Thesis

Synchronization

in

Dynamic Neural Networks

David Edward Cairns

Ph.D.

Department of Computing Science and Mathematics

University of Stirling

Submitted in partial fulfilment of

the degree of Doctor of Philosophy

September 1993

Declaration

I declare that this thesis has been composed by myself and that the research reported therein has been conducted by myself unless otherwise indicated.

Stirling, 20th September 1993

David E. Cairns

Acknowledgements

First and foremost I would like to thank Leslie Smith for his patience and faith in guiding me through my Ph.D. He has always been available for discussion and has been prompt in commenting on drafts of this thesis. Without his consistent support I think the years spent doing this Ph.D would have been considerably harder.

I would like to thank my friends here at Stirling, both from the Centre for Cognitive and Computational Neuroscience and from the Computing Science Department. Their company has made the process of completing this Ph.D. far more enjoyable than I think it otherwise would have been. I would particularly like to thank Roland Baddeley for useful discussion and comment on some of the work in Chapter 5.

I am grateful to Sam, Graham and Catherine for providing equipment and technical support. They have always been helpful and available to solve the variety of problems I seem to be able to create.

Finally, I thank Claire for her love, support, humour and patience. Without her I doubt that I would have got this far.

This work has been funded by a SERC post-graduate grant.

Abstract

This thesis is concerned with the function and implementation of synchronization in networks of oscillators. Evidence for the existence of synchronization in cortex is reviewed and a suitable architecture for exhibiting synchronization is defined. A number of factors which affect the performance of synchronization in networks of laterally coupled oscillators are investigated. It is shown that altering the strength of the lateral connections between nodes and altering the connective scope of a network can be used to improve synchronization performance. It is also shown that complete connective scope is not required for global synchrony to occur. The effects of noise on synchronization performance are also investigated and it is shown that where an oscillator network is able to synchronize effectively, it will also be robust to a moderate level of noise in the lateral connections. Where a particular oscillator model shows poor synchronization performance, it is shown that noise in the lateral connections is capable of improving synchronization performance.

A number of applications of synchronizing oscillator networks are investigated. The use of synchronized oscillations to encode global binding information is investigated and the relationship between the form of grouping obtained and connective scope is discussed. The potential for using learning in synchronizing oscillator networks is illustrated and an investigation is made into the possibility of maintaining multiple phases in a network of synchronizing oscillators. It is concluded from these investigations that it is difficult to maintain multiple phases in the network architecture used throughout this thesis and a modified architecture capable of producing the required behaviour is demonstrated.

Contents

1 Introduction **1**

1.1 Thesis structure 1

1.2 Thesis contribution 3

2 Synchronization in neural systems **5**

2.1 Introduction 5

2.2 Evidence for synchronization 6

2.2.1 Synchronization in response to a causal event 7

2.2.2 The binding problem 9

2.2.3 Further evidence for stimulus evoked synchronization 11

2.2.4 Synchronized activation or phase locked coherence? 13

2.3 Modelling neural activity 14

2.3.1 The integrate and fire oscillator 14

2.3.2 Coupled excitatory/inhibitory systems 17

2.3.3 Phase response equations 19

2.4 Synchronization in artificial neural networks 20

2.4.1 The phase comparator 21

2.4.2 Lateral connectivity 24

2.4.3 A simple architecture 26

3 Synchronization in neural networks **30**

3.1 Experimental format 31

3.1.1 Architecture 31

3.1.2 General method 34

3.1.3 Measuring network coherence 35

3.2 Connection strength 38

3.2.1 Method 39

3.2.2 Discussion 43

3.3 Connective scope and coherence 45

3.3.1 Method 46

3.3.2 Discussion 47

3.4 Connective scope and speed of synchrony 53

3.4.1 Method 53

3.4.2 Discussion 54

3.5 Noise in dynamic neural systems 58

3.5.1 Modelling noise 58

3.5.2 The effects of noise on network coherence 61

3.6 Summary 73

4 Computation using synchronizing oscillators **76**

4.1 The split line experiment 76

4.1.1 Architecture 77

4.1.2 Method 77

4.1.3 Results 78

4.1.4 Discussion 79

4.2 Adaptive synchronizing networks 85

4.2.1 Learning rules for synchronizing oscillator networks 85

4.2.2 An example of learning in a synchronizing oscillator network 89

4.3 Multiple phase networks 97

4.3.1 Architecture 98
4.3.2 Method 98
4.3.3 Results 99
4.3.4 Discussion 99
4.4 Summary 104

5 Constraints on multiple phases **106**
5.1 Phase response dynamics 107
5.2 A study on phase response graphs 110
5.2.1 Method 110
5.2.2 Discussion 111
5.2.3 Single or multiple phase systems? 113
5.3 Input modulation 115
5.4 Synchronized activation 120
5.5 Summary 121

6 Conclusion **123**
6.1 Summary of Results 123
6.1.1 The evidence for synchronization 123
6.1.2 Factors affecting synchronization 124
6.1.3 Computation using synchronizing networks 127
6.1.4 Constraints on the number of stable phases 128
6.2 Further work 129

A Model Oscillators **131**
A.1 Terminology 132
A.2 Oscillator models 133
A.2.1 Leaky integrator 133
A.2.2 Morris-Lecar 133

A.2.3 Wang et al oscillator 134

References **135**

Chapter 1

Introduction

This thesis is concerned with the function and application of synchronizing oscillator networks. The original inspiration for this work stems from observations of oscillations and synchronization in cat visual cortex by Gray *et al* and Eckhorn *et al* (Gray *et al.*, 1989a; Eckhorn *et al.*, 1988; Eckhorn *et al.*, 1989). These observations led to a number of hypotheses concerning the possible function and potential use of systems of synchronizing oscillators. This thesis has two main aims. The first is to provide some foundations for understanding the behaviour of synchronizing oscillator networks within a practical or applied framework. The second is to illustrate how networks of synchronizing oscillators can be made to perform useful computation and to examine the implications for theories of neuro-computation.

1.1 Thesis structure

In Chapter 2, we begin the thesis by discussing the neurophysiological studies which led to this work. A number of interpretations of this data are made and some of the implications of these hypotheses investigated. From an analysis of these requirements, two architectures capable of exhibiting synchronization behaviour are compared and one chosen as a suitable model to be used to investigate synchronization phenomena within this thesis. A definition of this architecture is given and an illustration of its behaviour is shown.

In Chapter 3 we ask how it is possible to modify the basic architecture defined in Chapter 2 and what effects this will have. The aim of these modifications are twofold. It is desirable to gain an understanding of the main factors which could affect coherence in synchronizing oscillator networks in order that we might make predictions about their general behaviour. The second aim is to determine the optimal setup to be used when constructing synchronizing oscillator networks. The studies performed in this chapter do not cover all the aspects affecting synchronization in these networks but they compliment other work already performed within the field and help to build toward a foundation of the understanding of synchronizing oscillator networks. These studies cover four areas. Three concern the principal means of synchronization in the networks studied within this thesis — the lateral connections between oscillating nodes. The fourth concerns the general robustness of synchronizing oscillators to noise.

In Chapter 4 we investigate how it is possible to perform computation using the networks defined and investigated in Chapters 2 and 3. One of the key hypotheses put forward as a result of the neurophysiological data discussed in Chapter 2 is that synchronization of nodes in a network can be used to encode binding information. This is related to the observation that cells responsible for detecting the same object in cat visual cortex are synchronized. A simplified replication of this type of behaviour is shown in the first part of Chapter 4. Having shown that it is possible to perform binding in the model network used in this thesis, we then illustrate how learning can be added to encode this information. We also illustrate how altering our learning rule can determine the type of binding information which is encoded. The last part of this chapter concerns one of the main hypotheses put forward as a result of the neurophysiological data, the principle of using multiple phases to bind multiple objects simultaneously. This is an extension of the binding principal from one phase to multiple phases. In this study we look at the behaviour of different oscillator model networks when initialised to a two phase state and find that none of the networks are able to maintain their initial state in a stable manner.

This led us to question the hypothesis of using multiple phases to perform simultaneous computation and this principal is investigated further in Chapter 5. Through the analysis of phase response graphs, we investigate the behaviour of the oscillators used in this thesis. The results of this analysis confirm that the standard network used in this thesis will not maintain sufficient multiple phases in a form capable of allowing computation to occur. An addition to the standard network model is suggested which will allow for the maintenance of multiple phases. An illustration of its operation is given and compared with results obtained in Chapter 4 showing that it is capable of maintaining at least two stable phases.

