Simulating vertical turbulent dispersal with finite volumes and binned random walks

Uffe Høgsbro Thygesen[1,*], Bjørn Ådlandsvik[2]

[1]**Danish Institute for Fisheries Research, Technical University of Denmark, Jægersborg Allé 1, 2920 Charlottenlund, Denmark**
[2]**Institute of Marine Research, PO Box 1870, Nordnes, 5817 Bergen, Norway**

ABSTRACT: Early life stages in fish are often modeled by individual-based models. The transport of individuals with ocean currents is addressed through particle-tracking techniques, which typically simulate vertical turbulent dispersal with random walk schemes. These schemes, however, perform poorly when the eddy diffusivity displays steep gradients, as in stratified water columns, and near the surface and sea floor. In the present paper, we advocate the use of a binned random walk, which keeps track, not of the exact vertical position of a tracer particle, but only of the layer in which the particle resides. The binned random walk is derived by discretizing the Eulerian equations governing the vertical dispersal with the finite-volume method. The scheme is easily implemented, even when layers are non-uniform and turbulence statistics originate from a circulation model, and, by construction, always satisfies the well-mixed criterion. We demonstrate the scheme and discuss its applicability.

KEY WORDS: Individual-based models · Dispersal · Random walk

INTRODUCTION

Many model studies in both physical and biological oceanography concern the transport and dispersal of tracers in the ocean. Examples of important processes are the transport of fish eggs or larvae between spawning grounds and nursery grounds (e.g. Christensen et al. 2007, this Theme Section), and the vertical distribution of fish eggs (e.g. Ådlandsvik et al. 2001, Boyra et al. 2003). Such studies can be done in the Eulerian framework, which models the density of tracers as it evolves in time, or in the Lagrangian or individual-based framework, which adds single particles or individuals to a numeric model of the ocean and tracks their motion. One advantage of the individual-based framework, which is particularly relevant for early life stages of fish, is the ease of adding internal states such as condition or developmental stages. The present paper concerns individual-based models (see also Thygesen et al. [2006] for comparisons and connections between the Eulerian and the individual-based approach).

In many situations, it is important to include turbulent dispersal in the model. Dispersal is typically mod-

eled with diffusion and simulated by means of random walk schemes (Visser 1997). This is a reasonable approach when the time scales of interest exceed the so-called Lagrangian time scale, which measures the persistency of the velocity of a passive tracer in turbulent flow (U. H. Thygesen & A. W. Visser unpubl.). For the vertical component, this time scale is on the order of minutes (Yamazaki al. 2002), whereas the horizontal component is on the order of days (Garrett 2006). It is often sufficient to model the vertical coordinate with a random walk. This is because the horizontal dispersal may be dominated by longitudinal, or shear, diffusion (Taylor 1954), i.e. the combined effect of vertical mixing and the horizontal velocities varying over the water column, or because horizontal eddies are resolved by the circulation model. In this paper we, too, will focus on vertical mixing.

One practical problem with random walk schemes is that they perform poorly when the eddy diffusivity profile is not smooth; then small time steps are required. This is problematic in an oceanographic context, where stratification may impose large local gradients and curvatures. Moreover, when the eddy diffusivity

profile is the output of a numerical model, the discrete data points must be interpolated sufficiently smoothly. Another issue concerns the boundaries at the surface and bottom. Here, the standard algorithms implement reflecting boundaries, but in this case the well-mixed criterion requires the diffusivity to be constant near the boundary (Ross & Sharples 2004), which does not agree with physics. A similar problem arises even with constant diffusivity when a vertical bias is present, whether due to buoyancy, sedimentation, or active vertical migration. Thus, the practical use of random walk schemes for turbulent dispersal in combination with circulation models is complicated by numerical considerations and statistical verifications that the well-mixed criterion (Thomson 1987) is met also in practice. See Brickman & Smith (2002) and references therein for a discussion of these difficulties.

On the other hand, in the literature of stochastic processes (Gardiner 1985, Grimmett & Stirzaker 1992; see also Csanady 1973, Okubo & Levin 2001), a common cartoon model of 1-dimensional diffusion is a random walk on a lattice in discrete time: at each time step, the particle moves 1 step 'up' or 'down' with equal probability. The standard construction goes on to demonstrate that if the spatial grid size k and the time step h go to zero such that $2D = k^2/h$ is constant, then the limiting process is diffusion with diffusivity D.

The aim of the present paper is to point out that this simple construction generalizes very easily to a situation in which the diffusivity is neither constant in space nor in time and the grid is not equidistant. In this situation, the transition probabilities may be obtained from a finite-volume discretization (Ferziger & Perić 2002) of the underlying diffusion equation. Thus, it is straightforward to implement the scheme using output from a circulation model that is discrete in space and time. It does not require this output to be smooth, so no vertical smoothing of diffusivity profiles is required. By construction, this scheme always satisfies the well-mixed criterion. This holds regardless of discontinuities and boundary behavior of the diffusivity profile, and does not require the time step to be infinitesimal, although the basic algorithm does have a maximum allowable time step.

We call the resulting algorithm a *binned random walk*, since it does not keep track of the exact position of the particle, but only models the bin (layer or cell) in which the particle resides. Thus, if one is willing to settle with this spatial resolution, then one may avoid the difficulties with verifying the well-mixed criterion. This allows the focus to shift towards the accuracy of the transients.

