From the microscale to the macroscale

Alan McKane

Theoretical Physics Division, University of Manchester

Brest, August 20, 2013

OUTLINE

- Introduction
- Creating stochastic models
- Analysis of stochastic models
- Some applications
- Concluding remarks

1. Introduction

Ecology was probably the first area of biology where quantitative models were constructed. These early models were formulated at the population level, frequently as differential equations which were practically the only tool available to solve them at the time.

The advent of electronic computers led to a huge increase in the variety and complexity of models which could be studied, most naturally formulated at the individual level. However the analytic study of ecological models is still very much influenced by the traditional approach.

In this lecture I'll discuss how individual-based models (IBMs) can be studied not only numerically, but analytically as well.

I'll also argue that they are simpler to construct, more intuitive and describe effects that population-level models (PLMs) miss.

Three approaches to ecological modelling

Population level models

PLMs are typically formulated as ordinary differential equations for the fraction of the population that is of a certain kind. For example, Levins' original metapopulation model is written as

$$\frac{dQ}{dt} = cQ(1 - Q) - eQ,$$

where Q is the fraction of patches which are occupied and c and e are effective parameters.

This is a phenomenological equation, i.e. written down consistently, but not derived from a more basic model. The variables, such as Q, are continuous—implicitly assuming an infinite population size—and thus the results are deterministic (non-random).

The advantage of this approach is that the well-developed mathematics of dynamical systems is applicable. This is still considered by many to be the dominant methodology, with many textbooks available.

If one starts with this approach then the question of how to include stochastic effects arises. Some authors discretise the equations while others add stochastic (noise) terms on to the equations. Both of these approaches are ill-defined in general.

Agent based models

Agent based models have a finite number of individuals in the system, each typically having many attributes or degrees of freedom. These are frequently defined in terms of computer algorithms ready for simulation.

For instance, individuals in the model may have an age, a sex, a location, and so on. The term IBM is occasionally used to describe such models, and more particularly to mean computer-based ecological models where individuals have many attributes, but we will not use this terminology here.

The advantages of models of this type are the potentially unlimited amounts of detail that can be included and the ease of programming. The main disadvantage is that algorithmic approach can only produce numerical results—in general they are too complex for mathematical analysis.

Individual based models

IBMs are also composed of a finite number of individuals, but typically with a smaller number of attributes. The rules for the time evolution of the individuals within the system are formulated probabilistically, usually as a Markov process.

State variables are therefore discrete instead of continuous and by definition the model is stochastic. In an IBM version of the metapopulation model above, n would be the number of patches and the dynamics of colonisation and extinction of patches are random events that depend on n

Usually the population is divided up into a small number of classes, reducing the degrees of freedom, thus all members of the same class are indistinguishable.

The probabilistic formulation of IBMs means they can be simulated exactly with Monte Carlo techniques, but most importantly they can be analysed analytically via the master equation. The corresponding PLM to a given IBM can also be analytically derived.

2. Creating stochastic models

Beginning with IBMs, and deducing PLMs from them, is the approach taken in many of the physical sciences. In chemistry, for example, reaction kinetics at the molecular level yield rate equations at the macroscale and in physics, statistical mechanical models at the microscale lead to a thermodynamical description at the macroscale.

In ecological IBMs, random events—for example: birth, death and predation—at the level of individuals (the microscale) give rise to macroscopic dynamics for large populations of individuals.

In the limit of an infinite population these dynamics become deterministic and can be described by a set of ordinary differential equations: a PLM.

As we will discuss later, when the population is large but still finite, these dynamics are often still strongly stochastic. Before we do this, we first discuss the construction of IBMs and some methods of analysis.

Suppose we wish to describe an ecological system that contains m different species and the only way we are going to label individuals is as belonging to a given species.

If the population size is assumed to be finite, equal to N, then the state of the system is given by a vector of integers $\mathbf{n} = (n_1, \ldots, n_m)$, where n_i is the number of individuals of species i at that time.