In Chapter 6, the conclusions drawn from the work in this thesis are summarised. There are many areas within the field that this thesis could have considered in greater detail and a short discussion of where it would have been useful to extend some of the lines of this work is also given.

1.2 Thesis contribution

There are two main contributions from this thesis. The first concerns the principle of synchronization in networks of laterally connected oscillators. It is shown that, where a model oscillator is suitable for operation in a synchronizing network, it will be capable of performing effectively with a range of lateral input values and levels of connectivity. Furthermore it will also be resistant to a reasonable degree of noise. It is concluded from these results that synchronization of oscillations within cortex would be a feasible and robust phenomena. Where an oscillator model is not particularly effective at synchronizing, it is shown that addition of noise to the lateral inputs of the system will improve its performance.

The second contribution concerns the issues of computation in synchronizing oscillator networks. It is shown that it is possible to perform useful computation with a network of synchronizing oscillators. However it is also shown that there are problems inherent within one hypothesis on the use of synchronizing oscillator networks. The principle of using multiple phases within a network of synchronizing oscillators relies on the network being capable of maintaining a stable set of independent phases. We show that with the configuration of networks considered here, this will not be the case. The normal tendency of the network will be to move toward a state of global synchrony. Failing this, it will move about chaotically in state space in an attempt to achieve synchrony. Neither of these behaviours are desirable for a system which is required to maintain multiple phases. Given this result, we produce a method for maintaining a small number of multiple phases and illustrate its function. It is difficult to determine the neurophysiological feasibility of this method however and we therefore question the hypothesis that multiple stable phases are present in cortex.

Chapter 2

Synchronization in neural systems

In this chapter, we give a brief overview of the phenomena of synchronization within the context of neurophysiological and artificial neural network systems. We begin by discussing the initial theories and studies which led to the development of research into synchronization phenomena. We then present an overview of the foundations of modelling synchronization in neural network models, discussing the basic structures required and some of the possible models of dynamic neural activity that can be used. We conclude with some examples of synchronization phenomena in small networks of coupled oscillators.

2.1 Introduction

The presence of oscillations in neural activity has been well observed (Poppel & Logothetis, 1986; Glass & Mackey, 1988; Gray & Singer, 1989; Steriade *et al.*, 1990). However, until recently little information has been available about the temporal structure which underlies these oscillations. In particular, it was not known how the oscillation of one group of cells related to another group. Studies of cortical activity using EEG and micro-electrodes had revealed the frequency bandwidth over which cortical oscillations were spread and studies of neural cell membrane had revealed some of the form of cellular oscillatory behaviour (Hodgkin & Huxley, 1952).

Studies of artificial neural network models suggested that temporal correlation could have a potentially useful function. Furthermore, theories on cognitive processing had suggested that temporal correlation could have a role in the process of information binding (Vidal & Haggerty, 1988; Singer, 1989). Theoretical studies of phase entrainment in systems of coupled oscillators had already shown that this type of phenomenon was possible in neural systems. Experimental work had also shown that phase entrainment of a neuron could be induced by external stimuli (Othmer, 1985; Glass & Mackey, 1988). The question that remained was whether or not this behaviour was present in cortex and if so, could it be accurately observed?

2.2 Evidence for synchronization

Synchronization of oscillations in physical systems is a commonly observed phenomenon. A system of oscillators is likely to exhibit synchronization behaviour if a feedback mechanism is present between two oscillators such that the state of one oscillator is able to affect the state of the other. Whether or not this will lead to tight synchronization or merely unstable and perhaps transient coherence will depend upon the dynamics of the system. It follows from these observations on physical systems of oscillators that synchronization phenomena are likely to occur in biological systems where oscillatory behaviour is present.

Studies of periodic stimulation of pace-maker cells revealed that it was possible to entrain the oscillatory rhythm of a pace-maker cell with an entraining input, proving that it was at least theoretically possible for synchronization to be present in neurophysiological systems (Cohen, 1987; Glass & Mackey, 1988; Rinzel & Ermentrout, 1989). As a result of this work and the observation that oscillations were present in the neural activity of the brain (Connors, 1984; Dumenko, 1988; Eckhorn *et al.*, 1988; Gray & Singer, 1989), researchers began to concentrate on searching for temporal synchronization in cortex.

2.2.1 Synchronization in response to a causal event

A key advance in research into the temporal structure of oscillatory behaviour in cortex occurred with the publication of the results of Gray *et al*'s study on oscillations and synchronization in cat visual cortex (Gray *et al.*, 1989a; Gray *et al.*, 1989b). A previous study by Gray and Singer had shown oscillatory behaviour in response to a visual stimulus (Gray & Singer, 1989). Using multi-unit recording equipment, Gray *et al* were able to improve on their previous findings by not only recording oscillatory behaviour but also by analysing the temporal relationships between oscillations present at different recording sites. They were then able to relate this temporal behaviour to a causal event, the visual stimulus of an oriented, moving bar of light.

A number of recording sites with similar orientation preferences were located in area 17 of cat visual cortex and recordings made of their response to differently oriented, moving bars. Multi-unit recordings were made from these sites in response to the presentation of three different stimuli, each stimulus presented over a number of trials. The first stimulus consisted of a single long bar moving with an orientation preference similar to that of the neurons at the recording sites. The second stimulus comprised of two short bars moving in the same direction but with a gap in the middle. The third stimulus comprised of two bars as above but moving in opposite directions to each other.

Responses were collected from cells with similar orientation preferences to the direction of the moving bars, recording sites being chosen up to 7mm apart from each other. Two sets of results were collected, the first using only the long bar stimuli and the second using all three stimuli. Auto-correlograms and cross-correlograms for each of the sets were then computed (figure 2.1).

From the first set of results, we can see that the auto-correlograms for the three cells in the figure show basic oscillatory behaviour. Although the data is noisy, it is possible to detect oscillatory behaviour in the 1-1 and 5-5 auto-correlograms. These results fell in line with the previous studies on oscillatory behaviour and were not in themselves particularly

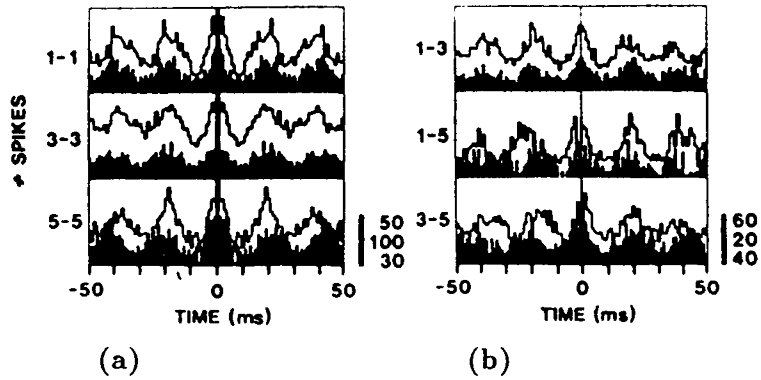

Figure 2.1: The auto-correlograms and cross-correlograms produced from the results of Gray *et al*'s study. (**a**) Single bar stimuli and three recording sites (**b**) Three different stimuli and two recording sites. (Gray *et al*, 1989a)

revealing; however the results of analysing the cross-correlograms were more important.

The correlograms show periodicity with areas of relatively high correlation between pairs of orientation detectors, despite the spatial differences between the sites. This is repeated in the second set of results (fig.2.1b), showing the responses of the cells to the three different stimuli. According to Gray *et al*, the results for the single long bar stimuli show evidence for synchronization of activation between the oriented feature detectors at the spatially separate sites. The level of this synchronization breaks down however as the two light bars are either split or moved in opposite directions. It is as though the split between the two parts of the bar effectively segregates them into different objects.

The importance of these results lies in the fact that they can be related to a causal event. Earlier studies had managed to show oscillations and even synchronization but these observations of neural activity were without definite proof of what specific stimuli were causing them (Skarda & Freeman, 1987). It could be said for example, that the results from the study of olfactory cortex related to various sniffing activities and the presence of certain odours, but it could not be determined what particular odour was causing the oscillatory behaviour. Because the neural activity picked up by Gray *et al* could be shown to be associated with the moving bar stimuli, it allowed them to draw more specific interpretations from their study.

