Lecture 1: Stirring and mixing

Jean-Luc Thiffeault

Notes by Sam Pegler and Amanda O'Rourke

21 June 2010

1 Introduction

Consider a patch of dye immersed in a flowing fluid, confined within a spatial domain Ω. The concentration of the dye $\theta(\boldsymbol{x}, t)$ can be modelled by the advection-diffusion equation

$$\frac{\partial \theta}{\partial t} + \boldsymbol{u} \cdot \boldsymbol{\nabla} \theta = \kappa \nabla^2 \theta, \tag{1}$$

where κ is a coefficient of diffusivity and $\boldsymbol{u}(\boldsymbol{x}, t)$ is a prescribed velocity field. By using equation (1), it is assumed that the concentration is advected precisely with the flow, and is also subject to molecular diffusion by which the concentration is allowed to gradually spread to adjacent fluid elements. The velocity field in equation (1) is assumed to have been determined a priori and satisfies the incompressibility condition

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0. \tag{2}$$

We also assume that the boundary conditions

$$\boldsymbol{n} \cdot \boldsymbol{u} = 0, \tag{3}$$

$$\boldsymbol{n} \cdot \boldsymbol{\nabla} \theta = 0, \tag{4}$$

apply on the boundary of the domain $\partial\Omega$, where $\boldsymbol{n}$ is the unit outward normal to the boundary. These conditions ensure that no fluid penetrates the boundary of the domain and that no concentration is allowed to diffuse into or out of the domain, respectively. (We can also use spatially-periodic boundary conditions, which ensure that there is no net flux in or out of the domain.)

We proceed to demonstrate some general properties of the system defined by equations (1)–(4) by deriving evolution equations for the spatial mean and variance of θ. To do this we begin by multiplying equation (1) by $m\theta^{m-1}$ and reordering the derivatives to give

$$\frac{\partial}{\partial t}(\theta^m) + \boldsymbol{\nabla} \cdot (\theta^m \boldsymbol{u}) = m\kappa \left[\boldsymbol{\nabla} \cdot (\theta^{m-1} \boldsymbol{\nabla} \theta) - (m-1)\theta^{m-2} |\boldsymbol{\nabla} \theta|^2 \right]. \tag{5}$$

Now we take the volume integral over the domain, defined by

$$\langle \phi \rangle \equiv \int_\Omega \phi \, \mathrm{d}V, \tag{6}$$

where $\phi(\boldsymbol{x}, t)$ is any function, which gives

$$\frac{d}{dt}\langle\theta^m\rangle = -m(m-1)\kappa\langle\theta^{m-2}|\boldsymbol{\nabla}\theta|^2\rangle, \tag{7}$$

on applying the divergence theorem to the second term on the left-hand side of equation (5) and the first term on the right-hand side, and applying boundary conditions (3)–(4) to make the resulting boundary terms vanish. Setting $m = 1, 2$ gives the two equations

$$\frac{d}{dt}\langle\theta\rangle = 0, \tag{8}$$

$$\frac{d}{dt}\langle\theta^2\rangle = -2\kappa\langle|\boldsymbol{\nabla}\theta|^2\rangle, \tag{9}$$

respectively. The first equation above represents conservation of total concentration, which is expected given that no concentration enters the domain as ensured by boundary conditions (3)–(4). The second equation is an evolution equation for the second moment of θ in terms of its gradient. This equation is not closed, i.e. it does not independently determine the time evolution of $\langle\theta^2\rangle$, because its right-hand side depends on $|\boldsymbol{\nabla}\theta|$, whose time evolution must be determined from a solution of the advection-diffusion equation (1). However, noting that the right-hand side of equation (9) is negative-definite if θ is not constant in space, we can deduce that the second moment of the concentration must decrease in time. Thus, at large times the second moment converges toward a constant value in which the right-hand side of (9) vanishes. This occurs if and only if the gradient of the concentration is zero everywhere. The system is driven toward a fully homogenized state in which the second moment takes the value

$$\langle\theta^2\rangle = \frac{\langle\theta\rangle^2}{\Omega}. \tag{10}$$

Note that although $\langle\theta^2\rangle$ never quite reaches this value, we say the system is *mixed* once it has fallen below a certain prescribed threshold.

As mentioned previously, equation (9) does not independently determine the rate at which $\langle\theta^2\rangle$ decreases in time. In fact, the rate at which $\langle\theta^2\rangle$ decreases usually depends strongly on the prescribed velocity field, which does not appear in equation (9). Importantly, the introduction of a velocity field, which we refer to as *stirring*, usually has the effect of increasing gradients in the concentration field $|\boldsymbol{\nabla}\theta|$ through advection and thus has the potential to significantly accelerate the onset of mixing.

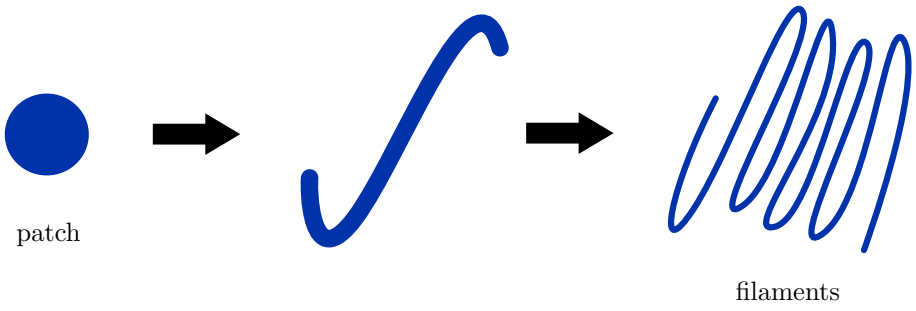

Figure 1: Typical mixing scenario in which a patch of dye is advected into filaments.

