Chapter 9

Networks of neurons

An essential component of the art of modelling is to carry out appropriate simplifications. This is particularly important when modelling networks of neurons. Generally, it is not possible to represent each neuron of the real system in the model, and so many design questions have to be asked. The principal questions concern the number of neurons in the model network, how each neuron should be modelled and how the neurons should interact. To illustrate how these questions are addressed, different types of model are described. These range from a series of network models of associative memory, in which both neurons and synapses are represented as simple binary or multistate devices, two different models of thalamocortical interactions, in which the neurons are represented either as multi-compartmental neurons or as spiking neurons, and multi-compartmental models of the basal ganglia and their use in understanding Parkinson's disease. The advantages and disadvantages of these different types of model are discussed.

In their implementation of Marr's influential theory of cerebellar cortex as a learning machine (Marr, 1969), Tyrrell and Willshaw (1992) constructed a simulation model of all the circuitry associated with a single Purkinje cell. With the limited computing resources available at the time, they did this by modelling each 3D layer of cells and connections in a 2D plane. To build the model they had to guess many parameter values about the geometry as these were not available. Their simulation results agreed broadly with the analysis carried out by Marr.

Two severe limitations prevent the modeller from constructing a model of a neural system in which each nerve cell is represented directly by a counterpart model neuron. One limitation is that there are so many neurons in the neural system that having a full-scale model is computationally infeasible. The second limitation is that usually only incomplete data is available about the functional and structural properties of the neurons, how they are arranged in space and how they interconnect.

The design issues that are most commonly addressed concern the numbers and types of model neurons and the topology of how they connect with each other. Another crucially important issue is how the cells are situated in 3D space. Since the embedding of network models in space is not normally attempted, this issue has not often been discussed, with some notable exceptions. An early attempt is the simulation model of Marr's theory of **cerebellar cortex** (Marr, 1969; Tyrrell and Willshaw, 1992). In Section 9.1 we consider these design issues.

The most common properties that are investigated in network models of the nervous system are the patterns of firing within the array of neurons and how such patterns are modified through specific synaptic learning rules. In this chapter, we examine these two properties in a variety of network

models in which the neurons are modelled to differing levels of detail. We start by looking at very simple networks where both neurons and synapses are modelled as two-state devices. In these networks most emphasis has been on the effects of synaptic modification. Accordingly, in Section 9.2 we describe three different approaches to constructing generic models of network associative memory in which the neuron is treated as a two-state device and the modifiable synapse as a simple binary or linear device. We show that one advantage of extreme simplification is that analytical results for how these networks can be used most efficiently can be obtained and the capacity of the system can be calculated.

For networks of more complex neurons, it is important to characterise their firing patterns before associative storage can be assessed. In Section 9.3 we examine an integrate-and-fire network model of a cortical column, and we explore associative storage and retrieval in this network. In Section 9.4 we look at network models of more complex, multi-compartmental model neurons, again looking at how associative memory can be embedded in them.

Section 9.5 contains three examples of modelling of thalamocortical connections using model neurons of different complexity. These large-scale network models are used to examine network phenomena such as oscillatory neural activity as recorded in electroencephalograms (EEGs).

Finally, we look at a clinically related application, in which the emphasis is on the patterns of activity under normal and abnormal conditions. In Section 9.6 we discuss how to model the effects of deep brain stimulation of the subthalamic nucleus in the basal ganglia, now used successfully for the relief of Parkinson's disease. We describe a multi-compartmental network model of the subthalamic nucleus and related structures and discuss the validation and predictions made from the model.

9.1 Network design and construction

In the preceding chapters we have seen that the construction of the model of a single neuron involves a vast range of choices concerning how to model components such as cell morphology, ion channels and synaptic contacts. Each choice involves a compromise over the level of biological detail to include. How to make useful simplifications is an important part of the modelling process.

The same is true if we want to build a network of neurons. A major decision is to choose at which level of detail to model the individual neurons. For a large-scale network with thousands, or hundreds of thousands of neurons, this may require using the simplified models introduced in the previous chapter. Other issues also arise with network models. How should we handle communication between neurons? Do we need to model axons and the propagation of action potentials along them? Do we need to model short-term dynamics and stochastic neurotransmitter release at synapses? Our network is almost certainly going to be smaller than real size in terms of the numbers of neurons. In which case, how should we scale the numbers of

Some neurobiological systems contain a small number of neurons enabling neurons to be represented one-to-one in the model. For examples see Abbott and Marder (1998).

neurons of different classes in the network? Finally, do we need to consider the location of neurons in space? In this section we explore possible answers to these questions.