Gray *et al* proposed two possible conclusions from their results. The first concerned the phenomenon of synchronized activation itself. They suggested that the synchronization of activation between different regions could be a mechanism by which separate cortical areas could establish relationships or associations between different features (where each cortical area processed a different part of the same feature). For example, in the case of the moving bar stimuli, the different cortical areas would be responsible for processing different regions of the bar but the complete synchronized group would represent the bar object.

The second conclusion that Gray *et al* drew concerned the possibility of combining the two observed phenomena, oscillations and synchronization, to produce a system which would allow for the parallel processing of multiple objects. This is a powerful concept with the potential to offer explanations for some important questions on object processing and short term memory. In principle, such a system would operate by using different phases of oscillatory cycles to store and manipulate different objects 'simultaneously' — the more phases, the more separate objects that could be dealt with at once.

2.2.2 The binding problem

It has long been known that there was a 'binding problem' concerning theories of cognitive processing (Malsburg & Schneider, 1986; Horn *et al.*, 1991). Although a great deal of evidence had been gathered on how the brain was able to segment patterns up into useful parts, it was not known how it associated or bound related features together. Given an object with a number of different features where each feature was processed and identified in a separate cortical area, how did the brain bind these parts together to form a representation of the whole object?

The theory of temporal synchrony discussed by Gray *et al* had previously been put forward as a potential solution to the binding problem but there had been little evidence to support it (Malsburg & Schneider, 1986; Damasio, 1989; Singer, 1989). Temporal synchrony would be able to solve the binding problem by using synchronization across spatially separate sites to indicate that the cells at each site were processing different parts of the same element. A direct path between the groups would not have to be traced. The fact that they were synchronized would be sufficient to indicate that they were related.

The principle of using multiple phases was an extension to the basic principle of temporal synchrony. Whereas temporal synchrony on its own would allow for the effective distinction of foreground (represented by synchronized activity) against background (incoherent activity), multiple phases would allow for a number of different features to be represented. Each feature would be independently bound with its own separate phase (Shastri & Ajjanagadde, 1993; Mozer *et al.*, 1991).

An example of this principle is shown in figure 2.2 where a pattern containing three objects has been segmented. Each segment in the input is represented by a group of nodes which are oscillating in a different phase of an oscillatory cycle to the other nodes. Only one group of nodes is maximally active at any one time and therefore only one pattern is 'attended' at any given point. Using this method, not only is the system able to bind the common parts of an object into one synchronized whole but it is theoretically capable of manipulating a number of objects at once.

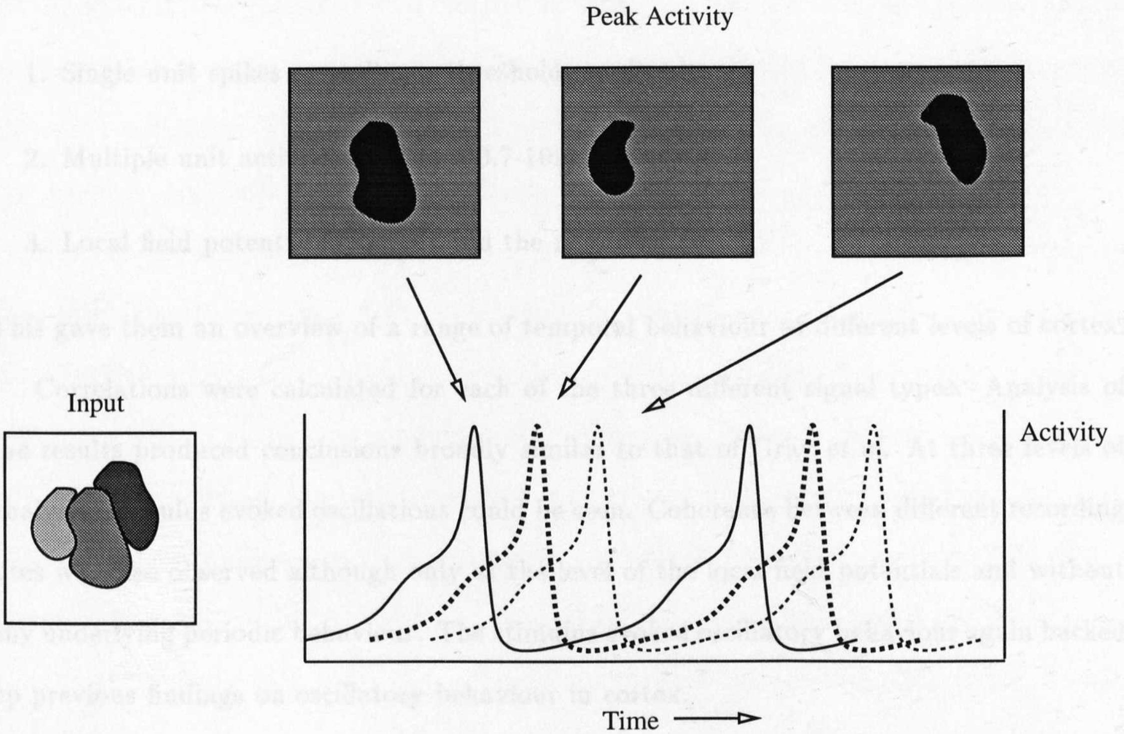

Figure 2.2: Temporal segmentation and binding using multiple phases: An initial input comprised of three separate objects is represented by three regions of activation. The activity within each region is coherent but the activity across regions is segregated over time, a different phase being used for each region.

2.2.3 Further evidence for stimulus evoked synchronization

Similar studies on oscillations and synchronization in cortex were also performed independently by Eckhorn *et al* (Eckhorn *et al.*, 1988; Eckhorn *et al.*, 1989; Eckhorn *et al.*, 1990). Their studies followed closely along the lines of the work by Gray *et al*, again concentrating on cat primary visual cortex and using multiple unit recording equipment to correlate behaviour over different recording sites.

Eckhorn *et al* located a number of different recording sites and experimentally determined their receptive fields. They took recordings from each site in response to optimal stimuli and separated them into three different types:

1. Single unit spikes exceeding a threshold amplitude.

2. Multiple unit activities within a 0.7-10kHz bandwith.

3. Local field potentials from around the recording site.

This gave them an overview of a range of temporal behaviour at different levels of cortex.

Correlations were calculated for each of the three different signal types. Analysis of the results produced conclusions broadly similar to that of Gray *et al.* At three levels of analysis, stimulus evoked oscillations could be seen. Coherence between different recording sites was also observed although only at the level of the local field potentials and without any underlying periodic behaviour. The stimulus evoked oscillatory behaviour again backed up previous findings on oscillatory behaviour in cortex.

Eckhorn *et al* concluded that the most important result of the study was the observation of synchronized activation in spatially separate sites. They ascribed this behaviour to an underlying feature linking process which worked in a similar manner to that proposed by Gray *et al.* Eckhorn *et al* also put forward suggestions as to the possible neural mechanisms which might be causing the synchronization to appear. Primary among these was the possibility that lateral inter-columnar and/or re-entrant cortical connections between the distant sites might be able to cause synchronization.

Evidence had already been collected for the existence of extensive intra-cortical connections in the cortex and previous studies on phase-locking behaviour in coupled cells had shown that this type of mechanism would allow phase-locking to occur (Gilbert & Wiesel, 1979; Mitchison & Crick, 1982; Rockland & Lund, 1982; Othmer, 1985). Could it be the case that this was a functional role for these connections?

2.2.4 Synchronized activation or phase locked coherence?

The results of both Gray *et al* and Eckhorn *et al* show strong evidence for underlying rhythmic behaviour in the neural activity of visual cortex. Combining the results with earlier studies on neural activity, it becomes clear that oscillatory behaviour also occurs at different resolutions. The rhythmic behaviour of spike trains from single cell analysis gives way to more stochastic macroscopic behaviour.

Both of the studies give support to synchronized activation across spatially separate sites although the purpose of the synchronization is left open to debate. The case for synchronized activation — a high correlation between activities of two or more systems within a relatively small time window, is relatively good. The case for phase locked activation — the occurrence of a high correlation of activation over a large time window between two or more more systems with a predictable, periodic behaviour is more constrained and therefore less certain. In the results of Gray *et al*, the cross-correlograms appear to show a cyclic coherence which would be indicative of phase locked oscillatory activity. In the case of Eckhorn *et al*, they concluded that synchronization was present but were unable to state if it was phase locked to the underlying oscillatory activity.