This binned random walk can also be used when a vertical bias is present, for example, when fish eggs are non-neutrally buoyant. The results will then be comparable to the Eulerian finite-volume approach used by Ådlandsvik et al. (2001) and Boyra et al. (2003), and has the advantage that it can be used as a component in a wider particle-tracking or individual-based framework.

The paper is organized as follows. In the section 'The diffusion equation and its finite-volume discretization', we consider the diffusion equation, which models unbiased dispersal and its discretization in space using finite volumes. In the section 'Lagrangian simulations', we discuss individual-based simulation of dispersal, using the discretized diffusion equation. The section 'An idealized example with stratification' uses a hypothetical example to illustrate dispersal when the rate of mixing changes abruptly; this serves as an extreme benchmark test for numerical schemes. The section 'Biased random walks in the vertical' includes vertical bias, e.g. due to the dispersing particles being non-neutrally buoyant. Finally, the 'Discussion' offers some conclusions and a discussion of the merits, limitations, and applicability of the scheme.

THE DIFFUSION EQUATION AND ITS FINITE-VOLUME DISCRETIZATION

The starting point for our analysis is the concentration field $C(z,t)$ giving the concentration C of a passive tracer substance, measured in mass or numbers per volume. Here, $t > 0$ is the time and $z \in [0, d]$ is the height over the ocean floor, so that $z = 0$ corresponds to the bottom and $z = d$ to the surface. It is governed by the diffusion equation:

$$\dot{C} = -J' = (DC')' \tag{1}$$

with no-flux boundary conditions, i.e. $J = -DC'$ vanishes at $z = 0$ and $z = d$. $\dot{C}$ indicates the time derivative $\partial C/\partial t$, whereas the prime in, e.g., C' denotes the spatial derivative $\partial/\partial z$. $J(z,t) = -D(z, t)C'(z,t)$ is the vertical diffusive flux. Note that the eddy diffusivity D may depend on position and time, i.e. we have $D = D(z,t)$.

To resolve this equation numerically, we pursue a finite-volume discretization (Ferziger & Perić 2002) of the vertical dimension, using a grid:

$$0 = z_0 < z_1 < \ldots < z_n = d \tag{2}$$

while we keep time a continuous variable. The discretized system uses n control volumes, with volume i containing the layer between $z = z_{i-1}$ and $z = z_i$. The system keeps track of the amount of material m_i in each control volume:

$$m_i(t) = \int_{z_{i-1}}^{z_i} C(z,t)\mathrm{d}z \tag{3}$$

with $i = 1,\ldots, n$, but not of how this material is distributed within the layer. (We will think of m_i as a biomass; it could equally well be measured in numbers of indi-

viduals). Note the difference between this approach and a pure finite-difference method, which will model the concentration of matter at grid points rather than the amount of matter within grid cells. To pose a dynamic equation for these amounts m_i, we first rewrite the diffusion equation (Eq. 1) in its integral form, which is the mass balance equation:

$$\dot{m}(t) = J(z_{i-1}, t) - J(z_i, t) \qquad (4)$$

To close the system, we must approximate the fluxes $J = -DC'$ in terms of the masses m_i. To this end, first note that the average concentration in control volume i is m_i/k_i, where $k_i = z_i - z_{i-1}$ is the width of layer i. Next, the distance from the center of layer i to that of the neighboring layer $i + 1$ is $(k_i + k_{i+1})/2$. Combining these, we obtain a first-order, finite-difference approximation of the spatial derivative C', which leads to the approximation:

$$J_i = D_i \frac{2}{k_i + k_{i+1}} \left(\frac{m_i}{k_i} - \frac{m_{i+1}}{k_{i+1}} \right) \qquad (5)$$

of $J(z_i, t)$. Here, D_i is the diffusivity at the interface between cells i and $i + 1$, i.e. at z_i. Note that J_i, D_i and m_i will depend on time, but that we suppress this for notational convenience. Combining with the mass balance of control volume i, we obtain:

$$\dot{m}_i = p_{i-1}m_{i-1} - (p_i + q_i)m_i + q_{i+1}m_{i+1} \qquad (6)$$

where the coefficients are:

$$p_i = D_i \frac{2}{(k_i + k_{i+1})k_i} \qquad (7)$$

$$q_i = D_{i-1} \frac{2}{(k_{i-1} + k_i)k_i} \qquad (8)$$

These equations apply to all inner boundaries between cells, but at the boundaries at the bottom and surface there is no flux. This is achieved by setting $p_0 = q_{n+1} = 0$.

Note that the scheme is guaranteed to conserve the total mass or number of tracers, since we are explicitly modeling the fluxes; this is the advantage of the finite-volume method. Also, the uniform concentration is necessarily stationary: if there are no concentration gradients, then there are no fluxes. These 2 desirable properties are independent of the grid and also of the diffusivities D_i applied to the interfaces; they may vary arbitrarily between adjacent cells, and they may be interpolated from the output of a numerical circulation model. So while interpolation requires some computational overhead and may give less accuracy, it does not jeopardize the well-mixed criterion.

