J Comput Neurosci (2006) 20: 5–23
DOI 10.1007/s10827-006-3843-z

On the application of "equation-free modelling" to neural systems

Carlo R. Laing

Received: 9 March 2005 / Revised: 20 July 2005 / Accepted: 25 July 2005 / Published online: 20 February 2006

Abstract "Equation-free modelling" is a recently-developed technique for bridging the gap between detailed, microscopic descriptions of systems and macroscopic descriptions of their collective behaviour. It uses short, repeated bursts of simulation of the microscopic dynamics to analyse the effective macroscopic equations, even though such equations are not directly available for evaluation. This paper demonstrates these techniques on a variety of networks of model neurons, and discusses the advantages and limitations of such an approach. New results include an understanding of the effects of including gap junctions in a model capable of sustaining spatially localised "bumps" of activity, and an investigation of a network of coupled bursting neurons.

Keywords equation-free · bifurcation · model · microscopic · macroscopic

1. Introduction

Many physical systems involve the interaction of many "units"—particles, cells, molecules, random walkers, etc.—and the equations governing the local dynamics of these units and the interactions between them are often known in some detail. Many of these systems have "emergent" macroscopic behaviour, such as a convective roll in a fluid heated from below. Often, it is this behaviour that is of interest to us, rather than the detailed behaviour of the many individual units (water molecules, in this example). For some systems such as fluid flow, we can construct approximate equations that govern the macroscopic dynamics and work directly with them, ignoring the microscopic details. This is appropriate if we are interested in phenomena that occur on much larger spatial scales, and much longer time-scales, than those associated with individual molecules.

However, for many systems of interest it is not possible to derive such macroscopic equations. Equation-free (EF) modelling, developed in the past few years by Kevrekidis et al. (Gear et al., 2002; Kevrekidis et al., 2003; Makeev et al., 2002; Möller et al., 2005; Runborg et al., 2002) is a way of analysing the macroscopic equations for such a system, even though the equations are not known explicitly, by using short bursts of appropriately-initialised simulations of the microscopic dynamics. Its success relies on there being a separation of time-scales in the system, with the macroscopic variables of interest changing on a much longer time-scale than most of the microscopic variables. It is computationally intensive, with repeated simulations of often highly-detailed models on a microscopic level being required; this makes them amenable to implementation on parallel computers. The idea is similar to that of approximate inertial manifolds in the analysis of partial differential equations (Garcia-Archilla et al., 1998), where the amplitudes of higher modes (which, for example, have rapid spatial variation) are "slaved" to (or are functions of) the amplitudes of a finite number of lower modes (which, for example, have slow spatial variation). Thus a simulation of the lower modes, together with knowledge of the slaving relationship, is sufficient to give an accurate description of the solutions of the whole PDE. For a broader overview

Action Editor: Barry J. Richmond

C. R. Laing
Institute of Information and Mathematical Sciences, Massey University,
Private Bag 102-904,
North Shore Mail Centre, Auckland, New Zealand e-mail:
c.r.laing@massey.ac.nz

of the problem of extracting macroscopic dynamics from a microscopic description of a system, see Givon et al. (2004).

The ideas involving in EF modelling have previously been implemented for a variety of different problems; for example, a Lattice-Boltzmann implementation of the FitzHugh-Nagumo PDE in one spatial dimension (Gear et al., 2002), and kinetic Monte Carlo models of simple chemical reactions on a surface (Makeev et al., 2002). The purpose of this paper is to demonstrate in detail the EF modelling of several different networks of model neurons. As will be seen, some of the results could have been obtained by different methods, but some new results that could not be obtained by other methods will also be shown. The emphasis here is on an exposition of the techniques rather than the novel results.

Neural systems are suitable for this type of analysis for several reasons. One is that neural systems normally have a wide range of time-scales, which is a crucial requirement for the ideas discussed here to work. For example, spike frequency adaptation normally occurs on a much longer time-scale than the processes involved in the generation of a single action potential. Another reason is the wide range in the level of description of various neural systems, from single ion channels to much larger ensembles of neurons (Koch, 1999) and ultimately, brain regions. While a specific system may be well understood at a particular level of description, it is often hard to integrate that system into a larger one without making drastic simplifications. For example, while a single neuron may be well-characterised in terms of the ion channels involved in its action potential generation, often that detail is thrown away when a network of such neurons is studied, and only the frequency of firing for a fixed input is considered. For the examples we study, we can include as much detail as is known in the description of the dynamics of individual neurons (keeping the single neuron model in a "black box" to be simulated when necessary) but we can still describe the behaviour of a network of neurons in terms of a small number of macroscopic variables.

The structure of the paper is as follows. We now give a brief introduction to EF modelling; much more detail can be found elsewhere (Gear et al., 2002; Kevrekidis et al., 2003; Makeev et al., 2002; Möller et al., 2005; Runborg et al., 2002). In Section 2 we discuss a simple network of model neurons, all-to-all coupled with slow excitatory coupling. In Section 3 we study two populations that mutually inhibit one another; this leads to bistability. In Section 4 we investigate a system capable of supporting spatial patterns, while Section 5 discusses a network of bursting neurons, and in Section 6 we study a noisy network. We conclude in Section 7.

1.1. Equation-free modelling

Consider numerical integration in time, or simulation, of a complex system whose microscopic dynamics one knows in detail. An example is a network of coupled neurons. Let U be the macroscopic variable that one hopes describes the system (for example, an average rate of firing), and let u be the vector-valued variable describing the microscopic state of the system (for example, a vector containing the voltages and gating variables of all of the neurons in the network). u is normally much higher-dimensional than U. We suppose that there exists a function $F(U)$ such that the dynamics of U are given by the differential equation

$$\frac{dU}{dt} = F(U) \tag{1}$$

We do not have an explicit formula for $F(U)$ and cannot evaluate it in the usual sense of a function evaluation, but we assume that $F(U)$ exists and is well-behaved (for example, that it is differentiable). We also assume that U changes slowly in time relative to the rate of change of most of the components of u. Suppose that we want to make an Euler step forward in time, from an initial condition U_0, i.e. we know the state of U now, and want to know it at a small time h in the future. The formula is

$$U_1 = U_0 + hF(U_0) \tag{2}$$

where U_1 is the value of U at time $t = h$, and h is the time-step we have chosen. To evaluate $F(U_0)$ we initialise the microscopic system (the full network) with an initial condition u_0 that is consistent with U having the value U_0. This is done with a "lifting" operator, M, such that $u_0 = MU_0$. This is a one-to-many operation, as there are generally an infinite number of initial conditions for the microscopic system that are consistent with U having a particular value. We typically do this by choosing u_0 from a conditional probability density function $p(u_0|U_0)$. We also have a "restricting" operator, m, which extracts the value of U, given u, i.e. $U = mu$. This is normally a many-to-one operator, and could be as simple as taking an average. The operators M and m satisfy $mM = I$, the identity, so that lifting and then immediately restricting should have no effect, within roundoff error.

We then run the microscopic system for a short amount of time, δ. During this time the quickly-changing components of u vary so that the probability density function of u becomes "slaved" to (or determined by) the current value of U (Kevrekidis et al., 2003). Let $U_\delta = mu|_{t=\delta}$, i.e. the restriction of u after a time δ. We then run the microscopic simulation for a further time Δt. This time is long enough for U to change appreciably, but not so long that nonlinear effects appear in its behaviour, i.e. over the time interval $[\delta, \delta + \Delta t]$, U should change approximately linearly with time. We then use a simple forward difference

$$\frac{U_{\delta+\Delta t} - U_\delta}{\Delta t} \tag{3}$$

as an approximation to

$$\left.\frac{dU}{dt}\right|_{U=U_0} \tag{4}$$

i.e. as an approximation to $F(U_0)$. This process is shown schematically in Fig. 1. We can then use this approximation to step forward in time by an amount h, obtaining U_1 from (2). To continue the process we need $F(U_1)$, which is evaluated in the same way as $F(U_0)$, and so on. Note that for some of the models we investigate, the probability density function used in lifting from U_0 to u_0 is known analytically, and the timestep δ can be set to zero.

In order to reduce the variance in our estimate of $F(U_0)$ we may want to do a number of "bursts" of computation, each with the same value of U_0 but with different (consistent) values of u_0, and then average the values of $F(U_0)$ obtained from these different bursts to obtain a more accurate value. This part of the simulation is easily parallelisable, as each burst can be done on a separate processor, with no communication required between them except at the end of the burst. In essence, we are numerically averaging over the high-dimensional dynamics of the microscopic variables, in the same way that we would analytically to explicitly derive an equation like (1) (if we could).

The above is an example of numerical integration, or simulation, which, if repeated evaluations of the microscopic system are required, may take longer than just integrating the microscopic system for the required amount of time. But other tasks such as finding fixed points of (1) (i.e. finding zeros of $F(U)$), or following these fixed points as a parameter is varied, all require evaluation of F, which can be done as outlined above. For some of these tasks, evaluation of derivatives of F with respect to variables or parameters are needed. These can be approximated through finite differences, i.e. through running the microscopic system with nearby initial conditions, or at slightly different parameter values. It is this ability to simulate the microscopic system at will with specific initial conditions (which is either very difficult or impossible for a physical system) that is exploited in the EF modelling paradigm.

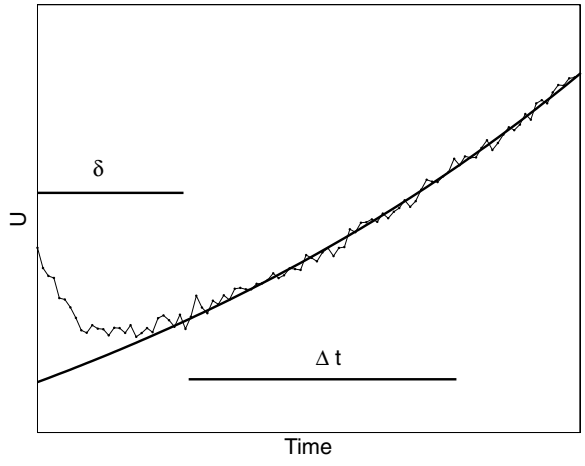

Fig. 1. A schematic diagram showing the time δ, during which the fast variables become "slaved" to the slow ones, and the time Δt, during which the change in U is determined. The thick line represents the true evolution of U, while the jagged line represents the restriction *mu* from one particular realisation of the microscopic dynamics. In this case, u was not initialised with a value consistent with the true value of $U(0)$, hence the difference in initial conditions. See Section 1 for more detail.

2. One population, positive feedback

Here we demonstrate some of these ideas using a simple all-to-all excitatorily coupled network of N integrate-and-fire neural oscillators. The equations are

$$\frac{dV_i}{dt} = I - V_i - \sum_k \delta(t - t_{ik}) + S + w_i \tag{5}$$

$$\tau \frac{ds_i}{dt} = A \sum_k \delta(t - t_{ik})(1 - s_i) - s_i \tag{6}$$

for $i = 1, \ldots N$, where t_{ik} is the kth firing time of neuron i (defined to be the times at which V_i reaches 1 from below), S is the average of the s_i,

$$S \equiv \frac{1}{N} \sum_{i=1}^{N} s_i \tag{7}$$

and the sums over k are over all past firing times. V_i is the voltage of neuron i and lies in $[0, 1)$, s_i is the strength of the synapses leaving neuron i, I is a constant input current and A controls the overall strength of synapses. The function $\delta(\cdot)$ is the Dirac delta function, used to reset the V_i to zero, and increment the s_i upon firing. The variables w_i are independent Gaussian white noise terms, with properties

$$\langle w_i(t) \rangle = 0 \qquad \langle w_i(t) w_i(s) \rangle = \sigma^2 \delta(t - s) \tag{8}$$

for each i, where the angled brackets indicate averages. The parameter σ controls the noise intensity. Noise was added to this system to smooth out the function describing the firing frequency of a single neuron as a function of its input current, which is known not to be differentiable everywhere for type I neurons, of which the integrate-and-fire neuron is an example (Gerstner and Kistler, 2002). This enables continuation algorithms, which rely on functions being sufficiently differentiable, to be used, and also makes the model more biologically realistic.