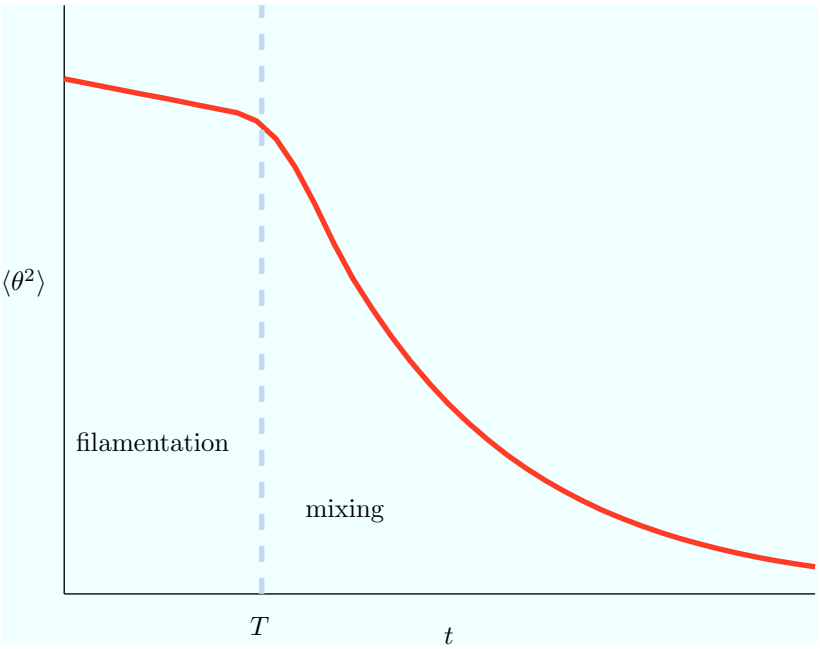

Figure 2: Decay of the second moment of the concentration field under the action of stirring.

A typical stirring scenario is illustrated in figure 1, in which an initially isolated patch of dye is advected into thin striations that exhibit a much greater concentration gradient than the initial patch. The evolution of the second moment, corresponding to this typical mixing scenario, is shown in figure 2. During the initial stirring phase, before thin filaments are formed, the concentration varies over relatively large length scales and the rate of mixing will be slow. After a short transient, by which point the filaments have become sufficiently thin, the rate of diffusion is large enough to balance the rate of thinning of the filaments and the rate of mixing $\partial\langle\theta^2\rangle/\partial t$ is no longer characterized by κ, but is instead characterized by the rate at which the velocity field is straining the fluid into filaments. Once this occurs the characteristic width of the filaments remains constant, and the rate of mixing reaches a maximum which persists for all time. The independence of the rate of diffusion on the parameter κ during the mixing phase implies the scaling

$$|\nabla\theta| \sim \kappa^{-1/2}, \tag{11}$$

which can be used as an indication of good mixing in practical applications.

The coefficient of diffusion κ is small in most applications. Hence, the characteristic timescale over which diffusion takes place in the absence of an advective flow ($\boldsymbol{u} = \boldsymbol{0}$) is typically very long. For example, in the context of heat diffusing through air, $\kappa \approx 2.2 \times 10^{-5}\mathrm{m}^2/\mathrm{s}$, so in a room with sides of length $L = 10\mathrm{m}$, a characteristic diffusion timescale is $L^2/\kappa \sim 53$ days. Thus, if one is aiming to mix the concentration field, relying on diffusion in the absence of advection is typically ineffective. In the following section we present an example that can be solved analytically to demonstrate the development of the mixing phase discussed above.

The filament solution

Consider the prescribed velocity field

$$\boldsymbol{u}(\boldsymbol{x}) = (\lambda x, -\lambda y), \tag{12}$$

where λ is a constant parameter, which drives fluid towards the x-axis in a straining motion with a hyperbolic point at the origin. In this case the advection-diffusion equation takes the form

$$\frac{\partial \theta}{\partial t} + \lambda x \frac{\partial \theta}{\partial x} - \lambda y \frac{\partial \theta}{\partial y} = \kappa \nabla^2 \theta. \tag{13}$$

This equation can be solved exactly, but for simplicity here we consider only x-independent solutions of the form

$$\theta(y,t) = e^{-\lambda t} f(y). \tag{14}$$

Substituting this into equation (13) gives the ordinary differential equation

$$-\lambda f - \lambda y f' = \kappa f''. \tag{15}$$

If we further assume that the concentration vanishes in the far field, $f \to 0$ as $y \to \pm\infty$, we can solve equation (15) and derive the solution

$$\theta = e^{-\lambda t} e^{-y^2/2l^2}, \tag{16}$$

where

$$l = \sqrt{\kappa/\lambda} \tag{17}$$

is called the Batchelor length.

The solution (16) represents a filament that is an attractor for typical initial conditions of the system. For any reasonable initial concentration field, the compression due to the straining velocity will attempt to confine the concentration into an ever thinner region about the x-axis until its gradient becomes sufficiently large for diffusion to balance the rate of compression. The width of the resulting filament is characterized by the Batchelor length (17), which contains the necessary scaling for (11) to apply. Furthermore, we see from the solution (16) that the rate of mixing λ during the mixing phase is indeed independent of κ. Note that equation (8) must be taken with a grain of salt, because the concentration does not vanish as $|x| \to \infty$. This is a consequence of the neglect of the x dependence.

