JOURNAL OF COMPUTATIONAL PHYSICS 106, 366–376 (1993)

On the Use of Random Walk Models with Spatially Variable Diffusivity

J. R. HUNTER, P. D. CRAIG, AND H. E. PHILLIPS

CSIRO Division of Oceanography, GPO Box 1538, Hobart, Tasmania, Australia

Received April 13, 1992; revised October 5, 1992

The random walk technique is commonly used to model diffusion in the environment. For a constant diffusivity K and model time-step δt, the random step should be chosen from a distribution with variance $2K\delta t$. However, if K varies spatially, this choice of step leads to the accumulation of particles in regions of low diffusivity. This problem may be overcome either by the incorporation of an apparent advection velocity, or by transforming to a coordinate system in which the diffusivity is constant. The latter technique requires no immediate approximations, is applicable to any reasonable diffusivity field and is therefore the preferred approach. In this case, as with constant K, the random step should be chosen from a normal distribution, for reasons of both theoretical accuracy and computational efficiency. Three important aspects of model design are discussed: the selection of the random number generator, the time step and the total number of particles.

1. INTRODUCTION

Particle-tracking models have been used quite extensively for the simulation of the transport of passive substances within fluids [e.g., 11]. The substance is modelled by an assemblage of discrete particles, which may be subjected to advection, diffusion, and decay by a variety of relatively simple deterministic or stochastic processes. For example, advection is simulated by a translation of each particle with a velocity derived from the local fluid velocity field. Diffusion is generally simulated using the random walk technique, where each particle position is periodically perturbed by a random vector, derived from a defined distribution function, P. The spatial size (specifically, the variance) of P is derived from the required effective diffusion coefficient(s). To simulate a spatially constant diffusion coefficient, K, the variance, or second moment, N_2^*, of P is given by

$$N_2^* = 2K\,\delta t, \tag{1}$$

where δt is the time step (see, for example, [11] or Eq. (11), later) and the asterisk is used specifically to indicate moments of P, rather than those of a general concentration distribution. Decay may be modelled deterministically (e.g., by assigning to each particle a certain quantity of substance which decreases with time) or stochastically (e.g., by the periodic random removal of individual particles).

Particle-tracking models have been used for the simulation of the transport of contaminants in turbulent boundary layers [e.g., 2], heat [e.g., 3, 9], oil [e.g., 1, 6, 8], biological entities [e.g., 7, 16] and hydrodynamically active particles [e.g., 5, 11].

Particle-tracking models offer advantages over finite-difference (FD) or finite-element (FE) solutions of the partial differential transport equations in the following two respects. First, they afford simpler and more accurate predictions of advection, which can only be modelled well by FD or FE methods if relatively complicated numerical schemes are used [e.g., 13]. If the velocity field can be locally described by an analytic function, then particles may be advected exactly through that field by simple integration. Second, for problems where the substance does not occupy the whole model domain, particle-tracking models may be significantly more computationally efficient than their FD or FE counterparts [9]. However, it should be noted that, in common with other stochastic methods, the accuracy of particle-tracking models which use random processes varies as the square root of the computational effort, whereas the accuracy of deterministic models generally varies in proportion to the computational effort.

In the presence of turbulence, the effective diffusivity is generally not a constant and varies spatially due to the presence of boundaries (e.g., "log layer" effects) and due to ambient density stratification, both factors tending to inhibit mixing. Transport models of turbulent systems should therefore be capable of simulating a spatially variable diffusion coefficient. The simplest solution to this problem (henceforth termed the "naive" method) would appear to be to define a spatially variable variance, equivalent to the value of N_2^* given by Eq. (1). However, it is easily found in practice that this leads to an apparent advection in directions of decreasing diffusivity and hence a

0021-9991/93 $5.00

concentration of particles in regions of low diffusivity. This paper indicates why this should be so and presents two ways of overcoming the problem.

2. THE MOMENTS OF THE CONCENTRATION DISTRIBUTION

At the beginning of a time step of a random walk model, each particle represents a delta function of concentration, being a finite quantity of substance (q), occupying an infinitesimally small space. After one time step, the particle has been stochastically "spread" since its new position is defined by a probability distribution, P, spanning a finite volume of space; P is equal to the new concentration distribution for each particle, normalised by q. This process may be conveniently described in terms of the moments of the concentration distribution, with respect to the particle position at the beginning of the time step. For example, for a conservative substance, the zeroth moment (the concentration integral) remains constant. Higher moments are zero at the beginning of the time step; advection moves the centroid of the distribution, changing the first moment; diffusion increases the variance, or second moment. In order for a random walk model to accurately simulate advection and diffusion, the moments of P should reflect the change of moments as described by the transport equation. In this section, we ignore the effect of advection and derive equations describing the evolution of the moments of the concentration distribution.

The forthcoming discussions will be confined to systems of one dimension, extension to higher dimensions being a straightforward exercise. The diffusion equation is:

$$\frac{\partial c}{\partial t} = \frac{\partial}{\partial x}\left(K\frac{\partial c}{\partial x}\right) \equiv (Kc')' \tag{2}$$

with x the spatial coordinate, t the time, c the concentration, and K the diffusivity. As indicated in Eq. (2), primes will be used to denote spatial derivatives.

We may define the unnormalised nth moment ($n \geqslant 0$) of c about the origin as

$$M_n = \int_{-\infty}^{\infty} cx^n \, dx \tag{3}$$

and the normalised nth moment as

$$N_n = M_n / M_0. \tag{4}$$

K may be expanded about the origin to yield the Taylor series:

$$K = K(0) + K'(0)x + \frac{K''(0)}{2}x^2 + \cdots + \frac{K^{(m)}(0)}{m!}x^m + \cdots = \sum_{m=0}^{\infty} \frac{K^{(m)}(0)}{m!}x^m. \tag{5}$$

Equations (2) (multiplied by x^n), (3), and (5), yield, after two integrations by parts and the assumptions that $Kx^n(\partial c/\partial x) \to 0$ and $nKx^{n-1}c \to 0$ as $x \to \pm\infty$ (i.e., that c is spatially limited),

$$\frac{d}{dt}M_n = \sum_{m=0}^{\infty} \frac{n(m+n-1)\,K^{(m)}(0)}{m!} M_{m+n-2}, \tag{6}$$

where the sum is only performed for $m + n - 2 \geqslant 0$.

