INT. J. NUM. METH. HEAT FLUID FLOW, VOL. 5, 549–573 (1995)

SIMULATION OF FLUID MIXING USING LEAST-SQUARES FINITE ELEMENTS & PARTICLE TRACING

G. F. CAREY AND Y. SHEN

Texas Institute for Computational Mechanics, The University of Texas at Austin

ABSTRACT

A least-squares finite element analysis of viscous fluid flow together with a trajectory integration technique for tracers is formulated and provides a mechanism for investigating mixing. Tracer integration is carried out using an improved Heun predictor-corrector. Results from our supporting numerical studies on the CRAY and Connection Machine (CM) closely resemble the patterns of mixing observed in experiments. A 'box-counting' scheme and other measures to characterize the level of mixing are developed and investigated. This measure is utilized in numerical experiments to determine an optimal forcing frequency for mixing by periodic boundary motion in a rectangular enclosure. Some details concerning the numerical schemes and vector-parallel implementation are also included.

KEY WORDS Mixing Particle tracing Parallel processing Least squares FEM

INTRODUCTION

Mixing is important in many industrial processes and in naturally occurring phenomena. There have been several experimental studies designed to afford a better understanding of the nature of mixing in fluids in the absence of diffusion. Of particular note in the context of the present work are experiments for laminar 'two-dimensional' mixing of tracers in viscous fluids under periodic boundary condition[19]. Complex patterns evolve from an initial 'blob' or 'blobs' of tracer as a result of 'stretching' and 'folding' with repeated periodic cycles. Moreover, different spatial domains in the flow field may exhibit very different mixing characteristics. For example, some regions will have convoluted striations with many long thin folds, whereas in other regions the initial 'blob' of dye tracer will be scarcely deformed indicating little mixing. 'Islands' in the fluid may exist that exterior tracer does not penetrate and from which tracer can not depart. Such issues may be particularly relevant in certain applications such as reaction or combustion processes.

Clearly, the extent of mixing is of considerable practical interest. Moreover, the configuration and size of 'islands' and of mixing patterns depend upon the forcing conditions such as the frequency of boundary motion and on material properties such as viscosity. A mechanism for quantifying mixing is highly desirable and can be applied in a control-feedback cycle to determine optimal mixing conditions (e.g. the best periodic boundary frequency for a given fluid). Simulation provides an opportunity to obtain precise quantitative data on mixing and has been a subject of several recent studies[1,11].

In the present work we develop a finite element formulation for laminar viscous flow in two dimensions and apply this to study mixing. The finite element formulation is non-standard and

0961–5539/95/060549–25$2.00

Received May 1994

is actually based on a least-squares mixed approach. However, this is incidental to the thrust of the present work and therefore the finite element scheme will be only briefly described here. Further details of the least-squares treatment are given elsewhere[6,8]. The flow field computed from the finite element analysis is used in a trajectory integration scheme for convective transport of sets of particles that simulate tracers. This trajectory scheme is implemented on a CRAY shared memory multiprocessor and on the CM2 distributed parallel processor.

Images of the tracer pattern provide the best indication of the nature and level of mixing. However, a simple scalar measure to quantify mixing or at least discriminate cases is highly desirable. Several scalar concepts related to area coverage, perimeter, tortuosity, spreading, and statistical measures are considered. Related fractal ideas and 'box-counting' are investigated and utilized in a set of numerical studies. In particular, we demonstrate how such measures of mixing might be used to improve the mixing characteristics of a system in which the periodic forcing frequency can be varied.

The outline of the paper is as follows: first, we briefly summarize the statement of the tracer mixing problem for a given velocity field. Next we introduce the least-squares mixed finite element scheme used in the flow field computations. The Heun predictor-corrector scheme for particle tracing is then presented. The problem of defining a simple measure of continuum mixing is subsequently described and this is followed by a treatment of related particle mixing computations. Numerical experiments give mixing patterns and compare several of the previous measures for test cases corresponding to steady flow and periodic flow. Finally, we give vector and parallel performance statistics for the simulations.

MIXING CONCEPTS

Mixing may be broadly categorized as of three types:

(i) *stretching and folding:* an initial designated material region stretches and folds by action of flow without diffusion and breakup;
(ii) *diffusion:* the stretching and folding in category (i) is accompanied by diffusion;
(iii) *breakup:* The initial region stretches, folds and breaks due to interfacial tension forces, producing smaller fragments which might in turn stretch and break.

In the present work we focus on the problem of convective mixing with negligible diffusion in category (i).

The actual tracer 'blob' can be modelled as either a continuum or as a discrete system of 'particles' and we discuss the relative merits of these respective approaches later. In the continuum model the transport of a species of concentration $c(x,t)$ is simulated using either the convective equation or a convection-diffusion equation in the case where the diffusion is not negligible. Since diffusion is small in the present work we can compute an approximation to the convection-dominated problem by a Taylor-Galerkin, SUPG or similar scheme. In fact a least-squares scheme can again be used[23]. The second approach is to model the tracer as a large set of representative 'particles' and this is the approach followed here. The mixing pattern can then be visualized and simple mixing measures can be investigated or tested. Moreover, this strategy lends itself to distributed parallel computing and image processing.

Let $u(x,t)$ be the velocity vector flow field for a fluid in motion. At some time $t=t_0$ (the initial time from the viewpoint of the mixing problem) we select a material 'blob' or 'blobs' in the fluid to correspond to a dye tracer. Thereafter, fluid mixing of this tracer in the absence of diffusion can be characterized by the stretching and repeated folding of the material in the 'blobs' (the tracer) under the action of the fluid velocity. That is, the position $x(t)$ of any point in the tracer blob satisfies (by definition) the pathline equation

$$\frac{dx}{dt}=u(x,t) \tag{1}$$

where at $t=t_0$, $x(t_0)$ is given. Since $u(x,t)$ is presumed known, it follows that the pathline of the tracer particle with initial position $x(t_0)$ can be traced out for $t>t_0$ by integrating (1).

In the 'particle' representation, the tracer 'blob' is approximated as a discrete set of particles where each particle corresponds to 'lumping' the tracer at the centre of a representative area. Clearly, to develop a reasonable representation of the mixing pattern, we need to take a sufficiently large sample of 'particles' distributed through the initial blob. (If there is extensive local stretching and folding, the relative motion of nearby particles is significant.) Problem (1) then reduces to integrating trajectories for the decoupled dynamical system,

$$\frac{dx_p}{dt}=u(x_p,t), \qquad p=1,2,\ldots,P \tag{2}$$

with initial vector $\{x_p(0)\}=X_0$ specified. In the numerical studies considered later, a modified Heun method is used to integrate trajectories in parallel on the Connection Machine.

The finite dimensional dynamical system (2) is deterministic for laminar flow and stochastic for turbulent flow. In the laminar case of interest here, (2) is autonomous for steady flow and non-autonomous for unsteady flow. Chaotic particle motion and mixing may be encountered in both the stochastic turbulent flow and in the unsteady laminar flow case. Although the driving velocity field in (2) can correspond to any fluid type, we shall, for convenience, consider the case of mixing in incompressible viscous fluid flow. This case is of considerable practical interest and there is also an abundance of experimental data in the literature.

LEAST-SQUARES FLOW SOLUTION

Incompressible laminar flow of a Newtonian fluid may be described by the Navier-Stokes equations,

$$\left.\begin{aligned}\frac{\partial u}{\partial t}+u\cdot\nabla u+\frac{1}{\rho}\nabla p-\nu\nabla^2 u=g\\ \nabla\cdot u=0\end{aligned}\right\} \tag{3}$$

for transport of momentum and mass conservation, respectively, where u is velocity, p is pressure, ρ is density, ν is kinematic viscosity and g is a prescribed body force. Initial conditions $u(x,0)=u_0(x)$ with $\nabla\cdot u_0=0$ and boundary conditions are specified (e.g. no slip $u=u_s$ on $\partial\Omega$ where u_s is the boundary velocity, is appropriate in the mixing studies later). These complete the mathematical statement of the viscous flow problem.

Except in a few simple cases, the velocity u satisfying the Navier-Stokes equations (3) is not known analytically. Numerical solution of the viscous flow problem provides instead an approximate velocity field $u_h(x,t)$ so that problem (2) may be modified by replacing u by u_h. Provided u_h is sufficiently close to u, the resulting tracer pathlines will approximate the actual tracer motion well.

As noted in the Introduction, an approximate solution of (3) can be obtained in a variety of ways. In a different context, we have been investigating the use of least-squares finite elements and have, for convenience, applied this approach here. For completeness, we briefly give the formulation since it is non-standard and an interesting topic in its own right. A more detailed treatment that also considers some limitations of the least-squares formulation is in preparation[9].

First, the vorticity ζ is introduced to recast the governing equations as a first-order system

for two-dimensional flow in domain Ω as

$$\begin{aligned}
&\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}=0\\
&\frac{\partial u}{\partial t}+u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}+\frac{1}{\rho}\frac{\partial p}{\partial x}+\nu\frac{\partial \zeta}{\partial y}=g_1\\
&\frac{\partial v}{\partial t}+u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}+\frac{1}{\rho}\frac{\partial p}{\partial y}-\nu\frac{\partial \zeta}{\partial x}=g_2\\
&\zeta-\frac{\partial v}{\partial x}+\frac{\partial u}{\partial y}=0
\end{aligned} \tag{4}$$

with appropriate boundary conditions specified on $\partial\Omega$. In the subsequent numerical studies, we consider the case with u, v specified on $\partial\Omega$ and reference pressure p_0 specified at an arbitrary point.