So far we have discussed the numerics of an Eulerian model. For our development of the particle-following algorithm, the following interpretation of the balance equation (Eq. 6) is key: focusing on the term $-p_i m_i$ in

the derivative, we see that in a short time interval h, there is an amount $p_i m_i h$ of matter that is initially in cell i and moves to cell $i + 1$. Thus, a fraction $p_i h$ of the material that is initially in cell i moves into cell $i + 1$. Taking a random tracer molecule that is initially in cell i, we see that the probability that this molecule moves to cell $i + 1$ during the time interval is $p_i h$. Similarly, the term $-q_i m_i$ means that the same random tracer molecule has probability $q_i h$ of moving into cell $i - 1$. These are exactly the probabilities that we need to know in order to simulate the random motion of a tracer molecule.

LAGRANGIAN SIMULATIONS

Lagrangian *random walk* simulations of vertical turbulent dispersal take as a mathematical starting point the *diffusion processes* Z_t (see Gardiner 1985 for background material). This is a stochastic process in continuous time that models the trajectory of a single particle; Z_t is the vertical position at time t. It is connected with the diffusion equation (Eq. 1), which governs the transition probabilities of Z_t. One characterization of this process Z_t is that it solves the stochastic differential equation:

$$dZ_t = D'(Z_t, t)dt + \sqrt{2D(Z_t, t)}dB_t \qquad (9)$$

This equation specifies the change dZ_t in the vertical position Z_t over an infinitesimal time interval dt. Here, B_t is Brownian motion, i.e. a stochastic process for which the increment $B_{t+h} - B_t$ is a Gaussian distributed random variable with mean 0 and variance h, for any positive t and h. Note that the physical unit of B_t is the square root of time, $s^{1/2}$. Equations such as Eq. (9) allow several different interpretations (Gardiner 1985), but to obtain pure diffusion and, in particular, a uniform steady-state concentration we must use what is known as the interpretation of Itô. That is to say that Z_t can be approximated in discrete time with the Euler scheme:

$$Z_{t+h} - Z_t = D'(Z_t, t)h + \sqrt{2D(Z_t, t)}(B_{t+h} - B_t) \qquad (10)$$

This is a stochastic recursion: given Z_t, we may use a random number generator to simulate $B_{t+h} - B_t$ from a Gaussian distribution with mean 0 and variance h, and thus compute Z_{t+h}. Note the term $D'(Z_t, t)h$; with heterogeneous turbulence, this is a biased random walk. Visser (1997) found that a more careful evaluation of the square root improved accuracy. It is a result from the theory for numerical analysis of stochastic differential equations (Kloeden & Platen 1995) that, as the time step h goes to 0, the transition probabilities of the discrete-time recursion converge to the solution of the diffusion equation (Eq. 1). Fig. 1 gives a schematic of the first few time points in such a simulation.

For accuracy, the Euler scheme (Eq. 10) requires that the relative change in diffusivity over a single time step is small. To assess this change, we use Itô's lemma (Gardiner 1985); this is a stochastic version of the chain rule of differentiation. We find that the change in diffusivity $D(Z_t)$ experienced by the tracer over a small time step h has mean $(|D'|^2 + D''D)h$ and standard error $|D'|\sqrt{(2Dh)}$, to the lowest order in h. Thus, the time step h must be small relative to both $D/(D')^2$ and $1/|D''|$ (compare Thomson 1987, Wilson & Flesch 1993). More powerful numerical schemes are available (Kloeden & Platen 1995), but in practical applications, where the eddy diffusivity is not a theoretical profile but the output from a circulation model, high-order schemes are not well suited for increasing performance. This is because the diffusivity is only known at discrete grid points or as averages over discrete cells. This makes it difficult to determine the derivatives of the diffusivity that are needed in the typical higher-order scheme.

Due to these obstacles, we may choose to pursue the less ambitious goal of just keeping track of which cell (bin, layer) the particle is in. Recall the interpretation of the finite-volume scheme (Eq. 6), that in a short time interval of duration h a fraction p_ih of the material initially in cell i moves into cell $i + 1$, and likewise for the other terms. For the Lagrangian simulation, this means that a particle in cell i should, with probability p_ih, move into cell $i + 1$ during the time interval. This leads to the following discrete-time algorithm:
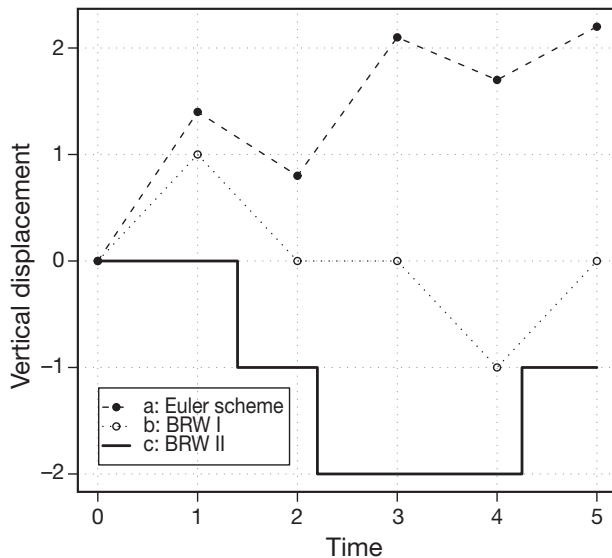