This difference affects the conclusions that both groups drew from the analysis of their results. The presence of synchronization led to the discussion of the possibility that they could have found evidence for a solution to the binding problem. If however the presence of the more powerful phenomenon of phase locking could be determined then it would be the starting point for a solution to the binding problem involving temporal segmentation and parallel feature linking. The evidence and conclusions drawn by Gray *et al* only show phase locking for one oscillatory group (a single phase or synchronized activation system). In order to show this mechanism at work in the cortex, evidence would have to be found for groups of oscillatory neurons in the same area of cortex which were oscillating out of phase with other groups. Evidence would also have to be collected to show that the cortex was able to maintain this system of multiple phases for a period of time sufficient to allow useful computation to be performed.

2.3 Modelling neural activity

Computer simulation of some of the processes involved in synchronization of oscillations within neurophysiology offers one method of investigating the feasibility of theories put forward for the role of synchronization in neuro-computation. Before attempting to design these simulations, it is important to choose reasonable models of cellular activity if we are to have any hope of relating our results to current theory. Given that we are interested in the dynamics of neural oscillations, it is important that the models we choose reflect some of the basic properties of neural activity. Static models of neural activity, commonly used in artificial neural network simulations, will not enable us to fully investigate the synchronization process. They generalise average neural activity over time into a single value and therefore do not accurately model the 'charge and fire' behaviour associated with low level neural activity. There are a wide variety of biological oscillators we could choose from to model neural activity (Chance *et al.*, 1973; Pavlidis, 1973; Abbott & Kepler, 1990). In this thesis we concentrate on three oscillators, the first due to its simplicity and the latter two because of their ability to model cellular and cell assembly neural activity respectively.

2.3.1 The integrate and fire oscillator

Perhaps the simplest activation function we can use to model dynamic activity is a linear integrate and fire function. The functionality of the integrate and fire oscillator is to add the current input I for a given time t to the current state A, provided A has not exceeded a given threshold α. If A equals or exceeds α then the node is deemed to have fired and the state A is reset to a resting state ϕ (equation 2.1).

$$A(t+1) = \begin{cases} A(t) + I(t+1) & A(t) \leq \alpha \\ \phi & A(t) > \alpha \end{cases} \tag{2.1}$$

where $\phi < \alpha$ and $I(t) \ll (\alpha - \phi)$

A_i	Activation for unit i	α	Threshold
I_i	Driving input for unit i	ϕ	Resting state

The behaviour of this function is simple and predictable. Given a constant input, a 'saw-tooth' pattern of activity is exhibited with increases in the level of input from 0 to α causing a directly proportional increase in the frequency of the model neuron (figure 2.3a).

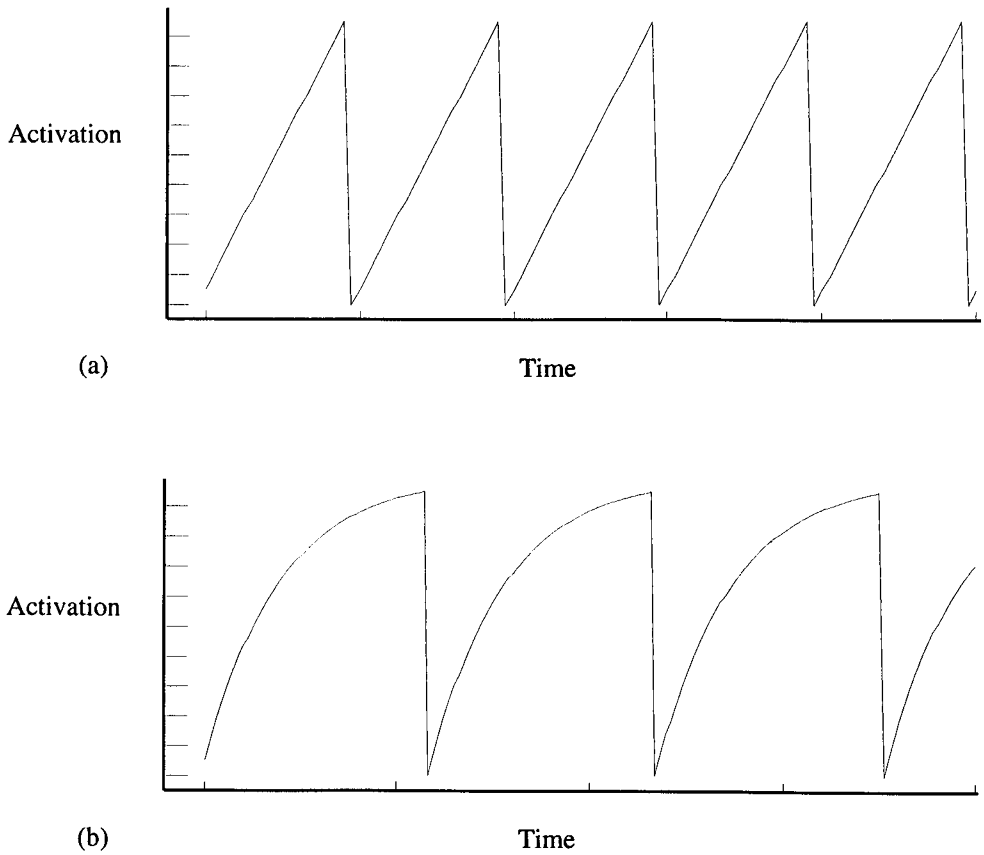

Figure 2.3: Oscillation in **(a)** a linear integrate and fire model and **(b)** a non-linear integrate and fire or 'leaky integrator' model.

As a model of the internal charging and firing behaviour of a neuron it is extremely crude. However it does show some of the macroscopic properties we would like in a model of neural behaviour. It has a regularly 'spiking' behaviour (conceptually the spike occurs when $A(t) > \alpha$) that is dependent upon input strength and it possesses a simple form of internal charging. Two other factors which make it a potentially suitable candidate for simplified modelling of dynamic processes are that, due to its simplicity, it is easier to understand interactions between a number of these oscillators and it allows for simple and fast simulations to be performed.

An improvement to this function which adds an important additional behaviour is to add a non-linear term κ to the previous equation for a linear integrate and fire function to give a non-linear integrate and fire or 'leaky integrator' function (equation 2.2) (Arbib & Amari, 1989). Depending upon the size of κ, the new function will exhibit one of three behaviours. With $\kappa = 1$, we maintain the same behaviour exhibited by the previous function. For $0 < \kappa < 1$, a capacitance charging effect will be exhibited resulting in what can be termed a 'leaky integrator' function. Given an initial value for A around 0 and α greater than 0, the value of A will rise relatively steeply at first. Provided the value for α is sufficiently large ($\alpha < \frac{I}{1-\kappa}$, if $I(t)$ constant), the rate of increase of A will begin to drop off as the effect of κ takes hold. Given that α has been set to a value just below the asymptote toward which A is rising, A will eventually cross the threshold and be reset to its resting value of ϕ. For these values of κ, we have now introduced a charge saturation effect into the activation model (figure 2.3b). This is a behaviour which is commonly observed in the activation traces of neurons just prior to spiking and is a useful addition in attempts to model neural activity.

$$A_i(t+1) = \begin{cases} \kappa A_i(t) + I_i(t) & A_i(t) \leq \alpha \\ \phi & A_i(t) > \alpha \end{cases} \tag{2.2}$$

where $\phi < \alpha$ and $I(t) \ll (\alpha - \phi)$

A_i	Activation for unit i	κ	Non-linearity describing leak of charge
I_i	Driving input for unit i	α	Threshold
ϕ	Resting state		

For $\kappa > 1$ an exponential increase in A is introduced. This produces a rapid, self-sustaining charging effect. This is an atypical behaviour for a model of cell activation as any small initial input received by the cell will always result in the cell firing. Since this is clearly undesirable and unlike neural behaviour, models for $\kappa > 1$ are not considered any further.