Next, semi-discretizing by finite differencing with respect to time at $t_\theta=\theta t_{n+1}+(1-\theta)t_n$, for time increment $\Delta t_n=t_{n+1}-t_n$ and $0\leqslant\theta\leqslant 1$

$$\begin{aligned}
&\theta\left(\frac{\partial u^{n+1}}{\partial x}+\frac{\partial v^{n+1}}{\partial y}\right)+(1-\theta)\left(\frac{\partial u^n}{\partial x}+\frac{\partial v^n}{\partial y}\right)=0\\
&\frac{u^{n+1}-u^n}{\Delta t}+\theta\left(u^{n+1}\frac{\partial u^{n+1}}{\partial x}+v^{n+1}\frac{\partial u^{n+1}}{\partial y}+\frac{1}{\rho}\frac{\partial p^{n+1}}{\partial x}+\nu\frac{\partial \zeta^{n+1}}{\partial y}\right)\\
&\quad+(1-\theta)\left(u^n\frac{\partial u^n}{\partial x}+v^n\frac{\partial u^n}{\partial y}+\frac{1}{\rho}\frac{\partial p^n}{\partial x}+\nu\frac{\partial \zeta^n}{\partial y}\right)=g_1(t_\theta)\\
&\frac{v^{n+1}-v^n}{\Delta t}+\theta\left(u^{n+1}\frac{\partial v^{n+1}}{\partial x}+v^{n+1}\frac{\partial v^{n+1}}{\partial y}+\frac{1}{\rho}\frac{\partial p^{n+1}}{\partial y}-\nu\frac{\partial \zeta^{n+1}}{\partial x}\right)\\
&\quad+(1-\theta)\left(u^n\frac{\partial v^n}{\partial x}+v^n\frac{\partial v^n}{\partial y}+\frac{1}{\rho}\frac{\partial p^n}{\partial y}-\nu\frac{\partial \zeta^n}{\partial x}\right)=g_2(t_\theta)\\
&\theta\left(\zeta^{n+1}-\frac{\partial v^{n+1}}{\partial x}+\frac{\partial u^{n+1}}{\partial y}\right)+(1-\theta)\left(\zeta^n-\frac{\partial v^n}{\partial x}+\frac{\partial u^n}{\partial y}\right)=0
\end{aligned} \tag{5}$$

The residual vector for admissible functions ζ, u, v, and p in (5) then may be conveniently expressed in terms of $z^T=(\zeta,u,v,p)$ as,

$$r^{n+1}=A(z^{n+1})z_x^{n+1}+B(z^{n+1})z_y^{n+1}+Cz^{n+1}-f^n \tag{6}$$

where f^n corresponds to all expressions in (5) evaluated at previous time level t_n as well as the contribution of the forcing vector g.

Introducing the successive approximation technique to iteratively linearize the residual expression in (6), at iteration $s+1$ for approximation z^{n+1} we have,

$$r^{n+1,s+1}=A(z^{n+1,s})z_x^{n+1,s+1}+B(z^{n+1,s})z_y^{n+1,s+1}+Cz^{n+1,s+1}-f^n \tag{7}$$

The least-squares residual formulation corresponds to finding the admissible solution $z^{n+1,s+1}$ minimizing the least-squares quadratic functional,

$$I=\frac{1}{2}\int_\Omega (r^{n+1,s+1})^T(r^{n+1,s+1})\,dx\,dy \tag{8}$$

In the present work a discretization Ω_h and finite element expansions ζ_h, u_h, v_h, p_h are introduced to construct an approximate least-squares finite element formulation[23]. An interesting feature

of the least-squares mixed method is that it is not subject to the LBB condition encountered in the mixed Galerkin method[3,8,10,14,21]. However, there are other open issues concerning the least-squares method that warrant study[9].

When the finite element expansion are substituted in (8), the stationary condition leads to a sparse symmetric linear algebraic system for the iterate vector $Z^{n+1,s+1}$ of nodal solution values,

$$K^{n+1,s}Z^{n+1,s+1}=F^{n+1,s} \tag{9}$$

System (9) is solved using sparse Gaussian elimination for iterate $s+1$. Nonlinear iteration ceases when specified convergence tolerances on successive iterates and on the residual (7) are met. The velocity field approximation $u_h(x,t^{n+1})$ can then be interpolated from the nodal velocity solution using the finite element basis. This velocity field provides the forcing function in the dynamical system (2) for tracing trajectories.

PARTICLE TRACING ALGORITHM

Mixing patterns for a collection of 'particles' lumped as discrete points can be constructed by integrating the dynamical system (2) and tracing the particle trajectories. Since we are interested in mixing over significant time periods, it is important that the time integration scheme be accurate, so that accumulated integration error does not excessively degrade the accuracy. Furthermore, many particle trajectories must be computed and the integration scheme should be efficient and easily vectorized or parallelized.

The improved Heun predictor-corrector scheme[27] is an $O(\Delta t)^2$ two-step integrator that has the desired attributes. The basic scheme for particle p at timestep n is:

Predictor:

$$\bar{x}_p^{n+1}=x_p^n+\Delta t\left[\frac{3}{2}u(x_p^n,t^n)-\frac{1}{2}u(x_p^{n-1},t^{n-1})\right] \tag{10}$$

Corrector:

$$x_k^{n+1}=x_p^n+\Delta t\left[\frac{1}{2}u(\bar{x}_p^{n+1},t^{n+1})+\frac{1}{2}u(x_p^n,t^n)\right] \tag{11}$$

In the predictor step (10) the velocities $u(x_p^{n-1},t^{n-1})$ and $u(x_p^n,t^n)$ are known (interpolated using the finite element basis) from the nodal solution at t^{n-1} and t^n, so the predictor calculation reduces to SAXPY vector operations. Using $\bar{x}_p^{n+1}$, the velocity vector u at the predicted particle locations can be interpolated similarly from the nodal velocity vector at t^{n+1} on the associated elements which are easily located. Then the corrector step (11) is applied and, similarly, involves SAXPY operations. Hence, the Heun method provides an efficient accurate scheme for trajectory integration on appropriate vector architectures. The tracing scheme is also readily parallelized since the motion of any tracer particle is independent of the motion of other particles. These issues are considered in more detail later.

CONTINUUM MIXING MEASURES

Before proceeding with numerical experiments and imaging the mixing patterns based on this trajectory scheme, we first consider the problem of quantifying mixing for the continuum and discrete cases. This is an open question and, as we show, it is doubtful that any simple measure can be derived. Nevertheless, some important features of the mixing problem are revealed.

The problem of reliably quantifying the level of mixing in either an actual experiment or in a numerical simulation is not straightforward. Moreover there are several subtle distinctions between these two respective situations which we discuss later. Let us first briefly consider the

experimental case[19]. Here a dye blob is introduced into the fluid at some location and the dye pattern is examined at later times. Dye is convected by the fluid and, to a small extent, also diffuses. The patterns shown in the experiments indicate the extent to which the dye is transported, but not the concentration (except in a qualitative sense based on the gradation in colour).

Thus it is important, *a priori*, to define what we mean by mixing. In the strictest possible sense we should intepret a 'completely mixed' case as one where the dye concentration is uniform throughout the domain. For example, in the absence of convection an initial dye blob would be diffused and diluted to a uniform constant steady state. The concentration as a function of position at any time would be determined as a solution of the diffusion equation. Similarly, if convection is involved, the solution to the convection-diffusion equation gives the evolving species concentration. The difference between the uniform steady state diffusion result and the solution to the convection-diffusion equation at any time represents the 'departure from a completely mixed configuration'. Hence, the more closely the constant uniformly-mixed steady behaviour is modelled, the better the mixing. Attempts to quantify mixing rigorously using a simple measure such as area coverage fail and counter-examples may be devised as seen later. Yet, qualitatively, these scalar measures may be useful and give insight into the mixing process or aid in design analysis. Accordingly, we now briefly consider some of the most natural choices and comment on their limitations. There are also some subtle distinctions between the physical situation and the mathematical model that are best considered via measure theory and which we comment upon later.

The simplest approach to measure the extent of two-dimensional mixing on a section or surface would be to attempt to calculate the area fraction of the dye pattern at times of interest during mixing[19]. This is relatively easy if the dye remains in a contiguous blob since the problem reduces to calculating the section area of the blob (assuming, of course, that the dye blob is essentially two-dimensional in nature). For example, a triangulation of the dye shape can be easily constructed from a set of points on the boundary contour of the blob[5] and the dye area is approximately the sum of element areas. Alternatively, a probabilistic (Monte Carlo) approach can be used with the area determined by the number of 'hits' on the blob relative to the number of 'misses' obtained using a random number generator with range corresponding to the fluid domain Ω.

If mixing is primarily due to diffusion, then this area measure is reasonable although still incomplete. In the numerical model, the diffusive transport for concentration $c(x,t)$ would need to be included and the fractional dye area with c greater than a small positive tolerance level would give the area coverage. (Obviously the area fraction also offers no information on the local concentration which may be of critical interest in applications.) The problem of determining the area becomes more complicated when mixing occurs by convection and the dye pattern is 'folded' back and forth with many thin lines. Even computing the area accurately is then difficult. This case also serves to pinpoint a fundamental weakness in using area to quantify mixing. Let us assume that the thin folds are regularly distributed over Ω with alternate thin strands separated by thin open strips. This distribution is conceptually consistent with the idea of a 'reasonably' well-mixed system. However, area alone can not describe mixing since we can not differentiate any other shape pattern (e.g., a square in one corner) that has the same area as this coverage by thin strands. Only if the pattern area approaches the domain area as in the diffusive mixing case, would this be viable.

The presence of folds suggests that the perimeter of the dye pattern or its tortuosity might also provide alternative descriptions of mixing. However, even in the diffusive case the perimeter may not grow monotonically as the dye spreads. The tortuosity in porous media applications is often characterized as the ratio of area to perimeter and can be similarly introduced here for the convective folded mixing mechanism. As folds degenerate to thin curves and merge, overlap, or cross, the difficulty of computing a tortuosity measure becomes prohibitive. Tortuosity also fails to describe the spread of dye as seen in later examples.

A complete detailed description (and image) of the concentration as a function of position and

time will characterize the transport solution but does not immediately offer a simple means for comparative measures of mixing. We can, however, use statistical procedures to compare a given concentration state with a uniform distribution by computing means, standard deviations and tests of fit. This will permit a more complete description of the state of mixing but will not discriminate pattern structures. For example, two patterns of very different structure would still give rise to the same primitive statistics. Furthermore, in many situations it is the structure, shape and extent of the pattern that is of interest: For instance, whether the initial dye blob is folded into many thin striations over most of the domain, over part of the domain, or remains a relatively undisturbed 'island' may be significant. This evidently is a more particular interpretation of 'mixing'.

Since mixing involves the extent of the pattern and its relative density these two concepts (pattern extent and pattern density) could be utilized to develop a statistically-based measure. For example, the maximum, minimum and mean diameters of the pattern relative to the domain together with standard deviation would provide a practical characterization and this could be utilized together with the area to produce a more complete description. The problem of assigning relative weighting to these features to determine which of several patterns is relatively more mixed must still be addressed.