To ensure a separation of timescales, τ is chosen to be 50, which is large relative to the timescale of changes in voltage (i.e. we have slow synapses). In the absence of coupling ($A = 0$), each model neuron behaves like an independent integrate-and-fire neuron, with a noisy input. When A is not zero, the neurons are coupled through the average of the s_i. Each time neuron i fires, s_i is increased by an amount $A(1 - s_i)/\tau$; in between firing times s_i undergoes exponential decay back to zero with time constant τ. For a range of values of I, the system (5) and (6) is attracted to a state where all neurons are firing approximately periodically and are not synchronised.

For this system the vector u describing the microscopic state of the network would be formed from all of the V_i and s_i. Using the EF modelling approach, we assume that exact values of all of these variables are not of interest, in terms of describing the state of the network, Instead, the assumption is that the single variable S characterises the dynamics of the system, and that there is a single equation governing the dynamics of S, say

$$\frac{dS}{dt} = F(S; I) \tag{9}$$

Thus for this system, the macroscopic variable U is just S. It should be stressed that it is an *assumption* that the behaviour of the network (5) and (6) can be accurately described by (9).

In order to perform numerical bifurcation analysis of (9) we need to be able to evaluate $F(S; I)$ and its partial derivatives with respect to both S and I. To evaluate $F(S;I)$ we run (5) and (6) for a short amount of time and monitor S during that integration. The integration must be long enough for S to change appreciably and for the V_i to redistribute (if necessary) so that their probability density function is appropriate for the current value of S (i.e. for the system to "heal" (Gear et al., 2002)), but not so long that nonlinear effects start to come into play, since nonlinear effects will reduce the validity of the approximation (3). The correct integration time can be determined by running a number of simulations and observing the different time-scales in the dynamics. Alternatively, for a system like (5) and (6), we can use the explicitly-given slow time-scale, τ, to choose the integration time, setting it to, say, $\tau/2$.

2.1. Initial conditions

The question of initial conditions must be addressed. From a particular value of S we need to generate $2N$ initial conditions, for the V_i and s_i. In principle, since S is the average of the s_i, we could choose the $s_i(0)$ from any distribution with mean S. For simplicity, we choose $s_i(0) = S$ for all i. (This is one component of our lifting operator, M; the restricting operator is just the averaging of the s_i.) Also in principle, the $V_i(0)$ could be chosen in any way from the interval $[0, 1)$. However, there are some points to be considered.

If all $V_i(0)$ are chosen to be equal then the neurons will start off synchronised, since all $s_i(0)$ are equal to one another as well. The nonzero noise intensity will break up this synchronisation, but we need this to occur well before the end of the short burst of simulation. A better choice would be to take the $V_i(0)$ from a uniform distribution over $[0, 1)$. This would remove the problem of synchronisation but it is still not the best solution. A moment's thought shows that during periodic firing the V_i are not spread uniformly through the interval $[0, 1)$ but are more likely to be near 1 than near 0. In fact, if $\sigma = 0$ (i.e. the noise intensity is zero), the probability density function (PDF) for the V_i is $p(V_i \mid I + S)$, where

$$p(V_i|I) = \begin{cases} \frac{1}{B(I-V_i)} & \text{if } I > 1 \\ \delta(V_i - I) & \text{if } I < 1 \end{cases} \tag{10}$$

where $B = \log[I/(I - 1)]$. When $\sigma \neq 0$ it is also possible to calculate the PDF for the V_i (Brunel and Hakim, 1999; Fourcaud and Brunel, 2002)

$$p(V_i|I) = \frac{2f_0}{\sigma} \exp\left(\frac{-(V_i - I)^2}{\sigma^2}\right) \times \begin{cases} \int_{-I/\sigma}^{(1-I)/\sigma} \exp(s^2)\, ds & \text{if} \quad V_i < 0 \\ \int_{(V_i-I)/\sigma}^{(1-I)/\sigma} \exp(s^2)\, ds & \text{if} \quad V_i > 0 \end{cases} \tag{11}$$

where f_0 is the stationary firing rate:

$$f_0 = \left(\sqrt{\pi} \int_{-I/\sigma}^{(1-I)/\sigma} \exp\left(x^2\right)[1 + \mathrm{erf}(x)]dx\right)^{-1} \tag{12}$$

and erf is the error function. We will use the PDF given by (10), rather than (11), for three reasons. Firstly, there are difficulties with numerically implementing (11) when $\sigma \neq 0$, due to overflow errors. Secondly, the difference between the PDFs (10) and (11) is only significant when $I + S \approx 1$. Thirdly, due to the difference in time-scales for the V_i and s_i, as mentioned above, the initial conditions for the V_i are largely irrelevant.

In Fig. 2 we show the effects of using different distributions of initial conditions, and also demonstrate how we estimate $F(S; I)$. We use a network of $N = 200$ neurons. Figure 2A shows a rastergram of 50 out of the 200 neurons, with the $V_i(0)$ chosen from a uniform distribution on $[0, 1)$. (The $s_i(0)$ were all taken equal to $S(0)$.) We see that there is an initial gap of approximately 1 time unit in which very few neurons fire. Figure 2B shows S as

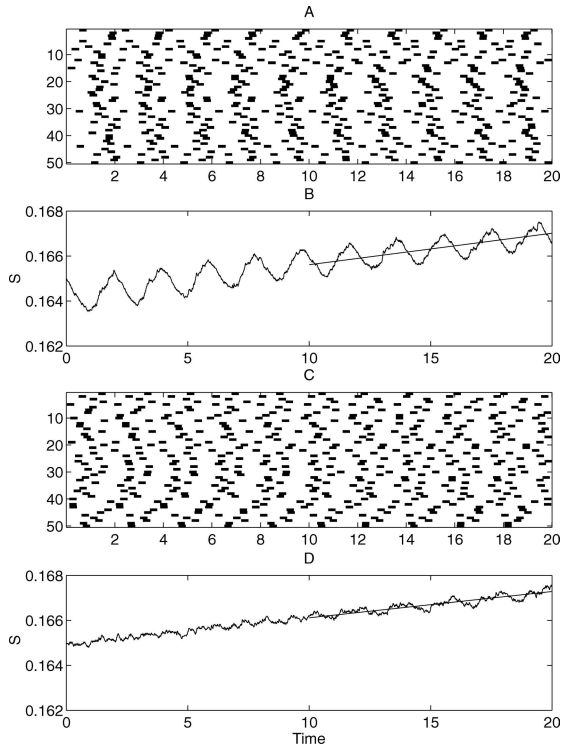

Fig. 2. Demonstration of the effects of different probability density functions for the $V_i(0)$, for the system (5) and (6). **A:** A rastergram showing the firing times of 50 out of 200 neurons (vertical scale: neuron index). The $V_i(0)$ were randomly chosen from a uniform distribution on [0, 1). **B:** S (the mean of the s_i) as a function of time for the simulation in **A**. Also shown is the straight line (slope $= 1.40 \times 10^{-4}$) fitted to the second half of the simulation. **C:** Same as **A**, but with the $V_i(0)$ taken from the PDF (10). **D:** Same as **B**, but for the simulation shown in **C** (slope $= 1.17 \times 10^{-4}$). Different realisations give qualitatively similar results (not shown). Parameters are $S(0) = 0.165$, $I = 1$. Other initial conditions are $s_i(0) = S(0)$ for $i = 1, \ldots, N$.

a function of time for the simulation shown in Fig. 2A. The small amount of synchrony induced by the initial conditions appears as oscillations superimposed on the slow growth in S, at the same frequency as that of the individual neurons. To calculate $F(S; I)$ we do a least-squares fit of a straight line to $S(t)$ for $t \in [10, 20]$ and take the slope of that line. We do this for N_{av} independent simulations and then average the results. Fitting a straight line and taking its slope as an estimate of $F(S(0); I)$, rather than simply using $[S(20) - S(10)]/10$, resulted in a smaller variance for the estimates, hence its use.

Figure 2C shows a rastergram for 50 out of the 200 neurons when the $V_i(0)$ are chosen from the PDF given by (10). The lack of any synchrony is clear in Fig. 2D, where we plot S as a function of time for the simulation in Fig. 2C. Any oscillations seen here are due to the noise and the fact that we do not have an infinite number of asynchronous neurons. Also shown in Fig. 2D is a least-squares fit of a straight line to S, as in Fig. 2B. The slope of this line differs from the slope in panel A by less than 20%. It is clear that other choices in determining $F(S; I)$ are possible. For example, one could fit a straight line over a different time interval, or use $S(20) - S(0)$, or fit a quadratic $q(t)$ to $S(t)$ and use $q(20) - q(0)$, or any number of other possibilities.

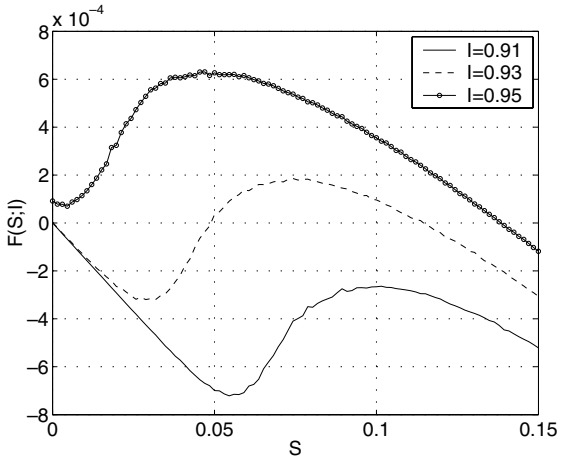

Fig. 3. The function $F(S; I)$ (Eq. (9)), as calculated from direct simulation of (5) and (6), for three different values of I. Macroscopic fixed points occur when $F(S; I) = 0$. We have used a network of $N = 200$ neurons, averaging over $N_{av} = 30$ realisations. Other parameters are $A = 0.4$, $\tau = 50$, $\sigma = 0.0245$.

This figure also shows that S changes on a time scale much slower than that of the V_i. Indeed, for this initial condition, the simulation must be run for 200 time units before S appears to saturate. This is expected, as each synaptic strength s_i evolves with a time constant $\tau = 50$.

In Fig. 3 we show $F(S; I)$ as a function of S, calculated as above, for three different values of I. We again used a network of $N = 200$ neurons, and averaged $N_{av} = 30$ times. We see that for $I = 0.91$, there is one stable fixed point very close to $S = 0$ (it is stable, since $F < 0$ for $S > 0$). There are three fixed points for $I = 0.93$ and only one, at a high value of S, for $I = 0.95$. We can see that there are some bifurcations of fixed points as I is varied. We now discuss this.