Fig. 1. The 3 schemes approximating diffusion, in the dimensionless case, with constant diffusivity $D = 0.3$. a, the discrete-time continuous-space Euler scheme (Eq. 10); b, BRW I: Binned Random Walk I, discrete in space and time, with a vertical grid size of 1 and a time step of 1; c, BRW II: continuous-time Binned Random Walk II, with a vertical grid size of 1. Note that these are 3 different realizations, so the actual sample paths are not expected to be identical

Binned Random Walk I:
(1) Let i be the number of the cell in which the particle resides at time t.
(2) Sample a random number U, uniformly distributed between 0 and 1.
(3) a. If $0 \le U < q_ih$, move the particle to cell $i - 1$.
  b. If $q_ih \le U < 1 - p_ih$, do not move the particle.
  c. If $1 - p_ih \le U < 1$, move the particle to cell $i + 1$.
(4) Advance time t to $t + h$ and go to Step 1.

This algorithm simulates a Markov chain, which approximates the diffusion process Z_t. The transition probabilities of this Markov chain converge to those of the diffusion process Z_t in the limit as the cell widths and the time step h go to zero. The scheme requires that the time step h is smaller than $\min_i 1/(p_i + q_i)$, because it should be impossible to remove more material from a cell than is available. As a note, this condition is also known in the context of numerical analysis of the Eulerian equation (Eq. 6), where it avoids negative concentrations and guarantees stability of the explicit time-marching scheme (see e.g. Ferziger & Perić 2002, p. 145).

The scheme has the advantage that it works robustly also when the jump rates p_i and q_i depend on time. Although we concentrate on the 1-dimensional case, this makes it possible to apply the scheme also in 3-dimensional situations; we comment briefly on this in the discussion section. Note that the uniform distribution, where the probability of the particle residing in each cell is proportional to the size of the cell, is necessarily stationary in time for any time step below this upper limit, by the construction of the finite-volume scheme. Thus, no extra effort is needed to assure or confirm that the well-mixed criterion is met.

With a uniform grid, constant diffusivity, and the appropriate time step, the algorithm reduces to a standard unbiased random walk on a lattice (Csanady 1973, Okubo & Levin 2001). Furthermore, with a uniform grid but varying diffusivity, the algorithm matches exactly the mean and the variance of the displacement of a particle over a short time interval (see Appendix 1).

When the jump rates p_i and q_i are constant in time, we can approximate the diffusion process Z_t with a continuous-time Markov chain, thus eliminating the fixed time steps and only sampling the process when it shifts from one cell to another. To this end, we exploit that the residence time in each cell is a random variable following the exponential distribution with mean $1/(p_i + q_i)$ (Grimmett & Stirzaker 1992). We arrive at the following continuous-time algorithm:

Binned Random Walk II:
(1) Let i be the number of the cell in which the particle resides at time t.

(2) Sample random numbers U and I, uniformly between 0 and 1.

(3) Compute the residence time H in cell i as $-(p_i + q_i)^{-1} \log U$.

(4) a. If $0 \leq I < q_i/(p_i + q_i)$, move the particle to cell $i - 1$.
    b. If $q_i/(p_i + q_i) \leq I < 1$, move the particle to cell $i + 1$.

(5) Advance time t to $t + H$ and go to Step 1.

This algorithm avoids errors associated with discretization in time, which is appealing. However, when many particles are simulated at once, it will result in a different sequence of time steps for each particle, which leads to a slightly more complicated implementation. When studying idealized flows, one will often prefer to output the sequence of time steps as well as the cell number, for each particle. In other situations, one may choose to simply output the position of each particle at fixed, regularly-spaced time steps.

Finally, when the jump rates p_i and q_i are constant, we can compute the transition probabilities over long time intervals as $\exp(Ah)$. Here, A is a tridiagonal matrix with elements p_i on the first superdiagonal, q_i on the first subdiagonal, and $-p_i - q_i$ on the diagonal. In the theory of Markov chains, A is known as the 'generator'. Algorithms for computing this matrix exponential are well known (Moler & Van Loan 2003) and are available in standard software packages. This allows us to take arbitrarily long time steps without any error associated with the time step, and thus also to verify the effect of using finite time steps.

AN IDEALIZED EXAMPLE WITH STRATIFICATION

We consider vertical dispersion in a stratified water column, with a total water depth of 50 m, and with perfect stratification so that the diffusivity is 0.1 m^2 s^{-1} above the pycnocline at 25 m and 0.02 m^2 s^{-1} below it. This is not meant to accurately represent the physics near a pycnocline, where we would expect a much reduced diffusivity at the very interface. Rather, it serves as an extreme benchmark for numerical schemes. The reason for this is that the gradient of the diffusivity D' is effectively infinite at the interface and zero everywhere else. Since this gradient term appears in the Euler scheme (Eq. 10), the scheme cannot be implemented directly; also, the maximum time step as computed in the section 'Lagrangian simulations' is zero. We will see, however, that the binned random walk performs well even in this extreme case.

We chose a non-regular grid with smaller grid cells near the interface. This demonstrates that the spacing in the finite-volume method needs not be regular, but also increases resolution in the initial phase when particles are concentrated near the interface. Sufficient resolution for the plots was obtained with 75 layers;

note that the well-mixed criterion will be met for any number of layers. The pycnocline itself is in the middle layer, Number 38.