2 Effective Diffusivity

Recall that filaments undergoing advection are strained along one direction at a rate λ and are compressed along the other to the Batchelor length $l = \sqrt{\kappa/\lambda}$, where κ is the molecular diffusivity and λ is the local rate of strain. In the previous section it was shown that the variance evolves according to

$$\langle \theta^2 \rangle \sim e^{-\lambda t},$$

where λ is a local rate of strain.

From simple arguments, molecular diffusivity alone is an exceedingly weak method for mixing scalar fields. We may then ask how stirring might improve tracer mixing and, more importantly, how we can quantify the effects of stirring by invoking an *effective* diffusivity.

2.1 A random walk model

The advection-diffusion equation for a passive scalar $\theta = \theta(\boldsymbol{x}, t)$ can be written as

$$\frac{\partial \theta}{\partial t} + \boldsymbol{u} \cdot \nabla \theta = \kappa \nabla^2 \theta, \tag{18}$$

where $\boldsymbol{u}$ is a given velocity field and κ is a constant of diffusivity. Even in the absence of stirring and the filamentation process, this diffusivity can act to mix the tracer field θ by monotonously decreasing the variance from equation (9).

The long timescales associated with diffusion alone indicates that the physical act of stirring, or the advective term $\boldsymbol{u} \cdot \nabla \theta$ of (18), is necessary to efficiently mix a tracer field. However, the advective term of (18) is more complex than the diffusivity. Ideally, one could write the effects of both the mixing of the tracer field by molecular diffusivity and by advective stirring as an effective diffusivity, such that the equations of motion can be approximated by

$$\frac{\partial \theta}{\partial t} = \kappa \nabla^2 \theta - \boldsymbol{u} \cdot \nabla \theta \approx \kappa_e \nabla^2 \theta. \tag{19}$$

We consider the displacement, x, of an *individual* particle stirred by a random flow field. We would like to compare the displacemnt of the particle to a true random walk,

$$x_n = x_{n-1} + \xi_n \tag{20}$$

where x_n is position of the particle at time n and where ξ_n is a random kick. Here we assume that all ξ_n are independent and identically distributed random variables.

At the nth step, the particle position x_n satisfies

$$x_n = x_0 + \sum_{k=1}^{n} \xi_k$$

where x_0 is the initial position. If the initial position of the particle is $x_0 = 0$ and the mean individual kick $\langle \xi_n \rangle = 0$, then the mean position of the particle at time n is

$$\langle x_n \rangle = \sum_{k=1}^{n} \langle \xi_k \rangle = 0$$

and the action of the random variations in the flow field given by ξ do not change the mean position of the particle.

Consider instead the mean-squared displacement of the particle:

$$\langle x_n^2 \rangle = \sum_{k=1}^{n} = n \langle \xi^2 \rangle = n\sigma^2, \tag{21}$$

where $\sigma^2 = \langle \xi^2 \rangle$. The mean-squared displacement of the particle then grows linearly with the number of steps, n, at a rate determined by the standard deviation of the random forcing field.

For a process with an average length of time between kicks, T, the number of kicks can be converted into a time by letting $t = nT$. The squared displacement in (21) then grows linearly in time as

$$\langle x^2(t) \rangle = 2\kappa t \tag{22}$$

where $\kappa = \sigma^2/2T$ is a diffusivity, or a measure of how quickly the squared displacement of the particle grows in time.

The above calculations are only in one dimension. In d dimensions, the mean-squared distance from the origin can be written as a sum of the mean-squared distance $\langle x_{k,n}^2 \rangle$ along each dimension k at timestep n,

$$\langle \sum_{k=1}^{d} x_{k,n}^2 \rangle = \langle x_{1,n}^2 \rangle + \langle x_{2,n}^2 \rangle + ... + \langle x_{d,n}^2 \rangle$$

which, assuming the random forcing is isotropic such that $\langle x_{k,n}^2 \rangle = \langle x_n^2 \rangle$, can also be written as

$$d\langle x_n^2 \rangle = dn\sigma^2 = 2d\kappa t \tag{23}$$

where $\kappa = \sigma^2/2T$ is the same as in the one dimensional case.

The diffusivity, κ, can be thought of as an effective diffusivity as it arises from the physical stirring generated by the stochastic process ξ_n. The effective diffusivity is again given by

$$\boxed{\kappa_{\text{eff}} = \frac{\sigma^2}{2T}} \tag{24}$$

where σ^2 is the standard deviation of the particle displacement due to the random kicks ξ_n and T is the average interval between kicks.

2.2 Particle concentration and diffusivity

The random walk description of a diffusive process above applies only to a single particle with position given by $\boldsymbol{x}_n$. We wish to consider the evolution of a cloud of particles. Let $\theta(\boldsymbol{x}, t)$ be the density of particles within a given area (or volume in 3D); then θ satisfies the diffusive equation

$$\frac{\partial \theta}{\partial t} = \kappa_e \nabla^2 \theta \tag{25}$$

if each point evolves independently according to a random walk $\boldsymbol{x}_{n+1} = \boldsymbol{x}_n + \xi_n$.