There are other issues that can complicate the situation. These concern concepts that we might consider to be local (micro-mixing) versus global (macro-mixing). To examine this further, let us now develop some related mathematical concepts. Consider, for example, a circular curve. This can be deformed into any other simple closed curve and all such curves are topologically equivalent or homeomorphic. That is, any two such curves can be transformed into one another by a simple continuous deformation. We remark that length need not be preserved in so far as this mathematical (continuum) argument is concerned. However, if the original curve is deformed to intersect itself, then the transformation is no longer continuous and the curves are no longer topologically equivalent. Since one aspect of mixing involves complex crossing of the tracer material, the topological characterization would provide part of a more complete description for mixing that other simpler approaches (e.g., length measures, etc.) alone can not. In this example, the number of curve intersections and their distribution would provide some characterization of the topological complexity of the pattern and this might be an important feature of the mixing from a practical standpoint. Similarly, ideas from 'knot theory' could be applied[15,17].

Let us briefly return to the area and length measures and consider a simple closed curve of length L. The shape of maximum area enclosed by the curve is a circle, and the ratio A/L of area $A=\pi r^2$ to perimeter $L=2\pi r$ is $L/4\pi$. If the circle is deformed to other simple curves then the area is decreased so the ratio $L/4\pi$ is an upper bound for the set of simple closed curves of length L. For example, the square of side $L/4$ has $A/L=L/16<L/4\pi$. The square has the maximum area of rectangles with given perimeter L and $A/L=wh/2(w+h)$ where $w=$width, and $h=$height. As w or $h\to 0$ with L fixed we have $A/L=wh/L\to 0$. That is, $0<A/L<L/4\pi$ for all simple closed curves of fixed length L.

Hence, if we have a simple closed tracer 'blob' with perimeter L, the maximum area would correspond to the circular shape ($A/L=L/4\pi$) and the more elongated shapes to smaller areas with $A/L\to 0$. We see that there is a tautology here when area and tortuosity measures are compared. The shape that has the greatest area would physically have the lowest local concentrations and in this sense has the greatest mixing. On the other hand the elongated (possibly folded) shapes characteristic of our tracer patterns imply greater mixing as indicated by greater tortuosity.

Next, assume that the simple closed curve is arbitrarily perturbed to an irregular curve. The problem of determining the length of the new curve leads us to consider ideas related to the theory of fractals. In particular, the fractal dimension of the irregular boundary is greater than the fractal dimension of a line. Therefore, fractal dimension offers a measure of local irregularity and hence some information regarding 'micro-mixing'. However it fails to provide information

on scale. The desire to estimate fractal dimension has also led to simple constructive approaches by 'box-counting methods' which fit naturally with the discrete particle tracing approach adopted here. Later, we develop a 'box-counting' algorithm motivated by this idea. This is also a consequence of the observation that we are really dealing with discrete particle tracing in the present treatment rather than the continuum mixing problem and this issue is taken up next.

PARTICLE MIXING MEASURES

We emphasize that the previous mathematical descriptions, and to some extent the discussion of experimental mixing, have relied on a continuum description. However, in reality the fluid and tracer physically consist of discrete molecules and therefore are an extremely large finite set. Similarly, computations are not on the real number system but assume finite precision arithmetic on a countable number set (which is, in fact, a set of measure zero). Hence it is appropriate to model mixing numerically using a finite particle set. However, practical considerations indicate that only a moderately large number of particles can be traced even on the massively parallel computer systems considered here.

The tracer patterns that arise in mixing may be quite complex as illustrated in the examples involving steady and periodic flows considered later. Moreover, it is important to quantify the level of tracer mixing from an initial configuration under different conditions such as periodic forcing frequencies and for different fluids. Since the mixing patterns may be convoluted with thin striations, many folds, isolated 'islands' and other complex structures the ideas associated with fractals may help quantify the level of mixing. In particular, the fractal dimension indicates how well an object (here the mixing pattern) covers the associated metric space and therefore can be used to construct a comparative measure for mixing that may be used to supplement other measures such as those mentioned previously.

Particle occupation estimate

Intuitively, we realize that the tracer blob is well-mixed if the tracer particles become uniformly distributed throughout the domain. Conversely, mixing will be low if the tracer particles are less uniformly distributed. This primitive concept leads to a simple contructive approach for determining fractal dimension by a 'box counting method'. The idea is similar to the probabilistic (Monte Carlo) technique mentioned earlier for determining areas or volumes based on object 'hits' by a random number generator. The basic idea is to cover the flow domain with a uniform mesh and to count the proportion of boxes containing the evolved tracer particles (the attractor). For simplicity let us consider the case of a unit square or cube domain. More formally we have[4]:

Let S be a compact non-empty set in $B=[0,1]^m \subset R^m$, $m=1,2,3$ with Euclidean metric. Cover B by a uniform mesh of boxes with side 2^{-n}. Let $N_n(S)$ be the number of boxes which intersect S. Then the fractal dimension D_f may be defined as,

$$D_f = \lim_{n\to\infty} \frac{\log(N_n(S))}{\log(2^n)} \tag{12}$$

It follows that, $D_f=0$ for a point, $D_f=1$ for a line segment, $D_f=2$ for a surface, and $D_f=3$ for a solid.

Now (12) alone can not describe mixing. This is clear since the fractal dimension of any surface 'blob' $S\subset\Omega$ is 2 independent of the size of S. For example if Ω is a unit square and S is the region $1/2<x<3/4$, $1/2<y<3/4$, the fractal dimension of S is 2. Similarly, the fractal dimension of S for $1/4<x<3/4$ and $1/4<y<3/4$ is also 2 even though the area is 4 times larger.

However, the definition of fractal dimension in (12) motivates the development of box-counting and similar discrete techniques as an additional measurement to help characterize the mixing of a large but finite set of tracer particles. Let S^* denote the corresponding particle pattern with P particles in domain $B=[0,1]^m$, $m=1,2,3$. Next, construct a box partition for specified n

such that $P \geqslant (2^n)^m$ hold. The number of boxes containing one or more particles is then $N_n(S^*)$ and can be easily computed. Following (12), we may define,

$$D_f^* = \frac{\log(N_n(S^*))}{\log(2^n)} \tag{13}$$

as a mixing estimate for S^*. It follows $D_f^* \in [0,2]$ for a 2-D mixing problem ($m=2$).

To illustrate the dependence on n, let us consider $P=16$ particles defining a pattern of particles S^* in $B=[0,1]^2$. Then for $n=1$ there are four boxes, and $D_f^* = \log N_1(S^*)/\log 2$. If all 16 particles are in one box then $N_1(S^*)=1$ so $D_f^*=0$. If the particles are in all four boxes then $D_f^*=2$. An intermediate value of D_f^* would indicate an intermediate measure of mixing. For example, if the particles are in two boxes we get $D_f^*=1$. If the 16 particles are uniformly distributed then $n=2$ would provide a uniform mesh of boxes with exactly one particle per box. For $n>2$, 16 particles would not be sufficient to characterize a uniform distribution. More generally, we require $P \geqslant 2^{mn}$. Setting $P=2^{mn}$ or equivalently $n=\log_2(P)/m$ will help discriminate patterns.

This estimate (13) for D_f^* is similar to an area measure for the continuum case where the area is computed approximately by accumulating box areas, and therefore has similar shortcomings. For example, D_f^* gives the same value irrespective of the precise distribution of particles among the box covering—most of the particles may be in one box and a single particle in each of the remaining boxes, etc. The mean, median, standard deviation and related statistical data on particle position can also be easily computed.

We have emphasized that the fractal dimension can not be used directly to measure particle mixing by box-counting. However, in chaotic mixing the tracer pattern can be interpreted as a strange attractor and the box-counting estimate D_f^* is an indication of the degree of chaotic mixing. Moreover, we anticipate that, in the event the mixing pattern is a fractal object, then box-counting should yield the fractal dimension. Let us briefy consider two well-known elementary patterns generated by simple iteration function systems—the uniform pattern on a square grid and the Sierpinski gasket for the triangular pattern[4]. As a mixing map, let a set of P particles in the initial square be mapped to $P/4$ particles in each of the four quadrants by the first iteration on the square and so on to $P/2^{2n}$ particles per box after n iterations and $n \leqslant 0.5 \log_2 P$. Then the number of boxes with particles is $N_n(S^*)=2^{2n}$, and (13) yields,

$$D_f^* = \frac{\log(N_n(S^*))}{\log(2^n)} = \frac{\log(2^{2n})}{\log(2^n)} = 2$$

which agrees with the expected fractal dimension for the square attractor. Similarly, for the Sierpinski map, P particles in the initial triangle are mapped to $P/3$ in each vertex subtriangle with none in a centre subtriangle and so on. At iteration n, there are $P/3^n$ particles in each 'active' subtriangle so that $N_n(S^*)=3^n$ and,

$$D_f^* = \frac{\log(N_n(S^*))}{\log(2^n)} = \frac{\log(3^n)}{\log(2^n)} = \frac{\log(3)}{\log(2)}$$

so we obtain the fractal dimension for the Sierpinski pattern.

Particle concentration estimates

As noted previously, a single measure can not provide a true characterization of the mixing pattern or process. However, several such parameters that complement one another may provide a useful description. We have also indicated that statistical information can also be incorporated in the parameterization. This approach can be utilized to estimate the particle concentration which is not addressed by the previous ideas.

Let $S^* \subset B=[0,1]^m$ with $m=2$. Let $P=2^{mn}$ be the number of particles in a given mixing simulation with integer n, and consider a uniform partition of the domain to $G=P/2^{2l}$ boxes

for integer $0 \leqslant l \leqslant n$. Number the boxes $g = 1,2,\ldots,G$ and let the local particle concentration c_i for box i be the fraction of particles in the box; i.e. $c_i = P_i/P$. Now a uniformly mixed particle pattern would correspond to a state with $P_i = P/G = 2^{ml}$ particles in each box and the concentration would be categorized accordingly as uniform with value $c_i = 1/G$. The other extreme case corresponds to all particles in a single box j so that $c_i = 0$ except for box j where $c_j = 1$. We may elect to plot the local values c_i as a piecewise-constant concentration surface (histogram) over the domain. Alternatively, we may fit the values in some related fashion and then construct particle concentration iso-contours.