2.2. Fixed points

Figure 4 shows the curve of macroscopic fixed points for (5) and (6) as I is varied. These are fixed points of S, not of the microscopic variables (the V_i and s_i). We have traced this curve using pseudo-arclength continuation (Doedel et al., 1991). Also shown in Fig. 4 are two curves that can be calculated analytically using a rate-based formulation, as we now discuss.

Taking the average of the term $\sum_k \delta(t - t_{ik})$ in Eq. 6 over an interval of length T, we obtain the number of times that neuron i has fired during that interval, divided by T, i.e. the average firing rate of that neuron during that time interval. The neurons are identical, each receiving an ef-

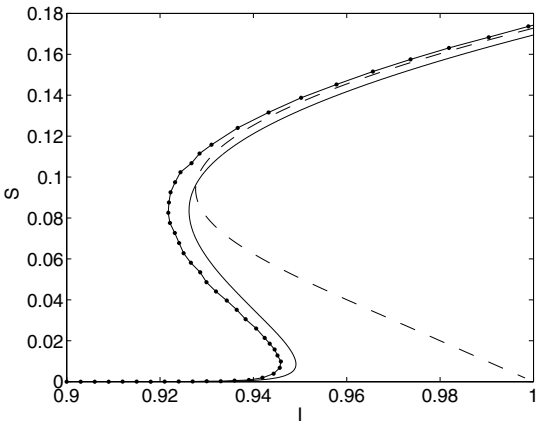

Fig. 4. The curve of macroscopic steady states of the integrate-and-fire network (5) and (6) (points joined by a line). Parameters are $A = 0.4, \tau = 50, \sigma = 0.0245$. A network of $N = 200$ was used, with averaging over $N_{av} = 30$ realisations. The dashed line shows the fixed points of the rate-based noise-free case [Eq. (15), with f given by (17)], while the solid line shows the solution of the rate-based system [Eq. (15), with f given by (18)] with $\sigma = 0.0245$.

fective input current of $I + S$. Thus each neuron will be firing at a rate $f(I + S)$, where $f(I)$ is the firing rate for a single integrate-and-fire neuron with input I. Thus we can approximate (6) by

$$\tau \frac{ds_i}{dt} = Af(I + S)(1 - s_i) - s_i \quad i = 1, \ldots, N \tag{13}$$

Taking the average of these N equations we obtain one equation for S:

$$\tau \frac{dS}{dt} = Af(I + S)(1 - S) - S \tag{14}$$

Fixed points satisfy

$$g(S, I) = 0 \tag{15}$$

where

$$g(S, I) = \frac{Af(I + S)(1 - S) - S}{\tau} \tag{16}$$

For a noise free neuron, i.e. when $\sigma = 0$,

$$f(I) = f_1(I) \equiv H(I - 1)\left(\log\left[\frac{I}{I - 1}\right]\right)^{-1} \tag{17}$$

where H is the Heaviside function. When $\sigma \neq 0$,

$$f(I) = f_2(I) \equiv \left(\sqrt{\pi}\int_a^b \exp\left(x^2\right)[1 + \mathrm{erf}(x)]dx\right)^{-1} \tag{18}$$

where $a \equiv -I/\sigma$, $b \equiv (1 - I)/\sigma$, and erf is the error function (Fourcaud and Brunel, 2002).

Plotted in Fig. 4 are solutions of (15) where $f(I)$ is given by (17) (dashed line), and where $f(I)$ is given by (18) (solid line). We see that for S greater than about 0.1, the three curves are close. However, below the turning point the curve of macroscopic fixed points is much better approximated by the curve for the rate-based model with the correct amount of noise, in the sense of the curves being closer to one another. The agreement between these curves is a useful confirmation that the method is working. We now discuss the stability of these fixed points.

2.3. Stability

The stability of a fixed point of (9) is given by the sign of $\partial F/\partial S$. A positive derivative indicates instability, whereas a negative derivative indicates stability. A similar remark holds for the rate-based Eq. (15). To find the partial derivative of $F(S; I)$ with respect to its arguments we use finite differences. For example,

$$\left.\frac{\partial F}{\partial S}\right|_{S=\overline{S}} \approx \frac{F(\overline{S} + \epsilon; I) - F(\overline{S}; I)}{\epsilon} \tag{19}$$

and similarly for $\partial F/\partial I$. In all following work we use a value of $\epsilon = 0.01$.

Figure 5 shows the stability of the fixed points found in Fig. 4, as indicated by the sign of $\partial F/\partial S$. We see that the upper and lower branches of fixed points are stable, while the "middle" branch is unstable. Using only numerical simulation of (5) and (6) until an attractor was reached, only the stable branches of this curve would have been found. It is only by using the EF approach of assuming that an equation describing all fixed points exists, and using short bursts of appropriate simulation to find roots of it, that we can trace out the unstable branch of macroscopic steady states and thus join the stable branches to one another.

2.4. Following bifurcations

We see in Fig. 4 that there are two saddle-node bifurcations as I is varied. These occur when the system has an eigenvalue of zero, i.e. when $\partial F/\partial S = 0$. If we allow another parameter, for example A, to vary, we can follow these bifurcations in a two-dimensional parameter space. We need to follow solutions of the system

$$F(S; I; A) = 0 \tag{20}$$

$$\frac{\partial F(S; I; A)}{\partial S} = 0 \tag{21}$$

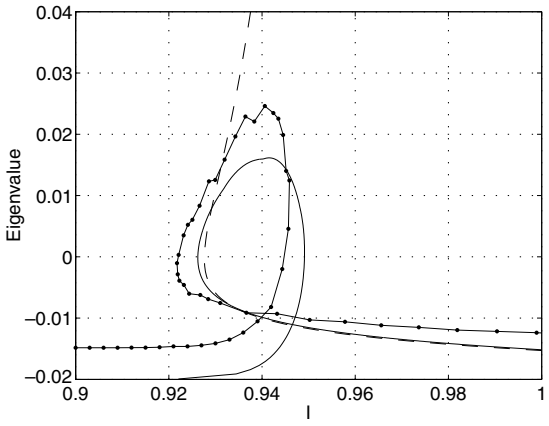

Fig. 5. Eigenvalues for the steady states shown in Fig. 4. Positive values correspond to instability, negative to stability. The points joined by a line are for macroscopic steady states of (5) and (6). The dashed line shows the stability for the rate-based noise-free case, while the solid line shows the stability for the rate-based system with $\sigma = 0.0245$.

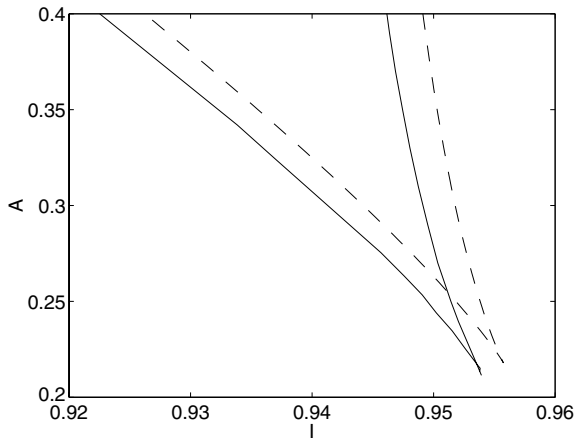

Fig. 6. Solid line: The curve of saddle-node bifurcations of macroscopic steady states of the integrate-and-fire network (5) and (6). A network of $N = 1000$ neurons was used, with averaging over $N_{av} = 50$ realisations. Other parameters are $\tau = 50, \sigma = 0.0245$. Dashed line: The curve of saddle-node bifurcations for the rate-based system with $\sigma = 0.0245$. Figure 4 shows a horizontal "slice" through this figure at $A = 0.4$.

where we have now explicitly included the dependence of F on A. Following solutions to these equations requires more averaging (i.e. a much larger value of N_{av} and/or N) to be successful, since we need to numerically estimate second derivatives of F during continuation, not just first derivatives (Makeev et al., 2002).

Results are shown in Fig. 6, where we plot the curves of saddle-node bifurcations [i.e. solutions of (20) and (21)] in the (I, A) plane. Also shown is the corresponding curve for the rate-based system (15) with f given by (18). For both systems, the curves meet in a cusp as A is decreased. For values of A below this point the system no longer has bistability for any values of I. Note that for the deterministic rate-based system, we expect the curve of saddle-node bifurcations to persist down to $A = 0$, terminating at $(I, A) = (1, 0)$ (Laing and Longtin, 2003).

2.5. Remarks

- In this modelling a separation of time-scales was created by choosing $\tau = 50$. We have successfully followed the curve of macroscopic fixed points of (5) and (6) using the ideas outlined above for values of τ down to $\tau = 10$, indicating that such a drastic separation of time-scales is not necessary, and can be relaxed somewhat. (Successfully followed means that the curve of fixed points lies on top of the curve calculated for $\tau = 50$.)
- There is no necessity for N, the number of neurons, to be greater than one, provided that N_{av}, the number of different simulations whose results are averaged, is sufficiently large. By running different simulations with different (appropriate) initial conditions, we are still effectively sampling the whole space of microscopic variables. Having a larger N makes S more smooth, as it is the average of N variables, which in turn reduces the value of N_{av} needed to obtain a reliable estimate of dS/dt.

  In particular, the results obtained for $N = 1$ are those that would be obtained if we had a network of N identical neurons that were *perfectly synchronised* during the entire burst of simulation. We return to this point in Section 7.
- We also note that during the averaging we are performing, we are actually simultaneously averaging over initial conditions for the V_i and over realisations of the Gaussian white noise in (5).
- As can be seen, there are a number of quantities that must be chosen, e.g. the number of bursts to be averaged over (N_{av}), and the lengths of the bursts. An element of trial and error may be involved in choosing these and it is possible that certain choices will cause the algorithms to break down. Verification of results through different means (for example, numerical simulation) may be prudent.

3. Two mutually inhibiting populations

We now discuss a more complicated network of two distinct populations of identical model integrate-and-fire neurons, each population inhibiting the other through the averaged synaptic activity.

3.1. The model

The equations are

$$\frac{dV_i^1}{dt} = I^1 - V_i^1 - \sum_k \delta(t - t_{ik}^1) - S^2 + w_i^1 \qquad (22)$$

$$\frac{dV_i^2}{dt} = I^2 - V_i^2 - \sum_k \delta(t - t_{ik}^2) - S^1 + w_i^2 \tag{23}$$

$$\tau \frac{ds_i^1}{dt} = A \sum_k \delta(t - t_{ik}^1)(1 - s_i^1) - s_i^1 \tag{24}$$

$$\tau \frac{ds_i^2}{dt} = A \sum_k \delta(t - t_{ik}^2)(1 - s_i^2) - s_i^2 \tag{25}$$

for $= 1, \ldots, N$, where the superscripts identify the population,

$$S^p = \frac{1}{N} \sum_{i=1}^{N} s_i^p \tag{26}$$

and the other terms have the same meanings as in the previous example. Note that all of the Gaussian white noise terms are independent.

For a fixed $I^1 = I^2 > 1$ and A small, both populations will fire at the same rate. But for A large enough it is possible for one population to fire strongly, completely suppressing the other in a "winner takes all" scenario. If $I^1 = I^2$, this system is symmetric with respect to interchanging the populations, and we expect this symmetry to manifest itself in terms of the possible bifurcations that can occur. We again take $\tau = 50$ to provide a separation of time scales.