The idea of a cloud of particles only works if the spatial and temporal resolution of the model are large enough such that the motion of individual particles are decorrelated. If we are zoomed in too closely, such that the grid scale is less than the standard deviation of the displacements ($L < \sigma$), then we see only individual particles rather than a cloud of particles with a given concentration, θ. If we choose sampling timescales that are too short, such that the timescale of the flow is less than the interval between kicks, T, then the kicks become correlated in time.

2.3 Numerical examples for sine flow

Consider, for example, the advection of a cloud of particles by the famous sine flow in a doubly-periodic domain [2]. The velocity field of the sine flow alternates between a horizontally-aligned shear flow and vertically-aligned shear flow given by

$$\begin{aligned} \boldsymbol{u}_H &= (U \sin(2\pi k y/L), 0) \quad \text{for } 0 \le t < \tfrac{1}{2}\tau; \\ \boldsymbol{u}_V &= (0, U \sin(2\pi k x/L)) \quad \text{for } \tfrac{1}{2}\tau \le t < \tau, \end{aligned}$$

where here U is a constant velocity, k is the wavenumber of the shear flow field, L is the size of the biperiodic domain, and τ is the period of the flow.

We then obtain a map of the position of the particle by solving

$$\dot{\boldsymbol{x}} = \boldsymbol{u}, \qquad \boldsymbol{x} = \boldsymbol{x}_0$$

exactly as a two-step process:

$$\begin{aligned} \text{Step 1}: \quad & x(\tau/2) = x_0 + U\tfrac{\tau}{2} \sin\left(\tfrac{2\pi k y_0}{L}\right) \\ & y(\tau/2) = y_0 \\ \text{Step 2}: \quad & x(\tau) = x(\tau/2) \\ & y(\tau) = y(\tau/2) + U\tfrac{\tau}{2} \sin\left(\tfrac{2\pi k x(\tau/2)}{L}\right). \end{aligned}$$

At the end of an interval of length τ, the new position (x', y') of a particle that started at (x, y) is then

$$\begin{aligned} x' &= x + T \sin\left(\frac{2\pi k y}{L}\right) \\ y' &= y + T \sin\left(\frac{2\pi k x'}{L}\right) \end{aligned} \tag{26}$$

where $T = U\tau/2$. For a constant velocity, U, T is proportional to τ and the intensity of stirring can be manipulated by varying T. Note that in the second step of (26), y', is a function of x' and not x. This is necessary to conserve area, as can be verified by computing the Jacobian of the map (26).

2.4 Diffusivity and the limit of small T

We can vary the period of the sine flow given in (26) by varying T. For small T, the period of the stirring is less than the timestep of integration and one can expect the flow to appear nondiffusive. Indeed, in the limit of small T, the sine map given in (26) approaches a symplectic integrator [3]. This result can be obtained by noting that

$$\frac{x' - x}{T} = \sin\left(\frac{2\pi k y}{L}\right) \qquad\qquad \frac{y' - y}{T} = \sin\left(\frac{2\pi k x'}{L}\right)$$

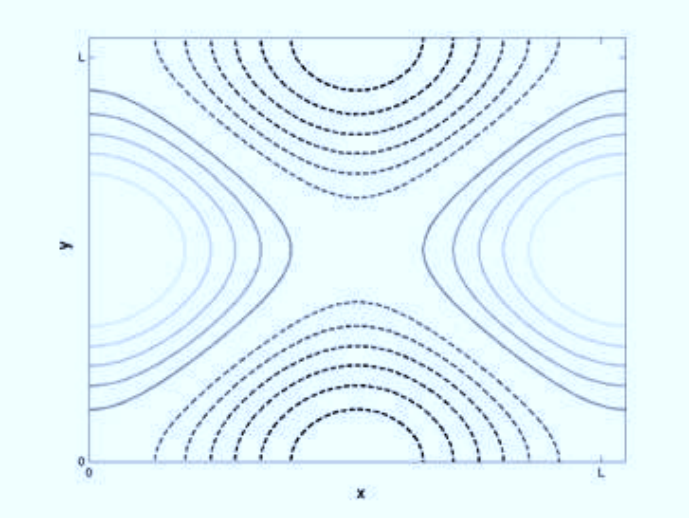

Figure 3: Streamfunction in the limit that $T \to 0$. The flow field then satisfies $\boldsymbol{u} = (\partial_y \psi, -\partial_x \psi)$ so that the flow is clockwise where ψ is negative (dashed) and counterclockwise where ψ is positive (solid).

and by taking the limit as $T \to 0$. In this limit, $\frac{x'-x}{T} \to \frac{dx}{dt}$ etc. such that

$$\frac{dx}{dt} = \sin\left(\frac{2\pi k y}{L}\right) \qquad\qquad \frac{dy}{dt} = \sin\left(\frac{2\pi k x}{L}\right). \tag{27}$$

Equation (27) is equivalent to saying the nondivergent velocity field, $\boldsymbol{u} = d\boldsymbol{x}/dt$, can be written in terms of a streamfunction $\boldsymbol{u} = (\partial_y \psi, -\partial_x \psi)$ where

$$\psi = \frac{L}{2\pi k}\left(\cos\left(\frac{2\pi k x}{L}\right) - \cos\left(\frac{2\pi k y}{L}\right)\right). \tag{28}$$

A contour plot of ψ is given in figure 3.

Consider now the advection of a cloud of particles displaced by the map (26). If we let $L = k = 1$ and use a relatively small value of T, $T = 0.1$, the flow follows nearly closed orbits as given in figure 4. As this flow closely approximates the streamfunction solution, the stirring given by (26) does not diffuse the concentration of particles. As a result, the use of an 'effective' diffusivity does not apply here.