From (6), it is clear that, for $n = 0$,

$$\frac{d}{dt}M_0 = 0, \qquad \text{yielding } M_0 = \text{a constant}, \tag{7}$$

so that

$$\frac{d}{dt}N_n = \sum_{m=0}^{\infty} \frac{n(m+n-1)\,K^{(m)}(0)}{m!} N_{m+n-2}. \tag{8}$$

We will now consider two cases, $K^{(m)} = 0$ for $m \geqslant 1$ (K is constant) and $K^{(m)} = 0$ for $m \geqslant 2$ (K varies linearly with x):

(a) K is a constant, so $K^{(m)} = 0$ for $m \geqslant 1$; Eq. (8) yields

$$\frac{d}{dt}N_n = n(n-1)\,KN_{n-2} \tag{9}$$

so that

$$\frac{d}{dt}N_1 = 0, \qquad \text{yielding } N_1 = \text{a constant} \tag{10}$$

and

$$\frac{d}{dt}N_2 = 2K, \qquad \text{yielding } N_2 = 2Kt, \tag{11}$$

if $N_2 = 0$ at $t = 0$.

Equations (10) and (11) indicate that the centroid of the distribution remains stationary while the variance increases at a rate equal to twice the diffusivity, K.

For the particular case where c has constant shape, in the sense that c can be described by a single (time-varying) length scale L, with $N_n = A_n L^n$ (where A_n are dimensionless constants), Eqs. (9) and (11) yield $N_n = (n-1)\,N_2 N_{n-2}$. This is the relationship satisfied by the moments of a normal

(Gaussian) distribution, with the origin of x chosen so that $N_1 = 0$. The relationship therefore indicates the well-known and important result that a spreading normal distribution is a fundamental solution of the diffusion equation (2), with constant diffusivity K.

(b) K varies linearly with x, so $K^{(m)} = 0$ for $m \geqslant 2$; Eq. (8) yields

$$\frac{d}{dt} N_n = n(n-1)\, K(0)\, N_{n-2} + n^2 K' N_{n-1} \tag{12}$$

so that

$$\frac{d}{dt} N_1 = K', \qquad \text{yielding } N_1 = K't, \tag{13}$$

if the origin of x is chosen appropriately.

Also

$$\frac{d}{dt} N_2 = 2K(0) + 4K' N_1,$$

yielding (using Eq. (13))

$$N_2 = 2K(N_1)\, t, \tag{14}$$

if $N_2 = 0$ at $t = 0$.

The second moment of the distribution of c relative to its centroid (to be used in Section 3) is, from Eqs. (13) and (14),

$$N_2 - N_1^2 = 2K(N_1/2)\, t. \tag{15}$$

Equation (13) indicates that, in the presence of a linear variation of diffusivity, K, the centroid of the distribution (N_1) moves with a velocity K'. The second moment of the distribution, N_2 relative to the origin, or $N_2 - N_1^2$ relative to the centroid, increases in a manner similar to Eq. (11), except that K must be evaluated at N_1 (Eq. 14)) or $N_1/2$ (Eq. (15)), respectively, rather than at the origin.

3. THE RANDOM WALK TECHNIQUE AND THE ADDITION OF AN APPARENT ADVECTION VELOCITY

The simplest random walk model moves each particle by an alternating sequence of "advective" and "diffusive" steps. The advective step may be computed from a first or higher order integration based on the local velocity field (e.g., [15]), or from analytic integration when the functional form of the velocity field is known (e.g., in a logarithmic layer). As discussed earlier, the diffusive step is generally implemented by adding, to the spatial position of each particle, a random vector derived from a defined distribution function, P. Since the diffusion equation (2) is linear in c, it is helpful to consider the concentration distribution at any one time as being composed of a set of delta functions, each corresponding to one particle. P should be chosen so that its moments approximate those derived from the integration of Eq. (8) over the range $t = (t_1, t_1 + \delta t)$, with initial conditions $N_m(t_1) = 0$ $(m \geqslant 1)$, t_1 being the start of the time step, δt.

Section 2 implies that, for the case of constant K, P should ideally be a normal distribution. However, the central limit theorem (e.g., [12, p. 351]) indicates that, for any general form of P, the concentration field resulting from an initial delta function approaches a normal distribution after a few time steps (where a "few" will be discussed more fully in Section 5). Since such a distribution is completely described by its first and second moments, the exact form of P is unimportant so long as its moments satisfy Eqs. (10) ($N_1^* =$ a constant) and (11) ($N_2^* = 2K\,\delta t$), where the asterisk denotes, specifically, the moments of P rather than those of the total concentration, c. Commonly used forms of P are therefore a pair of equal delta functions (the well-known "drunken man's walk" with random step $\pm(2K\,\delta t)^{1/2}$, henceforth termed the "two-step" distribution) and a uniform ("top hat") distribution, of extent $\pm(6K\,\delta t)^{1/2}$.

For the case of spatially varying K, we will assume that K is linear, so that the moments of P should be chosen to satisfy Eq. (12), integrated over the range $t = (t_1, t_1 + \delta t)$, with initial conditions $N_m(t_1) = 0$ $(m \geqslant 1)$. The first moment should therefore satisfy (from Eq. (13))

$$N_1^* = K'\,\delta t. \tag{16}$$

The second moment should satisfy (from Eqs. (14) and (15))

$$N_2^* = 2K(N_1^*)\,\delta t \tag{17a}$$

or

$$N_2^* - (N_1^*)^2 = 2K(N_1^*/2)\,\delta t. \tag{17b}$$

The centroid of P (defined by N_1^*) should therefore be *advected* with a velocity K' in the direction of *increasing* diffusivity. It is this correction that is necessary to prevent particles concentrating in regions of low diffusivity. The correct relationship for the second moment depends on the way in which the apparent advection velocity (K', above) is applied. If the new particle position is chosen from a single distribution function, P, which has first and second moments given by Eqs. (16) and (17), respectively, then Eq. (17a) should be used to derive N_2^*. However, it is generally simpler to move the particle in two steps, one

advective and the other diffusive. In this case the advective step should consist of a simple translation of the particle with a velocity K' (which itself increases the variance by $(N_1^*)^2 = (K'\,\delta t)^2$), and the diffusive step should involve a random walk using a distribution with variance $N_2^* - (N_1^*)^2$, given by Eq. (17b).

