Superprocesses and plankton dynamics

Robert Adler

Faculty of Industrial Engineering and Management, Technion, Haifa, Israel
Department of Statistics, University of North Carolina, Chapel Hill

Abstract. This paper is a first introduction to superprocesses for non-probabilists. They are presented so as to suggest natural models for the growth and motion of large plankton populations in either stable or flowing ocean environments.

November 30, 1997

Introduction

For the sake of academic honesty, I should start this paper with a disclaimer: Although the title contains both the keywords "superprocesses" and "plankton dynamics", and I do know something about the former, my knowledge of the latter, and of biological oceanography in general, starts and stops at the level of the elementary textbook Lalli and Parsons (1993). This would not generally be considered a level of expertise that would allow an author to start writing on a topic. Nevertheless, it does seem to me that there is a natural connection between the two topics, and that both are likely to benefit from interdisciplinary cross-fertilisation. This point was emphasised, for me, by a number of talks at the *Aha Huliko'a* meeting, and I shall return to this and some general philosophy on stochastic modelling in the closing, and what may be the most relevant, section of the paper.

Since cross-fertilisation must start somewhere, you should think of this paper as a mathematical σ in search of its matching oceanographic $\circleddash$, and forgive what may seem at first like an oversimplistic view of plankton dynamics.

The paper is structured as follows: In Section 1 we describe what are generally known as the "particle pictures" preceding and motivating most work in superprocesses. The particle pictures will, for us, descibe the reproduction and temporal-spatial spread of plankton communities. In Section 2 we look at what happens when the number of plankton generations, and the size of each generation, tend to infinity. Thus we shall define superprocesses, which will be the infinite limit of the particle pictures, and require the language of stochastic partial differential equations to be properly described.

In Section 3 we shall describe some basic geometric properties of superprocesses, and in Section 4 how all of this might relate to plankton dynamics. As one might expect, a theory developed by abstract probabilists is going to have to undergo a little work before it can be tailored to applications.

In the Section 5 we shall describe some variations on the basic model, including the way interactions between plankton can be modelled and, more importantly, how to place the entire superprocess structure within a randomly moving frame of reference such as an oceanographic flow. The final Section 6 has the promised general comments on stochastic modelling, motivated by the superprocess models of the remainder of the paper.

It does not seem to make too much sense in an introductory paper of this kind to worry too much about assigning detailed credits for the structures and results that will be described. Since the interested reader will have to go elsewhere for missing details, (s)he can search for references at the same time. Without doubt, the best place to start is with the Saint Flour lectures of Dawson (1993), which has a properly credit-assigned and almost encyclopædic treatment of superprocess accurate up to the time of its writing.

1. Particle pictures

We shall start by describing a model for generic particles, which, for the remainder of this paper, we shall think of as plankton, moving in the Euclidean space $\Re^d$, $d \geq 1$. The most interesting case for applications is obviously the three dimensional space $d = 3$, but it costs us nothing to work in wider generality. While moving, our plankton will also reproduce

(asexually) and die. The birth-death process will be such that the overall plankton population will remain of a more or less stable size, although overall linear death and growth rates will be permissable.

Thus, suppose that at time $t = 0$ we have $K > 1$ plankton, distributed over space according to a Poisson point process with control measure m, so that K is a Poisson random variable with mean $|m| = m(\Re^d)$, and the numbers of plankton in disjoint regions $A_1, \ldots, A_k$ are independent Poisson variables with means $m(A_1), \ldots, m(A_k)$. A special case that we shall use for the simulations shown below will be when m is a measure with all its mass concentrated at one point. In this case, we start with K plankton all at one point.

(It is important to note that this initial setup can be changed quite radically, without having much of an affect on the overall structure that we are describing. However, in order to make this exposition simple and short, we shall work throughout with specific, simplifying assumptions. We shall have more to say about these generalities in Section 5.)

Each of these K plankton follows the path of independent copies of a Markov process B, until time $t = \rho/\mu$.

