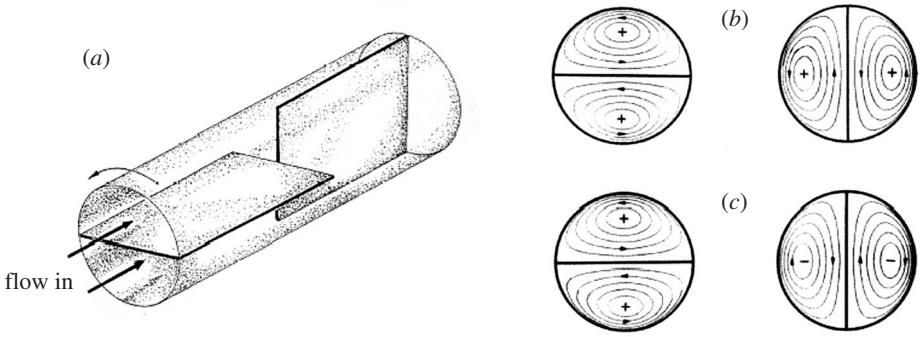

Figure 9. The partitioned pipe mixer. (*a*) The duct and the internal arrangement of plates. (*b*) The cross-sectional flows for the PPM at the indicated locations. (*c*) Cross-sectional flows that would occur for the Kenics mixer.

vortices. The other a flow resembles the motion in a Kenics static mixer: two co-rotational vortices followed by two counter-rotational vortices. As discussed earlier in this article, the PPM design fits the LTM formalism precisely.

Static mixers try to mimic the baker's transformation. Static mixers invariably involve *internal* surfaces; two regions have to be split and then reconnected (a method that tries to mimic cutting in two dimensions). This is an issue that makes these systems complicated to build at small scales. Some remarkably small mixers have been built (e.g. Bertsch *et al*. 2001), but these designs are scaled-down versions of designs

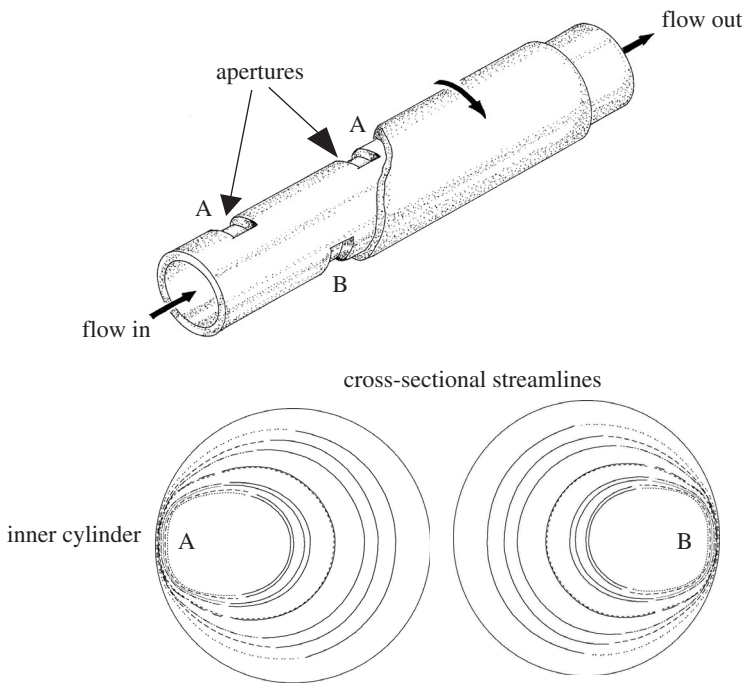

Figure 10. The rotated arc mixer.

commonly encountered in routine large-scale applications. However, this need not be the case and other design possibilities, more in line with current microfabrication technologies, should be explored.

Nevertheless, large-scale applications may be a source of inspiration. Consider the rotated arc mixer (Metcalfe *et al*. 2001, 2004). The design depends critically on a clever use of the cross-sectional flow. The system consists of two hollow cylinders with a very small gap between them; the outer cylinder rotates while the flow is driven axially by a pressure gradient. The inner cylinder has a strategically placed cut-off, exposing the flow contained in the inner cylinder to the drag of the moving outer cylinder (figure 10).