The departure of a given particle mixing state from the uniform state can also be used to characterize the level of mixing. For example, the standard deviation σ_c of box particle concentration c_i with respect to the uniform distribution of particle $\bar{c} = 1/G$ can be computed as,

$$\sigma_c = \sqrt{\frac{1}{G-1} \sum_{g=1}^{G} (c_i - \bar{c})^2} \tag{14}$$

Then $\sigma_c = 0$ for the uniform state; σ_c increases as the configuration departs from uniformity; and $\sigma_c = 1/\sqrt{G}$ for the most nonuniform state. Let the scaled deviation be defined by $D_c = \sqrt{G}\sigma_c$ so that $D_c \in [0,1]$, which gives a normalized statistic related to the distribution of local particle concentration.

Particle separation estimate

The 'spread' of the particles can also be characterized statistically in several ways, e.g. let g_{ij} be the distance separating particles i and j in the state. The matrix $D = (d_{ij})$, $i,j = 1,2,\ldots,P$, then can be introduced to describe the relative separation of particles in the system. Clearly, $d_{ii} = 0$ and the entries d_{ij} increase as particles move further apart. The groupings of small and large entries in matrix D gives an indication of particle spread and clustering. E.g. if all entries are small, then there is one tight cluster of particles. If matrix D has 2 sets of small entries and 2 sets of large entries, there are two tight particle clusters that are widely separated, and so on.

Various norms of this matrix can also be introduced to help quantify the particle spread. For example, we can compute the Euclidean length $\|y\|$ of vector $y = Dx$ where x ranges over the unit radial vector from an origin,

$$\|D\| = \max_{x \neq 0} \frac{\|Dx\|}{\|x\|}$$

or, equivalently,

$$\|D\| = \max_{\|x\| = 1} \|Dx\| \tag{15}$$

and then characterize particle separation using $\|D\|$.

Dynamical system interpretation

Variance in the particle trajectories can be used as a mechanism for further characterizing the tracer pattern behaviour and investigating the dynamical system for the tracer[2,13,25]. The approach is based on similar ideas to those for studying the dynamics of patterns in flow between concentric cylinders[26].

Let $(x_p(t_i), y_p(t_i))$ be the position of particle p at time t_i. Then the displacement of particle p during timestep $\Delta t = t_i - t_{i-1}$ is given by,

$$\Delta d_p(t_i) = \sqrt{[x_p(t_i) - x_p(t_{i-1})]^2 + [y_p(t_i) - y_p(t_{i-1})]^2}$$

The total distance traversed by particle p at time t_i is,

$$d_p(t_i)=d_p(t_{i-1})+\Delta d_p(t_i)$$

with $d_p(t_0)=0$. The mean distance traversed by all particles at time t_i follows as,

$$\bar{d}(t_i)=\frac{1}{P}\sum_{p=1}^{P} d_p(t_i)$$

and the variance in tracer particle distance traversed relative to the mean is given by,

$$\langle\sigma^2\rangle=\frac{1}{P}\sum_{p=1}^{P}[d_p(t_i)-\bar{d}(t_i)]^2 \tag{16}$$

For a chaotic Hamiltonian system,

$$\langle\sigma^2\rangle\sim t^{\mu}$$

and $\mu(\in[0,3])$ measures the diffusivity of the dynamic system[24]. Then μ provides a relevant property of the dynamical system.

NUMERICAL RESULTS AND ANALYSIS

Several physical experiments for dye tracing in glycerine for a cavity of width W and height H under both steady and periodic forcing motion of the walls have been reported[11] and are considered in the following numerical studies. The corresponding dimensionless variables are $x'=x/L, u'=u/V, p'=p/\rho V^2$, with $L=H^2/W$ the characteristic length scale and V a characteristic (wall) velocity. The dimensionless time is taken to be $t'=Vt/L$ for nonperiodic flow and $t'=t/T$ for periodic flow with period T. Omitting the prime notation and introducing Reynolds number $Re=VL/\nu$, and Strouhal number $Sr=fL/V$ with periodic flow at frequency $f=1/T$, we get,

$$\begin{aligned} Sr\frac{\partial \boldsymbol{u}}{\partial t}+\boldsymbol{u}\cdot\nabla\boldsymbol{u}+\nabla p-\frac{1}{Re}\nabla^2\boldsymbol{u}&=\boldsymbol{g}\\ \nabla\cdot\boldsymbol{u}&=0\\ Sr\frac{\mathrm{d}x_p}{\mathrm{d}t}&=\boldsymbol{u}(x_p,t),\qquad p=1,2,\ldots,P \end{aligned} \tag{17}$$

The initial particle distribution is specified by $\{x_p(0)\}=X_0$. The Strouhal number Sr in (17) can be viewed as a dimensionless forcing frequency for the periodic cavity flow. For the following numerical studies, the kinematic viscosity $\nu=6\,\mathrm{cm}^2/\mathrm{sec}$, Reynolds number $Re\sim O(1)$, and $Sr\ll 1$. In the mixing studies we display the computed patterns and where available compare with experiment. We also compute the particle occupation estimate D_f^* from (13), and the variance $\langle\sigma^2\rangle$ in the particle trajectory from (16).

Mixing patterns in steady flow

The mixing pattern is first simulated in a steady flow. For the steady flow, we take $Sr=1$ in (17) and first integrate the flow problem to a steady-state. Then this steady state flow field defines the forcing function for the particle trajectory equation. Mixing is induced by motion of the vertical side walls of a cavity with width $W=10\,\mathrm{cm}$ and height $H=6\,\mathrm{cm}$. The side walls move in opposite directions with velocities of magnitude $10/6\,\mathrm{cm/sec}$ ($Re=1$). The steady problem offers a good validation case for the flow solver and trajectory computation since the streamlines are fixed, and the tracer cannot cross a streamline. This problem also provides an accuracy check on the integration method for the class of problems of interest.

The flow field was integrated to steady state on a uniform 10×6 finite element mesh with 9-node biquadratic elements for all four field variables in the least-squares formulation described previously. A fixed time step $\Delta t = 0.6$ sec. was employed for $0 \leqslant t \leqslant 12$ sec. In the solution procedure, the velocities at a representative node vary with time and reach the steady state in about 3 sec. as indicated in *Figure 1*. The streamline, pressure and vorticity are shown in *Figure 2* for the solution at $t = 12$ sec., and are essentially steady state results. The streamfield plot reveals a hyperbolic point in the centre and two elliptic points symmetrically located on the horizontal centre line. This defines the steady state solution $u(x)$ in the trajectory calculation following.

Since the flow is steady in this case, we simply store the steady velocity field for the trajectory integration scheme with $u(x,t) = u(x)$ for use in (10), (11). A uniform band of (400×25) tracer particles is placed along a diagonal line extending from bottom left to top right as indicated in

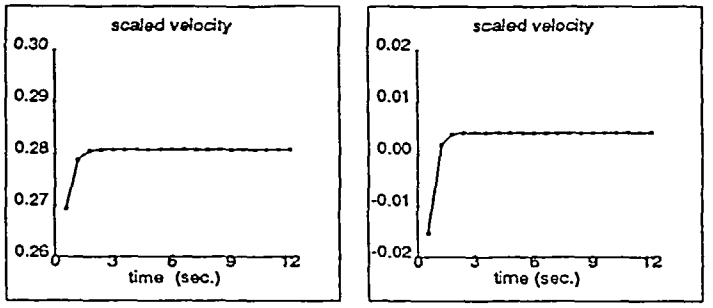

(a) velocity u/V (b) velocity v/V

Figure 1 Velocity history at (1 cm, 4.8 cm)

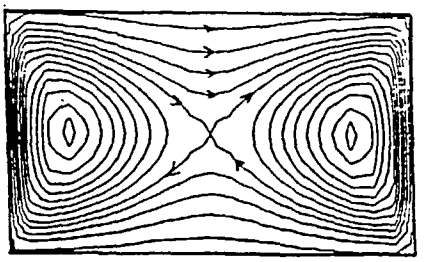

(a) streamline

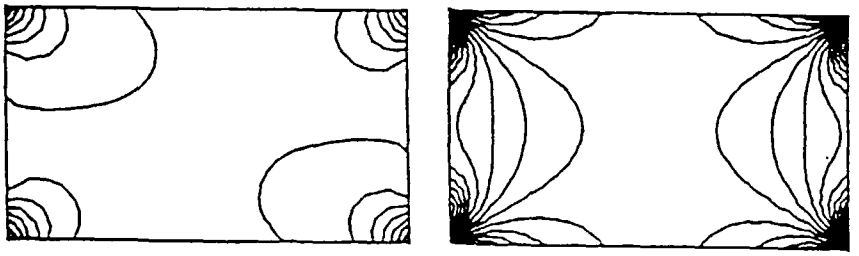

(b) pressure (c) vorticity

Figure 2 Quasi-steady flow solutions at $t = 12$ sec.

Figure 3(a). This defines the initial state and the system is integrated with fixed (conservatively small) timestep $\Delta t = 0.06$ sec. Following the evolution of the tracer from this initial state, we find that the narrow strip is first stretched and then folded by $t = 18$ sec. in *Figure 3(b)*. The main structure of the mixing pattern is evident at $t = 60$ sec. with some tracer confined to the elliptic interior zones and other tracer in the region exterior to the 'figure-eight' dividing streamline, see *Figure 3(c)*. Further stretching and folding occurs as the process continues, see *Figure 3(d)*, but the two 'islands' surrounding the elliptic points are essentially fixed in size from $t = 60$ sec. onwards. The spread of tracer is limited to the region defined by the streamlines containing the initial band. The simulated mixing pattern is similar to that obtained from laboratory experiment[18] as indicated in *Figure 3(e)*.

The calculation was repeated using larger time steps of $\Delta t = 0.3$ sec. and 3.3 sec. and the same pattern structure is maintained for significant time (3000 sec.) until accumulated integration error effects are evident. Hence, a time step of order $O(1$ sec.) is adequate to maintain accuracy in the range of times of interest in our study.