Our assumption is that the dynamics are governed by the equations

$$\frac{dS^1}{dt} = F(S^1, S^2; A) \tag{27}$$

$$\frac{dS^2}{dt} = F(S^2, S^1; A) \tag{28}$$

where the same function F is used, due to the symmetry. We investigate these equations in the same way as in the previous section. To evaluate $F(S^1, S^2; A)$, we initialise each $s_i^1(0) = S^1$ and $s_i^2(0) = S^2$, and initialise the $V_i^1(0)$ using the PDF $p(V_i^1 | I^1 - S^2)$ and the $V_i^2(0)$ using the PDF $p(V_i^2 | I^2 - S^1)$, where p is given by (10). We then integrate for 20 time units and fit a straight line to $S^1(t)$ for $10 < t < 20$ and take its slope—this is our estimate of $F(S^1, S^2; A)$. Fitting another straight line to S^2 over the same time interval gives us an estimate of $F(S^2, S^1; A)$. This procedure is repeated N_{av} times with different initial conditions for the voltages and different Gaussian white noise terms, and the results are averaged.

3.2. Rate equations

In a similar way to that in the previous section, we can write approximate rate equations for the system (22)–(25), under the assumptions that we have an infinite number of perfectly asynchronous neurons. The equations we obtain are

$$\tau \frac{dS^1}{dt} = Af(I^1 - S^2)(1 - S^1) - S^1 \tag{29}$$

$$\tau \frac{dS^2}{dt} = Af(I^2 - S^1)(1 - S^2) - S^2 \tag{30}$$

where $f(I)$ is the firing frequency of a neuron with input current I. We use the function given by (18), with the same value of σ as that used in the simulations of (22)–(25). From (29) and (30) we see that when $S^1 = S^2$, both of these values are given by the roots of

$$Af(I - S)(1 - S) - S = 0 \tag{31}$$

whereas when $S^1 \neq S^2$, all fixed points of (29) and (30) can be found by finding the roots of

$$Af\left(I - \frac{Af(I - S)}{1 + Af(I - S)}\right)(1 - S) - S = 0 \tag{32}$$

where $I^1 = I^2 = I$. The stability of these fixed points is determined by the eigenvalues of the Jacobian of (29) and (30), evaluated at the fixed points.

3.3. Numerical results

Numerical results are shown in Fig. 7, where we vary A, the strength of inhibition between the two networks, keeping $I^1 = I^2$. For low values of A we have only one steady state, for which $S^1 = S^2$. This is a stable state. As A is increased this symmetric state becomes unstable through a subcritical pitchfork bifurcation. Two unstable branches of steady states on which $S^1 \neq S^2$ are created in this bifurcation. These are destroyed in saddle-node bifurcations with stable branches of steady states which exist for large A. In these stable states, one network is firing and completely suppressing the other. Figure 7 shows results from both the analysis of (27) and (28), and of (29) and (30).

3.4. Forced symmetry-breaking

As mentioned, the system (22)–(25) has a reflectional symmetry involving the interchange of the two populations. This underlies the pitchfork bifurcation seen in Fig. 7. We can easily destroy this symmetry by making the two populations

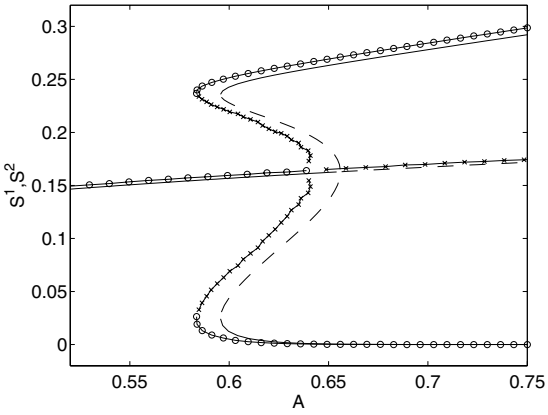

Fig. 7. Steady states for a pair of mutually inhibiting populations. The system is symmetric with respect to interchanging S^1 and S^2. Circles and crosses joined by lines indicate macroscopic steady states of (22)–(25). Circles indicate stable states and crosses unstable, as determined by the sign of the most positive eigenvalues of the Jacobian of (27) and (28). Along the "central" branch, $S^1 = S^2$, while along the other branches $S^1 \neq S^2$. We used two networks of $N = 1000$ neurons each, averaging over $N_{av} = 50$ realisations. Solid lines are stable fixed points of (29) and (30) while dashed lines are unstable fixed points of these equations. Other parameters are $I^1 = I^2 = 1.2$, $\tau = 50$, $\sigma = 0.0245$.

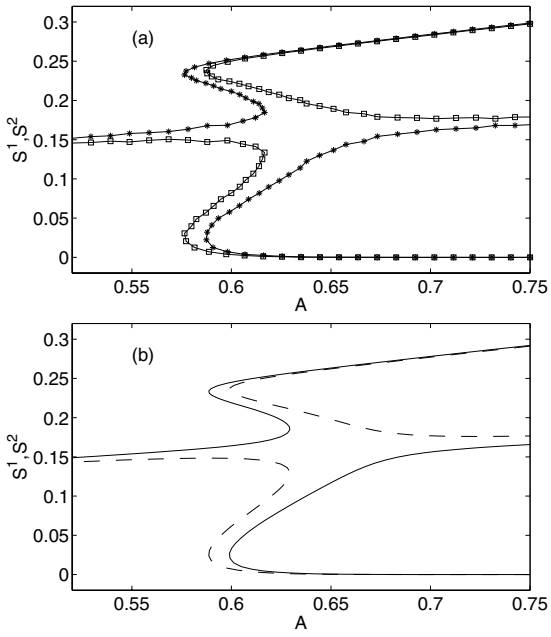

Fig. 8. Forced symmetry breaking. (a): macroscopic steady states of (22)–(25) with $I^1 = 1.198$ and $I^2 = 1.2$. Squares are S^1, stars are S^2. Two networks of $N = 200$ were used, with averaging over $N_{av} = 30$ realisations. (b): Fixed points of (29), (30), with f given by (18), and $I^1 = 1.198$ and $I^2 = 1.2$. The dashed line is S^1; the solid line, S^2. Stability is discussed in the text. Other parameters are as in Fig. 7.

non-identical. This will generically break the pitchfork bifurcation into two saddle-node bifurcations, and move apart the pairs of curves that are superimposed on one another in Fig. 7. This is shown in Fig. 8, where we have set $I^1 = 1.198$ and $I^2 = 1.2$ in both (22), (23) and (29), (30). Although not indicated in Fig. 8, the stability of the branches is as expected from Fig. 7: the fixed point for which S^1 and S^2 are almost equal is stable for small A and unstable for large A, and a stable and unstable solution are annihilated at each saddle-node bifurcation.

4. A spatially extended system

We now consider the problem of pattern formation in a spatially-extended neural system. There has been much recent interest in these problems, one area of interest being the formation of "bumps"—spatially-localised patches of active neurons (Compte et al., 2000; Gutkin et al., 2001; Laing and Chow, 2001; Laing et al., 2002). These are thought to be involved in working memory tasks and orientation tuning in the visual system. Work has been done on rate models, where the variable of interest is the average firing rates of the neurons (Laing et al., 2002), and also spiking neuron models (Compte et al., 2000; Gutkin et al., 2001; Laing and Chow, 2001). When studying a rate model one implicitly assumes that the neurons fire asynchronously, and that it is only their firing rate that is relevant. Spiking models are more realistic, but it is normally impossible to study unstable macroscopic states such as unstable bumps in these systems. However, we will now demonstrate how to do so, under the assumption that there is a separation of time-scales in the system. This will enable us to get a better picture of the bifurcations that such systems can undergo.

4.1. The model

The model we study is the following network of integrate-and-fire neurons.

$$\frac{dV_i}{dt} = I - V_i + \frac{2\pi}{N}\sum_{j=1}^{N} J_{ij}s_j - \sum_k \delta(t - t_{ik}) + w_i \qquad (33)$$

$$\tau\frac{ds_i}{dt} = A\sum_k \delta(t - t_{ik})(1 - s_i) - s_i \qquad (34)$$

where the w_i are the usual (uncorrelated) Gaussian white noise terms and

$$J_{ij} = J\left(\frac{2\pi|i - j|}{N}\right) \qquad (35)$$

where J is a difference of Gaussians:

$$J(x) = 5.25\sqrt{\frac{0.6}{\pi}}\exp\left(-0.6x^2\right) - 5\sqrt{\frac{0.5}{\pi}}\exp\left(-0.5x^2\right) \qquad (36)$$

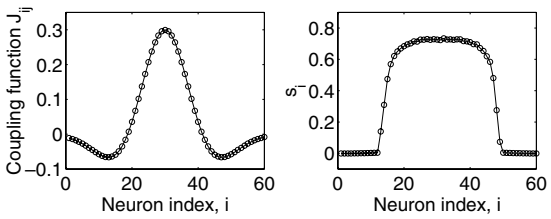

Fig. 9. Left: The coupling function J_{ij} (35) and (36) for $j = 30$, $N = 60$. Right: A stable stationary bump solution of (33) and (34) at $I = 0.9$. Other parameters are $A = 5$, $\tau = 50$, $\sigma = 0.0245$.

This coupling function is shown in Fig. 9 (left). We think of the neurons as lying on a ring of circumference 2π; their position could correlate with a feature that they are tuned for, e.g. orientation. We again take $\tau = 50$, so there is a separation of time scales between the s_i and the V_i. It is well-known that systems like (33) and (34) can support "bump" solutions, in which only a fraction of the neurons are firing (Laing and Chow, 2001). An example of such a bump state is shown in Fig. 9 (right).

4.2. Representation and implementation

For this system, there are a variety of ways of representing the macroscopic solutions. First note that the problem is spatially homogeneous, so we can assume that the bump is even about $x = \pi$, where x is the distance around the ring, with $0 \le x < 2\pi$. We choose the first two coefficients of an even Fourier series representation of the s_i, where we are thinking of the index i as a discrete spatial variable. We thus represent our bump as

$$S(x) = b_0 + b_1 \cos x \tag{37}$$

We suppose that there are functions F_1 and F_2 such that the dynamics of b_0 and b_1 are governed by

$$\frac{db_0}{dt} = F_1(b_0, b_1) \tag{38}$$

$$\frac{db_1}{dt} = F_2(b_0, b_1) \tag{39}$$

and that these can faithfully represent the dynamics of bumps in (33) and (34). See below for discussion about our choice of macroscopic variables.

Our restricting operator m takes the values $\{s_i\}$ and generates b_0 and b_1 via

$$b_0 = \frac{1}{N}\sum_{i=1}^{N} s_i \quad \text{and} \quad b_1 = \frac{2}{N}\sum_{i=1}^{N} s_i \cos\left(\frac{2\pi i}{N}\right) \tag{40}$$

(This is just the first two terms of a discrete cosine transform of the s_i.) Our lifting operator M generates initial conditions $\{s_i(0)\}$ from b_0 and b_1 in the obvious way:

$$s_i(0) = b_0 + b_1 \cos\left(\frac{2\pi i}{N}\right) \quad i = 1, \ldots, N \tag{41}$$

For simplicity we do not include any random components in the $\{s_i(0)\}$, although it would be consistent to add a random number from a distribution with mean zero to each $s_i(0)$, for example.