For moments higher than the second, and any general field of K, it is evident that the correct form of P is *not* a normal distribution. This is clear from the simple example of a field which includes points at which K approaches zero. A diffusing patch must always be bounded by such points and hence cannot be normal in form. We must therefore invoke the central limit theorem and assume that moments higher than the second are unimportant.

The requirement for an apparent advection velocity presents a paradox. Expansion of Eq. (2) yields

$$\frac{\partial c}{\partial t} = Kc'' + K'c'. \qquad (18)$$
$$\text{(i)} \quad \text{(ii)}$$

It could be considered (erroneously) that, if term (ii) were omitted from Eq. (18), the solution would be that simulated by using the "naive" method of Section 1, with a local diffusivity given by Eq. (1). It would then appear that term (ii) represents the adjustment required to yield the correct solution. However, term (ii) is equivalent to an advection velocity of $-K'$, in the direction of *decreasing* diffusivity—i.e., in the opposite direction to the required velocity determined above. The paradox may be resolved by considering the moments of the concentration distribution described by

$$\frac{\partial c}{\partial t} = Kc'', \qquad (19)$$

where K varies linearly in space.

A similar analysis to that described in Section 2 yields

$$\frac{d}{dt} N_1 = 2K', \qquad \text{yielding } N_1 = 2K't, \qquad (20)$$

if the origin of x is chosen appropriately, indicating that the centroid moves *twice as fast*, in the case of Eq. (19), as in the case of the "correct" diffusion equation (2). The apparent advection velocity contained in term (ii) in Eq. (18) is simply the difference between the centroid velocity, K', for Eq. (2) and the centroid velocity, $2K'$, for Eq. (19). The paradox is hence resolved by appreciating that Eq. (19), with spatially variable K, is *not* the equation that is simulated by the "naive" method of Section 1.

4. A TRANSFORMATION THAT YIELDS A CONSTANT DIFFUSIVITY

Equation (17) indicates one way of implementing a random walk model with spatially variable diffusivity. An alternative method is to remap the spatial coordinate to yield a transport equation with constant diffusivity. The transformation is shown here for the case in which the diffusivity varies in only one dimension; extension of the mapping to cover variability in the other dimensions is not simple. This present technique is therefore suitable for the modelling (in one, two, or three dimensions) of boundary layers in which the diffusivity varies in only one direction. Including the effect of advection, the transport equation is

$$\frac{\partial c}{\partial t} = \frac{\partial}{\partial x}\left(K\frac{\partial c}{\partial x} - uc\right), \qquad (21)$$

where u is an advection velocity and both K and u may vary with x.

We now define a new coordinate system, $X(x)$, and a transformed concentration, γ. Mass conservation on transformation is ensured if

$$\gamma\, dX = c\, dx. \qquad (22)$$

Equation (21) becomes

$$\frac{\partial \gamma}{\partial t} = \frac{\partial}{\partial X}\left(K(X')^2\frac{\partial \gamma}{\partial X} + KX''\gamma - uX'\gamma\right), \qquad (23)$$

where primes continue to indicate differentiation with respect to x.

In order to obtain a constant diffusivity in the transformed system we set $K(X')^2 = 1$ (the value of the constant on the right-hand side being unimportant). Hence, the transformed coordinate is given by

$$X = \int \frac{dx}{K^{1/2}}, \qquad (24)$$

so that Eq. (23) becomes

$$\frac{\partial \gamma}{\partial t} = \frac{\partial}{\partial X}\left(\frac{\partial \gamma}{\partial X} - \frac{\gamma}{K^{1/2}}\left(\frac{K'}{2} + u\right)\right), \qquad (25)$$

and γ is given by

$$\gamma = cK^{1/2}. \qquad (26)$$

The diffusive term in Eq. (25) may be solved by a random walk technique with constant (unit) diffusivity. The transformation yields an apparent advection velocity which is

$K'/(2K^{1/2})$ in X-space (equivalent to $K'/2$ in x-space or *half* the apparent velocity that was shown, in Section 3, to be required if the transport equation is solved in x-space). The advection velocity, u, is scaled by $K^{-1/2}$ in X-space.

The physical meaning of the above transformation may be understood by considering that, after an elapsed time δt, an initially infinitesimally small patch reaches a size, given by Eq. (17), of

$$N_2^{1/2} = (2K\,\delta t)^{1/2}, \tag{27}$$

where the asterisk has been dropped to indicate the properties of an actual concentration distribution. The transformation scales the spatial coordinate by $K^{-1/2}$, which is proportional to the above patch size, thereby effectively yielding a spatially constant diffusivity.

We will illustrate the technique with two examples in which the diffusivity varies linearly in space. The first case involves a continuous source and the second an instantaneous release, both at the "left hand" boundary ($x = 0$). For both examples, there are simple analytic solutions within the semi-infinite domain $x = (0, \infty)$. The numerical domain is limited to $x = (0, 4)$. The examples, in which the numerical predictions are compared with analytic solutions, are as follows:

(a) *Continuous release.* We assume that $K = ax + b$, where a and b are constants, $u = 0$, and the substance flux in the positive x-direction is F, so that

$$(ax + b)\frac{\partial c}{\partial x} + F = 0. \tag{28}$$

The steady-state solution to Eq. (28) (which requires that F is constant) is

$$c = c(0) - \frac{F}{a}\ln\left(\frac{ax}{b} + 1\right). \tag{29}$$

In the random walk technique, the transformation is, from Eq. (24),

$$X = \frac{2(ax + b)^{1/2}}{a}, \tag{30}$$

where the constant of integration has been arbitrarily chosen to be zero.