While the orbit given in figure 4 is not strictly confined to the streamfunction in (28) because of the use of a finite T, it is not space filling. We can add noise to the sinemap though

$$\begin{aligned} x' &= x + T\sin\left(\frac{2\pi k y}{L}\right) + \sqrt{2D}\, a \\ y' &= y + T\sin\left(\frac{2\pi k x'}{L}\right) + \sqrt{2D}\, b \end{aligned} \tag{29}$$

where a and b are Gaussian random variables with zero mean and unit variance. Here D is a constant that controls the amplitude of the noise and can be consider a molecular diffusivity.

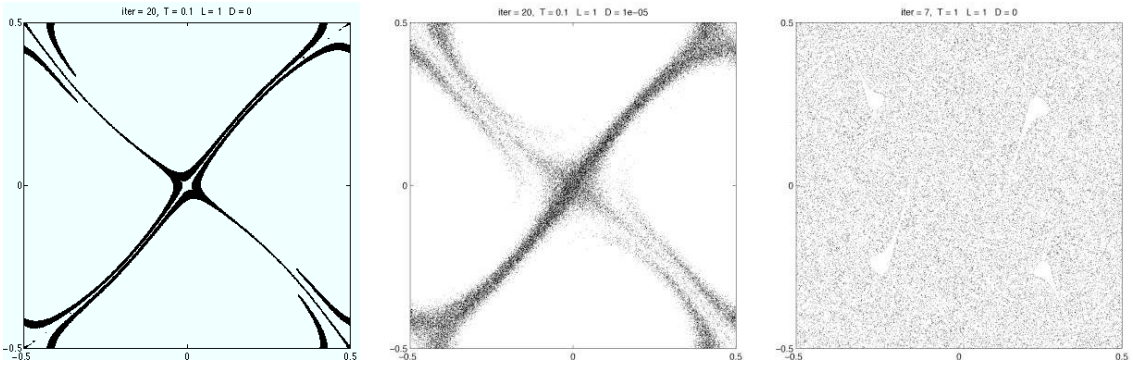

(a) $D = 0$, $T = 0.1$, and $it = 20$ (b) $D = 10^{-5}$, $T = 0.1$, and $it = 20$ (c) $D = 0$, $T = 1$, and $it = 7$

Figure 4: Particle concentration in a box of sides $L = 1$ and stirring timestep $T = 1$ after 20 iterations. There is no noise in (a) and the map is confined near the streamlines given in 3. With the addition of noise, where $D = 10^{-5}$, the orbit becomes diffusive. Figure 4(c) is the same as in 4(a) but with $T = 1$ and is shown only at the 7th iteration.

In the absence of stirring, (29) becomes the two-dimensional version of the random walk model given in (20). As we have chosen $\langle a^2 \rangle = \langle b^2 \rangle = 1$, the standard deviation σ^2 of the noise is simply $2D$.

If we let $D = 10^{-5}$ (figure 4(b)), the near-integrability of the map of figure 4(a) is broken, however the flow still remains constrained to near the streamlines in (28) for a moderate number of iterations. The addition of this molecular diffusivity, D, will cause the trajectory to eventually fill the domain. However, this process is quite slow.

The diffusive behavior captured by adding random noise in figure 4(b) can also be obtained by increasing the step size even in the absence of noise. For a box size $L = 1$ and with $D = 0$ and a timestep of $T = 1$, ten times longer than that of the examples given in figure 4, the flow no longer follows the streamfunction given in (28). The flow in figure 4(c) is chaotic and quite rapidly fills the entire domain with the tracer after only seven iterations.

2.5 Filament width and the Batchelor length

The Batchelor length, $l = \sqrt{\kappa/\lambda}$, is the approximate scale at which filaments, stretched by stirring, become so thin that the tendency for the filament to become narrower due to stirring is balanced by the diffusion of the patch. In the definition of the Batchelor length, λ is the local rate of strain. In the sine map given in (26), modified with the diffusion equation of (29), the rate of strain is approximately T and the coefficient of "molecular" diffusion is D. For this problem, then, the width of the stirred filaments should scale approximately as

$$l = \sqrt{D\,T} \tag{30}$$

such that the Batchelor length scales with $\sqrt{D}$ for fixed T.

Snapshot of the concentration fields for 'strong' diffusivity, with $D = 10^{-4}$, and 'weak' diffusivity, with $D = 10^{-6}$, are given in 5(a) and 5(b), respectively. The width of a typical