The initial conditions for the V_i are generated in a similar way to that in earlier sections. Given the set $\{s_i(0)\}$, for each neuron we calculate the initial effective drive current:

$$I_i = I + \frac{2\pi}{N}\sum_{j=1}^{N} J_{ij} s_j(0) \qquad i = 1, \ldots, N \tag{42}$$

We then choose each $V_i(0)$ from the PDF $p(V_i(0)|I_i)$, where p is given by (10).

We estimate $F_1(b_0(0), b_1(0))$ and $F_2(b_0(0), b_1(0))$ in the usual way. Given values of $b_0(0)$ and $b_1(0)$, lift them to initial conditions using Eq. (41) and generate the $V_i(0)$ as above, run the system (33) and (34) for 20 time steps, generating $b_0(t)$ and $b_1(t)$ using the time-dependent versions of (40), then do a least-squares fit to find the slopes of $b_0(t)$ and $b_1(t)$ as functions of time over the time interval $[10, 20]$. These are our estimates of $F_1(b_0(0), b_1(0))$ and $F_2(b_0(0), b_1(0))$, respectively. Actually, N_{av} simulations are run in parallel, with different initial conditions for the V_i and different realisations of the white noise, and the $s_i(t)$ are then averaged over realisations before $b_0(t)$ and $b_1(t)$ are generated and the slopes are fit. Another possibility would be to find the slopes for each realisation, and then average these slopes.

4.3. Results

In Fig. 10 we show the Fourier coefficients b_0 and b_1 as functions of I for the parameter values $A = 5$, $N = 60$, $\tau = 50$, $\sigma = 0.0245$. We average over $N_{av} = 50$ realisations. Two branches are shown. One, for which $b_1 = 0$, represents the spatially-uniform state, and forms an "S" shaped curve in the top panel. The other branch (on which $b_1 \neq 0$) represents the bump state, and is both created and destroyed in subcritical pitchfork bifurcations from the spatially-uniform state. This is a pitchfork bifurcation in our system because we have "pinned" the bump by forcing it to be even about $x = \pi$, thus destroying the continuum of solutions that exists, parametrised by the angular position of the maximum. Only one half of the family of bumps is represented in Fig. 10 (bottom): the family for which $b_1 < 0$. These have their maximum at $x = \pi$. The other half of the family, which have their maximum at $x = 0$ and which have $b_1 > 0$, can be obtained by reflecting the curve in Fig. 10 (bottom) about the I axis.

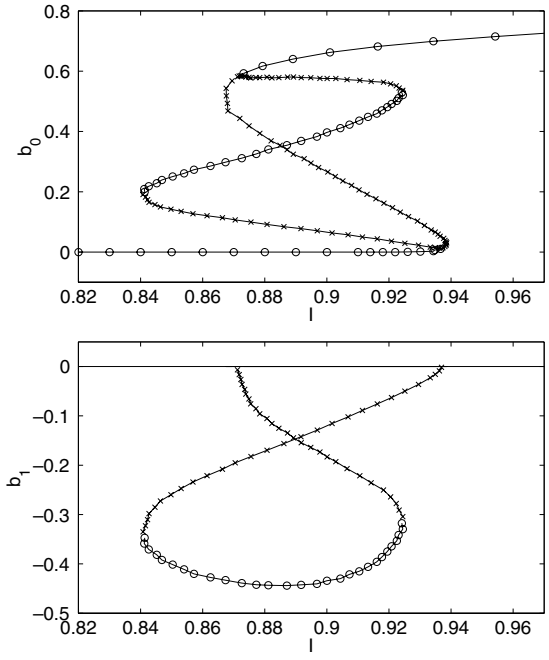

Fig. 10. b_0 (top) and b_1 (bottom) as functions of I, for macroscopic steady states of the model (33) and (34). Circles indicate stable solutions while crosses indicate unstable. The "S"-shaped curve in the top panel corresponds to spatially uniform states, for which $b_1 = 0$. The "Z"-shaped curve corresponds to bump solutions, for which $b_1 \neq 0$. Bump solutions are created/destroyed in pitchfork bifurcations from the spatially uniform state at $I \approx 0.87$ and $I \approx 0.94$ (where the curve in the bottom panel touches the I axis). Only one branch of bump solutions created/destroyed in the pitchfork bifurcations (the branch for which $b_1 < 0$) is shown. The other branch can be obtained by reflecting the curve in the bottom panel about the I axis. See text for more details.

Stability is also indicated in Fig. 10. This is determined by examining the eigenvalues of the 2×2 Jacobian of the associated differential equations for b_0 and b_1, at the steady states. On the "middle" branch of the spatially-uniform state, both of these eigenvalues are positive. On all other unstable branches, one eigenvalue is positive and the other negative. When following the spatially-uniform state through a pitchfork bifurcation, the eigenvector corresponding to the zero eigenvalue has a large entry in the b_1 component and a very small entry in the b_0 component, indicating that the instability acts to break the spatial uniformity. Note that there are regions of both bistability and tristability, between the "all-off" state, the bump state, and the "all-on" state.

By tracing out the unstable branches, we have been able to piece together the stable branches, which are the only branches we would have observed using straight-forward numerical integration.

4.4. A rate model

As with some previous examples, we can derive an effective rate model whose dynamics should closely mimic those of the spiking neural network (33) and (34). The sum over j in Eq. (33) is the discretised version of the convolution of J and S, where J is the continuous function (36). Thus, moving to a spatial continuum, the effective drive to a neuron at position x is $I + (J * S)(x)$, where the convolution of J and S is given by

$$(J * S)(x) = \int_0^{2\pi} J(x - y)S(y)\, dy \tag{43}$$

Replacing the sum of delta functions in (34) by the firing rate f, where f is given by (18), we obtain the nonlocal PDE

$$\tau \frac{\partial S(x,t)}{\partial t} = Af[I + (J * S)(x)][1 - S(x,t)] - S(x,t) \tag{44}$$

Writing S as the Fourier series (37), we can find the differential equations that b_0 and b_1 obey:

$$\tau \frac{db_0}{dt} = \frac{A(1-b_0)}{2\pi} \int_0^{2\pi} f(g(x))\, dx - \frac{Ab_1}{2\pi} \int_0^{2\pi} f(g(x)) \cos(x)\, dx - b_0 \tag{45}$$

$$\tau \frac{db_1}{dt} = \frac{A(1-b_0)}{\pi} \int_0^{2\pi} f(g(x)) \cos(x)\, dx - \frac{Ab_1}{\pi} \int_0^{2\pi} f(g(x)) \cos^2(x)\, dx - b_1 \tag{46}$$

where

$$g(x) = I + \mu b_0 + \nu b_1 \cos x \tag{47}$$

and

$$\mu = \int_0^{2\pi} J(y)\, dy \quad \text{and} \quad \nu = \int_0^{2\pi} J(y) \cos y\, dy \tag{48}$$

Note that the manifold defined by $b_1 = 0$ is invariant, and on this manifold the dynamics of b_0 are given by

$$\tau \frac{db_0}{dt} = Af\,(I + \mu b_0)(1 - b_0) - b_0 \tag{49}$$

Comparing this with Eq. (14), which is derived for a homogeneous population with positive feedback, we see the similarities, and also the importance of the quantity μ, equal to 2π times the mean of J. Fixed points of (49) are shown with a dashed line in Fig. 11 (top).

In Fig. 11 we show fixed points of (45) and (46) as I is varied, for both the bump state and the spatially-uniform state. We see very good agreement with the results obtained

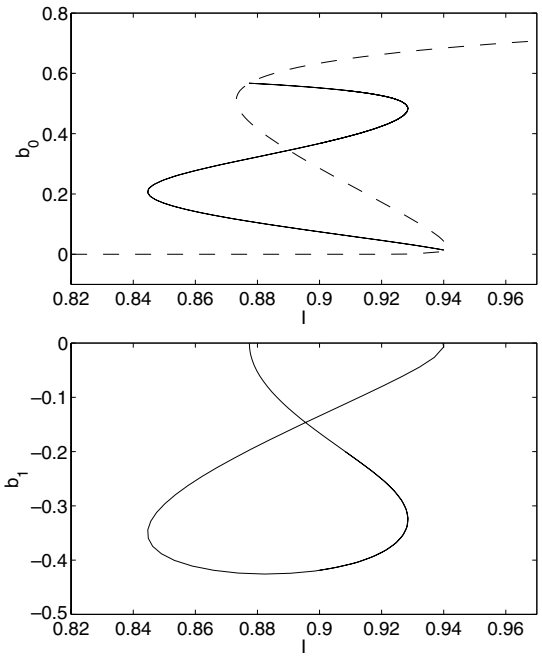

Fig. 11. b_0 (top) and b_1 (bottom) as functions of I, at steady states of the system (45) and (46). The dashed line corresponds to the spatially uniform state, while the solid line corresponds to the bump state. The same pattern of bifurcations occurs as occurred in Fig. 10. Note that the dashed curve describes the fixed points of (49). Parameters are $A = 5$, $\tau = 50$, $\sigma = 0.0245$, with J given by (36).

from the network of spiking neurons (Fig. 10). However, we now discuss a situation in which it is not possible to derive an equivalent rate formulation, for which EF modelling provides information that could not be derived any other way.

4.5. Including gap junctions

Although much of the communication between neurons occurs through synapses, there are often significant connections via gap junctions (Chow and Kopell, 2000; Kopell and Ermentrout, 2004). To investigate the effects of including such connections, we modify (33) and (34) to

$$\frac{dV_i}{dt} = I - V_i + \kappa(V_{i+1} - 2V_i + V_{i-1}) + \frac{2\pi}{N}\sum_{j=1}^{N} J_{ij}s_j - \sum_k \delta(t - t_{ik}) + w_i \tag{50}$$

$$\tau\frac{ds_i}{dt} = A\sum_k \delta(t - t_{ik})(1 - s_i) - s_i \tag{51}$$

where $\kappa > 0$ is a measure of gap junction conductivity, and $V_{-1} \equiv V_N$ and $V_{N+1} \equiv V_1$. All other terms have their previous meanings. Including a term like this acts to keep the voltages of neighbouring neurons more similar.

In Fig. 12 we show results of the same form as those in Fig. 10, but for (50), (51) with $\kappa = 0.5$. Including gap junctions in this way simultaneously decreases the size of the interval of I values for which the system can support a stable bump and shifts this interval to higher values of I. We further demonstrate this in Fig. 13 where we trace the four saddle-node bifurcations in Figs. 10 and 12 as κ is varied.

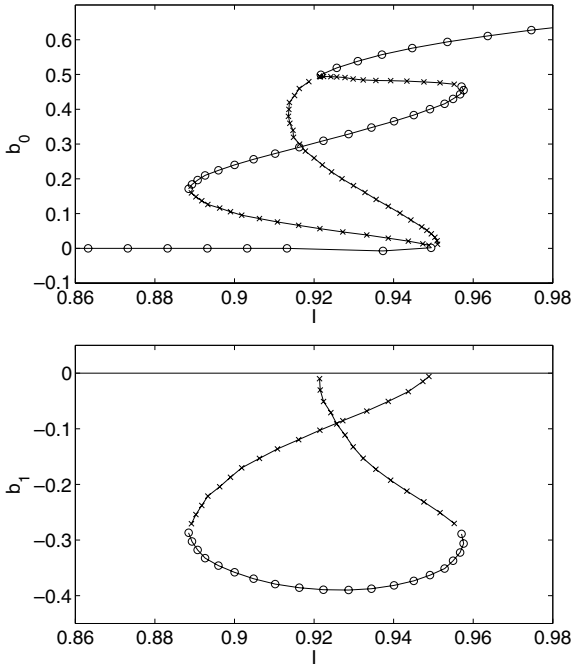

Fig. 12. b_0 (top) and b_1 (bottom) as functions of I, at steady states of (50) and (51) with $\kappa = 0.5$. Circles indicate stable solutions while crosses indicate unstable. Parameters are $N = 200$, $\tau = 50$, $A = 5$, $\sigma = 0.0245$. We used the coupling function given in (35) and (36), and averaged over $N_{av} = 100$ realisations. Compare with Fig. 10 and note the different axes.