The apparent advection velocity ($K'/(2K^{1/2})$ as defined by Eq. (25)) may be integrated analytically. Neglecting the stochastic diffusive step, a particle trajectory is given by

$$\frac{dX}{dt} = \frac{a}{2(ax + b)^{1/2}} = \frac{1}{X} \tag{31}$$

which, on integration, becomes

$$X_2^2 = X_1^2 + 2(t_2 - t_1), \tag{32}$$

where the subscripts, 1 and 2, denote the beginning and the end of the advective step, respectively.

Figure 1 shows a simulated solution, run to approximately steady state from initial conditions of zero concentration throughout the model. Particles were added at the left-hand boundary ($x = 0$) and removed if they reached the right-hand boundary ($x = 4$). P was chosen to be a top hat distribution. The following values were used:

$a = 1$, $b = 0.1$,

$F = 500$ particles per unit time,

elapsed time = 100,

time step = 0.1.

The particle concentration was computed by dividing the model into adjacent cells of equal size and counting the numbers of particles in each cell. The model contained 40 cells. The constant in Eq. (29) was evaluated by equating the analytic and simulated concentrations at the right-hand side of the model (the centre of the extreme right-hand cell, at approximately $x = 4$). If the expected particle concentration in cell i, of volume V_i, is $\langle c_i \rangle$, then the expected number of particles in the cell is $V_i\langle c_i \rangle$. The actual particle concentration in the cell follows a binomial distribution of standard deviation, s_i, given by

$$s_i = \langle c_i \rangle \left(\frac{1}{V_i\langle c_i \rangle} - \frac{1}{N}\right)^{1/2}, \tag{33}$$

where N is the total number of particles in the model (e.g.,

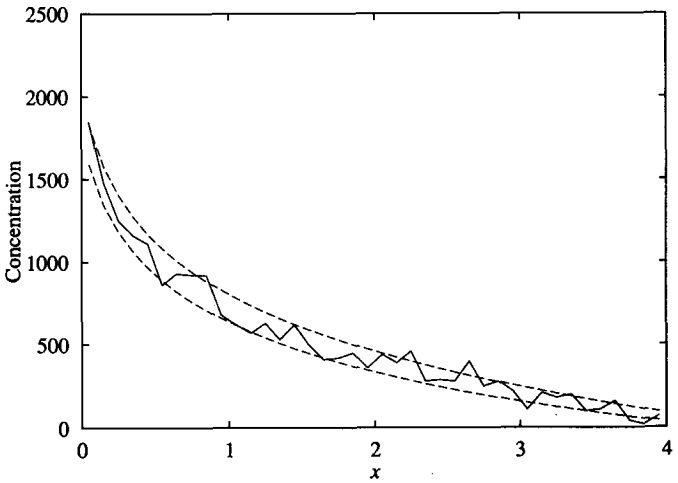

FIG. 1. The simulated and analytic concentration due to a continuous release. The dashed curves indicate one standard deviation on either side of the analytic solution. The continuous curve indicates the simulated concentration.

[9]). The dashed curves of Fig. 1 indicate one standard deviation of the expected concentration distribution on either side of the analytic solution.

The continuous curve indicates the simulated concentrations. It is evident that these are generally within one standard deviation of the analytic solution with only three values (7.5% of the total) falling outside two standard deviations (the expected number being approximately 1.8 or 4.6% of the total), indicating satisfactory agreement.

(b) *Instantaneous release.* As in the previous example, we assume that $K = ax + b$, where a and b are constants, and $u = 0$. A time-dependent solution of Eq. (2), for an instantaneous release at $x = -b/a$ and $t = 0$, is

$$c = \frac{Q}{at} \exp\left(-\frac{x + b/a}{at}\right), \qquad (34)$$

where Q is the total number of particles in the range $x = -b/a$ to $x = \infty$. The diffusive flux is zero at $x = -b/a$, where the diffusivity is zero.

In the random walk technique, the transformation is given by Eq. (30), and the advection step was evaluated analytically using Eq. (32). It is clear from Eqs. (25) and (34) that the numerical solution cannot include the point $x = -b/a$, where K is zero. A zero flux (reflective) condition was hence imposed at the "left-hand" boundary ($x = 0$) and b/a was chosen to be sufficiently small (0.001) to obtain a good approximation to the analytic solution. At the "right-hand" boundary ($x = 4$), a further zero flux condition was imposed. This latter condition (which is strictly incompatible with the solution given by Eq. (34)) had little effect on the numerical simulation, which was terminated after only a few particles had reached this boundary.

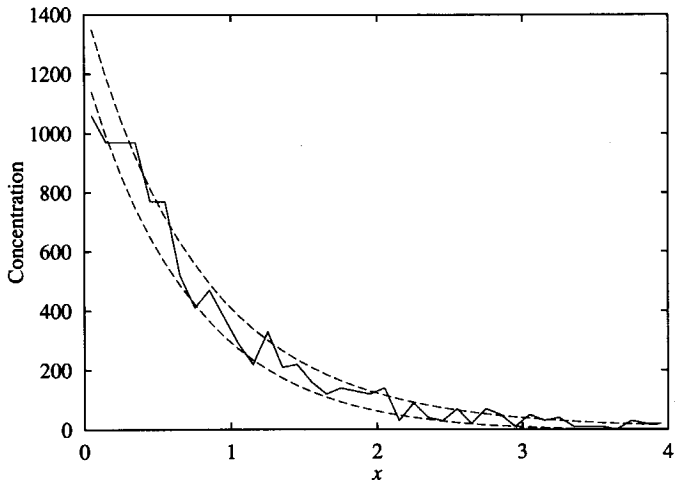

FIG. 2. The simulated and analytic concentration due to an instantaneous release. The dashed curves indicate one standard deviation on either side of the analytic solution. The continuous curve indicates the simulated concentration.

Figure 2 shows a simulated solution in which Q particles were initially released at the left-hand boundary ($x = 0$). P was chosen to be a top hat distribution.

The following values were used:

$a = 1$, $b = 0.001$,
$Q = 1000$ particles,
elapsed time $= 0.75$,
time step $= 0.01$.

The concentration was computed as in the previous example, and, again, the model contained 40 cells. The dashed curves indicate one standard deviation of the expected concentration distribution (as calculated in the previous example) on either side of the analytic solution.