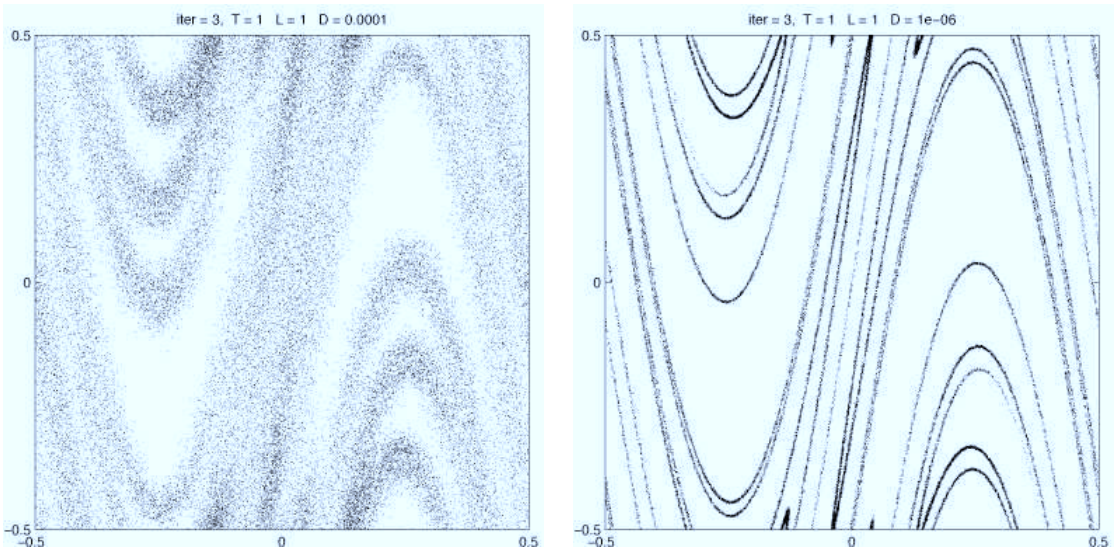

(a) $D = 10^{-4}$ (b) $D = 10^{-6}$

Figure 5: Particle concentration in a box of sides $L = 1$ with $T = 1$ and with varying D at iteration 3. The width of the filaments scale as the Batchelor length given in (30).

filament in the 'strong' diffusivity case is 0.109 nondimensional units while a 'weakly' diffusive filament is approximately 0.006 nondimensional units. The width of the filaments for these two cases do roughly scale with $\sqrt{D}$ as suggested from (30).

2.6 Effective diffusivity and the domain size, L

For the effective diffusivity scaling given in (25) to hold for a stirred fluid, the length scale of the stirring must be less than that of the domain. In the examples above, the length scale of the stirring was the same as that of the domain. If we increase the size of the domain to $L > 1$ while holding T at an intermediate value, $T = 1/2$, the concept of an 'effective' diffusivity becomes relevant.

Shown in figure 6 are snapshots of the particle concentration field after ten iterations for a domain size varying from $L = 1$ to $L = 25$ while the spatial correlations $\langle x^2 \rangle$, $\langle y^2 \rangle$, and $\langle X^2 \rangle = \langle x^2 + y^2 \rangle$ and cross-correlations $\langle xy \rangle$ are shown in figure 7. For small domain sizes, such as in 6(a) and 6(b), the sine flow stirring efficiently fills the domain with the tracer. This homogenization of the tracer field is evident in 6(c) and 6(d) by the leveling off of the correlation values with time.

In cases with a relatively small domain size, such as 6(a) and 6(b), the spatial variances for $\langle x^2 \rangle$, $\langle y^2$, and $\langle X^2 \rangle$ do not grow linearly with time as predicted from the effective diffusivity equations in one dimension, equations (22) and (23). The concept of an effective diffusivity fails in cases where the stirring is at approximately the same scale as the domain, as predicted.

Increasing the domain size to $L = 10$ and $L = 25$ while holding the scale of the stirring fixed, however, shows a more promising diffusive-like behavior. The tracer concentration in 6(c) and 6(d) spreads outward from an initial patch due to the stirring with spatial variances