Mixing patterns in periodic flow

The mixing pattern is next simulated for a periodic flow in a 10 cm × 6 cm cavity. Mixing is induced by the alternate motions in opposite directions of the respective top and bottom walls

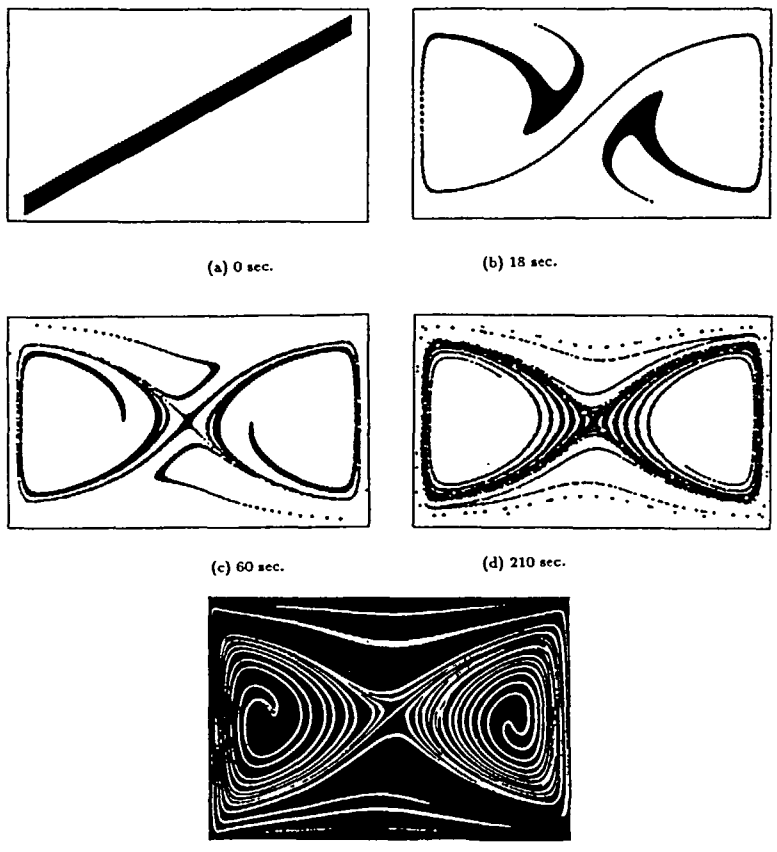

Figure 3 Mixing pattern in steady flow[18]

with velocity of magnitude $V=2$ cm/sec ($Re=1.2$) and period T, i.e. we have for top and bottom velocities, respectively:

$$u_T=\begin{cases} V, nT\leqslant t\leqslant(2n+1)T/2 \\ 0, \text{otherwise} \end{cases}$$

$$u_B=\begin{cases} 0, nT\leqslant t\leqslant(2n+1)T/2 \\ -V, \text{otherwise} \end{cases}$$

The viscous flow solution is computed on a 10×6 mesh of 9-node finite elements with timestep $\Delta t=0.1$ sec. The velocity history computed at representative interior points ($x=9$ cm, $y=1$ cm and $x=9$ cm, $y=2$ cm) during two periods is shown in *Figure 4* for a period of boundary motion $T=4$ sec. The small 'overshoot' at the point (9 cm, 1 cm) corresponds to the flow adjustment as one boundary velocity is reduced from V to zero and the other increased from zero to V. Point (9 cm, 2 cm) is further from the forcing boundary and no noticeable overshoot occurs.

The streamline, vorticity and pressure solutions at $t=\frac{3}{2}T$ and $2T$ are sketched in *Figure 5*. Comparing streamline patterns, there is a single elliptic point with position oscilating from the top to bottom location once during each period. This behaviour is similar to that of the two 'blinking' vortices described in Reference 1. We anticipate that the periodic forcing and adjustment of streamline patterns will enhance tracer mixing.

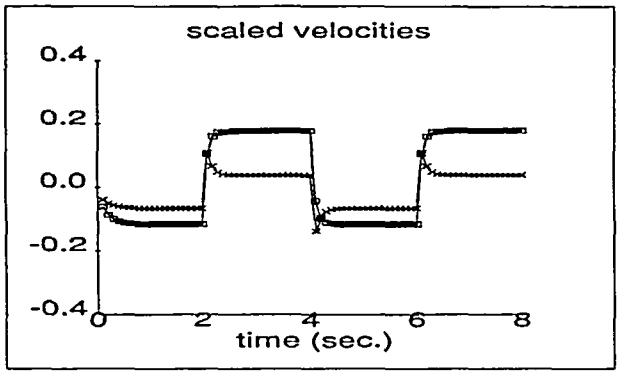

(a) velocities u/U

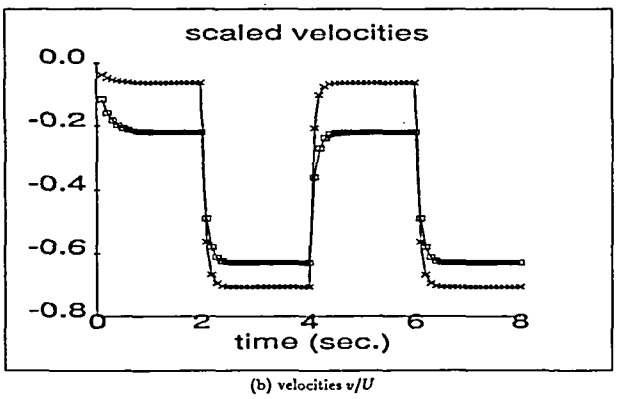

(b) velocities v/U

u and v at (9cm, 1cm) by x and (9cm, 2cm) by □

Figure 4 Velocity history for periodic flow

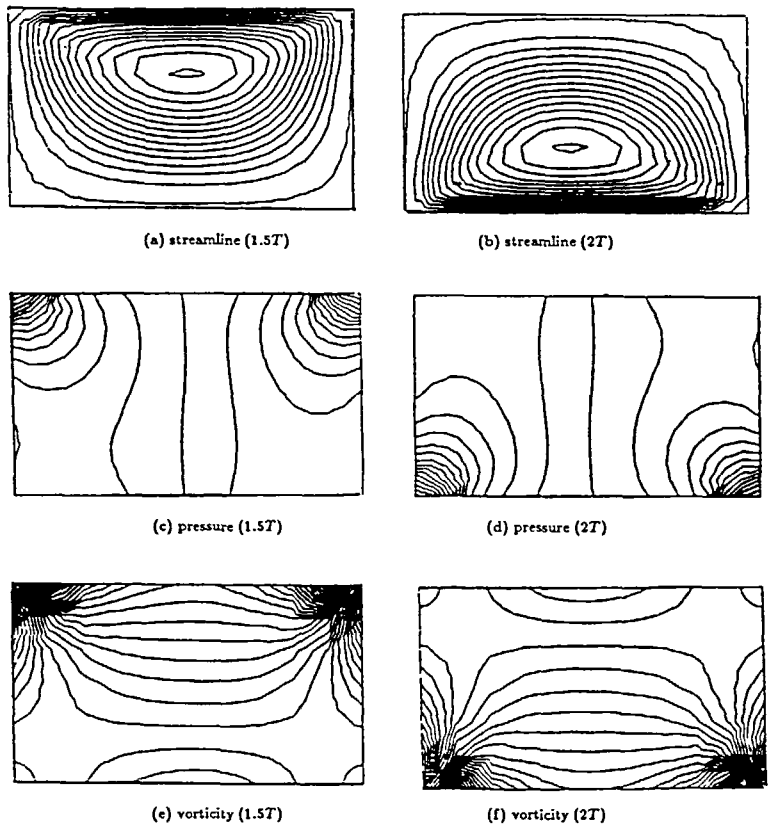

Figure 5 Periodic flow solutions

Experimental studies[19] have considered the mixing pattern with forcing period 48 sec. The corresponding mixing patterns after 10 periods are displayed in *Figure 6(e)* for a circular dye blob of diameter 0.3 cm initialy injected at position (2 cm, 3 cm). These results provide a basis for comparison with the present computational model. In our numerical simulation, a square tracer of size 0.3×0.3 cm^2 and containing 100×100 particles is initially centred at (2 cm, 3 cm) at the beginning of a periodic cycle as indicated in *Figure 6(a)*. The stretched and folded tracer after 5, 7, and 10 periods is shown in *Figures 6(b), (c), (d)*. The computed mixing pattern after 10 periods in *Figure 6(d)* is quite extensive and resembles closely the experimental result in *Figure 6(e)* for essentially the same situation. There are some minor differences since, for instance, our initial tracer is square and the experimental initial tracer is approximately circular.

In the above mixing pattern for periodic flow, an island without tracer is observed in the middle region. Laboratory experiments have also shown that a tracer located in this island mixes very little, while tracer outside the island is extensively mixed. There is no communication between them[18,19]. The experimental studies were duplicated in a set of simulations to explore this behaviour for an initial configuration consisting of three square 'blobs' of size 0.4×0.4 cm^2 and containing 60×60, 50×50 and 80×80 particles for the tracers centred at (1.6 cm, 3.2 cm), (3 cm, 3 cm) and (8.4 cm, 1.8 cm) respectively, see *Figure 7(a)*. The period was 35 sec. and the computed mixing pattern is shown in *Figures 7(b), (c)* and *(d)* after 3, 7 and 8 periods, respectively. The tracer on the left is initially located near the border of an island subregion, and part of the boundary of the island can be seen from that tracer in the final plot. The middle tracer inside

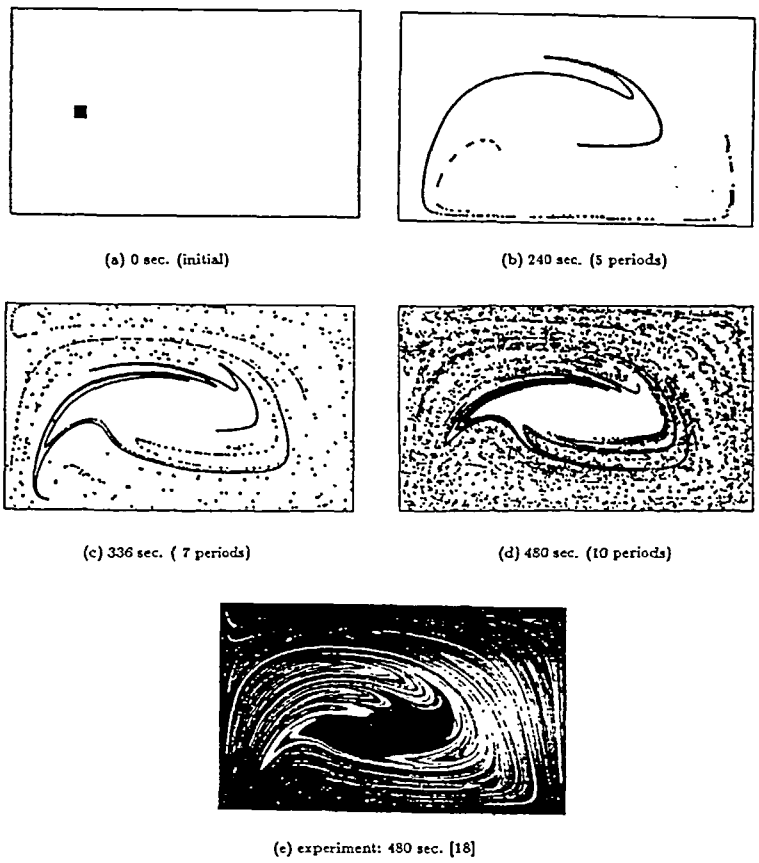

Figure 6 Mixing pattern in periodic flow

the island is much less mixed than the right tracer outside the island. The pattern in *Figure 7(d)* compares favourably with the experimental results in *Figure 7(e)* and *(f)*.