We release an ensemble of 10 000 tracers at time $t = 0$ at the pycnocline, i.e. in the central layer. Fig. 2 shows the vertical distribution during the initial phase of the dispersion (Fig. 2b), when the effects of bottom and surface are not yet noticeable. In this phase the solution is analytically available (Appendix 2); note that more than half the material moves into the zone with high diffusivity, where the tail is also longer. Fig. 2b also shows histograms of the vertical position of tracers, which move according to Binned Random Walk I (see the section 'Lagrangian simulations'). Fig. 2b demonstrates the good agreement between the stochastic simulation and the analytical solution, also during transients.

Fig. 2 also shows the final phase (Fig. 2d–f), when the material is approaching the final uniform distribution over the water column. Note that the time scale of the final transition to uniformity is between 10 000 and 30 000 s; this may be confirmed by computing the half-time of the slowest mode of the transition rate matrix A, which is 5790 s.

The dots in Fig. 2d–f are histograms of stochastic simulations with the naive random walk model $Z_{t+h} = Z_t + \sqrt{2D(Z_t)}(B_{t+h} - B_t)$, obtained by omitting the bias term $D'(Z_t)h$ in the Euler scheme (Eq. 10) (cf. Visser 1997). (Recall that D' in our case is zero everywhere except at the pycnocline where it is infinite). It is well known (Visser & Thygesen 2003) that this scheme produces an incorrect steady-state concentration of $C(z) \sim 1/D(z)$, compare Fig. 2f. Fig. 2d,e demonstrates that also the transients are qualitatively wrong, in that initially the majority of particles move down rather than up.

Finally, Fig. 2 contains sample paths obtained using Binned Random Walk I (see the section 'Lagrangian simulations'; Fig. 2c).

BIASED RANDOM WALKS IN THE VERTICAL

While the previous derivation was for pure diffusion, many applications have a vertical bias due to buoyancy, sedimentation, or active vertical migration of individuals. For example, Ådlandsvik et al. (2001) derive the dynamics of the vertical distribution of fish eggs and larvae from their buoyancy. Another example of such a bias is diel vertical migrations. The starting point for a simulation of vertical motion is the advection–diffusion equation:

$$\dot{C} = -(uC - DC')' \qquad (11)$$

Here, $u = u(z, t)$ is the bias, which mathematically appears as an advective term and which may again
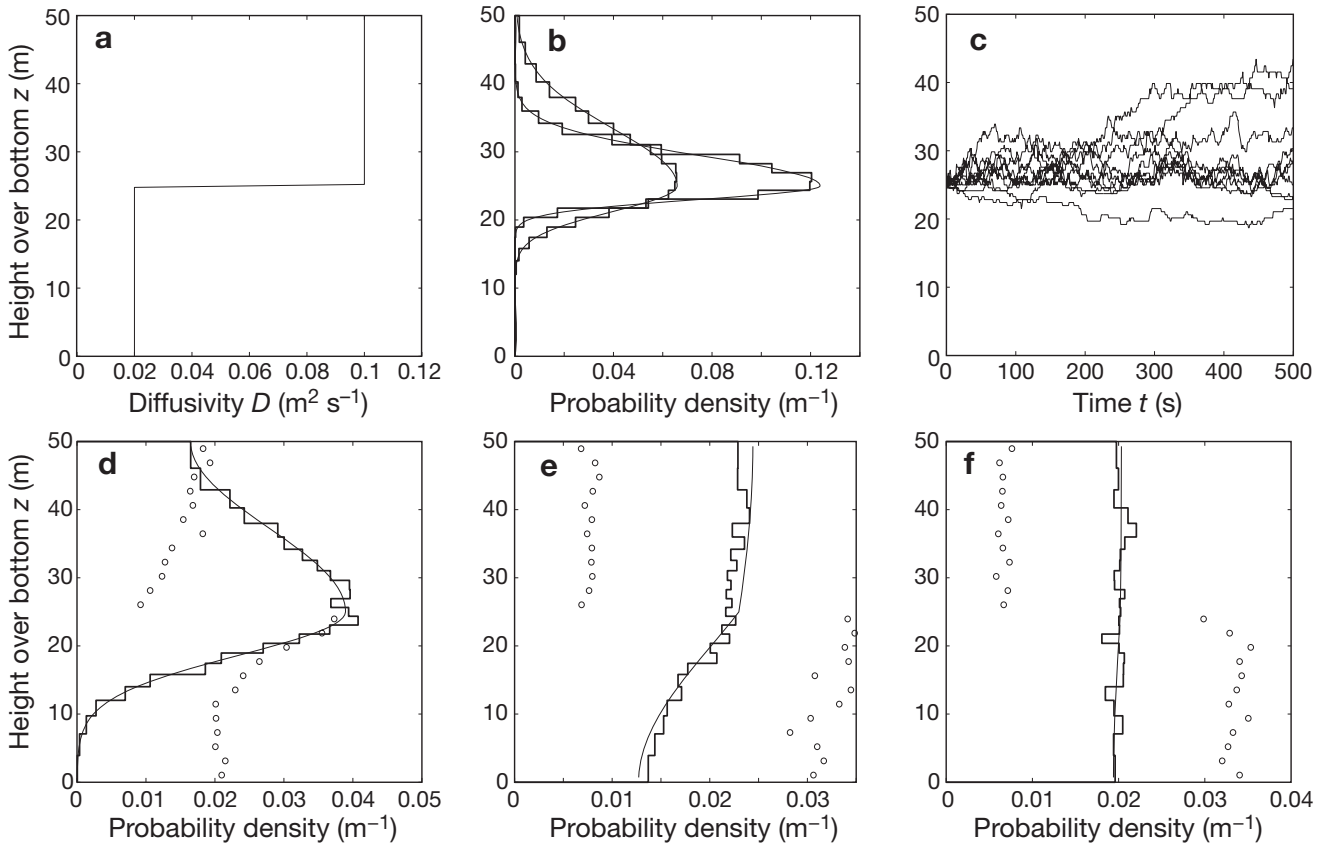
Fig. 2. Numerical solution of the diffusion equation with stratification. (a) The diffusivity profile. (b) The initial self-similar (scaling) phase; analytical solution (smooth curves) versus histograms (staircases) of the stochastic simulation of Binned Random Walk I. The more peaked distribution is for $t = 100$ s and the less peaked for $t = 350$ s. (c) Sample paths for Binned Random Walk I. (d) The final transition to a stationary state. Numerical solution using the final-volume method (smooth curve), histograms for Binned Random Walk I (staircase), and histogram for naive random walk (dots); time $t = 1000$ s. (e) As (d); time $t = 10\,000$ s. (f) As (d); time $t = 30\,000$ s