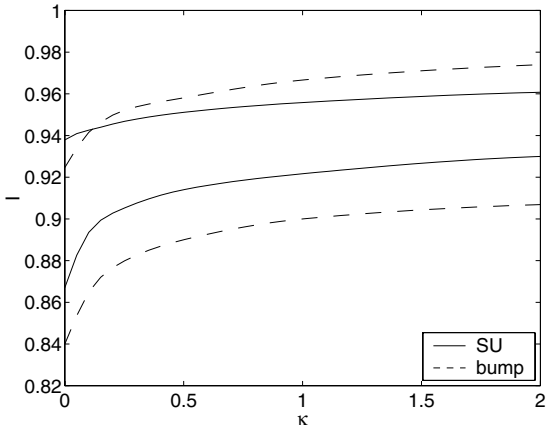

Fig. 13. Curves of saddle-node bifurcations of the spatially-uniform state ("SU") and the bump state ("bump") for (50) and (51). Figures 10 and 12 are vertical slices through this diagram at $\kappa = 0$ and $\kappa = 0.5$, respectively.

The inclusion of the term involving voltage differences in (50) means that deriving a rate model description of (50) and (51) [such as that in Section 4.4 for the system (33) and (34)] is not possible. Thus, as far as we are aware, these results are novel and could not be derived any other way. This demonstrates one of the limitations of rate models. Note that from the EF point of view, macroscopic stationary states of (50) and (51) are just as easy to analyse as those of (33) and (34). Indeed, further levels of complexity could be included and the model neurons in (33) could be replaced by models that are as realistic as one would like, and the process repeated.

4.6. Discussion

An obvious question relates to the appropriateness of representing a bump, an example of which is shown in Fig. 9 (right), by just the first two components of a Fourier series. The main justification for this is that we are specifically interested in bump solutions which have only one maximum on the domain. Previous work on models of this type shows that they can typically support only one bump, and extensive numerical simulation for the system (33) and (34) suggests that this is also the case here (not shown). This is an example of choosing the macroscopic variables (b_0 and b_1) so that questions of interest can be answered. In this case, the question is: what bifurcations are responsible for the creation and destruction of single bump solutions? If, for example, we were interested in two-bump solutions, for which there are two disjoint intervals on which neurons are active (Laing and Troy, 2003), we would clearly need more terms in our Fourier series (37) to be able to accurately describe these solutions.

As an aside, it is easy to show that if the coupling function $J(x)$ can be represented exactly by a finite Fourier series with $m+1$ terms (i.e. the last term being $c_m \cos(mx)$), then including more than $m+1$ terms in our Fourier representation of S (Eq. (37)) will not change the bifurcation structure. This can be seen by calculating the convolution between J and S that occurs in Eq. (44).

Several authors have recently discussed the mechanisms by which bumps of the form studied here can lose stability through synchronisation of the neurons by either a transient external stimulus being applied (Gutkin et al., 2001), or by speeding up the synapses to make the neurons more synchronous (Compte et al., 2000; Laing and Chow, 2001). Unfortunately, it seems that the EF approach will not be useful in studying these bifurcations, since by definition, they occur in the parameter regime (fast synapses) where the method is not valid. However, one could use these ideas to study a system like that of Compte et al. (2000) which includes both fast and slow synapses and investigate, for example, the effects of changing the strength of the fast synapses.

5. Bursting neurons

The neuron model used in the previous sections was the integrate-and-fire model, a simple example of a type I neuron (Gerstner and Kistler, 2002). For a fixed input current, one of these model neurons fires at a well-defined frequency, and that frequency is a continuous function of the input current. This allows us to derive rate models, for which we characterise a single neuron by a function giving the firing frequency in terms of the input current. However, many neurons do not fire periodically when stimulated with a constant input but instead fire bursts of action potentials (Doiron et al., 2002; Izhikevich, 2000). For these neurons, such a characterisation is not appropriate, and it may not be possible to derive a rate-based approximation.

In this section we analyse a network of model "ghostbursting" neurons (Doiron et al., 2002), coupled with slow excitatory synapses. These neurons are found in the electrosensory lateral line lobe of the weakly electric fish *Apteronotus leptorhynchus* and are thought to be involved in electrosensory processing (Doiron et al., 2002). Although these cells receive input from the outside world via electroreceptors, the majority of their input is via feedback loops (Berman and Maler, 1999). We have coupled them with synapses having a time-constant of 100 msec, where the longest time-constant associated with the intrinsic neuron dynamics is 5 msec. The equations are given in Appendix A, along with a discussion of the implementation of initial conditions.

The behaviour of a single isolated ghostbursting neuron [Eqs. (59)–(64) with $A = 0$] as the input current, I, is varied is shown in Fig. 14. The top panel shows the instantaneous firing frequencies (i.e. reciprocals of the interspike intervals) over a period of 400 msec, for different values of I. There are three different types of firing behaviour: (i) periodic firing ($5.6 < I < 8.6$), (ii) bursting ($8.6 < I < 18.6$) and (iii) "doublet" firing ($18.6 < I$), in which the interspike intervals are alternately short and long (Doiron et al., 2002). Note that there are many bifurcations as I is varied in the bursting regime. The bottom panel shows the average firing frequency as a function of I. The non-smoothness of this curve is due to the bifurcations that occur as I is varied, and the curve does not become smoother if the averaging is over more than 400 msec (data not shown).

The results for the coupled system [Eqs. (59)–(65) with $A = 0.2$] are shown in Fig. 15, where we plot Q, the average synaptic strength, as a function of the input current to the neuron somas (all equal). As expected, the excitatory nature of the coupling causes bistability to occur, with stable firing now possible at values of I below the threshold for the onset of firing in an isolated neuron ($I \approx 5.6$). (Quiescence is also stable for these values of I.) With reference to Fig. 14, we can divide the curve in Fig. 15 into three sections, depending

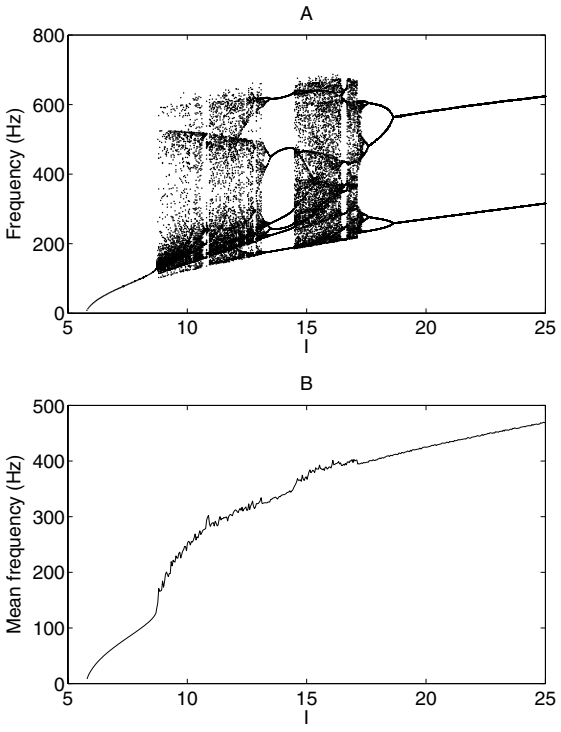

Fig. 14. **A:** Instantaneous firing frequencies during an interval of 400 msec, as a function of I, for a single ghostbursting neuron [Eqs. (59)–(64) with $A = 0$]. **B:** Mean firing frequency. Transients have been eliminated.

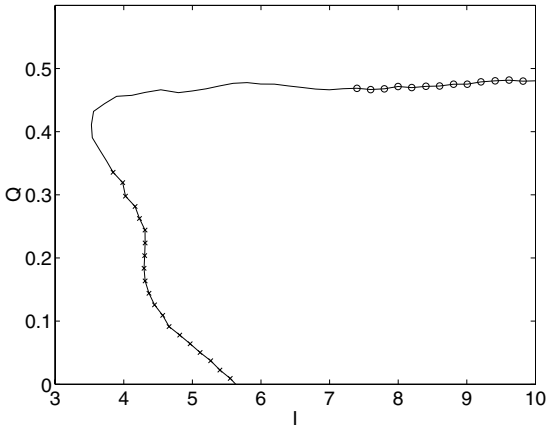

Fig. 15. Average synaptic strength, Q, as a function of input current, I, for a network of excitatorily coupled ghostbursting neurons (59)–(65). Crosses indicate periodic firing, a solid line bursting, and circles indicate where each neuron is firing "doublets". The upper branch ($Q > 0.4$) is stable, which the lower branch ($Q < 0.4$) is unstable. Coupling strength is $A = 0.2$, and we use a network of $N = 15$ neurons, averaging $N_{av} = 10$ times.

on whether the neurons are firing periodically, bursting, or firing doublets. We see that on the stable branch, the neurons are either bursting or firing doublets, while on the unstable branch they are either bursting or firing periodically.

The curve in Fig. 15 is relatively smooth, in comparison with the curve in Fig. 14 (B), presumably because coupling the neurons in this all-to-all fashion and measuring an average quantity "smears out" the fine structure seen in Fig. 14. While knowing the details of the fine structure may be of interest when one is analysing the behaviour of an individual neuron, it could be argued that the behaviour shown in Fig. 15 is a more appropriate representation of the macroscopic fixed points of the coupled system.

To attempt to derive a macroscopic equation governing the behaviour of this system, one could take the equation for the evolution of the slow variables (65)

$$100\frac{dq^i}{dt} = \sigma(V_s^i)(1 - q^i) - q^i \tag{52}$$

and replace it by

$$100\frac{dQ}{dt} = f(I, Q)(1 - Q) - Q \tag{53}$$

where

$$f(I, \overline{Q}) = \lim_{T\to\infty} \frac{1}{T} \int_0^T \sigma(V_s(t))\, dt \tag{54}$$

and $V_s(t)$ is taken from an isolated neuron (Eqs. (59)–(65)) with input I and Q set to $\overline{Q}$. Zeroes of (53) would then presumably be close to those plotted in Fig. 15. Near the onset of periodic firing f has a stereotypical shape, and the method of Ermentrout (Ermentrout, 1994) could be used, i.e., we could find an explicit formula for f, valid over some region in parameter space. However, in general (and particularly for a bursting neuron) $f(I, \overline{Q})$ can be determined only by extensive numerical simulation and, as seen in Fig. 14 (B), the average firing frequency is not necessarily a smooth function of I and Q, and thus could not be accurately approximated by, for example, a low-order polynomial. Also, (53) is only valid if the neurons are identical. If the neurons all had different values of the conductance $g_{Dr,d}$, for example, it is known that the bifurcations that they undergo as I is increased would differ (Doiron et al., 2002), and the function $f(I, \overline{Q})$ would then have to be calculated for each neuron separately and averaged to give an equation like (53).