2.3.2 Coupled excitatory/inhibitory systems

A more accurate model of dynamic neural activity can be gained by modelling a system as an antagonistic excitatory/inhibitory pair. The excitatory system produces a positive response in the overall state of the system when presented with an external input. This is offset by the inhibitory system which attempts to bring the state of the system down. The inhibitory system is usually driven indirectly by the excitatory system but with a delay or momentum term to enable the system to oscillate. Although the inhibitory system starts off at a lower activation than the excitatory system, it eventually becomes sufficiently active to cause the level of activation of the excitatory system to drop. A large family of biological oscillators are based on this principle (Hodgkin & Huxley, 1952; Wilson & Cowan, 1972; Abbott & Kepler, 1990). An example of this coupled excitatory/inhibitory process can be seen in figure 2.4 which shows the state of the excitatory and inhibitory systems in the Morris-Lecar model of cell membrane potential (Morris & Lecar, 1989; Rinzel & Ermentrout, 1989).

The level at which the activation function models the behaviour of neural activity varies widely. In the case of the Morris-Lecar model above (a simplified derivation of

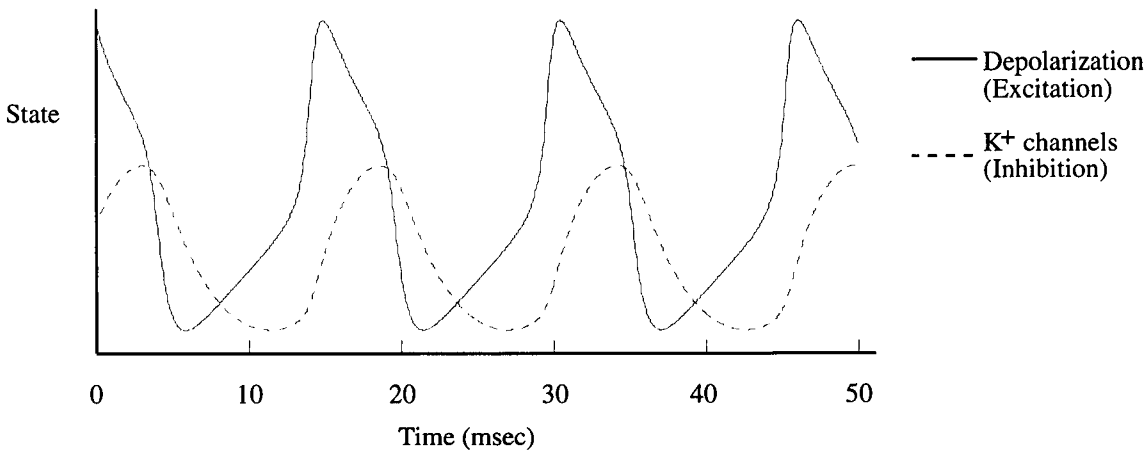

Figure 2.4: The antagonistic excitatory/inhibitory system in the Morris-Lecar model of cell membrane potential.

the Hodgkin-Huxley cell membrane model), the excitatory system models the fluctuation of the internal voltage of a piece of cell wall membrane in the axon of a neuron. The inhibitory system models the change in the proportion of K^+ gates or channels which are open as a consequence of the internal depolarization. The more channels which are open, the more inhibited the rise in the voltage becomes. Provided a reasonable choice of parameters are chosen, this model simulates the behaviour of a piece of neural cell membrane relatively accurately (Morris & Lecar, 1989; Rinzel & Ermentrout, 1989; Shepherd, 1990)(Appendix A.2.2).

At a higher level of neurophysiology, sets of excitatory/inhibitory systems can be used to model the interaction between excitatory and inhibitory networks in cortex (Wang *et al.*, 1990). The excitatory system is taken to represent the average state of a group of excitatory neurons and the inhibitory state, that of a group of inhibitory neurons. Networks which model this type of behaviour are used to illustrate how such systems might explain the observation of oscillations in measurements of local field potentials in neural activity. Despite the large number of different models of neural activity which are currently in use, the fundamental synchronization process is quite similar (Chapter 5). Differences between the behaviour of each model tend to reflect the specialization of the model to a particular level of modelling neural activity.

2.3.3 Phase response equations

As an alternative to modelling the oscillatory activity of a unit and observing how the frequency and phase of the oscillation is affected by external input, we can model the phase response of a hypothetical oscillator directly (Baldi & Meir, 1990; Lumer & Huberman, 1992). This allows us to concentrate on how the state of coherence in a network varies as a function of the type of input being provided, without the complexity of modelling the dynamics which produced these changes. We compute the changes in phase that would (or could) have been produced by a given oscillator and observe how this affects the state of the system.

In order to perform studies using this method the activation state of a given node is discarded and replaced by its phase state — a measure of the position of where its peak activity would be compared with the peak of a global pacemaker node. Two nodes with the same phase state value would correspond to a state of synchronization whereas a node which had a phase state difference of π with another node would be exactly out of phase.

To consider the change in the state of a node modelled using phase response equations we examine the effect a given magnitude of input has on the phase of the node receiving the input. Regular driving input to a node is taken to produce no change in the phase of the node, only inputs which are greater or less than normal (due to the effect of entraining input from other nodes) produce a change. The direction and magnitude of change that is produced by differences in input vary from model to model but a wide variety of behaviours can be produced, each modelling different aspects of oscillator models. A more in depth discussion on phase response requirements is found in Chapter 5.

One of the problems of using phase response equations in dynamic networks is that it is not always possible to create a model oscillator which produces the behaviour defined by a given phase response equation (it is of course possible to go in the other direction). A wide family of oscillators will produce fairly similar phase response behaviour. As a result of this it becomes difficult to judge whether or not the phase response that we are modelling is actually neurophysiologically feasible. The relevance of a network depends upon whether or not its phase response equations correspond to real oscillators.

2.4 Synchronization in artificial neural networks

In order to determine the likelihood of phase-locked coherence being present in the cortex, it is possible to construct simple simulations based on the foundations of what is already known about cortical structures. If it can be shown that synchronized activation or phase-locked coherence occurs in these simulations then, provided they remain within neurophysiological constraints, it would give further weight to the argument that they are present and form a key part of cortical processing. Before considering the methods which can be used to simulate synchronization in neural networks, it is important that we understand the process of synchronization. The following section reviews the basic requirements necessary to achieve synchronization between oscillators. It describes two approaches currently applied to achieve these requirements and the merits of each approach.

Given two nodes which are oscillating at a similar frequency but out of phase, how can they be made to synchronize? The two nodes must communicate by some mechanism so that the result of their combined interaction produces a convergence toward a common phase. Communication of the phase state of one node should be integrated into the equations governing the oscillatory behaviour of the second node resulting in a positive move toward that state. There are currently two main approaches to solving this problem. A solution proposed by Kammen *et al* uses feedback from a *phase comparator* to produce a phase adjustment in a group of connected nodes (Kammen *et al.*, 1990). An alternative method is the use of *lateral connectivity* which draws its approach from physiological evidence. Lateral communication of the activation of a node produces a change in the activity of neighbouring nodes and therefore a shift in their phase. Provided the dynamics of the system are correct, this results in a move toward a common stable phase state where the nodes remain in synchrony over subsequent oscillations.

Achieving synchrony through lateral connectivity has been used by a number of researchers studying synchronization phenomena in artificial neural networks (Vidal & Haggerty, 1988; Atiya & Baldi, 1989; Schillen & Konig, 1989; Sompolinsky *et al.*, 1990; Wang *et al.*, 1990; Wilson & Bower, 1990; Nischwitz *et al.*, 1991; Grossberg & Somers, 1991; Malsburg & Buhmann, 1992). Following on from the results of Gray *et al* a particularly large simulation was undertaken by Sporns *et al* to model some of the processes involved (Sporns *et al.*, 1989). Their simulation modelled a network of over 25,000 cells representing different areas of visual cortex (primarily those where the original neurophysiological results had been found). Using their model, Sporns *et al* were able to replicate some of the synchronization behaviour discovered in the study of Gray *et al.* Sporns *et al* indicated that they attributed the cause of the synchronization behaviour to the 'intra-cortical' and 'inter-cortical' connections modelled in their simulation and implied that this was probably one of the functions of these connections in cortex.

2.4.1 The phase comparator

In a comparison of architectures capable of producing synchronization effects, Kammen *et al* discussed the use of hierarchical feedback. The principal feature of this method is the use of a locally connected *phase comparator* (figure 2.5a). The comparator sums the phases of connected oscillators and feeds back an average. This average is integrated into the phase equations of the oscillators to produce a shift toward the average phase of the group connected to the comparator.