9.1.1 Connecting neurons together

Networks of neurons are formed predominantly through neurons connecting together via chemical synapses formed between axonal terminals from efferent neurons and the postsynaptic membrane of receiving neurons. The signal that passes from the efferent to receiving neuron is the action potential. A possibility for modelling these connection pathways is to include in each cell model a compartmental model of its axon along which action potentials propagate. This is computationally very expensive and arguably unnecessary. Action potentials are stereotypical and the information content of signals passing from one neuron to another is carried by the times of action potentials arriving at synapses, rather than the precise voltage waveform of the action potential.

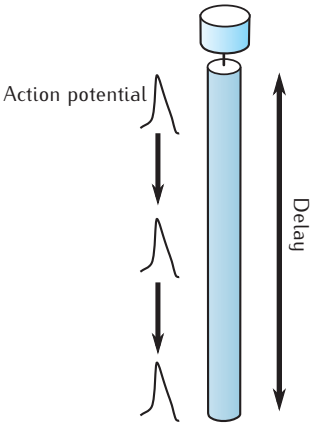

Fig. 9.1 An action potential is initiated in the axon initial segment and propagates along the axon. This can be modelled as a delay line, which specifies the time taken for the action potential to travel the length of the axon. The action potential itself is not modelled.

Consequently, the approach that is almost uniformly applied is to treat the signal that passes from one neuron to another to be the presence or absence of an action potential. Then the connection from one neuron to another is modelled as a **delay line** (Figure 9.1). The voltage in the soma or axon initial segment of the efferent cell is monitored continuously. If the voltage goes over a defined threshold (e.g. 0 mV), this signals the occurrence of an action potential. The delay line then signals this occurrence to the synaptic contact on the receiving neuron at a defined time later, corresponding to the expected transmission time of the action potential along the real axon. This approach is not only vastly cheaper computationally than compartmental modelling of axons, but it is also easily implemented on parallel computers, as only spike times need to be sent between processors (Brette *et al.*, 2007; Hines and Carnevale, 2008).

There are circumstances where it is necessary to model the detail of action potential propagation along axons. The delay line model assumes that action potential propagation is entirely reliable and is not modulated along the length of the axon. The possibility of action potential failure at branch points or due to presynaptic inhibition are ignored. These effects have been explored using compartmental models of isolated axons (Parnas and Segev, 1979; Segev, 1990; Manor *et al.*, 1991a, b; Graham and Redman, 1994; Walmsley *et al.*, 1995). They could certainly be expected to influence network dynamics and thus raise the challenge of modelling action potential propagation in a network model.

9.1.2 Scaling neuronal numbers

Many networks of interest contain thousands or even millions of neurons, which it is often not feasible to model. It is then necessary to model a scaled-down version of the actual network. This involves scaling both the numbers of neurons and the number of synapses between neurons.

Suppose our network is going to be one-tenth the size of the brain nucleus we are modelling. This nucleus contains three cell types – a principal excitatory neuron that makes up 80% of the cell population, and two types of inhibitory interneuron, each constituting about 10% of the

population. The obvious way to scale neuronal numbers is to retain the relative proportions of cells of different types (80:10:10) in our one-tenth-sized model. Provided this results in reasonable numbers of cells of each type in the model, then this could be an appropriate choice. What may constitute a reasonable number of cells is discussed below.

The principal use of the model is likely to be to study the population response of the excitatory cells. For this to be an accurate reflection of physiology, it is important that the excitatory and inhibitory synaptic input onto these cells represents realistic population activities. In our model network, inhibition from each population of inhibitory interneuron should be as close as possible to that experienced by a real excitatory neuron in vivo. Given that we have fewer interneurons in our model network than exist in vivo, there are two ways of achieving this:

(1) Scale up the maximum synaptic conductance of each connection from an inhibitory interneuron onto an excitatory cell by a factor of ten, in this example.
(2) Create ten times the number of synaptic contacts from each interneuron onto each excitatory cell than exist in vivo.

Neither approach is perfect. Scaling the synaptic conductances may give an equivalent magnitude of inhibition. However, this will be applied as a few large conductance changes at isolated points on the excitatory cell. As discussed in more detail in Section 9.3.3, the resulting voltage changes and integration with excitatory inputs will be distorted. Creating more synaptic contacts from each interneuron will result in a realistic spatial distribution of inhibitory inputs, but spikes arriving at these inputs may have unnatural correlations since groups of them are more likely to derive from the same interneuron. Unless it is actually possible to include physiological numbers of interneurons in the network model, one of these compromises is required. The same considerations apply to the inputs from excitatory neurons.