At time ρ/μ each plankton, independently of the others, either dies or splits into two, with probability $\frac{1}{2}$ for each event. The individual plankton in the new population then follow independent copies of B, starting at their place of birth, in the interval $[\rho/\mu, 2\rho/\mu)$, and the pattern of alternating critical branching and spatial spreading continues until there are no plankton left alive. (This happens with probability one, since the process which simply counts the total number of plankton alive is a critical branching process, which ultimately dies out.)

The process of interest to us is the measure valued Markov process

$$X_t^\mu(A) = \frac{\{\text{Number of plankton in } A \text{ at time } t\}}{\mu}, \quad (1)$$

where $A \in \mathcal{B}^d =$ Borel sets in $\Re^d$. Note that, for fixed t and μ, X_t^μ is an atomic measure. Note also that if $\rho = \infty$ there is no branching occurring.

Now suppose we set $K = \mu$ and $\rho = 1$, and send $\mu \to \infty$. Fixing ρ is simply a matter of scaling time. Sending $K \to \infty$ involves a high density limit, in which the number of plankton is high. Sending $\mu \to \infty$ means that X_t^μ measures the spatial dispersion of plankton after a large number (μ) of generations.

It is now well known that under very mild conditions on the motion process B the sequence $\{X^\mu\}_{\mu \geq 1}$ converges (weakly, on an appropriate Skorohod space) to a measure valued process which is called the *superprocess* for B. Before we look at this limit process, and attempt to characterise it, let us look at some simulations.

Figure 1 shows what happens to $K = 1,000$ plankton that begin life together at the center of a square and slowly spread out according to a Brownian motion. (Actually, since this is on a computer, the "Brownian motion" in this case must be discrete, so that it is really a random walk. However, we shall ignore this fact in the following.) There is no branching in this example, so that the spread, shown after each plankton has taken 50, 100 and 250 steps, is in the spherical fashion usually associated with random walks. The left hand graphs are the contour lines for the particle density plots in the right hand graphs.

In Figure 2 we have added the birth/death process that leads to the superprocess. That is, intermingled with the steps of the random walk are birth/death events, at which particles either split into two or die, as described above. The three sets of diagrams show the extent of the plankton spread after the same number of random walk steps as in the corresponding diagrams of Figure 1. Note however, that the spread is no longer purely spherical. Furthermore, note how disjoint, small "communities" of plankton develop, a phenomenon that we shall discuss in more detail later.

(More examples of the above kind are available, either as photographic stills, an interactive computer program, or a video movie, in Adler (1994a-c).)

Three dimensional versions of Figures 1 and 2 are easy to generate, and exhibit similar phenomena, but are are little less easy to appreciate on the two dimensional, uncoloured, page.

2. The superprocess

We now look a little more carefully at what happens in the high density, multiple generation situation, when the parameters K and μ of the particle picture tend to infinity together.

To make the notation easier, we shall restrict interest for the moment to the case in which B is a d-dimensional Brownian motion, and return to the more general case in Section 5. Here, then, is one way to describe the limiting process:

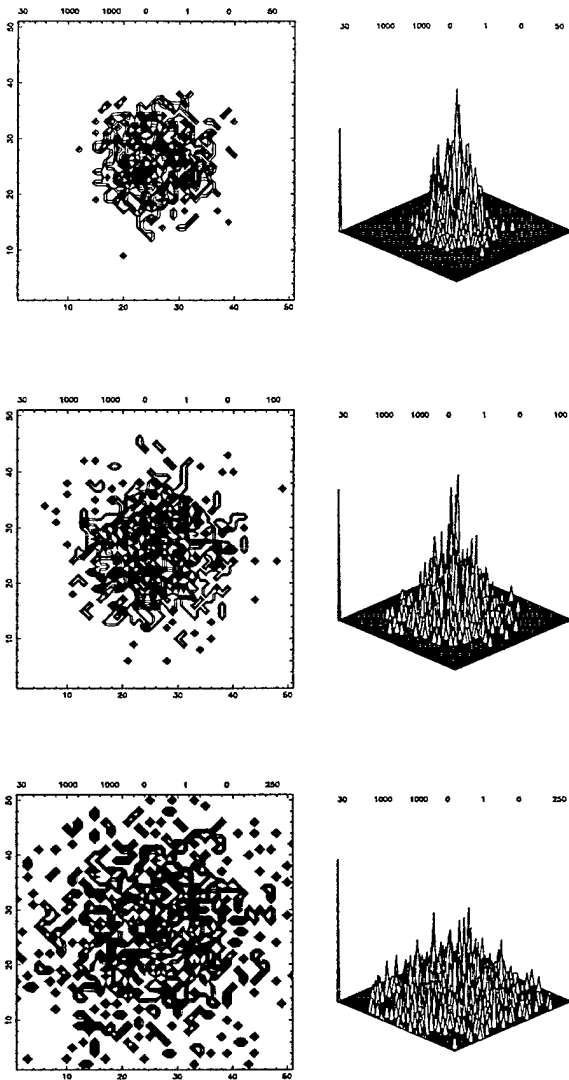Suppose that the limiting measure $X_t = \lim X_t^\mu$
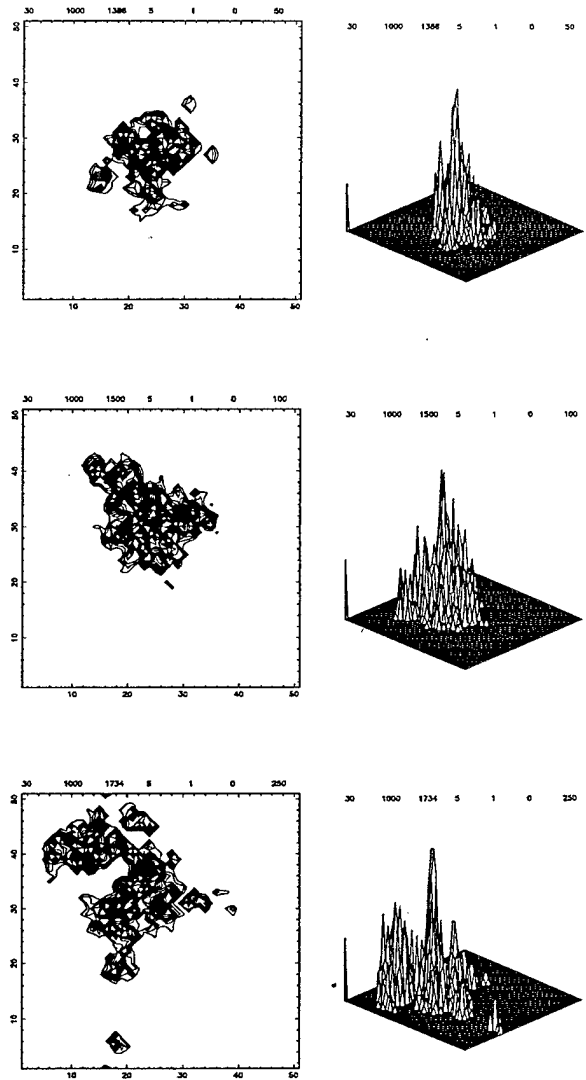
Figure 1. Random Walks

Figure 2. Superprocesses

has a density $f(x,t)$, $x \in \Re^d$, $t \geq 0$. Then f would solve the stochastic partial differential equation

$$\frac{\partial f(x,t)}{\partial t} = \sum_{i=1}^{d} \frac{\partial^2 f(x,t)}{\partial x_i^2} + \sqrt{f(x,t)}\xi(x,t)$$

$$= \Delta f(x,t) + \sqrt{f(x,t)}\xi(x,t), \qquad (2)$$

where ξ is a space-time Gaussian white noise.