This could also be the number of individuals of species i per unit area, or some other measure.

Since the basic idea of model specification can be illustrated on a system with just one species, let us focus on a system with n identical individuals of type A.

Suppose also that the only processes are death and (asexual) birth. We represent these as

$$A \xrightarrow{d} E \quad \text{and} \quad A + E \xrightarrow{b} A + A.$$

The first equation indicates that an individual of type A dies at a rate d to give a vacancy (E for ‘empty’) and the second, that if there is a vacancy, another individual of the same type is born at a rate b.

The rate at which the number of individuals reduces from n to $(n-1)$ is then

$$T(n-1|n) = dn,$$

and the rate at which it increases from n to $(n+1)$ is

$$T(n+1|n) = bn\frac{(N-n)}{N} = bn\left(1-\frac{n}{N}\right).$$

This last result assumes a well-mixed population and is proportional to the probability of an individual and a vacancy existing in the system. These are the fundamental ingredients that underpin the dynamics.

We should emphasise that these rates depend on N and thus so do the dynamics of the model. This is not true of a PLM. We now discuss a number of examples of IBMs of increasing complexity.

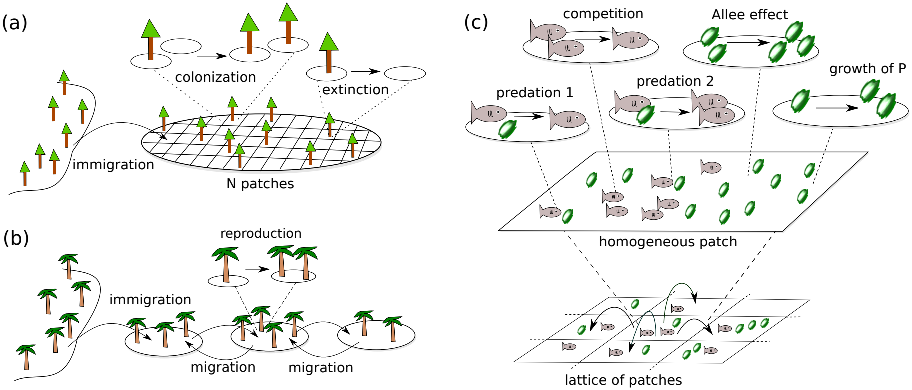

Three examples of ecological IBMs. (a) shows a simple metapopulation model. (b) shows a neutral island chain. (c) the Levin-Segal model, which describes the explicitly spatial plankton-herbivore dynamics.

Examples of ecological IBMs

An advantage of IBMs is that specifying them is simple and they are straightforward to biologically motivate, with additional features easily incorporated. Here we give a number of examples, three of which are illustrated in the figure.

Neutral models

The ‘classic’ neutral theory, introduced by Hubbell consists of a local community with birth/death processes of the Moran type and immigration from a metacommunity, which acts as a well-mixed source pool of potential immigrants.

Since all individuals are assumed to have the same birth and death rates, we may fix our attention on one species, A, and denote all other species as B. Then the transitions are:

(i) Death/Birth: $A + B \xrightarrow{(1-m)} A + A$; $A + B \xrightarrow{(1-m)} B + B$

(ii) Immigration: $A \xrightarrow{m} B$; $B \xrightarrow{m} A$

Here the rates do not include the combinatoric factor that gives the probability of choosing individuals of the desired species to interact.

The quantity m is the probability that the replacement of an individual is due to an immigration event, rather than by a birth/death event.

The probability of a particular immigrant being chosen is given by the relative abundance of that particular species in the metacommunity, which can also be described as an IBM with birth/death and speciation.

This model is only implicitly spatial, but it can be extended in many ways, for example by making it explicitly spatial, incorporating a chain of islands, or a network structure.

Metapopulation models

Metapopulations are sets of fragmented local populations connected by migration.

At the simplest level of description the patches can be thought of as being either occupied or unoccupied.