If it is likely that different sized network models will be tested, it is very useful to fix the number of afferent inputs that a given cell receives from the population of cells of each type in the model. For example, the number of inhibitory inputs that each excitatory cell receives from each of two populations of inhibitory interneurons should remain fixed regardless of the actual number of each cell type in the model. When the number of cells is changed, a cell of a particular type will provide fewer or more synaptic contacts onto a target cell, but the target cell will always have the same number of synaptic inputs from the efferent cell population (Orbán *et al.*, 2006).

Another effect of scaling the numbers of neurons is that the small populations of interneurons may be scaled to the point of having only one or a few cells representing these populations in the model. In this case the population activity in the model of these interneurons may be a serious distortion of the activity in vivo. Real activity may involve thousands of asynchronously firing cells, with the instantaneous population activity providing a good estimate of some modulating driving force, such as slowly changing sensory input (Section 8.2.2; Knight, 1972; Hospedales *et al.*, 2008). The small population in the model may only provide a poor representation of the modulating input.

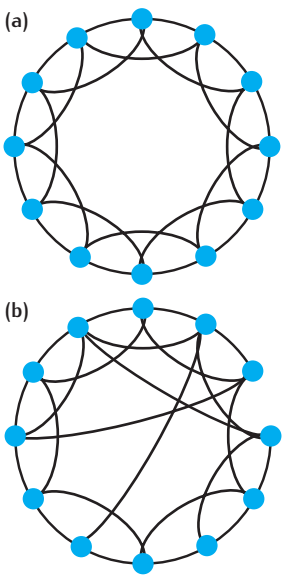

Fig. 9.2 **(a)** Local connectivity in which a neuron connects only to near neighbours. **(b)** Small-world connectivity in which some of the local connections are replaced by longer-range connections.

If this is the case, then it may be possible to scale each population of cells differently. If the excitatory cells are not strongly recurrently connected, then only a relatively small number of these cells are required in the model to allow a good study of their network activity (Orbán *et al.*, 2006). This then allows relatively larger populations of interneurons to be modelled so that both their population activity and their inhibitory effect on the excitatory cells are much more physiological. This approach was taken by Orbán *et al.* (2006) in a model of the CA1 area of hippocampus, where recurrent connections between pyramidal cells are sparse. Their network model of theta activity contained a small number (15–30) of detailed 256-compartment pyramidal cells, but with populations of up to 200 basket and 90 oriens lacunosum-moleculare cells, each cell modelled by a single compartment.

9.1.3 Positioning neurons in space

Real neurons have a particular location within a brain nucleus, and connectivity patterns between neurons are often distance-dependent. To capture these patterns it may be necessary to place our model neurons in virtual space.

For, say, a cortical column or other small part of a brain area, it may be reasonable to assume that connectivity is completely uniform (e.g. every neuron connects to every other neuron) or that there is a fixed probability that one neuron makes contact with another neuron. In this case the precise spatial location of a neuron is not relevant and can be ignored.

In general, though, we will need to lay our cells out in some 1D, 2D or 3D arrangement that reflects the physiological layout. Typically this is done with a regular spacing between cells. Then, when forming connections between cells, the probability that an efferent cell forms a connection onto a target cell can be a function of the distance between them. This function is often an exponential or Gaussian function so that the probability of connection decreases with distance (Figure 9.2a). This reflects the basic connection arrangement in many brain nuclei. More complex connection strategies can easily be implemented. So-called **small-world networks** (Watts and Strogatz, 1998; Netoff *et al.*, 2004; Földy *et al.*, 2005) can be generated by first creating a network with only local connections between cells (a cell connects to a few of its nearest neighbours) and then randomly reassigning a small proportion of the connections to be much longer-range connections (Figure 9.2b).

One problem to deal with is that of edge effects, in which cells at the edge of our spatial layout receive fewer connections than interior cells. This could be overcome by assuming the spatial arrangement actually wraps around, so that a cell at the end of the line is assumed to be a neighbour of the cell at the opposite end of the line (Netoff *et al.*, 2004; Wang *et al.*, 2004; Földy *et al.*, 2005; Santhakumar *et al.*, 2005), i.e. the line is actually a circle (Figure 9.2). A more biological solution might be to have a sufficiently large model network that the cells at the boundaries can be ignored.