In fact, this description does make sense in one dimension, but not when $d \geq 2$, for then the limiting superprocess does not have a density. However, an integrated version of (2) does make sense. So take a

nice test function ϕ and with the notation

$$\int \phi(x)f(x,t)\,dx \equiv \int \phi(x)\,X_t(dx) \equiv X_t(\phi)$$

rewrite (2) as

$$X_t(\phi) = \int_0^t X_s(\Delta\phi)\,ds + \int_0^t \int_{\Re^d} \phi(x)Z(dx,dt), \quad (3)$$

where Z is a martingale measure with $Z_0 = 0$ and quadratic variation given by

$$\langle Z.(\phi)\rangle_t = \int_0^t X_s(\phi^2)\,ds, \qquad (4)$$

all of which is a way of writing (2) when the density f does not exist. Below we shall write (3) in shorthand as

$$\dot{X}_t = \Delta X_t + \dot{Z}_t. \qquad (5)$$

This equation is, in part, familiar to all of us. If the birth/death branching phenomenon was not present in our plankton story, then the martingale term Z_t in (5) would disappear, and so the equation would "degenerate" to the standard heat equation. This, of course, is what is to be expected from the strong law normalisation in (1), which, in the limit, yields the measure describing where any one of our plankton can be expected to be at time t.

Details of how to properly formulate all the weak convergences used above can be found in Walsh (1986), where one can also find a precise formulation of the stochastic PDE's. The original result is due to Watanabe (1968), and a good general reference on related subjects is Ethier and Kurtz (1986). However, as mentioned in the Introduction, Dawson (1993) is an all-inclusive reference.

3. Geometry of superprocesses

What has made superprocesses interesting to probabilists has, in part, been their interesting and delicate geometric properties. We shall describe only the most basic of these here, which relate to the sets shown in the left hand columns of Figures 1 and 2.

Since, for each $t \geq 0$, $X_t(\cdot)$ is a random measure, it is carried by a random support set S_t: roughly speaking, S_t is defined as the smallest closed set for which

$$X_t(S_t) = \max_{A \in \mathcal{B}^d} X_t(A),$$

where $\mathcal{B}^d$ is the collection of Borel sets in $\Re^d$; i.e. S_t is the set of "sites" of the particles alive at time t.

Most of the known interesting geometrical results about superprocesses are related to the dimensional properties of S_t. For example, the following, in which "$\dim(A)$" refers to the Hausdorff dimension of the set A, holds with probability one, for each $t > 0$:

$$\dim(S_t) = \min(2, d). \qquad (6)$$

This result can also be made uniform in t.

Furthermore, if we define the "range" of X (which amounts to those points ever visited by the particles during the time interval $(0, t]$) by

$$R = \lim_{\epsilon \downarrow 0} \left\{ \text{closure} \left(\bigcup_{t > \epsilon} S(X_t) \right) \right\},$$

then the dimension of the range is $\min(4, d)$.

These two basic facts lead to an entire theory of interesting sample path behaviour for superprocesses. Unfortunately, however, very little of it is likely to be of use in the study of plankton dynamics.

4. Plankton dynamics

The main issue of this short paper is not, of course, to discuss superprocesses *per se*, but rather whether or not they have something to offer as models of plankton dynamics. For this, results about Hausdorff dimension, while elegant, are mere mathematical curiosities. What is far more important, in comparing theory with experiment or data, is knowing things like the distribution of the number and sizes of clusters in the support of a superprocess.

Unfortunately, these are random variables about which very little quantitative is known. The main reason behind this is the following very basic result, which gives the Laplace transform of the marginal distributions of a superprocess:

If X_t is the superprocess of (3,4) then

$$E \exp\{-X_t(\phi)\} = \exp\{-\int u(x, t) X_0(dx)\}, \qquad (7)$$

where X_0 is the (usually non-random) starting measure of the superprocess and $u(t, x)$ is the solution of the non-linear PDE

$$\frac{\partial u(x, t)}{\partial t} = \Delta u(x, t) - u^2(x, t) \qquad (8)$$

$$u(x, 0) = \phi(x). \qquad (9)$$

This is not an easy equation. No explicit solution is available, and only general scaling properties (as $t \to \infty$ or $\|x\| \to \infty$) of the solution are known. Since (8), even at its best, is information about a Laplace transform, direct, quantitative information about actual distributions is almost impossible to find.