In the example in figure 10 there are two cut-offs per period, but obviously the system can be generalized to any number of cut-offs. With two cut-offs per period, the system corresponds exactly to an LTM. The theory for the case of more than two cut-offs has not yet been developed. It is apparent that this design can be implemented, at least in theory, by means of EOFs.

Another variation on the cavity flow is the exploitation of time-dependent changes in geometry by adding a secondary baffle (figure 11). This idea goes back to Jana *et al*. (1994), and variations on this idea have been patented in the context of polymer-processing applications. In the original case the cross-flow was induced by an upper wall sliding diagonally. However, it is easy to see that the design will work as well when the driving is due to the patterned ridges, as developed by Stroock *et al*. (2002). In this case the portrait changes from one figure of eight to another, but the location of the central hyperbolic point has been shifted (Jana *et al*. 1994). In some sense this resembles two LTMs, but in fact the hyperbolic point changes the

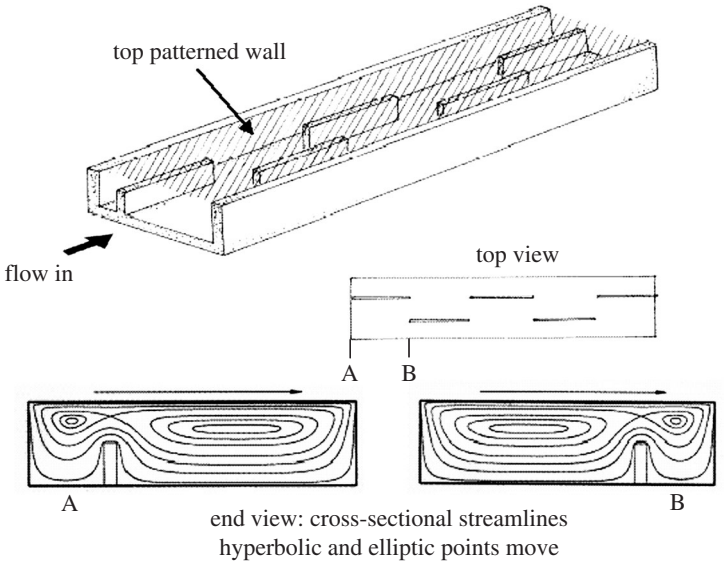

Figure 11. Micromixer with a periodic series of baffles driven by a grooved bottom wall. 'A' and 'B' denote the beginning of each half cell where the cross-sectional flow is shown.

mathematical structure. The theory for this case remains to be developed and the $g(r)$ has a double-hump structure. However, the same mathematical approach taken in the original LTM papers should apply.

All the above designs are spatially periodic. For example, vertical and horizontal elements in the case of the PPM mixer, right-handed (R) and left-handed (L) helices in a static mixer, or zigzagging to the right and to the left (R-L-R-L-R) in the Stroock mixer. Thus, in all the above mixers, when we assemble the mixer we have a sequence of R-L-R-L-R, and so on. This clearly works, but, if there are unmixed regions and tubes, resulting from the map, they will persist even with an infinite number of elements. In other words, it is possible for unmixed streams to pass through the mixer. In the past, the only way to investigate the presence of islands was to propose a design and to resort to computations. However, to the extent that a sequence of R-L can be viewed as a twist map, we can now be assured that mixing is effective in a region that can be calculated *a priori*. The task is therefore to mix these regions with the outside, unmixed regions. Symmetry manipulations provide a route by which to achieve this objective (Franjione & Ottino 1992; Ottino 1990).

5. Other mathematical directions

(*a*) *Pseudo-Anosov maps*

Recently, techniques from the study of the topology of maps of surfaces (Boyland 1994) have been applied to mixing in macroscopic flows by Boyland *et al*. (2000). These *pseudo-Anosov* maps also have the Bernoulli property, and LTMs can be viewed as an example of such maps where it is straightforward to control the region of the flow on which the Bernoulli property holds. While the mathematical language may appear formidable (and we will not go into it here), once it has been translated into the context of a mixing problem it becomes 'almost' intuitive. This is a line

of research that should be pursued more vigorously. MacKay (2001) has written a review of the subject.