If A labels the occupied patches and E the unoccupied ones, then the possible events are:

(i) Colonization: $A + E \xrightarrow{c} A + A$
(ii) Extinction: $A \xrightarrow{e} E$
(iii) Migration: $E \xrightarrow{m} A$

This leads to a simple one-dimensional master equation, where the state variable is the number of occupied patches. The same transition scheme applies to many models such as logistic growth.

Again this simple model has been elaborated on, for instance by making the number of colonisable patches into a dynamic variable.

Predator-prey model

The IBM consists of n individuals of species A (the predators) and m individuals of species B (the prey)

The processes are taken to be

$$BE \xrightarrow{b} BB \quad \text{(birth)}$$

$$\begin{array}{llll} AB \xrightarrow{p_1} AA & & AB \xrightarrow{p_2} AE & \text{(predation)} \\ A \xrightarrow{d_1} E & & B \xrightarrow{d_2} E & \text{(death)} \end{array}$$

From these we can use combinatoric arguments to construct the transition rates

$$\begin{aligned} T(n-1,m|n,m) &= d_1\, n \\ T(n,m+1|n,m) &= 2b\left(\frac{m}{N}\right)(N-n-m) \\ T(n,m-1|n,m) &= 2p_2\left(\frac{n}{N}\right)m + d_2\, m \\ T(n+1,m-1|n,m) &= 2p_1\left(\frac{n}{N}\right)m \end{aligned}$$

The introduction of spatial variation introduces no new points of principle in the modelling procedure, although the analysis typically becomes more complex.

3. Analysis of stochastic models

Once a model has been specified, the complete stochastic dynamics are encoded analytically by the master equation, or individual realisations can be obtained by Monte Carlo simulation, which outputs a stochastic time series.

These are the two main ways of investigating and solving these models (that is, obtaining a complete description of the dynamics).

The stochastic simulation algorithm (SSA) due to Gillespie provides a method for the exact simulation of a master equation. This type of simulation is the usual way of investigating these models, mainly because of its simplicity to program.

However, another important point is that the SSA and master equation are both derived from the same underlying Markovian assumptions, so that there is an exact correspondence between the two.

In general, simulations provide individual realisations of a stochastic process.

The master equation (also known as Kolmogorov's forward equation) describes analytically the temporal evolution of the probability density of being in a particular state n, thus encoding the full stochastic dynamics.

If the master equation could be solved (for the probability density) then one would have a complete description of the properties and the dynamics of the stochastic system.

However, it cannot be solved analytically for most cases of interest, and so a number of approximation methods have been developed.

One calculation that is straightforward and can always be carried out is to determine the macroscopic behaviour from the master equation, which is the $N \to \infty$ limit, and which yields the PLM corresponding to the original IBM.

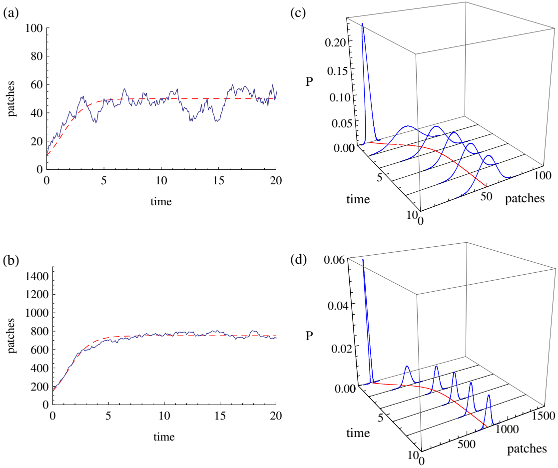

Emergence of macroscopic behaviour. The number of occupied patches for the metapopulation model for (a) $N = 100$, and (b) $N = 1500$ carrying capacity. The red dashed lines shows the deterministic result and the blue lines are from stochastic simulation.