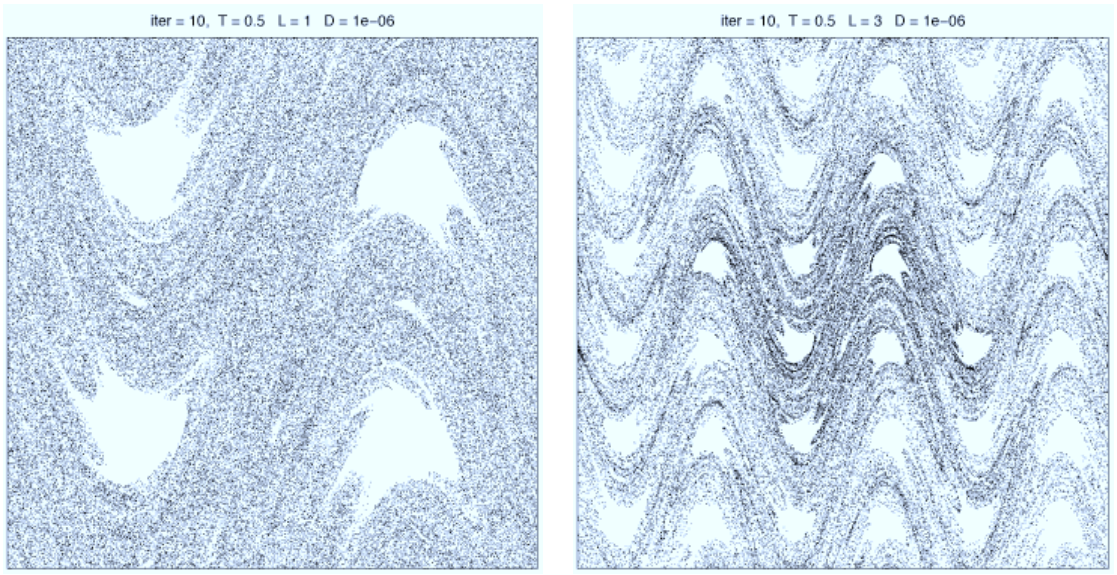

(a) $L = 1$ (b) $L = 3$

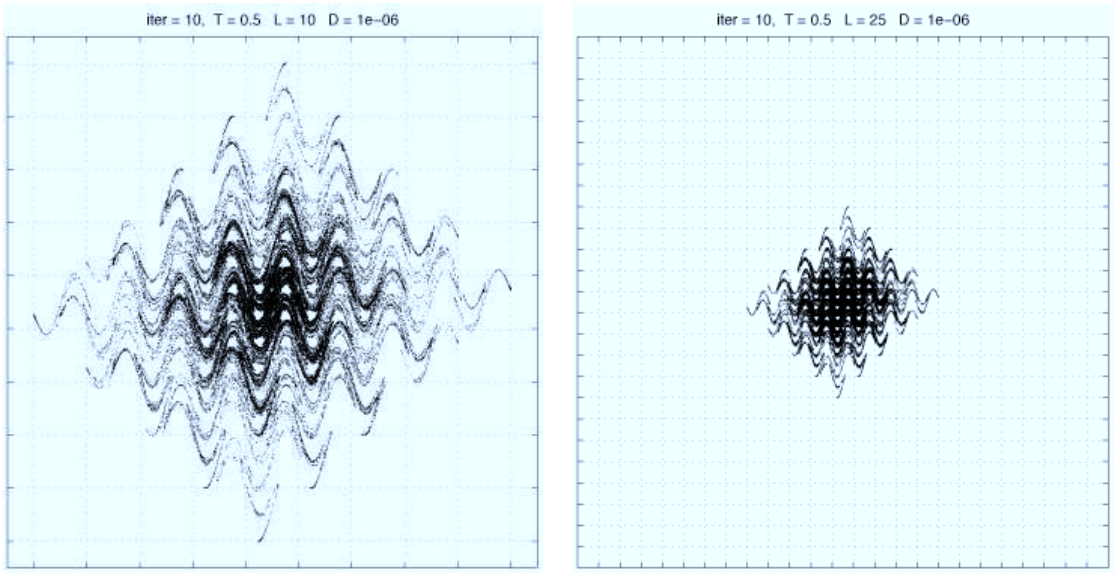

(c) $L = 10$ (d) $L = 25$

Figure 6: Particle concentration in a box of of domain size of $L = 1$ in 6(a), $L = 3$ in 6(b), $L = 10$ in 6(c), and $L = 25$ in 6(d) with $T = 1$ and with $D = 1 \times 10^{-6}$ at iteration 10.

growing linearly with time, as shown in figures 7(c) and 7(d).

Note that the variances in the x and y directions for all domain scales indicate isotropy, as shown by the similar behavior of $\langle x^2 \rangle$ and $\langle y^2 \rangle$ in figure 6. This agrees with the assumption that the variances in multiple dimensions have the same statistics/standard deviation as we assumed in the derivation of the multi-dimensional form of the effective diffusivity equation, (23). Additionally, the cross-correlation between the x and y spatial dimensions, $\langle xy \rangle$, is