(*b*) *Complex time dependence*

In this article we have been concerned with either 2D time-periodic flows, or spatially periodic steady flows. It is natural to ask if gains in mixing can be made by making the flow 'more unsteady' in the sense of making the time dependence more complex. Some early work, both theoretical and experimental, has been done in this area by Franjione *et al*. (1989).

However, it must be emphasized that the dynamical systems framework has not been fully developed so that the same types of analyses can be carried out as for time-periodic or 3D steady flows. Moreover, there are some fundamental differences. Nevertheless, this promises to be a rich and rewarding line of research.

In recent years an effort to develop the building blocks of a geometric theory for the dynamics generated by aperiodically time-dependent velocity fields has begun. Aspects that have been studied include bifurcation theory, normal form theory, shadowing lemmas, and chaos- and Smale-horseshoe-like constructions. Numerical methods for computing hyperbolic trajectories and stable and unstable manifolds have also been developed. An overview of all of these issues can be found in Mancho *et al*. (2003).

6. Concluding remarks

It is important to realize that there is a gap between the rigorous mathematical results on mixing in LTMs and the flows arising in mixers in the sense that the LTM results do not *immediately* apply to the entire flow domain in a way that will allow us to conclude that the entire domain has the Bernoulli property. Rather, the mathematical results of LTMs apply to a given pair of annuli, one in each half-cycle of the advection cycle, on which the hypotheses of the relevant theorems are satisfied. Nevertheless, we have seen that in simple situations two appropriately chosen annuli can occupy a significant fraction of the flow domain (*ca.* 78%). A particularly important general conclusion is that there is also a clear difference between the co- and counter-rotating cases. Thus, for a particular flow configuration we must determine the maximal region occupied by pairs of annuli from each half-cycle on which the LTM results hold. There are many possibilities here. For example, suppose the first half-cycle contains one recirculation cell (i.e. one elliptic point surrounded by closed streamlines) and the second half-cycle contains two (or more) recirculation cells. This will necessitate choosing some annuli that counter-rotate and some that co-rotate. Is this better, or worse, than having one recirculation cell in each half-cycle with the elliptic point displaced in the second half-cycle with respect to its position in the first half-cycle? What if there are hyperbolic points in each half-cycle (as in the case of figure 11, where $g(r)$ has two maxima)? Does this help, or inhibit, mixing? There are numerous open questions and more work is obviously necessary. It is apparent, however, that the way forward is clear. The theoretical tools and approach surrounding LTMs provide the framework for the development of a predictive theory for the design of mixing systems and, in particular, micromixers. In fact, the LTM framework will apply to any situation where the flow can be modelled as what

we refer to as a 'segmented flow'. Segmentation can occur in either space or time (or both). An example of a temporally segmented flow is a blinking flow or two 2D steady flows, each having regions of closed streamlines, that are temporally switched between each other. An example of a spatially segmented flow is a duct flow. This consists of a series of 3D spatially steady flows, each having closed streamlines in the cross-section, that are joined in space. The examples we have shown have been segments whose axes are all in a line. However, this is not necessary, and the serpentine channel (Liu *et al*. 2000) would fall into this classification, and therefore the LTM approach should apply.

We are grateful to Robert MacKay and Howard Stone for comments on an earlier version of this paper. The work of S.W. has been supported by Office of Naval Research Grant no. N00014-01-1-0769. The work of J.M.O. has been supported by DOE, Office of Basic Energy Sciences.

Appendix A. Proof that the baker's transformation is isomorphic to a Bernoulli shift

In our list of definitions above we described a property that was even stronger than mixing: the Bernoulli property. A map has the Bernoulli property if it is isomorphic to a Bernoulli shift. What does this mean, and why is it useful to know? We now want to explain this, and also to show that the baker's transformation has the Bernoulli property.

We begin by describing the most common example of a Bernoulli shift.

We define a map from Σ^2 into itself, called the *shift map*. The shift map, denoted by σ, acts on a bi-infinite sequence by shifting the period one place to the right, i.e.

$$\sigma(s) = \{\cdots s_{-n}s_{-n+1}\cdots s_{-2}s_{-1}s_0.s_1s_2\cdots s_{n-1}s_n\cdots\}.$$

Recall that an isomorphism is a one-to-one correspondence between points in the space of symbol sequences and points in the domain of the map that preserves the essential structures in both spaces. We construct it explicitly for the baker's transformation.