9.1.4 Variability in cell properties

The vast majority of neuronal network models contain populations of cells with completely uniform properties, including morphology and membrane

physiology. This does not reflect the variation seen within biological neurons and can lead to artifacts in network behaviour due to uniformity in cellular responses to synaptic input. A better approach is to introduce variance into one or more cellular properties, including membrane resistance, resting membrane potential and ion channel densities. Experimental estimates of these parameters may be available that indicate the magnitude of variance in a biological population. Variations in electrophysiological responses may indicate variability in membrane ion channel densities (Aradi and Soltesz, 2002).

Alternatively, some variation can be introduced into a population of otherwise identical cell models by either starting a simulation with different initial conditions for each cell; e.g. different starting membrane potentials, or providing a different background stimulus, in the form of a small depolarising or hyperpolarising current injection, to each cell (Orbán *et al.*, 2006).

Computational models and experiments have shown that signal integration in cells and collective network behaviour is strongly influenced by variability in individual cell characteristics (Aradi and Soltesz, 2002; Aradi *et al.*, 2004). Another consideration is the relative proportion of cells of different types within the network. Classification of cell types is an art form that is still evolving (Somogyi and Klausberger, 2005; Markram, 2006). Thus the number of cell populations and their relative sizes may be free variables in the network model. Simulations have shown that networks containing the same cell types, but in different proportions, can show significantly different behaviour (Földy *et al.*, 2003, 2005).

9.1.5 New network quantities

Modelling a network of spatially located neurons, as opposed to a single, isolated neuron, allows for the possibility of modelling new signals in addition to cellular voltages. These include pseudo EEGs and **extracellular field potentials**. A basic field potential equation is (Rall, 1962; Protopapas *et al.*, 1998; Bédard *et al.*, 2004):

$$\Phi(x,y,z,t) = \frac{1}{4\pi\sigma}\sum_{i=1}^{n}\frac{I_i(t)}{d_i}, \tag{9.1}$$

where Φ is the field potential at a particular recording site (x,y,z), and each of the n current sources I_i is a distance d_i from the recording site. The conductivity of brain tissue is assumed to have a uniform value, σ, throughout. For a network of spatially located compartmental cell models, according to Equation 9.1, the current sources correspond to the membrane current in each compartment in each cell. Figure 9.3 gives example traces of the extracellular membrane potential calculated in this way in the vicinity of a schematic compartmental neuron model.

There are two principal limitations to using Equation 9.1. The first is that uniform extracellular conductance is an approximation to reality, and the second is that the extracellular medium has capacitive properties as well as conductive ones (Ranck, 1963; Bédard *et al.*, 2004). In general, the problem