Mixing efficiency

The mixing patterns in steady and periodic flows show different degrees of mixing. If the patterns in *Figures 3* and *6* are compared with respect to time using the degree of mixing by the particle occupation estimate D_f^*, two kinds of variations are observed in *Figure 8*. Both mixing patterns are simulated by tracing 10 000 particles, and their degrees of mixing (D_f^*) are measured using 2^{mn} boxes ($m=2$, $n=6$). D_f^* is initially lower in periodic flow than in steady flow since a small square of initial tracer is compared with a long band initial tracer. However, D_f^* for periodic flow increases faster than that in steady flow because D_f^* exhibits an exponential growth in periodic flow. The low efficiency of mixing in steady flow results from the inability of the tracer to cross the streamlines. There are large regions that the tracer will not invade even if the initial tracer is distributed across the entire domain. Mixing is totally different for periodic flow, and the tracer extends over the domain. Therefore, periodic flow should be introduced to enhance mixing.

When periodic flow is considered further, we find that the degree of mixing depends upon the

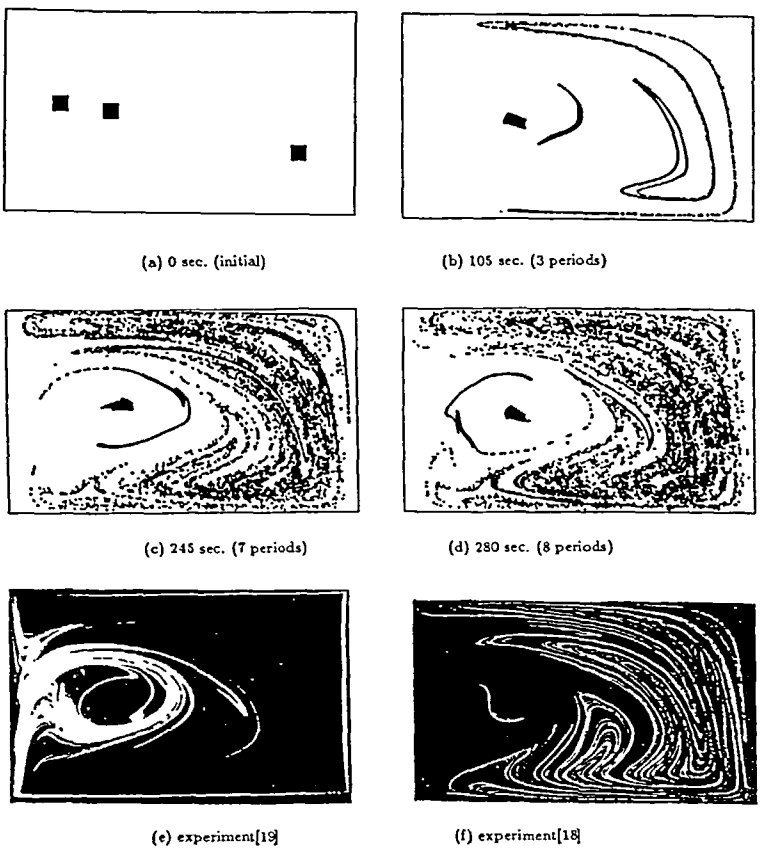

Figure 7 Island in mixing

forcing frequency and the initial position of the tracer. Our numerical scheme can be utilized to determine the optimal forcing frequency to achieve maximal mixing. To illustrate this, mixing was computed for forcing frequencies $Sr=0.0036$, 0.0090, 0.018, 0.0225, 0.0360, 0.0450, 0.0600 and 0.090, which correspond to periods $T=500$, 200, 100, 80, 60, 40, 30 and 20 sec. The initial tracer is centred at (8 cm, 3 cm). The mixing patterns are shown at $t=250$ sec. in *Figure 9* for each frequency. Relatively less mixing has been achieved at the extreme frequencies and the best mixing is observed at a period of 80 sec, where we see that almost uniform mixing has taken place by $t=250$ sec. The lower level of mixing in the extreme lower fequency range is because the tracer is stretched very slowly by an almost steady flow. The lower mixing in the extreme upper frequency range occurs because the tracer is hardly stretched by 'shaking' flow. Therefore, the optimal mixing would be expected under an optimal frequency near $Sr=0.225$, which corresponds to the period of 80 sec. in *Figure 9(d)*.

The mixing patterns in *Figure 9* show different size and shape islands. The island size decreases gradually as $Sr\rightarrow 0.0036\sim 0.0255$, and then increases as $Sr\rightarrow 0.0255\sim 0.090$. The second part of this correlation between the size of open regions and forcing frequency has also been noted in experiments. For example, when experiments were made at periods 20, 40, 60 and 600 sec. ($Sr=0.09$, 0.045, 0.036 and 0.003) and the best mixing is observed at period 60 sec.[11]. When comparing with our numerical simulation at periods 20, 30, 40, 60, 80, 100, 200, 500 sec., where mixing with period 80 sec. is better than that with period 60 sec.

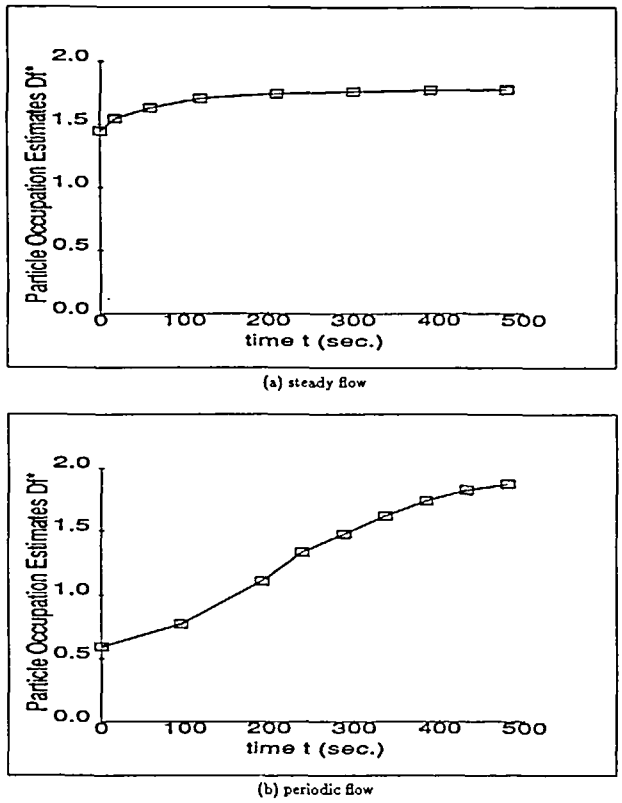

Figure 8 Degree of mixing vs. time

Another set of calculations with tracer located initially at centre of cavity (5 cm, 3 cm) in *Figure 10* with the same periods shows similar behaviour to that of *Figure 9* except at certain intermediate frequencies (e.g. at periods 40, 100 sec.) where the mixing is very poor. At periods 40 and 100 sec., the mixing patterns reveal that the tracers are in the islands, see *Figure 10(c)* and *(f)*. When the tracer is outside the island, the mixing for the centre tracer is only slightly worse than that for the right tracer (8 cm, 5 cm), see *Figure 9* and *10(a), (b), (d), (e), (g)* and *(h)*.

Next, the degree of mixing is measured at $t = 250$ sec. using the same box-counting as before ($2^{2\times 6}$ boxes), and D_f^* is graphed against frequency in *Figure 11* for tracers initially at (8 cm, 3 cm) indicated by the solid line at (5 cm, 3 cm) by the dashed line. Over the frequency range shown, the degree of mixing is found to depend on the forcing frequency and initial tracer location. When the tracer is initially located on the right, D_f^* attains its maximum of 1.9402 at period 80 sec. and has minima of 1.1775 and 1.6777 at periods 500 sec. and 20 sec., respectively. When the tracer is initially located at the centre, the D_f^* profile has local maxima near periods 30, 80 and 200 sec., with a global maximum of 1.8797 again occurring near period 80 sec. and local minima occurring near periods 40 and 100 sec. The D_f^* curves represent the degree of mixing reasonably for the mixing patterns in *Figures 9* and *10*. That is, for the centre tracer, the tracer is in an island near these local minimal frequencies. When the initial tracers are outside the island, the initial right tracer results in a slightly better mixing than the initial centre tracer. Hence, the forcing frequency is a key factor and initial position is the secondary factor in