We now estimate the relative costs of deriving a function like $f(I, \overline{Q})$ versus doing the computations to produce Fig. 15. Suppose that we estimate $f(I, \overline{Q})$ on a 100×100 grid in the $(I, \overline{Q})$ plane and then interpolate between these data points. We could simulate for 1 s at each point, meaning that T in (54) is at most 1 s (we discard transients). This is equivalent to simulating one neuron for 10,000 s. Alternatively, there are about 50 data points in Fig. 15, each of which took nine "bursts" of simulation to obtain (three Newton steps per point, each of which took three bursts to estimate two partial derivatives). We used 15 neurons, averaging 10 times, with bursts being 100 ms

long, resulting in effectively simulating one neuron for 6,750 s. Although obviously these numbers can be varied and do not represent the total cost they still indicate that for a system like this, without an analytic f–I curve for each neuron, the EF approach can compete with more traditional methods, particularly if the system is heterogeneous.

In terms of implementation, there is little difference between the network of bursting neurons studied here and the network of integrate-and-fire neurons studied in Section 2—it is only the complexity of the individual neuron model that has changed. Thus this method could easily be extended to networks of very detailed model neurons (for example, those of Doiron et al., 2001) with, for example, less-ordered connectivities.

Although this example is not particularly biologically realistic, since in practice the positive feedback is faster than that modelled here, it does demonstrate that the ideas put forward here can be used with realistic models that incorporate ion channel dynamics, rather than just integrate-and-fire models. It also demonstrates how to initialise the "fast" variables in the microscopic description when an explicit probability density function for them is not known (see Appendix).

6. A noisy network

In this section we discuss a network of coupled excitatory and inhibitory neurons that could be used to store a single bit of information, i.e. it is bistable for some range of parameters, as is the network in Section 2. We use integrate-and-fire neurons, as in Sections 2–4, but now all of the neurons are subject to high levels of noise, in the form of randomly occurring synaptic events. We include heterogeneity in the excitatory population, and assume that the excitatory synapses are slow and inhibitory ones fast. The system can be thought of as describing spatially uniform states of previous working memory models (Compte et al., 2000; Gutkin et al., 2001). The high levels of noise and the fact that the inhibitory synapses are not slow precludes the derivation of a rate equation for this model.

6.1. The model

We have N_e excitatory neurons and N_i inhibitory, with coupling both within and between populations. The equations are

$$\frac{dV_e^j}{dt} = I + \sigma^j - V_e^j + g_{ee}(E_e - V_e^j)(S_e + \nu_e^j(t)) + g_{ie}(E_i - V_e^j)(S_i + \nu_i^j(t)) - \sum_k \delta(t - t_e^{jk}) \tag{55}$$

$$\tau_e \frac{ds_e^j}{dt} = A \sum_k \delta(t - t_e^{jk})(1 - s_e^j) - s_e^j \tag{56}$$

$$\frac{dV_i^m}{dt} = I_i - V_i^m + g_{ei}(E_e - V_i^m)(S_e + \nu_e^m(t)) + g_{ii}(E_i - V_i^m)(S_i + \nu_i^m(t)) - \sum_n \delta(t - t_i^{mn}) \tag{57}$$

$$\tau_i \frac{ds_i^m}{dt} = A \sum_n \delta(t - t_i^{mn})(1 - s_i^m) - s_i^m \tag{58}$$

for $j = 1, \ldots, N_e$ and $m = 1, \ldots, N_i$. The subscripts on the variables and parameters label their type (excitatory/inhibitory), and the superscripts index them. t_e^{jk} is the kth firing time of the jth excitatory neuron, and similarly for t_i^{mn}. We have

$$S_e = \frac{1}{N_e}\sum_{j=1}^{N_e} s_e^j \qquad S_i = \frac{1}{N_i}\sum_{m=1}^{N_i} s_i^m$$

The functions $\nu_{e/i}(t)$ (mimicking randomly occuring excitatory/inhibitory synaptic activity) are formed from the sum of pulses of the form $0.1e^{-t/2}$ $(t > 0)$, whose arrival times are chosen from a Poisson process whose mean rate is 0.2. There are no correlations between arrival times for different neurons, nor between random inhibitory and excitatory input.

Reversal potentials are $E_e = 1.4$, $E_i = 0.5$. Other parameters are $A = 0.3$, $\tau_e = 30$, $\tau_i = 1$, $I_i = 0.9$, $g_{ee} = 5$, $g_{ei} = 2$, $g_{ie} = 1$. We vary g_{ii}. The heterogeneity in the excitatory population comes from choosing $\sigma^j = -0.1 + 0.2(j - 1)/(N_e - 1)$, i.e. the current offsets for the excitatory neurons range uniformly from -0.1 to 0.1. We use $N_e = 160$ and $N_i = 40$, to model the approximate observed ratio of cell types (Compte et al., 2000). Note that the excitatory synapses are the only slow variables in the network, mimicking, for example, NMDA-type transmission (Compte et al., 2000).

For this model, we assume that the macroscopic description of the system is given by $dS_e/dt = F(S_e, I)$, in the usual way. To evaluate $F(\widehat{S}_e, I)$, we choose to initialise the $V_e^j(0)$ and $V_i^m(0)$ from the distribution (10) with $I = 1.1$ (the choice of this value of I was arbitrary), we set $s_e^j(0) = \widehat{S}_e$ for each j, and we choose the $s_i^m(0)$ from a uniform distribution between 0 and A. A more accurate lifting operator, M, could be chosen, but the numerical results indicate that this one is sufficient, and it is easy to implement. This confirms that the initialisation of the fast variables need not be particularly

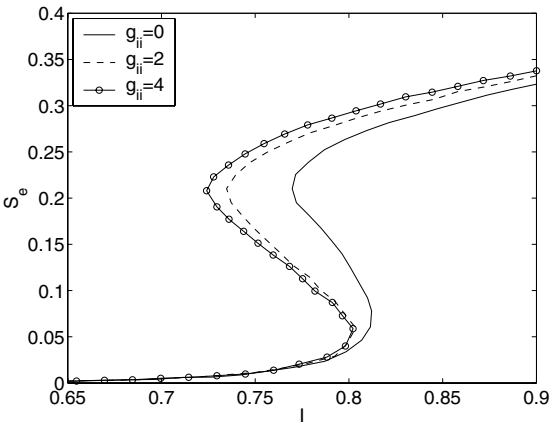

Fig. 16. Steady state values of S_e as a function of I for three different values of g_{ii}, for the system (55)–(58). For each curve, the upper and lower branches are stable, while the middle one is unstable. Parameters are $N_e = 160$, $N_i = 40$ and we average over 20 realisations. Other parameters are given in the text.

accurate, as their probability distribution is quickly slaved to the dynamics of the slow variables.

6.2. Results and discussion

An example of the sort of results we can obtain is shown in Fig. 16, where we plot the steady state values of S_e (the average strength of synaptic connections from the excitatory population) as a function of I (the average current injected into the excitatory population), for three different values of the inhibitory-to-inhibitory conductance (g_{ii}). We see that increasing g_{ii} increases the range of values of I for which the system is bistable, for the indicated values of the other parameters. Note that this parameter, g_{ii}, directly affects only the dynamics of some of the fast variables, in much the same way that κ does in the model in Section 4.5. Of course, much more can be determined about this network, in a similar way.

There are several reasons why a rate-based description of this system cannot be analytically derived. One is the high level of noise in the system, in the form of randomly arriving synaptic input. This is of such a high intensity and form that it is not meaningful to derive a firing rate function for an uncoupled individual neuron such that given in (18) (which is appropriate only when the noise appears as Gaussian white noise). Another reason is that while the excitatory synapses are slow the inhibitory ones are not. Thus it is not possible to treat S_i as approximately constant, so we cannot derive the firing rate of an individual neuron as a function of S_e and S_i, as Ermentrout does (Ermentrout, 1994).

7. Conclusion and discussion

We have presented here some applications of "equation-free" modelling to simple neural systems. For some of the systems, we have recovered results that could have been obtained more easily through the derivation of rate equations, which provides a useful check of their success. For other systems, we have found new behaviour that is obtainable through these methods.

Similar work relating to the derivation of rate models from spiking models has been done previously. Shriki et al., (2003) perform such a derivation without any assumption of a separation of timescales. However, these authors have to assume asynchronous firing, which we do not (see below), and also that each neuron is well-described by an *f–I* curve. This rules out the analysis of networks of bursting neurons (such as those studied in Section 5), very noisy networks (such as in Section 6) and neurons whose intrinsic dynamics are bistable (which we have not studied). In a similar vein, Ermentrout (1994) showed that for type I neurons (which have a continuous and single-valued *f–I* curve) coupled by slow synapses, approximate differential equations for the evolution of synaptic strengths could be derived. This is essentially the same technique as we used to derive the rate equations (14), (29) and (30). However, as Ermentrout states, these ideas cannot be used when the neurons intrinsically burst. Also, his method still results in one equation per neuron, which makes studying even moderately-sized networks difficult (assuming that the neurons are not identical). In contrast, the EF approach assumes the existence of a small number of relevant macroscopic variables, independent of the network size, and copes well with heterogeneous, bursting, and noisy networks.

A crucial requirement for the success of EF methods is the separation of time-scales (although see the remark in Section 2.5). Without this, by the time the fast variables have become "slaved" (i.e. after the time δ in Fig. 1) the slow variable(s) will have also changed significantly, so that a linear approximation to their time derivative will no longer be valid, and the method is no longer successful. We have only considered slow synapses, but one hallmark of neural systems is their wide range of time-scales, from the submillisecond opening of ion channels to synaptic plasticity that may occur over seconds or longer (Varela et al., 1997), and the ideas presented here should be more widely applicable than just to slow synapses. At a more technical level, for the EF approach to work one must assume that there is a low-dimensional macroscopic description of the state of the system which is complete in the sense that the dynamics of the variables in this description depend only on the values of those variables, i.e. the system is closed (Makeev et al., 2002). In practice this is difficult, if not impossible, to show, but the usefulness of this method lies in the results that can be obtained simply by applying the methods, without too much concern for their justification.

One issue we have not addressed is the analysis of periodic orbits. The ideas presented here can be used to follow both

stable and unstable periodic orbits in the phase space of the macroscopic variables (and detect their bifurcations) by treating them as fixed points of a return map on a Poincaré section. However, one must now consider the relationship between the period of the periodic orbit and the intrinsic time-scales of the underlying system. For example, it may be necessary to "lift-evolve-restrict" the macroscopic variables several times in succession in order to complete one period of oscillation (Makeev et al., 2002).

Another issue not addressed here is synchronisation (although see the remark in Section 2.5). It is well known that populations of coupled model neurons can synchronise, and such synchronisation is often thought to be important (Gerstner and Kistler, 2002). In the work discussed here we have always assumed that the populations of neurons were not synchronised, by choosing microscopic initial conditions for each neuron from probability density functions of isolated neurons. This, combined with short simulation times and the addition of noise, ensured that our simulations accurately described states of asynchronous firing. (In the bump states discussed in Section 4, synchronous firing is not possible (Laing and Chow, 2001).) However, we can still use the ideas presented here even if neurons become synchronised at the spike-to-spike level. In this case, we effectively have a population of $N = 1$ neuron. As discussed at the end of Section 2, this just means that we need to increase N_{av} in order to obtain accurate estimates of the derivatives of the slow variable(s). Conceptually, we can see that since our macroscopic variables change on a slow time-scale relative to interspike intervals, whether neurons are synchronised or not on a spike-to-spike level will not affect the rate of change of those variables.