vary with space and time. The appropriate boundary conditions are *no flux*, i.e. that $J = uC - DC'$ vanishes at the boundaries $z = 0$, d. To simulate the motion of a particle, one may modify the Euler scheme (Eq. 10) to include the bias term, but, once again, the boundary behavior causes difficulties. In fact, a simple reflection scheme with constant diffusivity corresponds to a boundary condition $C' = 0$ rather than $uC - DC' = 0$. This provides another motivation for performing the simulation using a discrete-space Markov chain rather than the Euler scheme (Eq. 10).

For a finite-volume discretization of Eq. (11), it is always possible to use a first-order upwind scheme, where the net advective flux from cell i to $i + 1$ is approximated as:

$$\frac{m_i}{k_i}(u_i \vee 0) + \frac{m_{i+1}}{k_{i+1}}(u_i \wedge 0) \qquad (12)$$

Here, we use the shorthand $u_i = u(z_i)$, i.e. the velocity is evaluated at the cell boundary. The notation $a \wedge b$ means min(a, b) and $a \vee b$ means max(a, b). From the Lagrangian point of view, it means that the rate of

transition to cell $i + 1$ for a particle in cell i is now:

$$p_i = D_i \frac{2}{(k_{i+1} + k_i)k_i} + \frac{u_i \vee 0}{k_i} \qquad (13)$$

and likewise the rate of transition to cell $i - 1$ for the same particle is:

$$q_i = D_{i-1} \frac{2}{(k_i + k_{i-1})k_i} - \frac{u_{i-1} \wedge 0}{k_i} \qquad (14)$$

With these modified transition rates, the simulation proceeds according to the algorithms of the section 'Lagrangian simulations'.

Since this is only a first-order scheme, it introduces numerical diffusion, which leads to poor accuracy. It is therefore tempting to use a second-order central scheme, in which the net advective flux from cell i to cell $i + 1$ is approximated as:

$$u_i \frac{k_i^2 m_{i+1} + k_{i+1}^2 m_i}{k_i k_{i+1}(k_i + k_{i+1})} \qquad (15)$$

This expression is derived by linear interpolation of the concentrations m_i/k_i and m_{i+1}/k_{i+1} from the cell centers to their interface. This leads to transition rates:

$$p_i = \frac{2D_i + u_i k_{i+1}}{k_i(k_i + k_{i+1})} \quad \text{and} \quad q_i = \frac{2D_{i-1} - u_{i-1} k_{i-1}}{(k_i + k_{i-1})k_i} \quad (16)$$

This will improve the accuracy, so that fewer levels may be needed, although the requirement of positive transition rates leads, in the uniform case, to the constraint $P < 2$, where $P = |u|k/D$ is the cell Peclet number. This condition is also well known in the field of numerical analysis of partial differential equations; it avoids oscillations (Vreugdenhil & Koren 1993) and guarantees boundedness (Ferziger & Perić 2002, p. 145] in the corresponding second-order numerical scheme for the Eulerian advection–diffusion equation. As for the time step, whether we use the first-order upwind scheme or the second-order central scheme, the requirement $(p_i + q_i)h < 1$ in this case incorporates the Courant-Friedrichs-Levy stability condition for the advection part. In the uniform case this simplifies to:

$$\left(\frac{2D}{k^2} + \frac{|u|}{k}\right)h < 1 \quad (17)$$

In summary, seen from a numerical perspective, it is easy to include vertical bias, whether this is due to non-neutral buoyancy or active vertical behavior: given an advection–diffusion equation (Eq. 11), which governs the concentration, we may use the technique of this section to simulate the motion of an individual. From a modeling perspective, however, it may not be trivial to establish the right bias and diffusivity in Eq. (11). Even for passive but non-neutrally buoyant particles like fish eggs, the dispersal will differ from that of a fluid element. When active vertical migration is present, as may be the case for fish larvae, dispersal is the result of both turbulence and unpredictable individual motion. Although statistics may be obtained from careful observations of individuals (Grünbaum 1999, Visser & Thygesen 2003), the variability in behavior is likely to limit the fidelity of the model. In this situation, there is a limit to the effort that should be invested in the accuracy of the numeric algorithms.