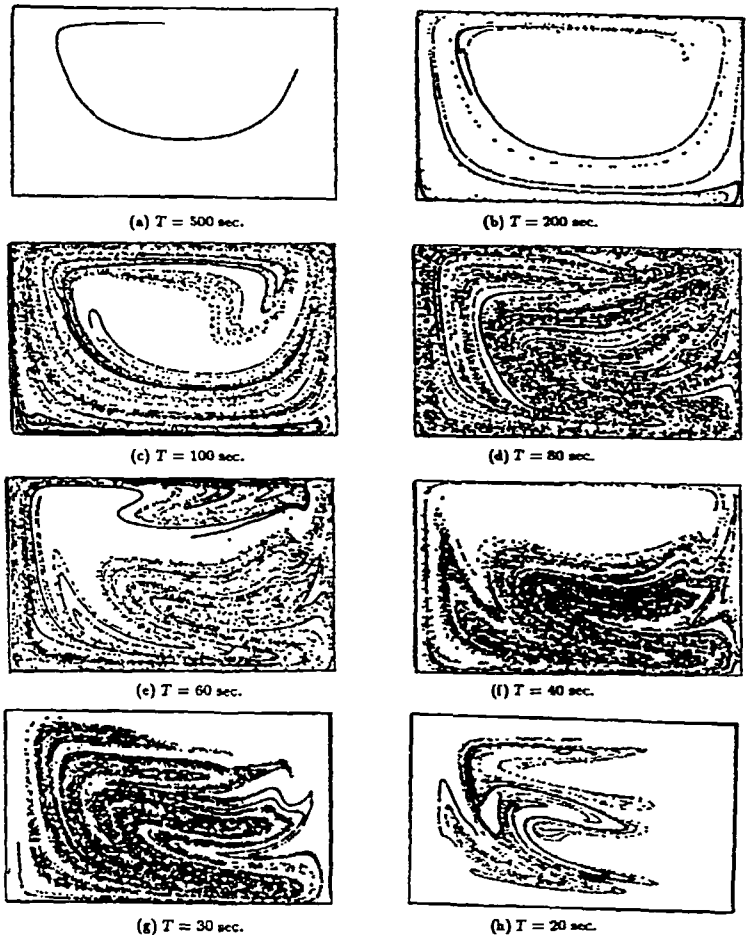

Figure 9 Mixing under different frequencies

obtaining an efficient mixing as long as the initial tracer is carefully placed outside the island. Our numerical results provide a more detailed identification of this behaviour than previous experimental work.

We also measure the variance of the particle position for several frequencies. Let a set of 50×30 tracer particles be distributed uniformly in the cavity at time $t=0$. In *Figure 12*, the variances are plotted on a log-log scale against time for periodic forcing with periods $T=20$, 35, 48, 80, 500 sec. For asymptotically increasing t, we have linear behaviour in the log-log plot.

VECTOR AND PARALLEL ASPECTS

In the present work we have focused on 2D mixing but the treatment of 3D mixing follows directly and is obviously more computationally intensive. Particle tracing is a straightforward procedure and can be carried out conveniently as a post-processing operation on a graphics workstation. There are situations, however, where rapid tracing is desirable during the computation rather than as a post-process. The desire for real-time simulation and real time

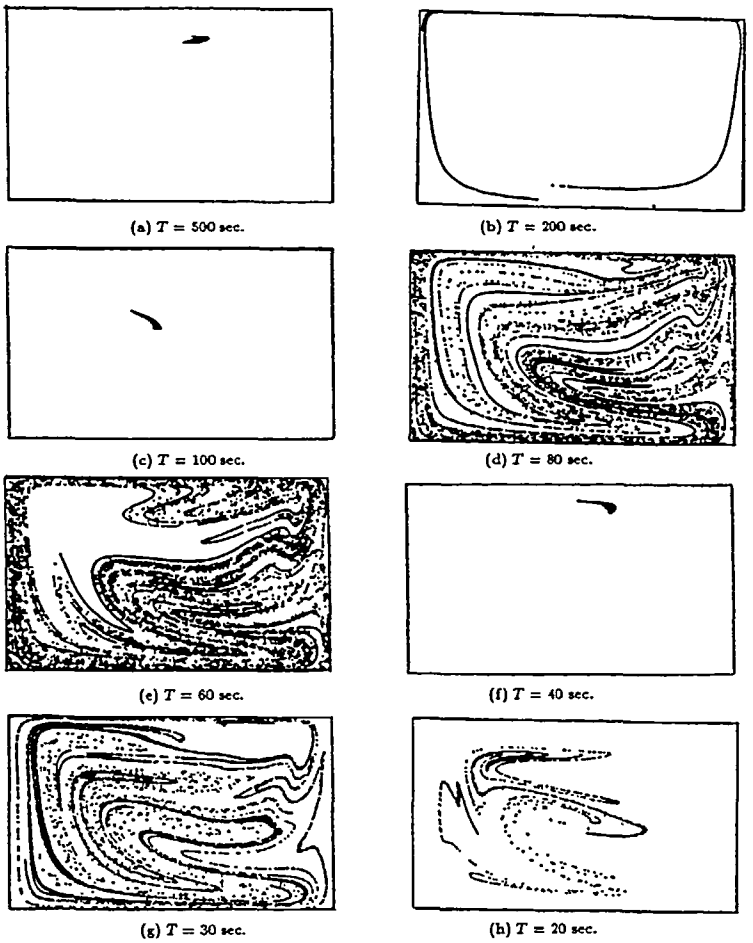

Figure 10 Mixing under different frequencies

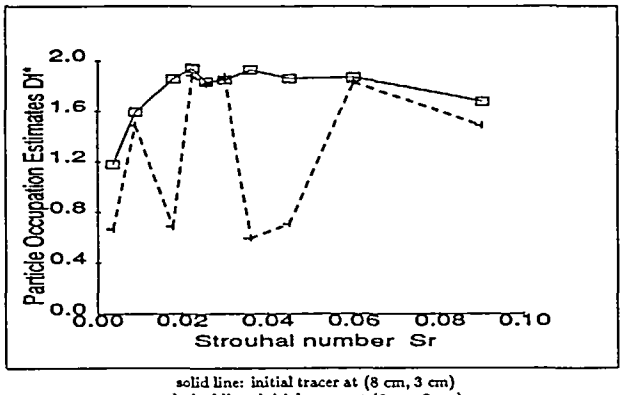

Figure 11 Degree of mixing vs. frequency

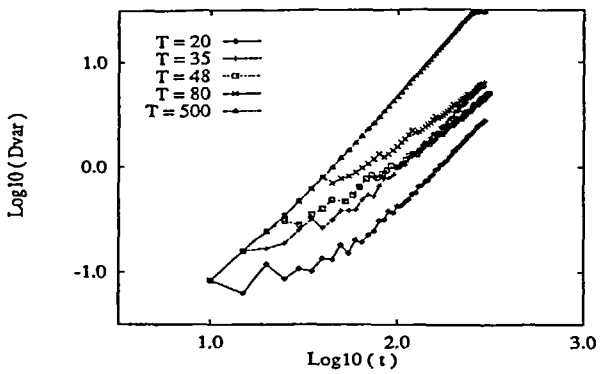

Figure 12 Variances vs. time for different frequencies

tracing is a further incentive. Finally, current generation computer hardware may offer vector and parallel capabilities that can be exploited to achieve these goals and yield very efficient tracing schemes[22]. In this section we briefly examine some aspects of both the vector and parallel appoaches.

Let us first consider vector implementation. The Heun predictor-corrector scheme in (10)-(11) has the general form,

$$x_{new} = x_{old} + \alpha u_a + \beta u_b \tag{18}$$

where x_{old} is the known position vector, α and β are given scalars and u_a, u_b are particle velocity vectors interpolated from the nodal solutions. Hence calculation of the new tracer position vector x_{new} consists of two main steps: interpolation of velocities and then the evaluation of (18). The calculation (18) is seen to involve multiplication of vectors by scalars (αu_a and βu_b) followed by vector addition and can be accomplished efficiently by two SAXPY operations. e.g. On the CRAY X-MP/14se employed for the present work a SAXPY runs at 166 Mflops[12].

In the present tracing studies we have considered a Cartesian grid. Consequently, it is straightforward and efficient to determine which element contains a given particle. The situation for unstructured grids is more complicated and necessitates both an expanded data structure including neighbour information and a trajectory test for side crossings[5]. Having determined the element containing a particle at (x_p, y_p), the particle velocity can be interpolated from the nodal velocity. Recall that the interpolation process is local with,

$$u_h(x_p, y_p) = \sum_{j=1}^{N_e} u_j^e \psi_j^e(x_p, y_p), \qquad p = 1,2,\ldots,P \tag{19}$$

where $\psi_j^e(x,y)$ are the element basis functions and P is the total number of particles. Therefore this calculation also vectorizes over the particles.

Vectorization

We now provide megaflop (Mflop) rates on the CRAY-XMP for the update in (18) and the velocity interpolation in (19). The calculations correspond to tracing two initial 0.4 cm square blobs, each with 70 × 70 particles distributed uniformly and centred respectively at (1.8 cm, 3.5 cm) and (8.2 cm, 3.5 cm) in a 10 × 10 biquadratic element mesh. The rate of 96 Mflops for the position update in (18) is reasonable but the poor performance of 33 Mflops reveals a difficulty for the velocity interpolation in (19). More detailed examination of this part of the calculation indicates that the source of the performance degradation is a memory bank conflict in extracting u_j^e for

(19). Since particles are numbered in the natural ordering sequence mainly contiguously-numbered particles initially lie in the same elements. Since each particle references memory for u_j^e, there are multiple close sequenced references to the same memory address. This results in a severe memory bank conflict. As mixing progresses and the tracer is dispersed over many elements, the bank conflicts subside and performance improves.

The initial poor performance can be improved by numbering the particles in a non-consecutive fashion so that common memory requests are staged further apart. Let the particles be renumbered in the sequence 1, P, 2, $P-1$, ..., $P/2$, $P/2+1$ so that the velocities are now picked alternately from the two elements. This renumbering reduces the bank conflict by 23.6% and the rate of interpolating velocities increases to 40.34 Mflops. If all 9800 particles are distributed element-by-element on a mesh with E element so that each element has approximately P/E particles and the particles in element e are numbered e, $e+p$, $e+2P$, etc. Then the bank conflict is reduced by 99.8%, and the rate for velocity interpolation increases to 78.9 Mflops. Hence we see that appropriately ordering the initial particles can substantially reduce bank conflicts and improve the performance.

Parallelization and visualization

Tracer motion and visualization can also be carried out effectively on parallel computers. In the present work we examine tracer mixing on the CM2 a parallel SIMD system. The CM2 also permits parallel I/O to a framebuffer which is a high-speed graphics display system for visualization. In the parallel algorithm, the position vector for the particles is distributed by components over the processors so that (18) is implemented component-wise in parallel. That is, the particles are distributed over the processors. Similarly, the velocity field is distributed by a **TABLE-SET** over the nodes in parallel. The velocities are retrieved from memory by an indirect addressing hardware system **TABLE-LOOK**. The patterns of particles at the desired times are displayed on the attached CM2 framebuffer. The algorithm was implemented primarily in CM FORTRAN with **TABLE-SET** and **TABLE-LOOK** in Paris (assembler) to obtain efficiency in indirect addressing. Without the indirect address hardware, the kernel for interpolation of particle velocities (19) is as low as 5 Mflops on the CM2-8K for a set of 8192 tracer particles in a steady flow field. The indirect address hardware permits the kernel (19) to run at 111 Mflops.