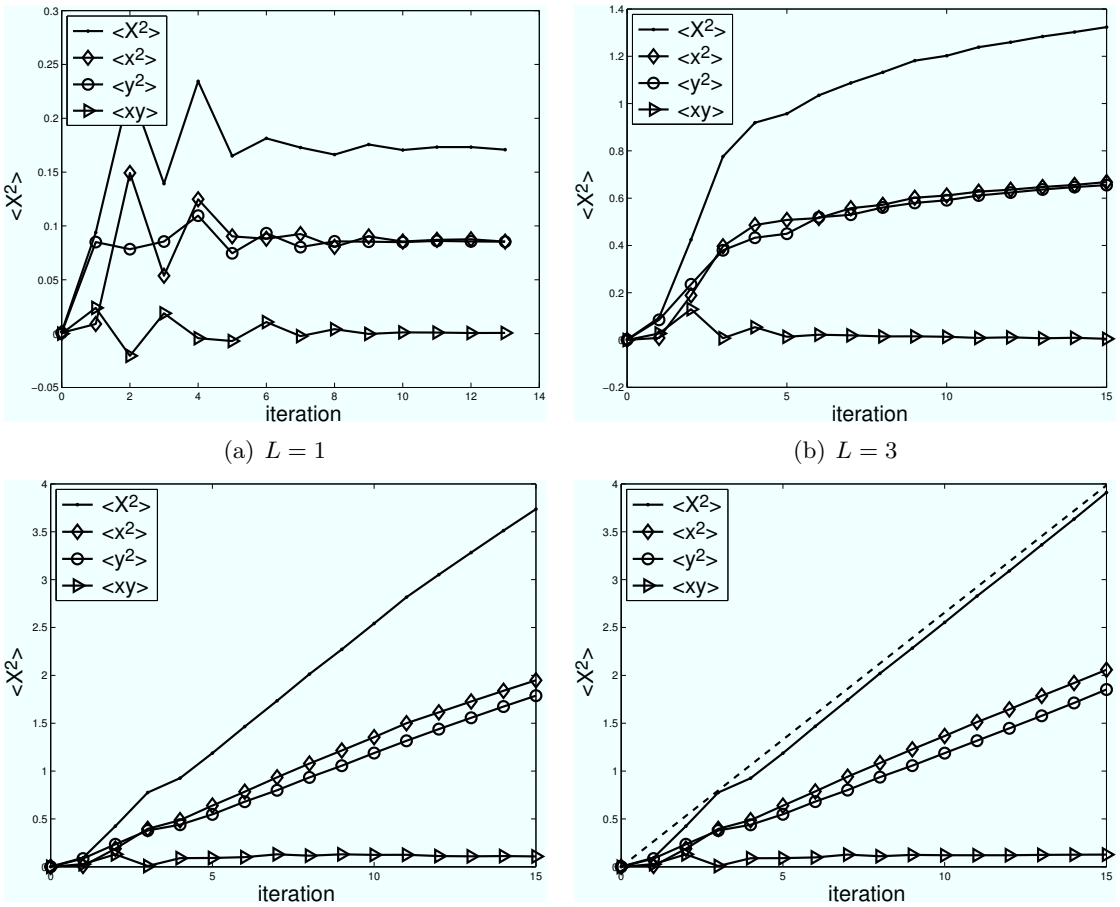

(a) $L = 1$ (b) $L = 3$

(c) $L = 10$ (d) $L = 25$

Figure 7: Calculations of the spatial variance for the examples shown in 6 as a function of iteration and increasing domain size.

zero at long times.

From (23), we calculate the approximate effective diffusivity for the $L = 25$ example as $\kappa = 0.068$, four orders of magnitude larger than the molecular diffusivity, $D = 10^{-6}$ — the stirring in this case more efficiently disperses the tracer field than molecular diffusivity alone.

3 Matlab Codes

3.1 `run_sinemap.m`

Listing 1: run_sinemap.m

```
function [varargout] = run_sinemap(T,L,D,Np,sqw)

if nargin < 1, T = 1; end
if nargin < 2, L = 1; end
```

```
if nargin < 3
  % Diffusivity (noise)
  D = 0;
end

if nargin < 4
  % Points initial square
  Np = 100;
end
Npx = Np; Npy = Np; Np = Npx*Npy;

if nargin < 5
  % Size of initial square
  sqw = .2;
end

f = @(X) sinemap(X,T/L,1/L);

x = linspace(-sqw/2,sqw/2,Npx)';
y = linspace(-sqw/2,sqw/2,Npy)';
X = [repmat(x,[Npy 1]) kron(y,ones(Npx,1))]/L;

dX2 = mean(sum(mymod(L*X,L).^2,2));
dx2 = mean(mymod(L*X(:,1),L).^2);
dy2 = mean(mymod(L*X(:,2),L).^2);
dxy = mean(mymod(L*X(:,1),L).*mymod(L*X(:,2),L));

it = 0;
figure(1)
plotmap(L*X,T,L,D,it)  % Stretch output.

while 1
  inp = input('Press a key (Q to quit)','s');
  if strcmp(lower(inp),'q'), break; end
  X = f(X); it = it + 1;
  dX2 = [dX2; mean(sum(mymod(L*X,L).^2,2))];
  dx2 = [dx2; mean(mymod(L*X(:,1),L).^2)];
  dy2 = [dy2; mean(mymod(L*X(:,2),L).^2)];
  dxy = [dxy; mean(mymod(L*X(:,1),L).*mymod(L*X(:,2),L))];
  % Add some noise
  if D > 0, X = X + sqrt(2*D)/L*randn(size(X)); end
  figure(1)
  plotmap(L*X,T,L,D,it)  % Stretch output.
end

hold off

if nargout > 0
  varargout{1} = dX2;
  figure(2)
  plot(0:it,dX2,'k.-')
  hold on
  plot(0:it,dx2,'r.--')
  plot(0:it,dy2,'b.--')
  plot(0:it,dxy,'g.--')
  hold off
  xlabel('iteration')
  ylabel('<X^2>')
  legend('<X^2>','<x^2>','<y^2>','<xy>','Location','NorthWest')
end

% --------------------------------------------------
function Xm = mymod(X,L)

Xm = mod(X + L/2,L) - L/2;

% --------------------------------------------------
function plotmap(X,T,L,D,it)

plot(mymod(X(:,1),L),mymod(X(:,2),L),'.','MarkerSize',1)
axis([-L/2 L/2 -L/2 L/2])
```

```
axis square
title(sprintf('iter = %d,  T = %g   L = %g   D = %g',it,T,L,D))

if L > 1
  gr = (0:L-1) - floor(L/2-.5) - .5;
  set(gca,'XTick',gr)
  set(gca,'YTick',gr)
  set(gca,'XTickLabel',[])
  set(gca,'YTickLabel',[])
  grid on
end
```

3.2 sinemap.m

Listing 2: sinemap.m

```
function Xp = sinemap(X,T,L)

if nargin < 2, T = 1; end
if nargin < 3, L = 1; end

Xp = zeros(size(X));

Xp(:,1) = X(:,1) + T*sin(2*pi*X(:,2)/L);
Xp(:,2) = X(:,2) + T*sin(2*pi*Xp(:,1)/L);
```

References

[1] G. K. Batchelor, Small-scale variation of convected quantities like temperature in turbulent fluid: Part 1. General discussion and the case of small conductivity, J. Fluid Mech., 5 (1959), pp. 113–133.

[2] R. T. Pierrehumbert, Tracer microstructure in the large-eddy dominated regime, Chaos Solitons Fractals, 4 (1994), pp. 1091–1110.

[3] W. R. Young, Stirring and mixing, in Proceedings of the 1999 Summer Program in Geophysical Fluid Dynamics, J.-L. Thiffeault and C. Pasquero, eds., Woods Hole, MA, 1999, Woods Hole Oceanographic Institution. http://gfd.whoi.edu/proceedings/1999/PDFvol1999.html.