Every number in the unit interval $[0, 1]$ can be represented by its binary expansion. We can use each 'half' of a bi-infinite sequence to form the binary expansion of a number in the unit interval as follows:

$$x = \sum_{i=0}^{\infty} \frac{s_i}{2^{i+1}}, \qquad y = \sum_{i=1}^{\infty} \frac{s_{-i}}{2^i}.$$

This observation leads to a natural way for defining a map from Σ^2 to points in the unit square R as follows:

$$\phi(\{s_{-n}s_{-n+1}\cdots s_{-2}s_{-1}.s_0s_1s_2\cdots s_{n-1}s_n\cdots\}) = (x, y) \equiv \left(\sum_{i=0}^{\infty} \frac{s_i}{2^{i+1}}, \sum_{i=1}^{\infty} \frac{s_{-i}}{2^i}\right).$$

The goal now is to relate the shift dynamics on Σ^2 to the baker's map on R through the map ϕ. This is related to what we mentioned in our (informal) definition of an isomorphism concerning 'preserving the essential structures'. Mathematically, this is

accomplished by 'proving that the following diagram commutes':

$$\begin{array}{ccc} \Sigma^2 & \xrightarrow{\sigma} & \Sigma^2 \\ \downarrow \phi & & \downarrow \phi \\ D & \xrightarrow{S} & D \end{array}$$

In other words, if we take any bi-infinite sequence of zeros and ones,

$$s = \{\cdots s_{-n}s_{-n+1}\cdots s_{-2}s_{-1}.s_0s_1s_2\cdots s_{n-1}s_n\cdots\},$$

and we start in the upper left-hand corner of the diagram, then map to R by first going across the top of the diagram, then down, we will get the same thing if we first go down, then across the bottom of the diagram. In other words, we want to show that for any sequence $s \in \Sigma^2$, we have $\phi \circ \sigma(s) = S \circ \phi(s)$. We will work out each side of the equality individually and, hence, show that they are equal.

From the definition of the map ϕ we have:

$$\phi \circ \sigma(s) = \left(\sum_{i=0}^{\infty} \frac{s_{i+1}}{2^{i+1}}, \sum_{i=0}^{\infty} \frac{s_{-i+1}}{2^i} \right).$$

This is the easy part. Showing the other side of the equality sign requires a bit more consideration. Note that the map is linear and diagonal on R. Therefore, we can consider the baker's map acting on the x- and y-components individually.

For

$$x = \sum_{i=0}^{\infty} \frac{s_i}{2^{i+1}} \quad \text{we have } 2x = \sum_{i=0}^{\infty} \frac{s_{+1i}}{2^{i+1}},$$
$$y = \sum_{i=1}^{\infty} \frac{s_{-i}}{2^i} \quad \text{we have } \tfrac{1}{2}y = \sum_{i=1}^{\infty} \frac{s_{-i}}{2^{i+1}}$$

(which is smaller than $\frac{1}{2}$), but this is not the end of the story. With respect to the y-component, we need to take into account the 'cutting and stacking' (i.e. the 'mod 1' part of the definition of the baker's transformation). Note that in the expressions for $2x$ and $\frac{1}{2}y$ the symbol s_0 no longer appears. Somehow it needs to reinstated, and taking proper account of the cutting and stacking will do that. Now, if $s_0 = 1$, the x-component is greater than or equal to $\frac{1}{2}$. This means that after mapping by the baker's transformation this point is outside the square, and therefore it is in the part that is cut off and stacked. Stacking means its y-component is increased by $\frac{1}{2}$, i.e.

$$\tfrac{1}{2}y = \sum_{i=1}^{\infty} \frac{s_{-i}}{2^{i+1}}.$$

Now if $s_0 = 0$, this additional term contributes nothing to the sum. Combining these facts, we have shown

$$S \circ \phi(s) = S\left(\sum_{i=0}^{\infty} \frac{s_i}{2^{i+1}}, \sum_{i=1}^{\infty} \frac{s_{-i}}{2^i} \right) = \left(\sum_{i=0}^{\infty} \frac{s_{i+1}}{2^{i+1}}, \sum_{i=1}^{\infty} \frac{s_{-i+1}}{2^i} \right).$$

Now since ϕ is invertible, $\phi \circ \sigma(s) = S \circ \phi(s)$ is equivalent to $S = \phi \circ \sigma \circ \phi^{-1}$. This looks very much like the formula for a similarity transformation for matrices. We know that if two matrices are similar, they share many basic properties. The same is true in this more general setting. In particular, by composing this expression for S with itself n times we have $S^n = \phi \circ \sigma^n \circ \phi^{-1}$. From this we can conclude that there is a one-to-one correspondence between orbits of S and orbits of σ. In particular, we can then immediately conclude that S has an infinite number of (saddle-type) periodic orbits of all periods.