Some general, qualitative results are known however. For example, the clustering phenomenon can be described by a compound Poisson process, by which we mean that the one can represent a superprocess as a superposition of clusters. The number of clusters in any given size range is Poisson distributed, but the parameter of this Poisson variable is not easy to compute. Furthermore, the properties of a "typical cluster" are known only in the most general of terms. There is very little that could be computed explicitly so as to compare with a real plankton flow.

Of course, all of this is a problem only for pure mathematicians. A natural, practical way to compare theory with the real world is easily available precisely the way that Figures 1 and 2 were generated; *viz.* simulation and Monte Carlo.

If, as we believe might well be the case, a marriage between superprocesses and plankton dynamics is in the offing, Monte Carlo estimation of superprocess distributions will undoubtedly play the rôle of matchmaker.

5. More general models

Since the previous sections all concentrated on plankton that either die or split in two with equal probabilities, and then follow Brownian motions, if the suggested project of applying superprocesses to plankton dynamics is to succeed, it will be necessary to entertain far more sophisticated models of motion and reproduction. We shall now descibe some of these, and the way in which they will affect the basic stochastic partial differential equation (SPDE)

$$X_t(\phi) = \int_0^t X_s(\Delta\phi)\,ds + \int_0^t \int_{\Re^d} \phi(x)Z(dx,dt), \quad (10)$$

with its associated martingale quadratic variation

$$\langle Z.(\phi)\rangle_t = \int_0^t X_s(\phi^2)\,ds. \quad (11)$$

The easiest of all model assumptions to change is that of the branching structure, although, given the asexual nature of plankton reproduction, this is one of the least worrisome to us. For example, instead of giving each plankton a 50/50 chance of disappearing or splitting in two at each generation, we can allow each plankton to replaced by a random number, say N, as long as N has mean one and finite variance σ^2. The only affect on the limiting SPDE would be to place a factor of σ^2 before the integral on the right hand side of (11). (In fact, even infinite variance, which would allow for the occassional very large family size, can be incorporated into the particle picture, but then the structure of the basic SPDE (10) changes quite significantly in that the limiting martingale becomes discontinuous.)

One can also move away from a mean family size of one for the branching, albeit not too far away. For example, there are no serious difficulties if the mean number of offspring is such that $EN = 1 + a/\mu$, where μ was the parameter tending to infinity in the basic

story and a is a finite – positive or negative – constant. In this case (10) changes in that a term of the form

$$a \int_0^t X_s(\phi)\,ds \quad (12)$$

must be added to the right hand side. In fact, the "constant" a can even be taken to be time and spacially varying, in which case the additional term (12) becomes

$$\int_0^t \int a(s,x)\phi(x)\,X_s(dx)ds \quad (13)$$

One application of such a space/time varying reproduction rate would be to allow for models in which certain geographical regions, or seasons, impact on the rate and/or success of reproduction

Of course, these changes are really very small at the level of the total population, since the increase or decrease in mean offspring size is very small for each plankton. This, however, is essential. For while this small change per plankton can be quite significant at the level of the total population (allowing, for example, the solution to (10) to grow in a linear fashion in time) a larger change, such as allowing each plankton to produce an average of two offspring per generation, would lead to immediate explosion of the limiting model, and so not provide a model that makes sense for a (roughly) stable population.

The model also allows for immigration and emigration. The effect here, which can, again, be time/space dependent, eventually amounts to adding a term of the form $\int_0^t \int a(s,x)\phi(x)\,dxds$ to (10), where a now describes the migration.

Finally, we can introduce interaction between the plankton in a way that will affect the branching rate. There are a number of ways to do this: plankton may increase their mean rate (again proportionally to $1/\mu$) in the presence or absence of neighbours. A model of this kind will lead to additional terms in (10) of the form

$$\int_0^t \int \int \phi(x)a(x,y)\,X_s(dx)X_s(dy)ds \quad (14)$$

where $a(x,y) > 0$ would imply an increase in the mean offspring number at x due to particles at y, etc. Interacting models in which the variance, rather than the mean, number of offspring can also be handled. In this case the change appears in the quadratic variation (11).