Although we have largely studied stochastic systems here, this is not a requirement for the EF method to work, as is demonstrated in Section 5 where we study a small network of deterministic bursting neurons. Adding noise to the neuron models brings them closer to reality (Koch, 1999), and has the added advantage of smoothing out some of the functions that determine the curves of fixed points, making the numerical work easier.

The ideas presented here have also been applied to continuum models with heterogeneous domains (Runborg et al., 2002), and it would be interesting to apply them to neural field models of the form (44) with, for example, a periodic spatial modulation of the current, I. Bressloff has studied a similar situation in the context of travelling waves, using homogenization and averaging (Bressloff, 2001), and it would be interesting to compare the results of the two different approaches. Another area of interest would be to study a spatially-extended system that is explicitly discrete (i.e. not resulting from discretizing a continuous system) from the point of view of there being an effective continuum description of, for example, travelling waves (Möller et al., 2005).

There is nothing in this type of analysis that restricts the model neurons to be identical—indeed, the excitatory population in Section 6 is heterogeneous—or even of the same type or described at the same level of complexity. In principle, neurons in different populations could be described with different levels of complexity. For example, excitatory neurons could be modelled using the Hodgkin-Huxley-type formalism, while inhibitory neurons could be of a more generic nature, e.g. integrate-and-fire.

In summary, we have demonstrated for several neural models a new technique for the analysis of systems for which a microscopic description is available but for which macroscopic equations cannot or have not been derived. Many of the results presented here could have been derived by using a rate model description with much less computational effort. Specifically, the results derived using EF modelling in Sections 2.3 and some of Section 4 were well-approximated by results derived from a rate description. However, new results have been found (in Section 4.5 where we investigate the effects of including gap junctions in a model for "bump" formation, in Section 5 describing the coupled ghostbursters, and in Section 6, where we study a noisy network with fast inhibitory synapses). We believe that the ideas presented here are widely applicable and will be valuable in bridging the gap between detailed models of single neurons and a macroscopic "coarse" description of the dynamics of networks of them.

Appendix A. Ghostbursting neuron model

We now give the equations for the network of "ghostbursting" model neurons discussed in Section 5. The single neuron model we use has been presented and analysed elsewhere (Doiron et al., 2002). For a network of N all-to-all excitatorily coupled neurons the equations are

$$\begin{aligned} C\frac{dV_s^i}{dt} &= I + g_{Na,s}m_{\infty,s}^2(V_s^i)(1-n_s^i)(V_{Na}-V_s^i) \\ &\quad + g_{Dr,s}(n_s^i)^2(V_K - V_s^i) + g_c(V_d^i - V_s^i)/\kappa \\ &\quad + g_{leak}(V_l - V_s^i) \end{aligned} \tag{59}$$

$$\tau_{n,s}\frac{dn_s^i}{dt} = n_{\infty,s}(V_s^i) - n_s^i \tag{60}$$

$$\begin{aligned} C\frac{dV_d^i}{dt} &= -AQ(V_d^i - V_e) + g_{Na,d}m_{\infty,d}^2(V_d^i)h_d^i \\ &\quad \times (V_{Na} - V_d^i) + g_{Dr,d}(n_d^i)^2 p_d^i(V_K - V_d^i) \\ &\quad + g_c(V_s^i - V_d^i)/(1-\kappa) + g_{leak}(V_l - V_d^i) \end{aligned} \tag{61}$$

$$\tau_{h,d}\frac{dh_d^i}{dt} = h_{\infty,d}(V_d^i) - h_d^i \tag{62}$$

$$\tau_{n,d}\frac{dn_d^i}{dt} = n_{\infty,d}(V_d^i) - n_d^i \tag{63}$$

$$\tau_{p,d}\frac{dp_d^i}{dt} = p_{\infty,d}(V_d^i) - p_d^i \tag{64}$$

$$100\frac{dq^i}{dt} = \sigma(V_s^i)(1-q^i) - q^i \tag{65}$$

for $i = 1, \ldots, N$, where a subscript "s/d" denotes a somatic/dendritic variable, the superscript labels the neuron, and

$$Q \equiv \frac{1}{N}\sum_{i=1}^{N} q^i \tag{66}$$

Parameters are $g_{Na,s} = 55$, $g_{Dr,s} = 20$, $g_c = 1$, $g_{leak} = 0.18$, $g_{Na,d} = 5$, $g_{Dr,d} = 15$, (in units of mS/cm^2) $\tau_{n,s} = 0.39$, $\tau_{h,d} = 1$, $\tau_{n,d} = 0.9$, $\tau_{p,d} = 5$ (in units of ms), $V_{Na} = 40\,\text{mV}$, $V_K = -88.5\,\text{mV}$, $V_l = -70\,\text{mV}$, $V_e = 0\,\text{mV}$, $C = 1\mu\text{F/cm}^2$. The feedback strength A is specified in the text.

Functions are

$$m_{\infty,s}(V) = \{1 + \exp[-(V+40)/3]\}^{-1} \tag{67}$$

$$n_{\infty,s}(V) = \{1 + \exp[-(V+40)/3]\}^{-1} \tag{68}$$

$$m_{\infty,d}(V) = \{1 + \exp[-(V+40)/5]\}^{-1} \tag{69}$$

$$h_{\infty,d}(V) = \{1 + \exp[(V+52)/5]\}^{-1} \tag{70}$$

$$n_{\infty,d}(V) = \{1 + \exp[-(V+40)/5]\}^{-1} \tag{71}$$

$$p_{\infty,d}(V) = \{1 + \exp[(V+65)/6]\}^{-1} \tag{72}$$

$$\sigma(V) = 5/\{1 + \exp[-(V+20)/5]\} \tag{73}$$

The macroscopic variable is Q. To implement the microscopic initial conditions, we chose $q^i(0) = Q(0)$ for $i = 1, \ldots, N$. To choose the initial conditions for the other variables, we first ran a simulation of an isolated neuron for 70,000 time steps, during which I was linearly increased from $I = 7.75$ to $I = 18.25$, thus sweeping through a range of possible states in which a neuron could be. To initialise N neurons we randomly choose N integers from a uniform distribution on $[1, 70{,}000]$, say $k_1, \ldots, k_N$. The six initial variables of the ith neuron, $V_s^i(0), n_s^i(0), \ldots, p_d^i(0)$ are then set to the state that the "test" neuron was in at the k_ith timestep of the sweep. The integers $k_1, \ldots, k_N$ are rechosen at the start of each simulation. In this way we hope to initialise each neuron "close" to a state consistent with the current value of Q.

Acknowledgments I thank Alona Ben-Tal for carefully reading the manuscript, and the referees for their useful comments. This work was supported by the Marsden Fund administered by the Royal Society of New Zealand.

References

Berman NJ, Maler L (1999) Neural architecture of the electrosensory lateral line lobe: Adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. J. Exp. Biol. 202: 1243–1253.

Bressloff PC (2001) Traveling fronts and wave propagation failure in an inhomogeneous neural network. Physica D 155: 83–100.

Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. Neural Comput. 11: 1621–1671.

Chow CC, Kopell N (2000) Dynamics of spiking neurons with electrical coupling. Neural Comput. 12: 1643–1678.

Compte A, Brunel N, Goldman-Rakic PS, Wang X-J (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex 10: 910–923.

Doedel E, Keller HB, Kernevez JP (1991) Numerical analysis and control of bifurcation problems. (I) Bifurcation in finite dimensions. Int. J. Bifn. Chaos. 1: 493–520.

Doiron B, Chacron MJ, Maler L, Longtin A, Bastian J (2003) Inhibitory feedback required for network oscillatory responses to communication but not prey stimuli. Nature. 421: 539–543.

Doiron B, Laing CR, Longtin A, Maler L (2002) Ghostbursting: A novel neuronal burst mechanism. J. Comput. Neurosci. 12: 5–25.

Doiron B, Longtin A, Turner RW, Maler L (2001) Model of gamma frequency burst discharge generated by conditional backpropagation. J. Neurophysiol. 86: 1523–1545.

Ermentrout B (1994) Reduction of conductance-based models with slow synapses to neural nets. Neural Comput. 6: 679–695.

Fourcaud N, Brunel N (2002) Dynamics of the firing probability of noisy integrate-and-fire neurons. Neural Comput. 14: 2057–2110.

Garcia-Archilla B, Novo J, Titi, ES (1998) Postprocessing the Galerkin method: A novel approach to approximate inertial manifolds. SIAM J. Numer. Anal. 35: 941–972.

Gear CW, Kevrekidis IG, Theodoropoulos C (2002) "Coarse" integration/bifurcation analysis via microscopic simulators: Micro-Galerkin methods. Comp. Chem. Eng. 26: 941–963.

Gerstner W, Kistler W (2002) Spiking neuron models: Single neurons, populations, plasticity. Cambridge University Press.

Givon D, Kupferman R, Stuart A (2004) Extracting macroscopic dynamics: Model problems and algorithms. Nonlinearity 17: R55–R127.

Gutkin BS, Laing CR, Colby CL, Chow CC, Ermentrout GB (2001). Turning on and off with excitation: The role of spike-timing asynchrony and synchrony in sustained neural activity. J. Comput. Neurosci. 11: 121–134.

Izhikevich EM (2000) Neural excitability, spiking and bursting. Int. J. Bifn. Chaos. 10: 1171–1266.

Kevrekidis IG, Gear CW, Hyman JM, Kevrekidis PG, Runborg O, Theodoropoulos C (2003) Equation-free, coarse-grained multiscale computation: Enabling microscopic simulators to perform system-level analysis. Comm. Math. Sci. 1: 715–762.

Koch C (1999) Biophysics of computation: Information processing in single neurons. Oxford University Press.

Kopell N, Ermentrout B (2004) Chemical and electrical synapses perform complementary roles in the synchronization of interneuronal networks. Proc. Nat. Acad. Sci. USA 101: 15482-1-5487.

Laing CR, Chow CC (2001) Stationary bumps in networks of spiking neurons. Neural Comput. 13: 1473–1494.

Laing CR, Longtin A (2003) Dynamics of deterministic and stochastic paired excitatory-inhibitory delayed feedback. Neural Comput. 15: 2779–2822.

Laing CR, Troy WC (2003) Two-bump solutions of Amari-type models of neuronal pattern formation. Physica D 187: 190–218.

Laing CR, Troy WC, Gutkin B, Ermentrout GB (2002) Multiple bumps in a neuronal model of working memory. SIAM J. Appl. Math. 63: 62–97.

Makeev AG, Maroudas D, Kevrekedis IG (2002) "Coarse" stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples. J. Chem. Phys. 116: 10083–10091.

Möller J, Runborg O, Kevrekidis PG, Lust K, Kevrekidis IG (2005) Equation-free, effective computation for discrete systems: A time stepper based approach. Int. J. Bifn. Chaos. 15: 975–996.

Runborg O, Theodoropoulos C, Kevrekidis IG (2002) Effective bifurcation analysis: A time-stepper-based approach. Nonlinearity. 15: 491–511.

Shriki O, Hansel D, Sompolinsky H (2003) Rate models for conductance-based cortical neuronal networks. Neural Comput. 15: 1809–1841.

Varela JA, Sen K, Gibson J, Fost J, Abbott LF, Nelson SB (1997) A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. J. Neurosci. 17: 7926–7940.