DISCUSSION

Although the theory of random walk models for turbulent dispersal is fairly complete, at least in the single-particle setting, the practical use remains impaired by difficulties concerning the choice of time step, the interpolation of eddy diffusivity profiles that are only available in discrete points, and reflection laws at the boundaries. These issues may lead to accumulation of particles in disagreement with the laws of physics, as may plain omission of the corrective bias term, and require statistical efforts to verify that the well-mixed criterion is not seriously violated. In this situation, we believe that the binned random walk discussed in this paper is a worthy alternative, because it satisfies the well-mixed criterion, by construction, and is substantially simpler to implement than any other scheme to our knowledge.

The binned random walk will not deliver vertical resolution beyond the width of the layers used. This in itself is no disaster, as it can be argued that the vertical resolution will, in any case, be bounded by the width of the layers in the underlying circulation model. Nevertheless, a practical question is how many layers to include in the binned random walk. One should at least use the same number as in the underlying circulation model, but one can easily use more, although the number of layers is naturally limited by computational resources. The computational burden per time step is largely independent of the number of layers, but the required time step to get accurate transients will decrease with the number of layers squared. Thus, the number of layers is a trade-off between computational resources and the need for vertical accuracy imposed by the variability of other fields, e.g. the horizontal flow velocity. These trade-offs are similar to the choice of time step for a standard random displacement model; although we have not done a detailed comparison of numerical performance, the 2 methods appear similar in terms of computing time.

With a fixed discretization of the vertical, the choice stands between the 2 algorithms described in the section 'Lagrangian simulations'. Binned Random Walk II is the more efficient in ideal case studies where the fields do not change with time, whereas Binned Random Walk I is able to handle the time-varying fields of typical realistic cases. With this algorithm, the choice of time step h still remains. The absolute requirement for the scheme to be well defined is that $h < 1/(p_i + q_i)$ for all i, so that all probabilities are between 0 and 1. For accuracy, the time step must be small enough so that >1 state transition in the continuous process is unlikely. This, it may be shown, requires that $h^2(p_i + q_i)^2/2$ is small. While this gives some guidance, the quantitative effect of discretizing time (and, for that matter, space) is easily analyzed using the Eulerian counterpart, under frozen conditions, or Binned Random Walk II. This should be contrasted with, for example, the random walk scheme (Eq. 10), where it is substantially more difficult to determine even how much the stationary distribution differs from the uniform one for a given time step, let alone errors in transients. Note, however, that the Euler scheme (Eq. 10) does not have absolute requirements for the time step, only requirements for accuracy.

One disadvantage of a binned random walk for vertical dispersal is that the vertical position is necessarily

discrete and, thus, cannot be a continuous function of underlying parameters. This makes it more difficult to base sensitivity analysis on individual trajectories rather than on ensemble statistics, as may be done with continuous-space, random walk schemes (U. H. Thygesen & A. W. Visser unpubl.).

We have focused on the vertical dimension. Binned random walk methods can easily be extended to 2- or 3-dimensional situations. However, for advection-dominated horizontal flow these methods will suffer from the same kind of numerical diffusion and dispersion problems as the finite-volume/differences methods from which they are derived. It is therefore recommended to use ordinary particle tracking in the horizontal, maintaining the continuous particle position. This can be combined with a binned random walk in the vertical for approximate treatment of mixing, buoyancy, and/or biological behavior. In this situation, the vertical jump rates (i.e. the probabilities p_i and q_i) will typically vary with both time and horizontal position.

The oceanographic community generally agrees that it is of primary importance that Lagrangian simulations of turbulent dispersal do not display aggregation of particles due to artifacts of models or implementations. In this regard, the formal statement of the well-mixed criterion (Thomson 1987) was seminal in that it gave a continuous-time, random-flight process governing the motion of a tracer, which displayed the correct steady-state statistics and the right inertial subrange. Unfortunately, the numerical issues regarding discretization of this process, including interpolation of gridded data and boundary behavior, mean that the well-mixed criterion remains a hurdle for practitioners, also when using random walks. In the present paper, we have directed attention to the fact that the well-mixed criterion can be satisfied in practice, also in discrete time, by using a binned random walk. Thus, we hope that the focus can shift towards the accuracy of the transients, which in many applications is of greater importance than the steady state.