The VP ratio is defined as the number of virtual processing elements emulated by each physical processing element. For our tracing problem on the CM2-16K, the VP ratio is $P/2^{14}$. The calculation rates are shown in *Figure 13* for both steady flow and for periodic flow. The higher the VP ratio, the higher the rate, but the increase in rate slows down with further increase of the VP ratio. The top rate is 295 Mflops for steady flow and 224 Mflops for periodic flow when $P=2^{18}$. In the above test problems, the transition steps, where **TABLE-SET** is needed, comprise 80% of all marching steps. In most cases for periodic flow, the percentage of transition steps is significantly less than this percentage, e.g. from 0.8% to 20% for flow with period 500 sec. to 20 sec respectively. Hence the actual rate will be between the steady flow and periodic flow cases in the above test problem. When compared with the peak rate of 800 Mflops on the CM2-16K for a SAXPY operation[16], our rates in the application code are about 28% (224 Mflops) to 37% (295 Mflops) of the peak rate.

In the visualization stage, the initial pattern of particles is input from a data vault and displayed on the frame buffer. Subsequent patterns are traced on the CM2 according to (18), (19) and then displayed at specified time intervals on the framebuffer with a 325×195 pixel map. The pattern can also be written to a data file for videotape. Timing and related data for sample calculations are given in *Table I*. Case 1 is the mixing pattern study for steady flow (see *Figure 3*), case 2 is the mixing pattern study for periodic flow (see *Figure 6*).

The performance for visualization is summarized in *Table I* as follows: One frame of the picture is displayed after a few time steps of integration (see 'Others-March/frame') so that the

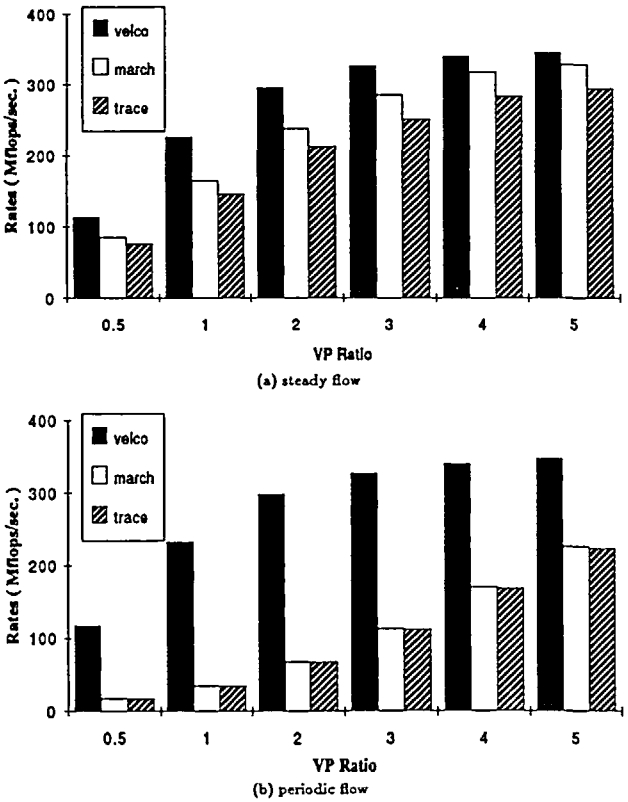

Figure 13 Performance on CM2-16K for velocity computation (velco), integration (march) and overall tracing (trace)

Table I Timing for visualizing mixing processes on CM2-16K

		Case 1	Case 2
Time	CM2	360	756
(sec.)	(time/frame)	(0.72)	(0.36)
	SUN	2074	3780
	(time/frame)	(5.76)	(4.99)
	Tape	16	31
	(time/frame)	(0.032)	(0.032)
	Experiment	120	480
Others	Particles	2^{14}	2^{14}
	Time steps	2000	4800
	Frames	501	961
	Marches/frame	4	5
	Data (Kbytes)	16040	30776

time for displaying 1 frame is about 1 sec. (see 'Time-CM2-(time/frame)'). The sequential time to write a data file for a videotape (see 'Time-SUN') is about 5 sec. per frame (see 'Time-SUN-(time/frame)'). The continuous stretching and folding of the tracer is represented on a corresponding videotape with about 0.05 sec. per frame (see 'Time-Tape-(time/frame)'). Finally the actual time taken for the corresponding laboratory experiments is given as 'Time-Experiment'.

CONCLUDING REMARKS

A least-squares finite element solution for viscous incompressible flow is employed in conjunction with a particle tracing algorithm to determine tracer patterns for mixing studies. The Heun particle tracing scheme is implemented on the CRAY-XMP shared-memory and the CM2 distributed memory system and mixing patterns are compared with experiment. Several measures of mixing are introduced and discussed for both the continuum and discrete cases. The question of 'tuning' the forcing frequency to enhance mixing in an experiment is examined. The attributes and limitations of these measures are examined and counter-examples constructed. We emphasize that there is no simple quantitative measure (such area) and that a more complete characterization can be made using several measures. Of course, one measure may be adequate for a specific purpose. Results of these numerical studies are very similar to published experimental work, but do differ in some minor respects. For instance, in the steady flow problem the simulation results show that unoccupied regions interior to closed streamlines are not invaded which is consistent with the mathematical model and assumptions. In the experiments these regions are slowly invaded, probably due to diffusion of dye and some small unsteady effects in the experiment. Furthermore, the numerical experiments permit a more complete investigation of forcing frequencies and hence are better able to approximate the optimal mixing frequency.

ACKNOWLEDGEMENTS

This research has been supported by grants from the Department of Energy and ARPA. The CM calculation were made at Schlumberger Laboratory for Computer Science. We would like to express our appreciation to Alfred Lorber and Wayne Joubert for their assistance with the vector-parallel implementation.

REFERENCES

1 Aref, H. Stirring by chaotic advection, *J. Fluid Mech.*, **143**, 1–21 (1984)
2 Aris, R., Aronson, D. G. and Swinney, H. L. *Patterns and Dynamics in Reactive Media*, New York, Springer-Verlag (1991)
3 Babuska, I. and Aziz, A. K. Survey lectures on the mathematical foundations of the finite element method, A. K. Aziz, edited *The Mathematical Foundations of the Finite Element Method with Applications to Partial Differential Equations*, 3–359, Academic Press, New York (1972)
4 Barnsley, M. *Fractals Everywhere*, Academic Press Inc., San Diego (1988)
5 Carey, G. F. *Computational Grids*, Wiley, U.K. (in preparation) (1994)
6 Carey, G. F. and Jiang, B. N. Least-squares finite elements for first-order hyperbolic systems, *Int. J. for Num. Meth. in Eng.*, **26**, 81–93 (1988)
7 Carey, G. F. and Oden, J. T. *Finite Elements: Fluid Mechanics*, Prentice Hall, New Jersey (1986)
8 Carey, G. F., Pehlivanov, A. and Shen, Y. Least-squares mixed finite elements, *Proc. of the Finite Element Method: Fifty Years of the Courant Element*, Jyväskylä, Finland (1993)
9 Carey, G. F., Pehlivanov, A., Shen, Y. and Base, A. Some questions concerning the least-squares mixed Navier-Stokes problem, *TICOM Report, 1994* (in press)
10 Carey, G. F., Pehlivanov, A. and Vassilevski, P. S. Least-squares mixed finite element methods for non-selfadjoint elliptic problems: II. Performance of block-ILU factorization methods, to appear in SIAM J. Scientific Computing
11 Chien, W. L., Rising, H. and Ottino, J. M. Laminar mixing and chaotic mixing in several cavity flows, *J. Fluid Mech.*, **170**, 355–377 (1986)
12 Dongarra, J. J., Duff, I. S., Sorensen, D. C. and van der Vorst, H. A. Solving linear systems on vector and shared memory computers, *SIAM*, Philadelphia, 63, 213 (1991)
13 Field, M. and Golubitsky, M. *Symmetry in Chaos: A Search for Patterns in Mathematics, Art and Nature*, Oxford, New York (1992)
14 Fix, G. J. and Gunsburger, M. D. On least squares approximations to indefinite problems of mixed type, *Int. J. Num. Meth. Eng.*, **12**, 453–489 (1978)
15 Hemion, G. *The Classification of Knots and 3-Dimensional Spaces*, Oxford, New York (1992)
16 Highnam, P. The connection machine tutorial, Schlumberger Software Conference (1990)
17 Kauffman, L. H. *On Knots*, Princeton, N.J. (1987)
18 Ottino, J. M. The mixing of fluids, *Scientific American* (1989)

19 Ottino, J. M. *The Kinematics of Mixing*, Cambridge University Press (1989)
20 Pearson, C. E. *Num. Meth. in Eng. and Science*, Van Nostrand Blaisdell, Reinhold, New York (1986)
21 Pehlivanov, A. I., Carey, G. F. and Vassilevski, P. S. Least-squares mixed finite element methods for non-selfadjoint elliptic problems: I. Error estimates, in press
22 Rodrique, G. and Shah, S. *Parallel Supercomputing: Methods, Algorithms and Applications*, edited by G. F. Carey, Wiley Series in Parallel Computing (1989)
23 Shen, Y. *Least-Squares Finite Element Solution and Mixing Simulation*, Ph.D. Dissertation, University of Texas at Austin (1993)
24 Shlesinger, M. F., Zaslavsky, G. M. and Klafter, J. Strange kinetics, *Nature*, **3363**, 31–37 (1993)
25 Skinner, G. S. and Swinney, H. L. Periodic to quasiperiodic transition of chemical spiral rotation, *Physica D*, **48**, 1, 1–16 (1991)
26 Solomon, T. H. and Swinney, H. L. Observations of anomalous diffusion and Levy flights in two-dimensional rotating flows, Report, Center for Nonlinear Dynamics, The University of Texas at Austin (1993)
27 Young, D. M. and Gregory, R. T. *A Survey of Numerical Mathematics*, **I, II**, Addison Wesley Inc. (1972)