The master equation and the emergence of macroscopic dynamics

The transition rates $T(n+1|n)$ and $T(n-1|n)$ define the model. From them the probability of finding n individuals in the system at time t, $P(n,t)$ can be calculated from the equation,

$$\begin{aligned} \frac{dP(n,t)}{dt} &= [T(n|n+1)P(n+1,t) + T(n|n-1)P(n-1,t)] \\ &- [T(n-1|n)P(n,t) + T(n+1|n)P(n,t)] \ . \end{aligned}$$

This is the equation that governs the dynamics of the system. It is the equation for a Markov chain in continuous time, and is called the master equation.

Although we have only discussed the case of a single species, all the formalism naturally generalises. We can replace n by $\mathbf{n}$ everywhere in the master equation where $\mathbf{n}$ is a vector with components which can represent any set of species in a large collection of spatial patches and having any number of other attributes.

The master equation contains far more information than the usual differential equations written down phenomenologically in terms of the population density.

The latter can be found simply taking the average of n/N; all the detail of the stochastic fluctuations is lost. Doing this for the metapopulation model described above gives the well-known logistic equation.

In the deterministic description the probability distribution, $P(n,t)$ is a spike at $n = \langle n \rangle$. To go beyond this we could simply replace the spike by a Gaussian.

The system would now be stochastic, being described by a probability distribution function, although of the simplest kind. The width of the Gaussian would be expected to scale like $1/\sqrt{N}$, and so we would effectively be writing

$$\frac{n}{N} = x + \frac{\xi}{\sqrt{N}},$$

where $x = \langle n \rangle / N$ and ξ is a stochastic variable.

This is the content of the van Kampen system-size expansion. It turns out to be an excellent approximation (far better than might naively be expected) away from boundaries.

Furthermore, since Gaussian probability distributions are derived from linear stochastic differential equations, the whole theory is linear and so can be treated exactly, even though the underlying stochastic processes may be highly non-linear.

For most interesting applications we need more than one degree of freedom, but the whole theory generalises in this case to a set of stochastic differential equations for ξ_i:

$$\frac{d\xi_i}{dt} = \sum_j \alpha_{ij}\xi_j(t) + \eta_i(t),$$

where $\eta_i(t)$ is a Gaussian white noise with zero mean, which is the only remnant of the demographic stochasticity of the system.

It has a correlation function given by a matrix β_{ij}. So the whole theory is given by two matrices α_{ij} and β_{ij} which can be systematically calculated from the transition rates which define the model.

Population level descriptions follow from individual level descriptions

Far more insight can be found by analytically deriving the corresponding PLM—a set of differential equations—from the IBM in the limit of infinite N, than postulating it on phenomenological grounds.

It is also helpful in writing down a PLM when the ecology of the system is complicated. In general it is easier to consider the events that define a system, write down an IBM, and then derive the corresponding PLM.

Unfortunately, often this procedure is reversed: a well-studied PLM is discretised, interpreted as an IBM, and then simulated to investigate various stochastic effects.

This is exactly the wrong way round: in general there will be many IBMs (sometimes infinitely many) giving a particular PLM, and so the PLM should always be derived from the IBM.

Many IBMs give the same PLM because IBMs define a stochastic process and PLMs are equations for the averages of this process, but many stochastic processes have the same averages.

4. Some applications

The modelling approach we have described has been applied to a number of problems over the last few years, and there is a clear trend towards the increasing use of this methodology.

For example, in the case of the neutral model discussed earlier, the quantity that has been studied most intensively is the species abundance distribution.

This is a plot of the number of species with a given number of individuals in a local community. It can be calculated exactly in the classic model starting from the master equation, and fits well with empirical data.

We will now discuss a number of other applications. For simplicity, we will restrict ourselves to cases where the van Kampen system-size expansion (also known as the linear noise approximation) is applicable.

Quasi-cycles

A common feature of some IBMs (with two or more degrees of freedom) is the ability of the demographic stochasticity to excite macroscopic scale coherent oscillations, known as quasi-cycles.