LITERATURE CITED

Ådlandsvik B, Coombs S, Sundby S, Temple G (2001) Bouyancy and vertical distribution of eggs and larvae of blue whiting (*Micromesistius poutassou*): observations and modelling. Fish Res 50:59–72

Boyra G, Rueda L, Coombs SH, Sundby S, Ådlandsvik B, Uriarte A (2003) Modelling the vertical distribution of eggs of anchovy (*Engraulis encrasicolus*) and sardine (*Sardina pilchardus*). Fish Oceanogr 12:381–395

Brickman D, Smith PC (2002) Lagrangian stochastic modeling in coastal oceanography. J Atmos Ocean Technol 19:83–99

Christensen A, Daewel U, Jensen H, Mosegaard H, St. John M, Schrum C (2007) Hydrodynamic backtracking of fish larvae by individual-based modelling. Mar Ecol Prog Ser 347:221–232

Csanady GT (1973) Turbulent diffusion in the environment. Springer, Heidelberg

Ferziger JH, Perić (2002) Computational methods for fluid dynamics, 3rd edn. Springer, Heidelberg

Gardiner CW (1985) Handbook of stochastic models, 2nd edn. Springer, Heidelberg

Garrett C (2006) Turbulent dispersion in the ocean. Prog Oceanogr 70:113–125

Grimmett GR, Stirzaker DR (1992) Probability and random processes, 2nd edn. Oxford University Press, Oxford

Grünbaum D (1999) Advection–diffusion equations for generalized tactic searching behaviors. J Math Biol 38:169–194

Kloeden PE, Platen E (1995) Numerical solution of stochastic differential equations. Springer, Heidelberg

Moler C, Van Loan C (2003) Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. SIAM (Soc Ind Appl Math) Rev 45(1):3–49

Okubo A, Levin S (2001) Diffusion and ecological problems: modern perspectives. Springer, Heidelberg

Ross ON, Sharples J (2004) Recipe for 1-D Lagrangian particle tracking models in space-varying diffusivity. Limnol Oceanogr Methods 2:289–302

Taylor GI (1954) The dispersion of matter in turbulent flow through a tube. Proc R Soc Lond A 223:446–468

Thomson DJ (1987) Criteria for the selection of stochastic models of particle trajectories in turbulent flows. J Fluid Mech 180:529–556

Thygesen UH, Nilsson AFN, Andersen KH (2006) Eulerian techniques for individual-based models based on additive processes. J Mar Syst doi 10.1016/j.marsys.2006.10.005

Visser AW (1997) Using random walk models to simulate the vertical distribution of particles in a turbulent water column. Mar Ecol Prog Ser 158:275–281

Visser AW, Thygesen UH (2003) Random motility of plankton: diffusive and aggregative contributions. J Plankton Res 25(9):1157–1168

Vreugdenhil CB, Koren B (eds) (1993) Numerical methods for advection–diffusion problems, Vol 45. Notes on numerical fluid mechanics. Vieweg, Wiesbaden

Wilson JD, Flesch TK (1993) Flow boundaries in random flight dispersion models: enforcing the well-mixed condition. J Appl Meteorol 32:1695–1707

Yamazaki H, Mackas DL, Denman KL (2002) Coupling small-scale physical process with biology. In: Robinson AR, McCarthy JJ, Rothschild BJ (eds) The sea, Vol 12, Chap 3. John Wiley & Sons, New York, p 51–112

Appendix 1. Incremental mean and variance in random walks on grids

Consider first the Euler scheme (Eq. 10) and assume Z_t is given. The mean displacement is $\mathbf{E}\{Z_{t+h}|Z_t\} = Z_t + D'(Z_t,t)h$, whereas the variance of the displacement is $\mathbf{V}\{Z_{t+h}|Z_t\} = 2D(Z_t,t)h$. For the continuous-time diffusion process given by the stochastic differential equation (Eq. 9), these expressions are correct to first order in the time increment h (Kloeden & Platen 1995). Note also that the mean-square of the increment coincides with the variance to first order in h, since the mean is of first order in h.

Next, consider Binned Random Walk I (see section 'Lagrangian simulations'), assume that the grid is uniform so that $z_i - z_{i-1} = k$ for all i, and assume that the particle at time t is in cell i. The mean displacement in the interval $(t, t + h)$ is:

$$-kq_ih + kp_ih \;=\; \frac{D_i - D_{i-1}}{k}\,h$$

Since the probability of a displacement of $-k$ is q_ih, just as th probability of a displacement of $+k$ is p_ih. Using the usual finite-difference approximation of D', this agrees with $D'h$ where D' is evaluated at the midpoint of cell i.

Next, the mean square displacement is:

$$k^2q_ih + k^2p_ih \;=\; (D_i + D_{i-1})h$$

Again, using $(D_i + D_{i-1})/2$ as an approximation for the diffusivity at the midpoint of cell i, this agrees with $2Dh$. In summary, the mean and variance of the increment in the Binned Random Walk I agree with the underlying diffusion process, to first order.

Appendix 2. Diffusion with piecewise constant diffusivity

We consider the concentration field $C(z,t)$ of a tracer in an infinite 1-dimensional space, subject to pure diffusion with diffusivity D_+ for $z > 0$ and diffusivity D_- for $z < 0$. If a unit quantity of a tracer substance is released at the interface at $t = 0$, then its density at time t is:

$$C(z,t) \;=\; \frac{1}{\sqrt{\pi t}}\,\frac{1}{\sqrt{D_+} + \sqrt{D_-}}\exp\!\left(-\frac{z^2}{4D(z)t}\right)$$

To see that this is the right solution in the presence of the discontinuity of D at the interface, note that this C satisfies the integral form of the conservation equation, from which the diffusion equation (Eq. 1) is derived:

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_a^b C(z,t)\mathrm{d}z \;=\; J(a,t) - J(b,t)$$

where $a < b$ are arbitrary and $J = -DC'$ is the diffusive flux.

Notice that the tracer substance is not equally divided between the 2 regions; more material will be present in the region where the diffusivity is higher.