Appendix B. Construction of an LTM from a duct flow

In this appendix we show how to rigorously obtain a linked twist map (LTM) for a spatially periodic duct flow. We refer to each spatial period as a *cell*, and the flow is described by a mapping from the beginning of a cell to the end of a cell. This mapping is the composition of two mappings, each a twist map. The first twist map is the mapping of particles from the beginning of the cell to the half cell. The second twist map is the mapping of particles from the end of the half cell to the end of the cell.

The cell-to-cell LTM obtained in this way will be discontinuous in the sense that each twist map is computed separately and continuity at the half cell is not enforced. In this case the requirement that the flow patterns at the beginning of consecutive half cells are related by a rigid rotation is rigorously true.

(*a*) *Construction of the first twist map*

We assume that we have a duct flow in the first half of the cell having the following form:

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \frac{\partial \psi_1}{\partial y}(x,y), \qquad \frac{\mathrm{d}y}{\mathrm{d}t} = -\frac{\partial \psi_1}{\partial x}(x,y), \qquad \frac{\mathrm{d}z}{\mathrm{d}t} = k_1(x,y).$$

Since the x- and y-components of the velocity field (i.e. the cross-flow) *do not* depend on the axial coordinate z we can consider transforming this 2D velocity field into standard action-angle variables. This requires the assumption that in some region in the cross-flow there is a pattern of closed streamlines given by $\psi_1 = c$.

(*b*) *Action-angle transformation in the cross flow*

If this assumption holds then it is well known from classical mechanics that in this region there is a transformation $(x,y) \to (r,\theta)$ satisfying the following conditions:

(i) $r = r(c)$, i.e. r is constant on the closed streamlines in the cross-flow,

(ii) $\displaystyle\oint_{\psi_1=c} \mathrm{d}\theta = 2\pi,$

(iii) $\displaystyle\frac{\mathrm{d}\theta}{\mathrm{d}t} = \Omega_1(r).$

The action variable is given by

$$r = \frac{1}{2\pi}\int_{\psi_1=c} y\,\mathrm{d}x,$$

and the angle variable is given by

$$\theta = \frac{2\pi}{T(\psi_1)}t,$$

where $T(\psi_1)$ is the period of the orbit in the cross-flow (which is a level set of ψ_1), and t denotes the time along the streamline measured from a certain starting point on the streamline.

We assume that this action-angle transformation on the x, y-component of the velocity field has been carried out so that these equations subsequently take the form

$$\frac{\mathrm{d}r}{\mathrm{d}t} = 0, \qquad \frac{\mathrm{d}\theta}{\mathrm{d}t} = \Omega_1(r), \qquad \frac{\mathrm{d}z}{\mathrm{d}t} = w_1(r,\theta), \quad w_1(r,\theta) = k_1(x(r,\theta), y(r,\theta)).$$

We remark that there may be multiple regions of closed streamlines in the cross-flow separated by homoclinic or heteroclinic orbits, or even solid boundaries. In general, a separate action-angle transformation is required for each such region of closed streamlines (and the different transformations may not be simply related).

(*c*) *Action-angle-axial transformation in the half cell*

Now we introduce a final change of coordinates, which leaves the action-angle variables in the cross-flow unchanged, but modifies the axial coordinate so that all particles on a given streamline in the cross-flow take the same time to travel the length of the half cell. We refer to these coordinates as action-angle-axial coordinates (Mezić & Wiggins 1994).