A pre-requisite for quasi-cycles to occur in a given IBM is that there is a stable fixed point in the corresponding PLM that is approached in an oscillatory manner.

Expressed mathematically, this means that the stability matrix at this fixed point must have at least one pair of complex eigenvalues.

However while the oscillations in the deterministic system are damped, and so die away, in the full (stochastic) system the demographic stochasticity acts as a forcing term and sustains the oscillations.

A physical analogy is a pendulum that is lightly damped, so eventually any oscillations will die away. If however it is randomly bombarded by weak white noise, these oscillations will be sustained.

In addition the amplitude of these oscillations will be much bigger than might naively be expected, because in amongst all the frequencies contained in the white noise will be one that corresponds to the natural frequency of the pendulum.

This resonance effect has been called stochastic amplification.

Quasi-cycles differ from limit-cycles in that they are a stochastic phenomena, and provide a simple robust mechanism that generates cyclic behaviour without the need for additional complexities, as are often assumed, to generate cyclic behaviour in deterministic PLMs.

One of the most powerful tools for analysing such oscillations is the power spectrum. This shows how the different frequencies that make up the time-series are distributed.

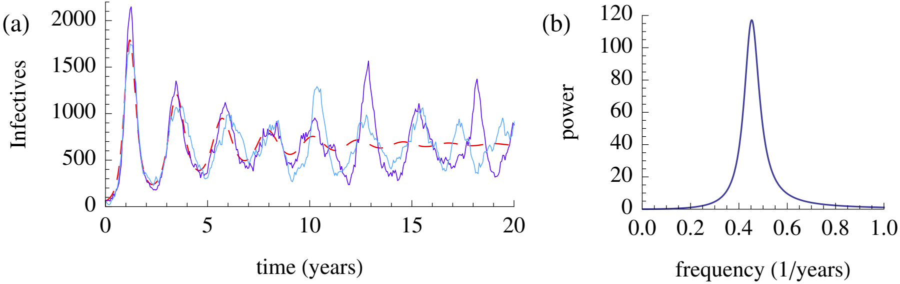

Quasi-cycles (a), and corresponding power spectrum (b). Two stochastic realisations of an SIR model are shown in (a) along with corresponding deterministic result (red dashed line). The cycles are maintained in the stochastic realisations, but can go out of phase.

The power spectrum

Quasi-cycles are stochastic in origin, and so do not have a single period, but a distribution of periods centred about an average value. This, together with the fluctuations in the amplitude of the cycles, mean that when the time series from a large number of realisations are averaged over, they average out to zero.

Yet a single realisation corresponds to what will be seen in an experiment, and so some averaging mechanism is needed which will not wipe out the cycles. The simplest way to achieve this is to take the Fourier transform of the time series:

$$\tilde{\xi}_i(\omega) = \int_{-\infty}^{\infty} \xi_i(t) e^{-i\omega t} dt$$

The frequencies centred around that corresponding to the average period will have the largest magnitude, and so taking the modulus squared of the Fourier transform and then averaging: $\langle |\tilde{\xi}_i(\omega)|^2 \rangle$, should give a smooth function spread about the characteristic frequency of the system.

This is the power spectrum of the fluctuations. The spread of the power spectrum will tell us how coherent the fluctuations are. If they were deterministic and sinusoidal, the power spectrum would be a spike.

The height of the peaks of the power spectra (reflecting the amplitude of the oscillations) are much larger than might be expected due to a resonance effect: all frequencies are present in the white noise and this resonates with the characteristic frequency of the system.

This characteristic frequency is not the same as the frequency of the decay of perturbations to the deterministic system. The latter is given by $|\mathrm{Im}\lambda|$, where λ is a complex eigenvalue of the Jacobian of fluctuations about the stationary state.

Instead the peak of the power spectrum is approximately at $\sqrt{(\mathrm{Im}\lambda)^2 - (\mathrm{Re}\lambda)^2}$, and the precise value also depends on the β_{ij}.