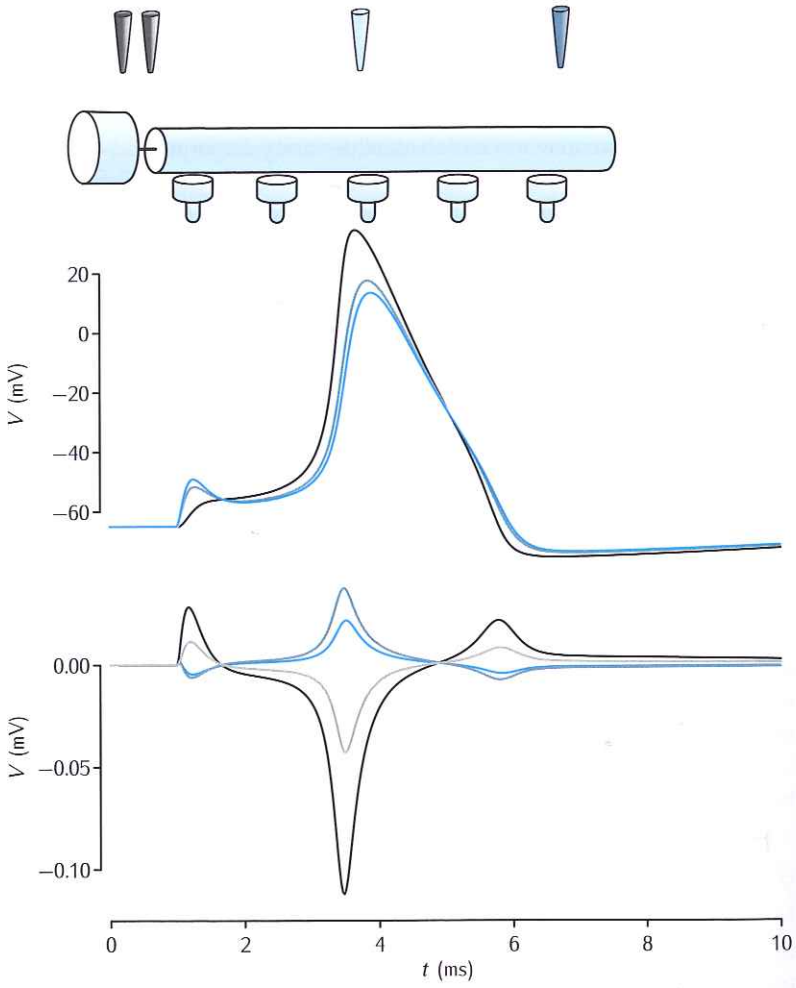

Fig. 9.3 Simulation of extracellular field potentials. Extracellular electrodes (black, grey, blue and dark-blue) are placed close to a ball-and-stick model neuron with an active soma and synapses on its dendrite. The soma is 40 μm long and 40 μm in diameter. The single dendritic cable is 200 μm long and 4 μm in diameter. The top traces show intracellular recordings when the synapses are activated enough to cause an action potential to be fired. Traces are from the soma (black), halfway down the dendrite (blue) and in the distal dendrite (dark-blue). The initial synaptic stimulation can be seen in the dendritic traces. The lower traces show the extracellular recordings corresponding to the electrodes of the same colour. During the synaptic stimulation, the dendrites act as a sink of extracellular current and the soma acts as a source. This can be seen in the negative deflection of the extracellular potential in medial and distal dendrites and the positive deflection of the extracellular potential close to the soma. During the action potential, the soma is a sink of current and the dendrites are current sources; this is reflected in the large negative deflection of the extracellular potential close to the soma and the smaller deflections of the extracellular potential near the dendrites. As the neuron repolarises, the roles of the soma and dendrites are again reversed.

of inhomogeneous media can be addressed using a finite-element model of the extracellular medium (Box 9.1). The second problem can be addressed by deriving a version of Equation 9.1 that incorporates capacitance (Bédard *et al.*, 2004). With capacitive properties included, the extracellular potential has its own time-dependent dynamics. The response to a periodic membrane current signal of a particular frequency can be computed. High-frequency signals are expected to attenuate more than low ones, meaning that the action potentials, which are very sharp and hence contain a lot of high-frequency components, will be highly attenuated at relatively small distances from the neuron. Fourier analysis can be used to predict the response to any particular time-varying signal (Bédard *et al.*, 2004).

Extracellular spatial concentration gradients of neuromodulators and volume transmitters, such as nitric oxide, can also be modelled using diffusion equations (Philippides *et al.*, 2000, 2005) or reaction–diffusion equations (Feng *et al.*, 2005b).

Box 9.1 Finite-element models of electric fields

The voltage spread in brain tissue due to an extracellular electrode can be modelled in a similar way to that described already for compartmental models of neurons. The extracellular space is modelled as a uniform conductance through which current flows, represented as a network of resistors:

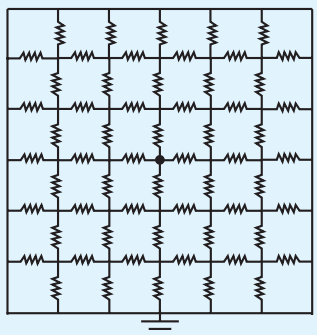

The electrode is the black disc and the return current to the electrode passes through the rim of the network. Here there is no capacitance, so the current balance equation for each internal node with no electrode is:

$$0 = \sum_{j \in \mathcal{N}_i} \frac{V_j - V_i}{R_{ij}},$$

where $\mathcal{N}_i$ is the set of nodes connected to node i. For the node containing the electrode, the zero on the left-hand side of the equation is replaced with the electrode current. For the nodes on the boundary, V_i is set to 0 mV, the potential of the return electrode. These equations can be formulated as a matrix equation and the steady state potentials can be computed (Butson and McIntyre, 2005).

In this example, there is no capacitance, and hence no dynamics, as there are no terms involving $\mathrm{d}V_i/\mathrm{d}t$. Capacitance can be incorporated in the model by replacing the real-valued resistance with a complex-valued impedance, in which the imaginary component represents capacitance. The amplitude and phase of the voltage at each node in response to an oscillating electrode current of any frequency can then be computed. Using Fourier analysis, this can be used to compute the response of each node to a periodic, time-varying electrode current (Butson and McIntyre, 2005).

In place of the regular square mesh used here, an irregular mesh may be used, in which the regions in which the voltage varies most have a finer mesh (Butson and McIntyre, 2005). Meshes can span 2D or 3D space. Software packages such as CALCULIX (`http://www.calculix.de`) are able to carry out finite-element analysis.