Kammen *et al* report that the comparator method of achieving phase locking is reliable and robust in the presence of noise. The effects of small differences in input are smoothed out by the averaging behaviour of the phase comparator. The sensitivity or granularity of a net using this architecture depends upon the extent of the connections from the oscillating nodes to the comparator — the more nodes connected to a comparator, the less sensitive the behaviour.

However there are some drawbacks to using this method of achieving synchronization. Each comparator affects the phase of nodes connected to it (figure 2.5b). As a result of this constraint, the allowable phase regions of a comparator based network are strongly dependent upon the position of the comparator. In effect, each comparator defines a phase region where connected nodes are encouraged to fall into phase with the average phase of that region. This limits the ability of the comparator architecture to dynamically adjust to the inherent regions defined by the input. If there is a discrepancy between the regions defined by the input and the comparator regions then problems arise. This will result in a conflict between the granularity of the input (the average size of pattern segments) and the granularity defined by the comparator regions. The connectivity schema will attempt to force the phase behaviour of the network into one which maps onto the local averages dictated by the comparator positions.

This problem can be offset to some degree by reducing the number of feedback connections from the comparator although there will be a tradeoff between this and the robustness of the net to the presence of noise and synchronization performance. By reducing the extent of the connections to a comparator and therefore decreasing its field of influence, less conflict will arise between regions defined by the input and regions defined by the comparators. If there is a discrepancy between the above regions, the disruptive influence of each comparator will be smaller leading to less overall conflict. Unfortunately, the reduction in connections also leads to a deterioration in the robustness and synchronization performance of the system. If each comparator has a relatively large area of influence then the effects of noise are smoothed out as the signal to noise ratio is high. In the opposite case however the signal to noise ratio is low and noise present in one node will have a more significant effect on the comparator behaviour.

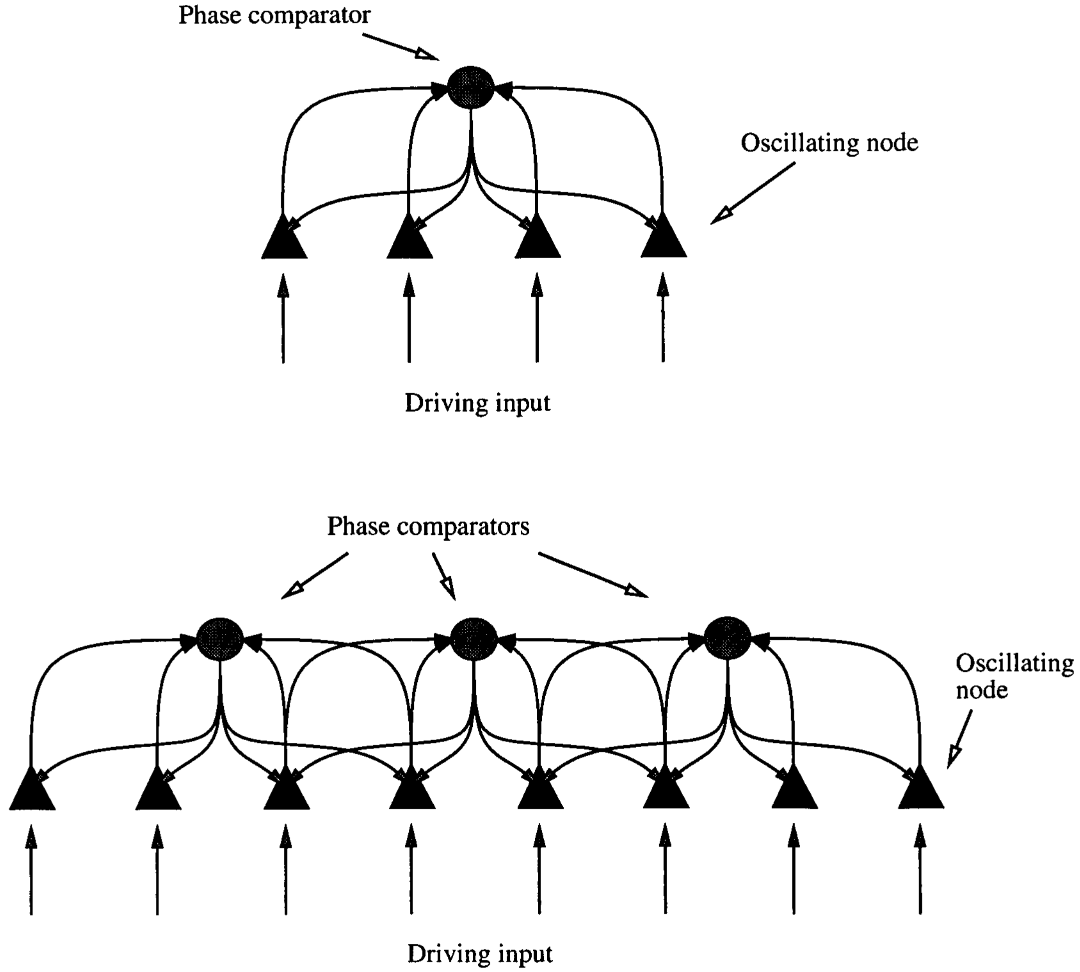

(a) (b)

Figure 2.5: **(a)** A single phase comparator with connected nodes. **(b)** A network of nodes connected to a phase comparator.

2.4.2 Lateral connectivity

An alternative to using phase comparators to correlate local phase is the simpler architecture of lateral connectivity (figure 2.6). The architecture of this model is biologically inspired, following on from observations on the structure of the visual cortex where excitatory inter-columnar fibres are found between specific orientation detectors of one column and identical detectors in another (Gilbert & Wiesel, 1979; Mitchison & Crick, 1982; Rockland & Lund, 1982; Gilbert & Weisel, 1983; Kisvarday, 1986; Chagnac-Amitai & Connors, 1989). Although it is possible that the large number of intra-cortical connections present in the brain could be responsible for producing synchronization behaviour due to effects analogous with that of the phase comparator architecture, lateral connectivity relies upon simpler principles and requires a relatively simple architecture to achieve its effect. The simplicity of this solution also produces a more flexible system, capable of being easily adapted to perform higher tasks such as learning (Chapter 4, Section 4.2).

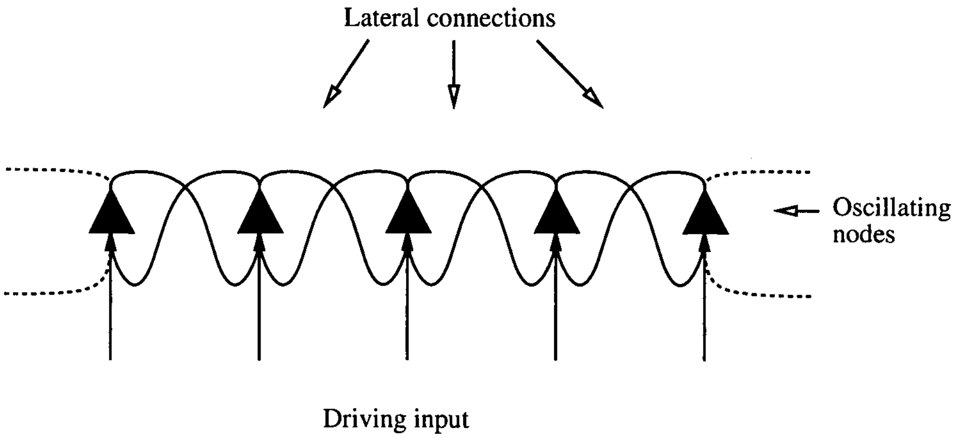

Figure 2.6: Lateral connectivity architecture.

Lateral connectivity relies upon local lateral connections between neighbouring nodes to cause adjustments in local phase behaviour. When a node generates an output spike a burst of excitation is passed out via the lateral connections to neighbouring nodes. This has the effect of increasing the probability of neighbouring nodes to fire. The increase in the probability of firing is in effect a phase shift toward the phase of the central node.

To illustrate, given two nodes A and B where A has just fired and B has shifted phase toward A as a result, we can observe the reciprocal effects of node B on node A. Node B will eventually reach peak activation and produce a burst of excitation. This will increase the probability of node A firing and in effect produce a phase shift toward node B in a similar manner to the above. In the first case the phase shift was in the direction of the phase of node A but in the reciprocal case it was in the direction of the phase of node B. This is a convergent process. Through a series of these communications, two neighbouring nodes move towards each other and into phase.