The connection of quasi-cycles with limit cycles, which are the conventional way of describing oscillations in nonlinear dynamical systems has been clarified. In addition there can be quasi-cycles about limit cycles.

Spatial patterns

In 1952 Turing predicted that systems which consisted of agents of at least two different species which reacted together and spatially diffused could, under certain circumstances, give rise to patterns, now known as Turing patterns.

However just as demographic stochasticity leads to stochastic oscillations where a deterministic analysis finds only static behaviour, it also leads to stochastic Turing patterns in spatial systems, where reaction-diffusion equations would indicate that none should exist.

This has been shown in a variety of systems, moreover the range of parameters for which stochastic Turing patterns exist is usually much larger than for conventional Turing patterns, which have restrictions such as the diffusion constants of the two species being at least an order of magnitude different to each other.

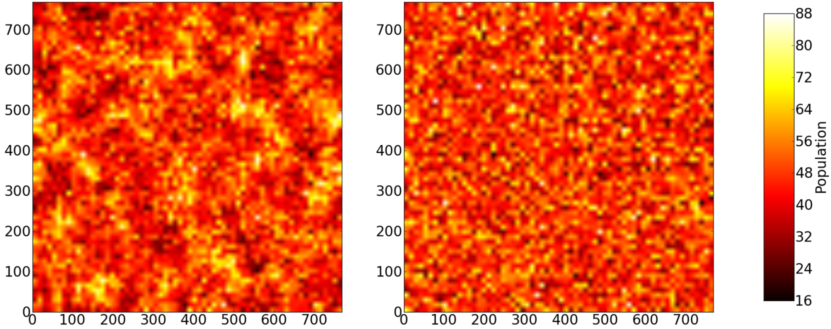

Quasi-pattern formation in the stochastic Levin-Segal model. The left panel shows a population map of the model in the quasi-pattern phase. The right panel show the same data, but randomised. If the patchiness were a statistical artefact then it would be preserved in the randomised version.

Recurrent epidemics

Individual based models have a long history in epidemiological modelling, beginning with Bartlett's pioneering work in the late 1950s.

The dynamics of childhood diseases such as measles, whooping cough and rubella provide particularly interesting case studies because they are subject to both external forcing, due to the aggregation of children in schools, and demographic stochasticity.

Previously, we discussed quasi-cycles about a deterministic fixed point. Here the forcing leads to a limit-cycle in the deterministic dynamics.

Now the noise excites the transient dynamics about the limit cycle, so the resulting macroscopic time series are a superposition of these and the deterministic limit cycle.

The relative proportions of the two periodicities reflected in the power spectrum then depend on the stability of the limit-cycle and the size of the system. Such an analysis can encompass and explain the differing periodicities of whooping cough and measles before and after mass vaccination.

5. Concluding remarks

Our aim in this lecture has been to highlight the trend towards models in ecology being individual based. There are several advantages in using stochastic models of this kind: they are easier to construct than PLMs, they are more intuitive, and they predict phenomena which deterministic models miss.

The importance of stochastic effects has long been known, going back to the earliest Monte-Carlo simulations of epidemic spreading and fade-out, but computational investigations have out-paced the theoretical understanding of such models.

Approaches based on the master equation are starting to remedy this and we have tried here to give an overview of the philosophy and methodology of these. We expect that future research will continue to develop techniques based on this starting point.

We would also expect that in future the trend will be towards whole classes of models which interpolate between these different traditions of model building.

Thus agent-based models would be simplified by construction of a chain of related models that would terminate with a relatively simple IBM, which in turn could be studied analytically and its deterministic limit compared with the differential equations that are generally postulated phenomenologically.

This lecture was based on the following article:

"Stochastic formulation of ecological models and their applications". Trends Ecol. Evol. **27**, 337 (2012).

It contains a large number of references and can be obtained from my website: `http://www.theory.physics.manchester.ac.uk/~ajm/`