Suppose $\Omega_1 \neq 0$. Then the transformation of variables $(r, \theta, z) \to (r, \theta, a)$, defined by

$$r = r, \qquad \theta = \theta, \qquad a = z + \frac{\Delta z_1(r)}{2\pi}\theta - \int \frac{w_1(r,\theta)}{\Omega_1(r)}\,\mathrm{d}\theta,$$

where

$$\Delta z_1(r) = \frac{1}{\Omega_1(r)}\int_0^{2\pi} w_1(r,\theta)\,\mathrm{d}\theta$$

transforms the velocity field in (r, θ, z)-coordinates to the form

$$\frac{\mathrm{d}r}{\mathrm{d}t} = 0, \qquad \frac{\mathrm{d}\theta}{\mathrm{d}t} = \Omega_1(r), \qquad \frac{\mathrm{d}a}{\mathrm{d}t} = A_1(r),$$

$$A_1(r) = \frac{\Delta z_1(r)}{2\pi}\Omega_1(r) = \frac{1}{2\pi}\int_0^{2\pi} w_1(r,\theta)\,\mathrm{d}\theta.$$

(*d*) *Explicit expression for the twist map*

Now an LTM can be constructed from the velocity field in (r, θ, a)-coordinates. Let $z = 0$ be the starting point of the channel. From the coordinate transformation above we see that the beginning of the first half cell, $z = 0$, corresponds to

$$a = \frac{\Delta z_1(r)}{2\pi}\theta - \int \frac{w_1(r,\theta)}{\Omega_1(r)}\,\mathrm{d}\theta.$$

The end of the first half cell (beginning of the second half cell) is at $z = \frac{1}{2}L$, which corresponds to

$$a = \tfrac{1}{2}L + \frac{\Delta z_1(r)}{2\pi}\theta - \int \frac{w_1(r,\theta)}{\Omega_1(r)}\,\mathrm{d}\theta.$$

The time of flight, T, from the beginning of the first half cell to the end is given by solving $a(T) = A_1(r)T + a(0)$, where

$$a(0) = \frac{\Delta z_1(r)}{2\pi}\theta - \int \frac{w_1(r,\theta)}{\Omega_1(r)}\,\mathrm{d}\theta \quad \text{and} \quad a(T) = \tfrac{1}{2}L + \frac{\Delta z_1(r)}{2\pi}\theta - \int \frac{w_1(r,\theta)}{\Omega_1(r)}\,\mathrm{d}\theta.$$

After some simple algebra we easily find that

$$T = \frac{L}{2A_1(r)}.$$

Therefore, the twist map for the first half cell is given by

$$r \to r, \qquad \theta \to \theta + \frac{L\Omega_1(r)}{2A_1(r)},$$

or

$$r \to r, \qquad \theta \to \theta + \frac{\pi L}{\Delta z_1(r)}.$$

(*e*) *Construction of the second twist map*

The construction of the twist map in the second half of the cell proceeds in exactly the same way as the construction for the first half. We suppose the flow in the second half is given by

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \frac{\partial \psi_2}{\partial y}(x,y), \qquad \frac{\mathrm{d}y}{\mathrm{d}t} = -\frac{\partial \psi_2}{\partial x}(x,y), \qquad \frac{\mathrm{d}z}{\mathrm{d}t} = k_2(x,y),$$

where $\psi_2(x,y)$ describes the streamline pattern in the cross-flow, which we choose to be a rigid rotation of the flow pattern defined by $\psi_2(x,y)$, and $k_2(x,y)$ describes the axial velocity. We then proceed with the same series of transformations as above.

The first and second twist maps are composed to give the LTM describing the mapping of fluid particles from the beginning to the end of a cell.

References

Anderson, J. L. & Idol, W. K. 1985 Electroosmosis through pores with nonuniformly charged walls. *Chem. Engng Commun.* **38**, 93–108.

Aref, H. 1984 Stirring by chaotic advection. *J. Fluid Mech.* **86**, 1–21.

Arnold, V. I. & Avez, A. 1968 *Ergodic problems of classical mechanics.* Addison-Wesley.

Baladi, V. 1999 Decay of correlations. In *Smooth Ergodic Theory and its Applications. Proc. Symp. Pure Mathematics, Seattle, WA, 1999*, vol. 69, pp. 297–325. Providence, RI: American Mathematical Society.

Beigie, D., Leonard, A. & Wiggins, S. 1994 Invariant manifold templates for chaotic advection. *Chaos Solitons Fractals* **4**, 749–868.