The continuous curve indicates the simulated concentrations. Again, there is good agreement, with the simulated concentration being generally within one standard deviation of the analytic solution with only one value (2.5% of the total) falling outside two standard deviations (the expected number being approximately 1.8 or 4.6% of the total).

5. THE CHOICE OF DISTRIBUTION PRESCRIBING THE RANDOM STEP

As was discussed in Sections 2 and 3, for the case of constant diffusivity K (or of the constant transformed diffusivity of Section 4), P should ideally be a normal distribution. However, the central limit theorem ensures that, after a certain number, R, of time steps, δt, the exact form of P is unimportant, so long as the first and second moments (N_1^* and N_2^*) satisfy certain relationships (Eqs. (10) and (11)), as discussed in Section 3. Commonly used forms are hence the two-step and the top hat distributions. However, it should be noted that this does not mean simply (as is often assumed) that the predictions are satisfactory after a time $R\,\delta t$ from the model start time, but rather that the model can only accurately resolve time scales greater than or equal to $R\,\delta t$, instead of δt. The minimum resolvable time scale has therefore been increased by a factor, R. The use of a simple two-step or top hat distribution is therefore only reasonable if it involves a saving in computation time by at least a factor of R, relative to that required by a normal distribution (the "correct" form for the case of constant diffusivity, or if the transformation of Section 4 is used). We here estimate the value of R and the computational overhead involved in the use of a normal distribution, rather than the simple two-step or top hat forms.

The implementation of a single time step of a random walk model, with constant diffusivity and hence a constant

distribution function, $P(\delta x, \delta t)$, is equivalent to the convolution of the concentration, c, by P (e.g., [4, p. 39]):

$$c_{i+1}(x) = \int_{-\infty}^{\infty} c_i(x+\delta x)\, P(\delta x, \delta t)\, d\,\delta x, \qquad (35)$$

where the subscript i denotes the time level $i\,\delta t$. The central limit theorem indicates that, if the convolution kernel, P, is applied R times to a delta function, the result is approximately a normal distribution, $G(x)$, with a variance equal to RN_2, where N_2 is the variance of P. The difference between the resultant and a normal distribution may be described by the normalised standard deviation, s_R, given by

$$s_R = \left(\frac{\int_{-\infty}^{\infty} (c_R(x) - G(x))^2\, dx}{\int_{-\infty}^{\infty} G^2(x)\, dx} \right)^{1/2}, \qquad (36)$$

where c_0 is a delta function. The value of s_R may therefore be determined, given a required level of agreement, using Eq. (36).

Equation (36) was evaluated numerically for both two-step and top hat distributions and for a range of R. The results are summarised in Table I. It was found that agreement with a normal distribution at about the 10% level requires R values of 3 and 2, for two-step and top hat distributions, respectively. At the 1% level, the R values are 21 and 13, respectively.

The computational overhead on using a normal distribution, instead of a two-step or top hat distribution, was estimated for a system-supplied random number generator and others described by Press *et al.* [14]. The subroutines (see Table II) were timed on a Sun Microsystems SPARCstation 1. The random number routines RAND, RAN, RAN0, RAN1, RAN2, and RAN3 (henceforth termed the "uniform" routines) return a value from a uniform distribution in the range 0 to 1. Derivation of the required random step using a two-step distribution requires a single call of one of these routines and a conditional statement; a top hat distribution requires one call of the random number routine, one subtraction and one multiplication. The normal distribution algorithm (GASDEV), from [14]) employs (generally) a single call of one of the above routines followed by a transformation to convert from a uniform to a normal distribution; a further multiplication yields the required random walk step. The time for the generation of a single random step in a model is therefore dominated by the execution time of the random number routine used to generate the uniform or normal variate. Timing tests of GASDEV, using each of the six uniform routines, has indicated that the overhead factor, defined by

$$\frac{\text{(Time for one call of normal algorithm)}}{\text{(Time for one call of uniform algorithm)}}$$

is 2.0 ± 0.4 (Table II). Hence it only takes about twice as long to generate a random number from a normal distribution as it does from a uniform distribution. The overhead on generating a model step from a normal distribution, rather than from a two-step or top hat distribution, would involve a similar ratio.

The R values of Table I indicate the ratio (typically 10) by which the time step may be increased when using a normal distribution, and since the computational overhead factor is only about two, we conclude that it is significantly more efficient to use a normal (rather than two-step or top hat) distribution in random walk models.

TABLE I

Normalised Standard Deviation, s_R, between Normal Distribution and Distributions Formed by R Applications of Two-Step or Top Hat Convolutions

R	s_R (two-step)	s_R (top hat)
2	0.144	0.0887
5	0.0449	0.0270
10	0.0218	0.0132
20	0.0108	0.00651
50	0.00428	0.00258

TABLE II

Random Number Generators Used in Tests

Name	Distribution	Source	Method	Time (μs)
RAND	Uniform	Sun Micro-sytems	Non-linear additive feedback	9.8
RAN	Uniform	Ref. [14]	Single linear congruential	15.5
RAN0	Uniform	Ref. [14]	Random shuffle of sequence of random numbers	32.6
RAN1	Uniform	Ref. [14]	Triple linear congruential	43.8
RAN2	Uniform	Ref. [14]	Similar to RAN0	20.9
RAN3	Uniform	Ref. [14]	Subtractive method, after [10]	9.2
GASDEV	Normal	Ref. [14]	Transformation of uniform distribution	23.8 (using RAND) 31.2 (using RAN) 53.0 (using RAN0) 67.1 (using RAN1) 37.9 (using RAN2) 23.0 (using RAN3)

6. THE CHOICE OF RANDOM NUMBER GENERATOR

Not all numerical random number generators are suitable for use in random walk models. The choice of a satisfactory random number generator should involve two main criteria, first that it should be computationally efficient and second that it should return values that are sufficiently random in the sense that diffusive processes may be simulated accurately. We will consider each of these criteria in turn.

Table II indicates that the computation time to generate a random number can vary widely (over factors of 4.8 and 2.9 for the uniform and normal distribution routines, respectively). Since the calculation of random numbers can occupy a significant portion of the total computation time of a random walk model (especially if processes other than diffusion (e.g., decay) are also modelled stochastically), it is worthwhile finding an efficient algorithm.