This covers, roughly, the type of changes that can be made to the branching structure of the basic model. The other factor that can be changed is the motion of the plankton. In general, it is trivial to replace the Brownian motion with that of any Markov process. The effect that this has on (10) is to replace the Laplacian Δ appearing there by the "infinitesimal generator" of the Markov process describing the motion. If, for example, the Brownian motion is replaced by an elliptic diffusion, then Δ will replaced by a second order differential operator.

Another interesting interaction involves allowing the particle motions some dependence. There are two ways to go about this: One involves interactions between the particles themselves, which would require the plankton having some "knowledge" of their neighbours' whereabouts, and modifying their motion according. This is quite difficult to handle mathematically, and so we shall not attempt to go into details here. The effect on (10) however involves the addition of a new, non-linear, term.

A more recent model, under study by Adler and Skolakis (1997), is somewhat more interesting in terms of plankton dynamics, and looks at the basic model placed within a random flow. In other words, the plankton can move and reproduce according to any of the above perturbations of the basic model, but in addition they live within a random flow, such as the ocean, that accounts for global motion. Some simple models of this form have been studied by Çinlar and Kao (1992) and Zirbel and Çinlar (1996).

Although the above describes new models with gay abandon, it is important to note that while new models are easy to generate at the particle level, they are generally at least an order of magnitude harder to handle at the limiting stochastic PDE level than is the basic superprocess. This is particularly true if any interactions are involved. Whether or not this is important from the point of view of plankton dynamics is not clear. Since, as pointed out in the previous section, very few of the variables that are of interest from this viewpoint are likely to be amenable to direct mathematical attack even in the simplest of cases, the additional mathematical complexity of the more sophisticated and more realistic models should not be an additional drawback.

6. On stochastic modelling

Most of the material in the preceding five sections was written in preparation for, and thus before, the *Aha Huliko'a* meeting at which it was presented. Questions that were raised during the talk found their way into the text via some rewriting. Thus the presentation was much the same as the preceding sections, although much less detailed.

However, as a response to some of the earlier talks presented at the meeting, a lively discussion arose as to what one might call "the rôle of white noise" in stochastic modelling, which I now want to address, using superprocesses as an example of a much wider class of phenomena.

The discussion arose after one speaker presented an oceanographic food chain model of some complexity. We all know the basics behind these models: A system is described by a collection of linked ordinary differential equations, in which (at the risk of gross oversimplification on my part) birth and death rates are combined into one (birth minus death) rate term, and predator-prey relationships by matching terms, but with opposite signs, in the equations for the overall size of each subpopulation in the chain.

The solution to a system of this kind is, of course, deterministic, since all stochastic variation has been replaced by mean behaviour in the ODE's. Since real systems of this kind are known to exhibit what seem to be random (or, at the very least, model-unpredictable) variations, the modeller's way of solving this problem was to add random noise terms to each of the ODE's, *without, in any way, tailoring the noise to the model.*

Furthermore, I should point out that this type of modelling was far from uncommon, and figured in a number of different models, of very different phenomena, throughout the meeting. Here, again grossly oversimplified, is a typical justification of this approach: "The ocean is really just Navier-Stokes equation, but we have no idea what the boundary conditions really are, and even if we did the equations would be too difficult to solve with available (and even foreseeable) computing power. Thus we approximate Navier-Stokes somehow, we approximate the boundary conditions, and then we add some white (or coloured) noise to indicate this approximation/lack of knowledge".

My claim is that it is *imperative* in this kind of (unavoidable) situation to give a great deal more thought to how the random noise itself should be modelled, and any relation it might have to the basic model.

As an example, consider the simplest model of

the superprocess. (The same kind of argument will apply to all the others.) Suppose that instead of building the model as we had, we used the following kind of argument: At splitting times, each particle either dies or splits into two with equal proababilities. Thus there is no change in the mean population size. So, for the moment, we might as well forget the branching phenomenon. Thus, we have only motion. Since we are averaging this over many particles, ultimately we will replace the superprocess stochastic PDE (2)

$$\frac{\partial f(x,t)}{\partial t} \;=\; \Delta f(x,t) + \sqrt{f(x,t)}\xi(x,t), \quad (15)$$

by the simple heat equation:

$$\frac{\partial f(x,t)}{\partial t} \;=\; \Delta f(x,t). \quad (16)$$

In other words, the stochastic term will have been lost, and we will have an equation with an extremely smooth solution. Since we know that in the real world, whether "real" means the ocean or the numerical experiment, there is stochastic variation, the kind of philosophy described above of merely adding a stochastic noise term will lead to an equation of the form

$$\frac{\partial f(x,t)}{\partial t} \;=\; \Delta f(x,t) + \xi(x,t), \quad (17)$$

where ξ is the same space-time white noise appearing in (15), but, as a result of the way randomness was ascribed to the model, which does not have the multiplying factor of $\sqrt{f}$. This is really quite different from the correct equation (15), and the differences are significant. Among them are the facts that:

- Whereas the solution to (15) is always positive, as one would expect of a particle density, the solution to (17) will often be negative.

- Whereas the solution to (15) will exhibit the kind of clustering phenomena that, for mathematicians, are what makes the process challenging, and, for biological oceanographers, may make the process interesting, the solution to (17) exhibits no such behaviour.

- The solution to (17) is Gaussian, which makes it readily amenable to mathematical analysis, and which therefore gives it a certain advantage over the solution of (15). On the other hand, the mathematical complexity resulting from the

non-Gaussian nature of the solution of (15) allows for the phenomena of larger than Gaussian variation in subpopulation (cluster) sizes that mimics what is known to happen in real branching phenomena.

The bottom line that comes out of this example is that the mere addition of random noise at a late stage to correct a model that has been oversimplified at an earlier stage cannot generally be expected to reproduce missing phenomena, neither quantitatively nor, more importantly, qualitatively.

Exactly how should this affect the existing stochastic models of physical and biological oceanography? I do not know, but, if I did, that would make me an oceanographer, rather than merely a meddling mathematician...

References

A word on the references: As mentioned in the Introduction, the key reference is Dawson (1993), which contains more than a practitioner will ever want to know about superprocesses and other measure valued diffusions. The other references appear either because they contain more recent material that is useful in the context of this exposition or, particularly in the case of references to my own work, because as pictures, or movies, or computer programs, they might be useful aids to a novice wanting to get a idea of whether or not superprocesses are the "right" modelling tool for him/her.

Adler, R.J. (1994a) Superprocesses – the particle picture, *Measure-Valued Processes, Stochastic Partial Differential Equations, and Interacting Systems, CRM Proceedings and Lecture Notes*, 5, ed. D. A. Dawson, Centre de Recherches Mathématiques, Université de Montreal, 1-15.

Adler, R.J. and Gross, A. (1994b) *SUPER*, an interactive C/GL graphics package for the visualisation of superprocesses in dimensions one through five. Available from the programmers.

Adler, R.J. (1994c) *Superprocesses – The Movie.* A 15 minute video of superprocess visualisations. Available from the producer.

Adler, R.J. and Skolakis, G. (1997) Superprocesses over random flows, in preparation.

Çinlar, E. and Kao, J.S. (1992) Birth and death on a flow. *Diffusion processes and related problems in*

analysis, Vol. II (Charlotte, NC, 1990), 121–137, *Progr. Probab.*, **27** Birkhäuser Boston, Boston.

Dawson, D. (1993) Measure valued Markov processes, in *Ecole d'Eté de Probabilités de Saint-Flour, XXI. Lect. Notes Math.* **1541** Springer, 1–260.

Ethier, S. and Kurtz, T. (1986), *Markov Processes*, Wiley, New York.

Lalli, C.M. and Parsons, T.R. (1993) *Biological Oceanography: An Introduction.* Butterworth-Heinemann, Oxford.

Walsh, J. (1986) An introduction to stochastic partial differential equations, *Lecture Notes in Math.* **1180** 265–437, Springer Verlag.

Watanabe, S. (1968), A limit theorem for branching processes and continuous state branching processes, *J. Math. Kyoto Univ.* **8** 141–167.

Zirbel, C.L. and Çinlar, E. (1996) Dispersion of particle systems in Brownian flows, *Adv. Appl. Probab.* **28** 53–74.