Bertsch, A., Heimgartner, S., Cousseau, P. & Renaud, P. 2001 Static micromixers based on large-scale industrial geometry. *Lab Chip* **1**, 56–60.

Boffeta, G., Lacorata, G., Redaelli, G. & Vulpiani, A. 2001 Detecting barriers to transport: a review of different techniques. *Physica* D **159**, 58–70.

Boyland, P. 1994 Topological methods in surface dynamics. *Topology Applic.* **58**, 223–298.

Boyland, P., Aref, H. & Stremler, M. A. 2000 Topological fluid mechanics of stirring. *J. Fluid Mech.* **403**, 277–304.

Burns, K. & Weiss, H. A. 1995 Geometric criterion for positive topological entropy. *Commun. Math. Phys.* **172**, 95–118.

Burton, R. & Easton, R. 1980 Ergodicity of linked twist mappings. In *Global Theory of Dynamical Systems. Proc. Int. Conf., Northwestern University, Evanston, IL, 1979*, pp. 35–49. Lecture Notes in Mathematics, vol. 819. Springer.

Camassa, R. & Wiggins, S. 1991 Chaotic advection in a Rayleigh–Benard flow. *Phys. Rev.* A **43**, 774–797.

Chien, W. L., Rising, H. & Ottino, J. M. 1986 Laminar and chaotic mixing in several cavity flows. *J. Fluid Mech.* **170**, 355–377.

Devaney, R. L. 1978 Subshifts of finite type in linked twist mappings. *Proc. Am. Math. Soc.* **71**, 334–338.

Devaney, R. L. 1980 Linked twist mappings are almost Anosov. In *Global Theory Of Dynamical Systems. Proc. Int. Conf., Northwestern University, Evanston, IL, 1979*, pp. 121–145. Lecture Notes in Mathematics, vol. 819. Springer.

Fountain, G. O., Khakhar, D. V. & Ottino, J. M. 1998 Visualization of three-dimensional chaos. *Science* **281**, 683–686.

Franjione, J. G. & Ottino, J. M. 1991 Stretching in duct flows. *Phys. Fluids* A **3**, 2819–2821. (Erratum: 1991 *Phys. Fluids* **6**, 3501.)

Franjione, J. G. & Ottino, J. M. 1992 Symmetry concepts for the geometric analysis of mixing flows. *Phil. Trans. R. Soc. Lond.* A **338**, 301–323.

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

Horner, M., Metcalfe, G., Wiggins, S. & Ottino, J. M. 2002 Transport mechanisms in open cavities: effects of transient and periodic boundary flows. *J. Fluid Mech.* **452**, 199–229.

Jana, S. C., Tjahjadi, M. & Ottino, J. M. 1994 Chaotic mixing of viscous fluids by periodic changes of geometry: the baffle-cavity system. *AIChE J.* **40**, 1769–1781.

Jensen, K. F. 1999 Micromechanical systems: status, challenges and opportunities. *AIChE J.* **45**, 2051–2054.

Jones, S. W., Thomas, O. M. & Aref, H. 1989 Chaotic advection by laminar-flow in a twisted pipe. *J. Fluid Mech.* **209**, 335–357.

Khakhar, D. V., Rising, H. & Ottino, J. M. 1986 An analysis of chaotic mixing in two model systems. *J. Fluid Mech.* **172**, 419–451.

Khakhar, D. V., Franjione, J. G. & Ottino, J. M. 1987 A case study of chaotic mixing in deterministic flows: the partitioned pipe mixer. *Chem. Engng Sci.* **42**, 2909–2926.

Kusch, H. A. & Ottino, J. M. 1992 Experiments on mixing in continuous chaotic flows. *J. Fluid Mech.* **236**, 319–348.

Lapeyre, G. 2002 Characterization of finite-time Lyapunov exponents and vectors in two-dimensional turbulence. *Chaos* **12**, 688–698.

Liu, R. H., Stremler, M. A., Sharp, K. V., Olsen, M. G., Santiago, J. G., Adrian, R. J., Aref, H. & Beebe, D. J. 2000 Passive mixing in a three-dimensional serpentine microchannel. *IEEE J. Microelectromech. Syst.* **9**, 190–197.