Algorithmic random number generators never generate truly random numbers. Since they involve numbers containing a finite number of digits, one pattern of output values is exactly repeated over a finite cycle length (typically 2^{32} sequential values). In addition, patterns exist over significantly smaller "scales." For example, the least significant bit of the output of some random number generators alternates between 0 and 1 over sequential values. Such patterns are of importance in random walk models, which typically involve a minimum of 10^3 particles and 10^3 time steps and hence at least 10^6 calls of a random number routine. The potential problem is illustrated by considering the case of the simulation of a single particle over a number of steps equal to the cycle length. At the end of the simulation, if the distribution is truly uniform, there will be *exactly* the same number of steps of $+D$ (say) as steps of $-D$, so that the sum of all the steps will be zero and diffusion is effectively

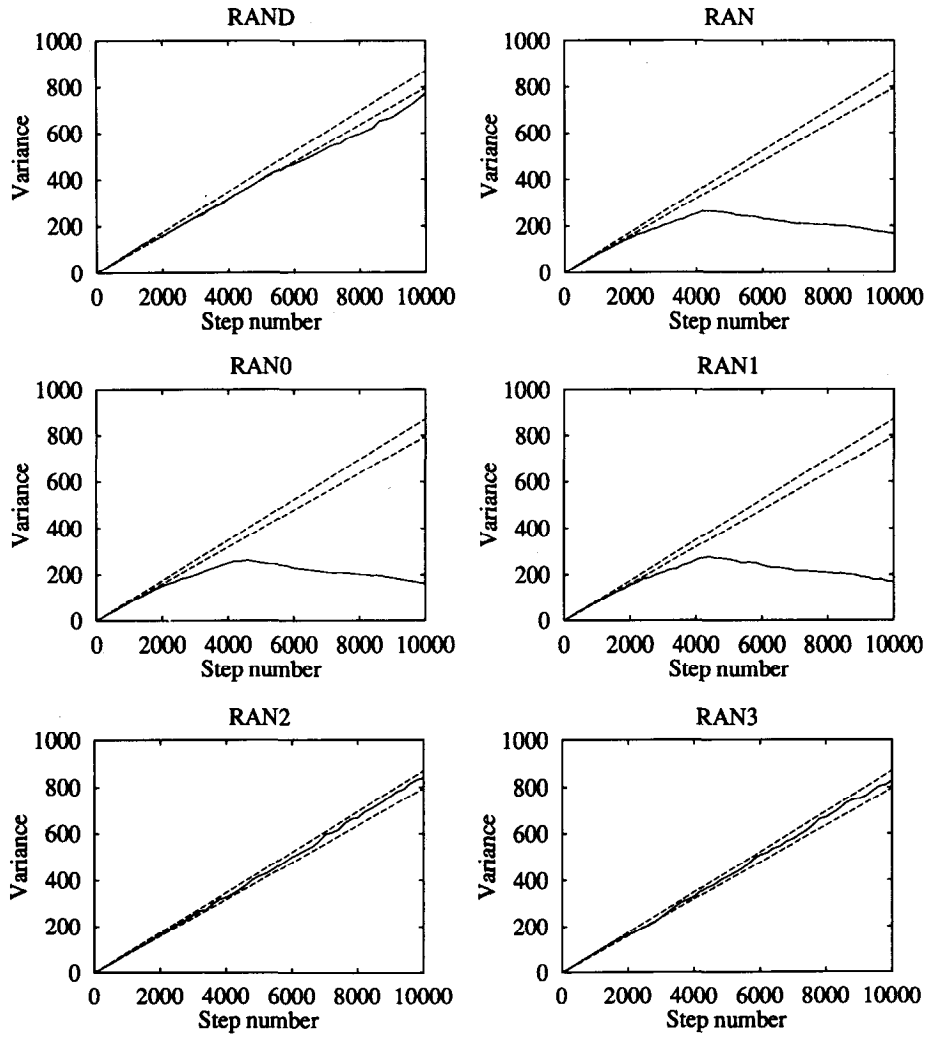

FIG. 3. Results of the tests of random number generators, Configuration A (time selected in inner program loop). The dashed lines indicate one standard deviation on either side of the theoretical variance. The continuous curves indicate the results of the various simulations.

absent over this period. In a model that utilises a sequence of random numbers that may even be considerably shorter than the cycle length, the effect of correlations is sometimes manifest by a patch of particles initially spreading at approximately the expected rate, but then shrinking in size at later times. A simple test of the suitability of a random number routine is therefore to compare the rate of increase of variance of a patch of modelled particles with the theoretical relationship (e.g., Eq. (11), for the case of constant diffusivity). Any distribution function may be used, since the exact form of the distribution does not affect the rate of increase of variance and does not mask the performance of the underlying random number generator.

Figures 3 and 4 show the patch variance as a function of time for simulations involving 10^3 particles, 10^4 time steps, no advection, $K = \frac{1}{24}$ and $\delta t = 1$. The simulation therefore involves 10^7 calls of the random number generator. From Eq. (1), the required distribution should have a variance equal to $\frac{1}{12}$, which is provided by a top hat distribution with limits ± 0.5. The variance should theoretically increase at a rate equal to $\frac{1}{12}$ per time step. The random number generators tested were those listed in Table II.