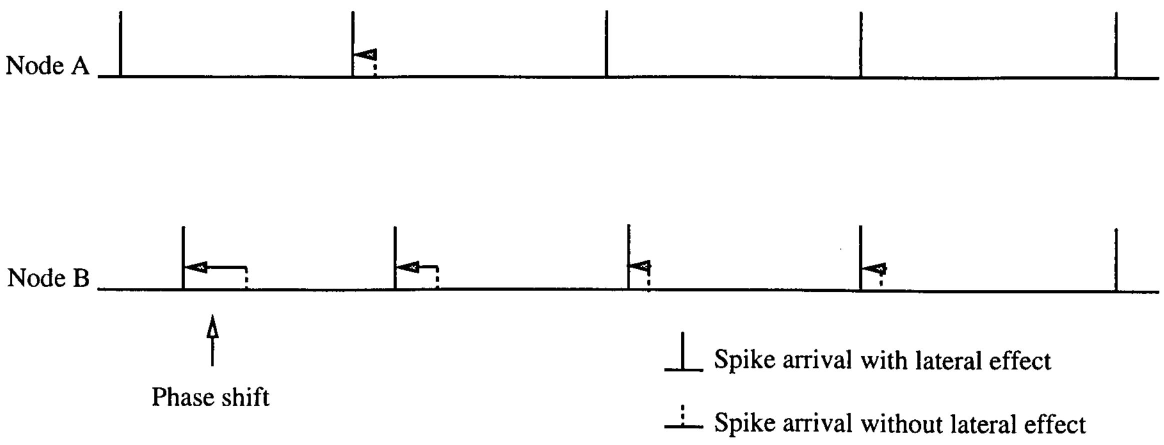

Figure 2.7: The process of synchronization.

A simplified example of this behaviour can be seen in figure 2.7. In the first instance, node A fires and causes node B to fire sooner than it originally would have done (the original firing point is indicated by the short dashed line). Due to the phase position of node B to node A, the phase shift that is induced is relatively large. Node B then fires and causes a phase shift in node A. The spike from node B to node A arrives closer to the firing point of node A and this causes a smaller phase shift than in the previous case. This is due to the non-linearity in the phase response behaviour of the oscillating node (Chapter 5). This process continues over a number of cycles until the nodes eventually synchronize.

2.4.3 A simple architecture

A minimal architecture which is capable of exhibiting synchronization behaviour is shown in figure 2.8a. Each of the two nodes receives a driving input combined with a lateral input from its neighbour to produce a net input for each node. This net input is integrated into the activation state of each node to produce a new state. As a result of the underlying dynamics of the activation model, this produces oscillatory behaviour over time where the net input affects the amplitude and frequency of the oscillation. The activation of each node is thresholded with an output function and the result transferred to its neighbour as the lateral input for the next cycle (Equations 2.3...2.5).

$$I_i(t+1) = \rho_i + w_{ji}.O_j(t) \tag{2.3}$$

$$A_i(t+1) = f(A_i(t), I_i(t+1)) \tag{2.4}$$

$$O_i(t+1) = \begin{cases} \gamma & A_i(t+1) \geq \theta \\ 0 & A_i(t+1) < \theta \end{cases} \tag{2.5}$$

Discrete model for two node architecture.

I(t)	Input	ρ_i	Driving input for node i
A(t)	Activation	θ	Output threshold
f(A,I)	Activation function	γ	Constant for output
O(t)	Output	$i,j = 1..2$	
w_{ji}	Weight from node j to i		

The synchronization behaviour of a network which uses the above architecture and a Morris-Lecar oscillator for the activation function is illustrated in figure 2.8b. The two nodes initially start almost 180^o out of phase but very rapidly begin a convergent process toward synchronization. It can be seen that within two cycles they almost completely phase-lock.

It is a useful feature of this type of system that the behaviour shown for the two node case also applies to larger groups of nodes. Each sub-group of nodes behaves like a single node. A separate group with a different phase will be pulled toward this common phase and vice versa. This can be seen in figure 2.9, where the synchronization behaviour of an 8 node cluster of Morris-Lecar oscillators is illustrated.

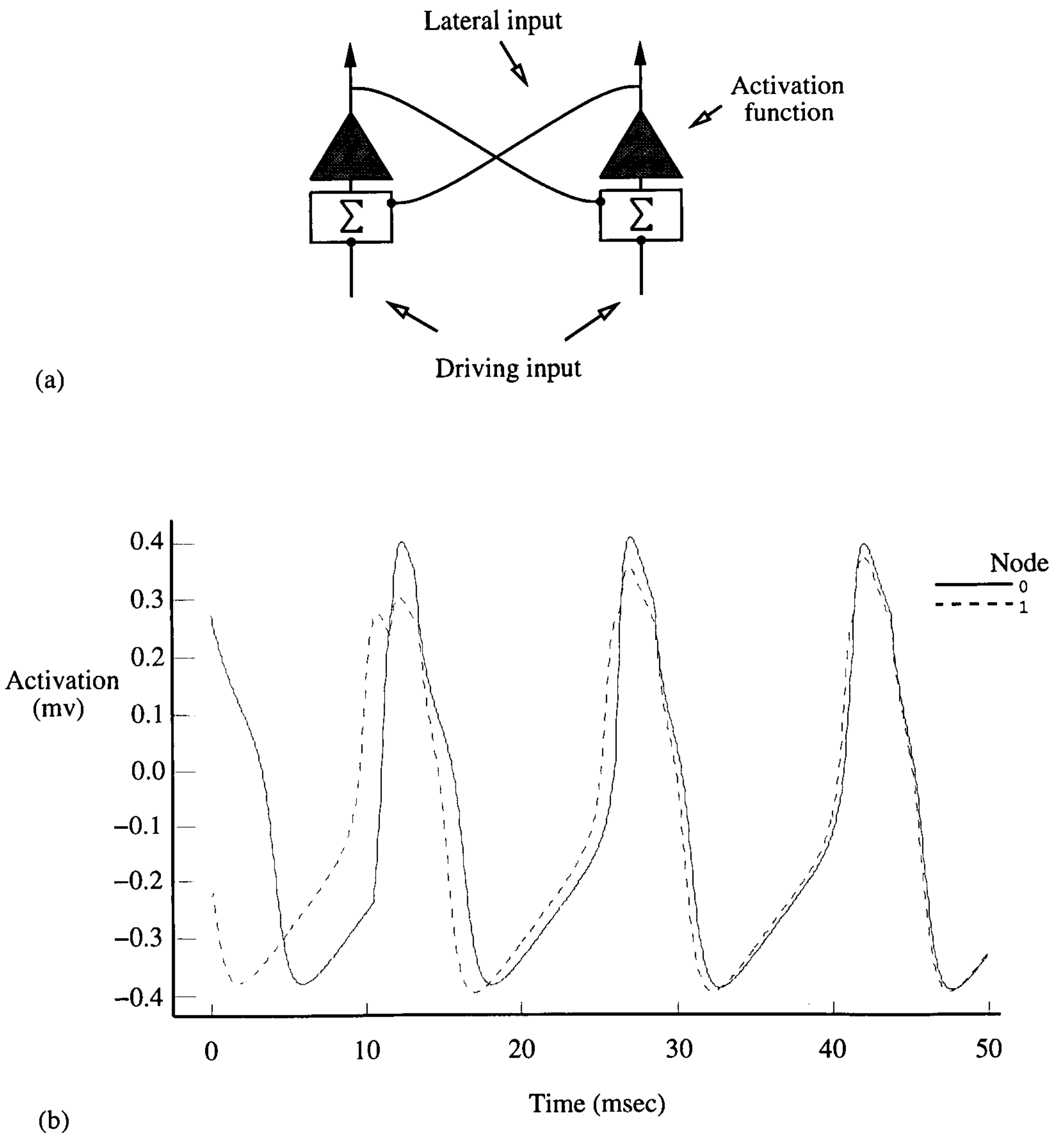

Figure 2.8: (**a**) Architecture of a 2 node laterally connected network (**b**) Synchronization in a 2 node laterally connected network. Simulation parameters : Oscillator - Morris-Lecar, Nodes 2, Time steps 1000, Driving input 0.3, Output threshold 0.26, Lateral input 0.33. All other parameters are identical to that of appendix A.2.2.