Losey, M. W., Jackman, R. J., Firebaugh, S. L., Schmidt, M. A. & Jensen, K. F. 2002 Design and fabrication of microfluidic devices for multiphase mixing and reaction. *IEEE J. Microelectromech. Syst.* **11**, 709–717.

MacKay, R. S. 2001 Complicated dynamics from simple topological hypotheses. *Phil. Trans. R. Soc. Lond.* A **359**, 1479–1496.

Mancho, A. M., Small, D., Wiggins, S. & Ide, K. 2003 Computation of stable and unstable manifolds of hyperbolic trajectories in two-dimensional, aperiodically time-dependent vector fields. *Physica* D **182**, 188–222.

Metcalfe, G., Rudman, M., Brydon, A. & Graham, L. 2001 Numerical mixing experiments in the rotated arc mixer. In *Proc. 6th World Cong. Chemical Engineering, Melbourne, Australia*.

Metcalfe, G., Rudman, M. & Graham, L. 2004 Chaotic structures from continuous blending. In *Proc. 19th Annual Meeting of the Polymer Processing Society, Melbourne, Australia*. (In the press.)

Mezić, I. & Wiggins, S. 1994 On the integrability and perturbation of three-dimensional fluid flows with symmetry. *J. Nonlin. Sci.* **4**, 157–194.

Mezić, I. & Wiggins, S. 1999 A method of visualization of invariant sets of dynamical systems based on the ergodic partition. *Chaos* **9**, 213–218.

Mezić, I., Wiggins, S. & Betz, D. 1999 Residence-time distributions for chaotic flows in pipes. *Chaos* **9**, 173–182.

Newhouse, S. E. 1988 Entropy and volume. *Ergod. Theory Dynam. Syst.* **8**, 283–299.

Newhouse, S. & Pignataro, T. 1993 On the estimation of topological entropy. *J. Stat. Phys.* **72**, 1331–1352.

Oseledec, V. I. 1968 A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems. *Trans. Mosc. Math. Soc.* **19**, 197–231.

Ottino, J. M. 1989 *The kinematics of mixing: stretching, chaos, and transport*. Cambridge University Press.

Ottino, J. M. 1990 Mixing, chaotic advection, and turbulence. *A. Rev. Fluid Mech.* **22**, 207–54.

Ottino, J. M. & Wiggins, S. 2004 Introduction. *Phil. Trans. R. Soc. Lond.* A **362**, 923–935.

Perry, A. D. & Wiggins, S. 1994 KAM tori are very sticky: rigorous lower bounds on the time to move from an invariant Lagrangian torus with linear flow. *Physica* D **71**, 102–121.

Pesin, Ya. B. 1977 Characteristic Lyapunov exponents and smooth ergodic theory. *Russ. Math. Surv.* **32**, 55–114.

Przytycki, F. 1983 Ergodicity of toral linked twist mappings. *Annls Sci. Ec. Norm. Super.* **16**, 345–354.

Qian, S. & Bau, H. H. 2002 A chaotic electroosmotic stirrer. *Analyt. Chem.* **74**, 3616–3625.

Rom-Kedar, V. & Wiggins, S. 1990 Transport in two-dimensional maps. *Arch. Ration. Mech. Analysis* **109**, 239–298.

Sinai, Ya. G. 1994 *Topics in ergodic theory*. Princeton University Press.

Stroock, A. D., Dertinger, S. K. W., Ajdari, A., Mezić, I., Stone, H. A. & Whitesides, G. M. 2002 Chaotic mixer for microchannels. *Science* **295**, 647–651.

Wiggins, S. 1990 *Introduction to nonlinear applied dynamical systems and chaos*. Springer.

Wiggins, S. 1992 *Chaotic transport in dynamical systems*. Springer.

Wiggins, S. 1999 Chaos in the dynamics generated by sequences of maps, with applications to chaotic advection in flows with aperiodic time dependence. *Z. Angew. Math. Phys.* **50**, 585–616.

Wojtkowski, M. 1980 Linked twist mappings have the K-property. In *Nonlinear Dynamics Int. Conf., New York*, pp. 65–76. Annual Report, vol. 357. New York Academy of Science.

Yomdin, Y. 1987 Volume growth and entropy. *Isr. J. Math.* **57**, 285–300.