The estimated patch variance is given by

$$\sum_{i=1}^{n} x_i^2,$$

where x_i is the coordinate of the ith particle and there are n particles. The expected spread of the patch variance may be derived from the properties of the χ^2 distribution, which has a variance of $2n$ and a mean of n (n being the number of degrees of freedom). The ratio of the standard deviation of the patch variance to the mean patch variance is therefore $(2/n)^{1/2}$. The dashed lines in Figures 3 and 4 indicate plus

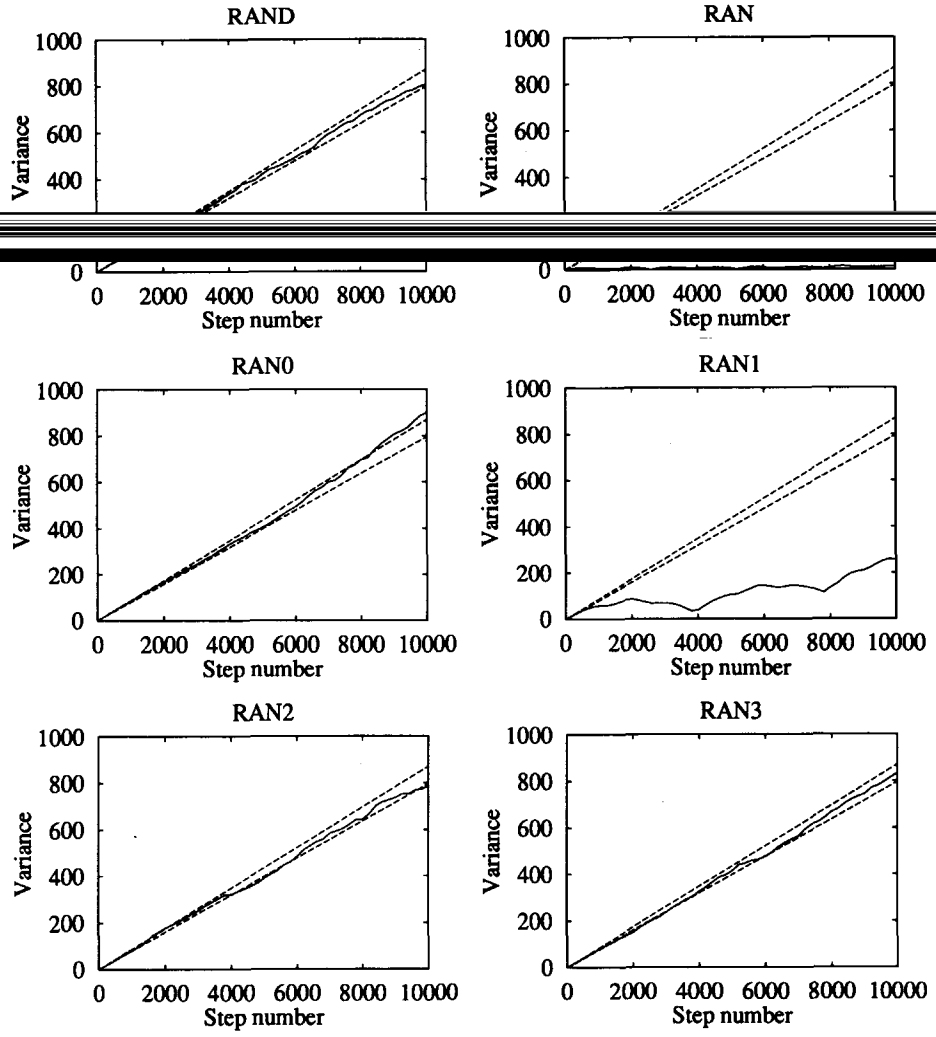

FIG. 4. Results of the tests of random number generators, Configuration B (time selected in outer program loop). The dashed lines indicate one standard deviation on either side of the theoretical variance. The continuous curves indicate the results of the various simulations.

and minus one standard deviation from the theoretical variance, or

$$\frac{t}{12}\left(1 \pm \left(\frac{2}{n}\right)^{1/2}\right),$$

where t is the time, or time step number. The continuous curves indicate the results of the various simulations.

Figure 3 shows results for a model in which time is stepped forward in the inner program loop and the particles are selected in the outer loop. Hence the sequence of operations (termed "configuration A") is:

Step forward particle (1) over duration of simulation,
step forward particle (2) over duration of simulation,
step forward particle (3) over duration of simulation,
etc.

Figure 4 shows results for a model in which time is stepped forward in the outer program loop and the particles are selected in the inner loop. Hence the sequence of operations (termed "configuration B") is:

Step forward all particles over time step (1),
step forward all particles over time step (2),
step forward all particles over time step (3),
etc.

According to this test, the performance of random number generators can be very poor and may depend on which configuration is chosen (A or B). In particular, routines RAN, RAN0, and RAN1 behave poorly with configuration A, while only RAN and RAN1 behave poorly with configuration B. Both RAND in configuration A and RAN0 in configuration B produce results which differ by more than a standard deviation from the expected variance.

Of the random generators tested, RAN3 is the most computationally efficient and performs satisfactorily over simulations involving 10^7 random numbers in either configuration A or B.

7. THE CHOICES OF TIME STEP AND OF TOTAL NUMBER OF PARTICLES

The optimal time step is closely related to the dominant length and time scales of the modelled system. In the transformed coordinate system, the total effective velocity may be defined by

$$U = K^{-1/2}\left(\frac{K'}{2} + u\right). \tag{37}$$

Three length scales can be defined (in the transformed coordinate system):

(a) the required spatial resolution, L_R,

(b) the length scale of the spatial variability of the total effective velocity, $L_U = U/(\partial U/\partial X)$, and

(c) the length scale of the spatial variability of the decay law, L_D.

The following time scales can then be defined:

(i) the time scale of the variability of the prescribed concentration boundary conditions,

(ii) the Eulerian time scale of the total effective velocity field, $U(\partial U/\partial t)^{-1}$,

(iii) the Lagrangian time scale related to the required spatial resolution, L_R/U,

(iv) the Lagrangian time scale related to the spatial variability of the total effective velocity, L_U/U,

(v) the Lagrangian time scale related to the spatial variability of the decay law, L_D/U, and

(vi) the diffusive time scale related to required spatial resolution, L_R^2 (noting that the effective diffusivity in the transformed coordinate system is unity).

Two examples of the relevance of time scale (iii) are as follows: First, if particles are injected into a current of velocity U, then at least one particle must be added every L_R/U time units in order to preserve a spatial resolution of L_R. Second, if the inter-particle spacing is appproximately equal to the required spatial resolution, L_R, then a particle will be advected, by a velocity U, across an open boundary every L_R/U time units. In each case, the addition of particles or the application of a boundary condition is required with a time scale of L_R/U.

If advection in the total effective velocity field can be described by an analytic solution, then the time scale (iv) is no longer a restriction. Alternatively, the effect of (iv) may be partially removed by defining variables, such as the diffusivity, as analytic functions within limited regions (e.g., in a piecewise linear fashion), such that analytic solutions exist for the particle trajectories in each region. It is common to employ a constant decay law, in which case the time scale (v) becomes infinite, therefore presenting no constraint. The time scale (vi) relates mainly to the time between boundary collisions caused by diffusion, requiring the application of some boundary condition (e.g., reflection).

The choice of time step is hence dominated by the time scales of the forcing of the system by the boundary conditions (i) and the velocity field (ii), and by time scales defined by the required spatial resolution (iii) and (vi). The constraints afforded by the spatial variabilities of the total effective velocity (iv) and of the decay law (v) may often be overcome.

The power of the transformation of Section 4, combined with the use of a normal distribution for step generation, may be illustrated by the following example. Consider a patch of substance released instantaneously into spatially

varying, but steady, fields of velocity and diffusivity. The boundaries are far from the patch and hence do not influence its spreading. The decay law is spatially constant and advection may be simulated by an analytic solution. In this simplistic, but not unreasonable, example none of the time scales (i) to (vi) limit the model time step; the model can employ an arbitrarily long single time step to predict *exactly* (subject to stochastic variations) the concentration field at any time in the future.

The choice of number of particles is closely related to the spread of simulated concentrations given by Eq. (33), which may be recast in terms of the proportional error for cell i,

$$\frac{s_i}{\langle c_i \rangle} = \left(\frac{1}{N} \left(\frac{1}{\alpha_i} - 1 \right) \right)^{1/2} \simeq (\alpha_i N)^{-1/2}, \qquad (38)$$

where $\alpha_i N$ is the expected number of particles in cell i and α_i are a set of (small) dimensionless constants that may be estimated in retrospect, or from prior approximate solutions to the problem in hand. Such solutions may be obtained from a simplified "back-of-the-envelope" model. Once the maximum allowable value of $s_i/\langle c_i \rangle$ (for any cell) is chosen, the minimum value of N, the total number of particles, may be chosen from Eq. (38).

8. CONCLUSIONS

In random walk models where the diffusivity varies spatially, the naive method (employing only a random step with a spatially variable variance, equivalent to the value of N_2 given by Eq. (1)), yields incorrect simulations (Section 3). The error is equivalent to the omission of an apparent advection velocity, in the direction of increasing diffusivity and of magnitude equal to the diffusivity gradient. The problem may be rectified in one of two ways, either by the addition of this apparent velocity or through the use of a spatial transformation that yields a constant diffusivity (and another apparent velocity) in the transformed coordinates (Section 4).

With the former method, the first and higher moments are modified by the finite diffusivity gradient (see, for example, Eq. (12)). The modification of the first moment is simply equivalent to the presence of the apparent advection velocity. Neglect of errors in the third, and higher, moments is only justified by recourse to the central limit theorem, which has been shown to involve a penalty in temporal resolution (Section 5).

The latter method, however, involves no approximation; further, it is applicable to any diffusivity field, provided that the reciprocal of the square root of the diffusivity is integrable (Eq. (24)) and that the diffusivity varies in only one direction. It is therefore the preferred method for the modelling of boundary layers.

In the case of a constant diffusvity (or the use of the transformation of Section 4), the normal distribution is the preferred distribution from which to choose the random steps (Section 5). The large number of model steps involved in a single simulation will frequently exceed the full cycle of a particular random number generator, in extreme cases leading to negligible diffusion. The random number generator should therefore be tested over a large number of diffusive steps, in the simple manner described in Section 6.

Two important aspects of model design are the selection criteria for the time step and the total number of particles (Section 7). The choice of time step is often dominated by the time scales of the forcing of the system and by the required spatial resolution. The proportional error of the predicted concentration is given by the square root of the inverse of the expected number of particles in a model cell (Eq. (38)). A "back-of-the-envelope" forecast, or a retrospective check, will indicate whether the number of particles is sufficient to achieve the required accuracy.

ACKNOWLEDGMENTS

We thank Peter McIntosh and Stephen Walker for their helpful comments on this manuscript. The project was partially funded by a grant from the Australian Research Council.

REFERENCES

1. S. W. Ahlstrom, *A Mathematical Model for Predicting the Transport of Oil Slicks in Marine Waters*, (Pacific Northwest Laboratories, Battelle, 1975).
2. C. M. Allen, *Proc. R. Soc. London A* **381**, 179 (1982).
3. I. Bork and E. Maier-Reimer, *Adv. Water Res.* **1** (3), 161 (1978).
4. G. T. Csanady, *Turbulent Diffusion in the Environment* (Reidel, Dordrecht, 1973).
5. J. W. Dippner, *Continental Shelf Res.* **10** (1), 87 (1990).
6. A. J. Elliott, *Dtsche. Hydrogr. Z.* **39**, *H* **3**, 113 (1986).
7. A. E. Hill, *Mar. Ecol. Prog. Ser.* **75**, 39 (1991).
8. J. R. Hunter, in *Oceanology International 80, Brighton, U.K., 1980*, p. 42.
9. J. R. Hunter, in *Numerical Modelling: Applications to Marine Systems*, edited by J. Noye (Elsevier Science, North Holland, New York 1987), p. 257.
10. D. E. Knuth, *The Art of Computer Programming*, Vol. 2 (Addison–Wesley, Reading, MA, 1981), Sections 3.2–3.3.
11. E. Maier-Reimer and J. Sundermann, in *Engineering Application of Computation Hydraulics*, edited by M. B. Abott and J. A. Cunge (Pitman, London, 1982), p. 198.
12. F. Mosteller, R. E. K. Rourke, and G. B. Thomas, *Probability with Statistical Applications* (Addison–Wesley, Reading, MA, 1970).
13. J. Noye, in *Numerical Modelling: Applications to Marine Systems*, edited by Noye (Elsevier Science, North Holland, New York, 1987), p. 231.
14. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes* (Cambridge Univ. Press, Cambridge, UK, 1986).
15. D. Ramsden and G. Holloway, *J. Comput. Phys.* **95**, 101 (1991).
16. P. C. Rothlisberg, J. A. Church, and A. M. G. Forbes, *J. Mar. Res.* **41**, 511 (1983).