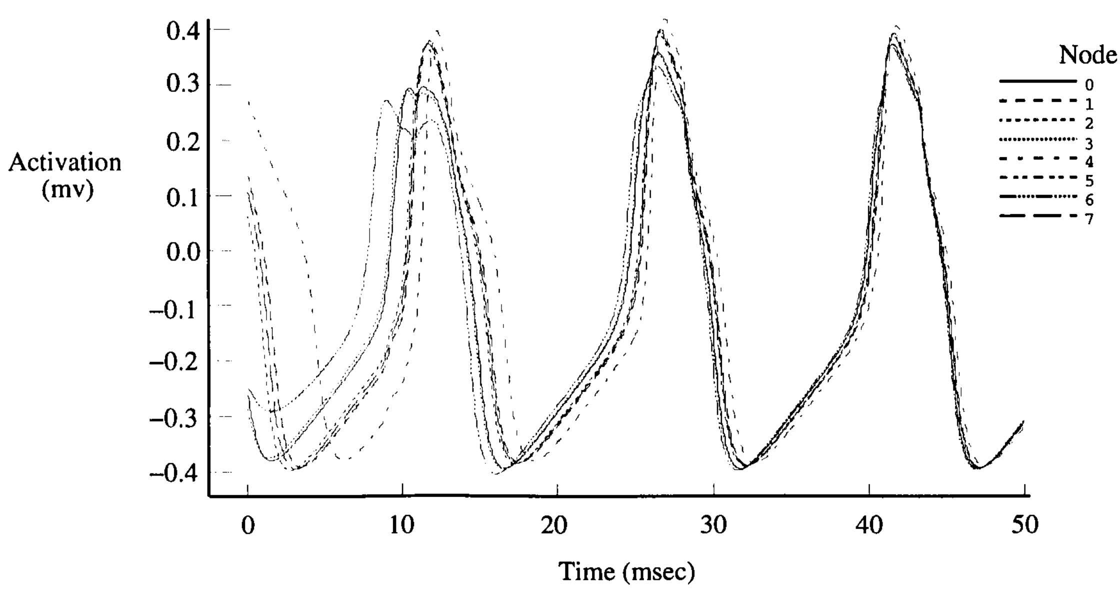

Figure 2.9: Synchronization in an 8 node laterally connected network. Simulation parameters: Oscillator — Morris-Lecar, Nodes 8, Time steps 1000, Driving input 0.3, Output threshold 0.26, Lateral input 0.33. All other parameters are identical to that of appendix A.2.2.

Chapter 3

Synchronization in neural networks

In chapter 2 we established that synchronization of laterally coupled oscillators is a reliable phenomenon. Given this fact, it would be useful to know what factors in networks of laterally coupled oscillators affect their synchronization performance. Knowledge of the effects of these factors would allow us to improve the performance of networks which use synchronization as a computational tool. Furthermore, it might allow us to make useful comparisons with neurophysiological data. If we find certain features improve synchronization performance then we could check for their presence in the neurophysiology of cortex. The following chapter investigates a number of factors which affect synchronization in networks of laterally coupled oscillators. There are a large number of factors which we could choose to investigate but we have decided to concentrate on those which we think have the most significant effect on synchronization performance. These are as follows:

1. The effect of altering connection strength on network coherence.

2. The effect of altering connective scope on network coherence.

3. The relationship between connective scope and network synchronization speed.

4. The effects of noise on network coherence.

3.1 Experimental format

In each of the studies in this chapter a common experimental format is used. This approach is taken in order to maintain consistency between the different experiments, hopefully making it easier to make comparisons between the results of the different sections. The following section outlines this general format and the architecture used. Except where mentioned in a specific study, the following details apply to all the studies in this chapter.

3.1.1 Architecture

The architecture used for these studies consists of laterally connected network of 20 nodes with wrap-around connections at the ends (effectively a ring). Each node in the network receives a driving input responsible for making it oscillate and a normalised lateral input (figure 3.1). The lateral input to a node comes from *all* the other nodes in the network apart from the network used in sections 3.3 and 3.4 where the connection scope is altered. The activation model for each node in the network is varied across different studies. However, for any individual study the same activation function is used for all the nodes in the network. The activation functions used in this chapter are those discussed in chapter 2, namely the 'leaky integrator' model, the Morris-Lecar model and the Wang *et al* model (Appendices A.2.1...A.2.3).

The synapse where the lateral connection from one node meets the dendritic system of another node is modelled using a threshold function which determines whether or not a spike has occurred (equation 2.5). For the interval that the pre-synaptic activation is greater than or equal to the threshold value, a constant size post-synaptic potential is generated (equation 3.1). The post-synaptic potential for a single synapse is combined with all the other post-synaptic potentials arriving at the node and normalized by dividing by the maximum number of lateral spikes a node is able to receive (equation 3.2). The

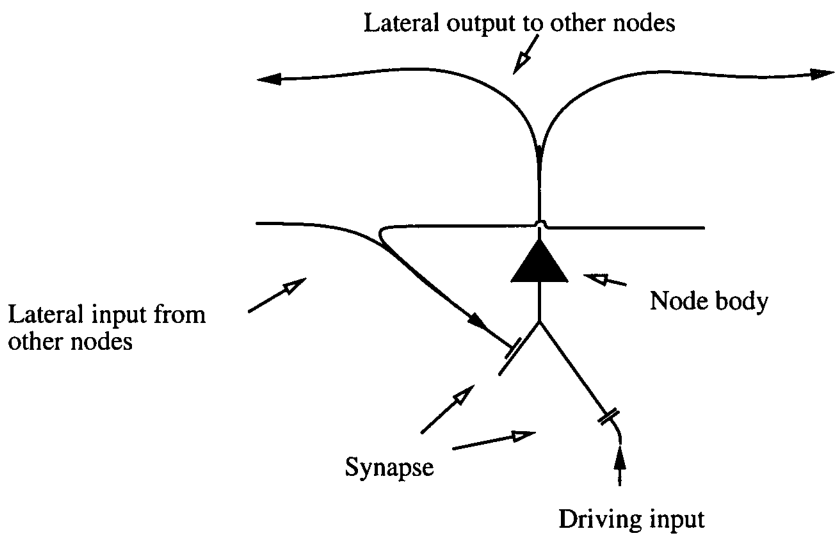

Figure 3.1: Connectivity for a single node.

driving input for a node is a constant value and is identical for each node in the network. The driving input is added to the normalized lateral input to give a total value for the input to a node for a given time step (equation 3.3). The new state (activation level) for each node is then calculated by applying the input and current state to the particular activation function (equation 3.4).

$$l_{ij}(t+1) = \begin{cases} \gamma & a_j(t) \geq \alpha \\ 0 & a_j(t) < \alpha \end{cases} \tag{3.1}$$

$$L_i(t+1) = \frac{1}{N-1} \sum_{j=1, i \neq j}^{j=N} l_{ij}(t+1) \tag{3.2}$$

$$I_i(t+1) = T_i(t+1) + L_i(t+1) \tag{3.3}$$

$$A_i(t+1) = f(A_i(t), I_i(t+1)) \tag{3.4}$$

N	Number of nodes in the network
l_{ij}	Lateral input at synapse from node j to node i
γ	Constant for post-synaptic spike
α	Constant for synaptic threshold
a_j	Pre-synaptic activation for node j
$L_i(t)$	Total lateral input for node i at time t
$T_i(t)$	Driving input for node i at time t
$I_i(t)$	Total input for node i at time t
$f(A,I)$	Oscillatory activation function
$A_i(t)$	State of node i at time t

It can be seen from these equations that there is only one discrete time delay modelled in these equations. Axonic and synaptic delays are not modelled directly. The combined effects of these delays are grouped into a transmission delay of one time step. This is a rather simple approach and does not reflect the real neurophysiology, specifically in the effects of signal transmission delays over long distances. A signal from a node separated by as much as 20 other nodes from the recipient only takes one time step to cross the gap. While this is reasonable within the scale of the small networks used in these studies, it becomes more unrealistic as the networks get larger.

Modelling synaptic and axonic delays introduce further complexity to the synchronization process which makes the understanding of the synchronization process more difficult.

It has been shown by Konig and Schillen that signal delays can be used to de-synchronize nodes(Konig & Schillen, 1991). Furthermore studies by Glunder and Nischwitz and also Smith *et al* have also shown that signal delays combined with inhibition can produce synchronization or desynchronization in a network depending upon the length of the time delays(Nischwitz *et al.*, 1992; Glunder & Nischwitz, 1993; Smith *et al.*, 1993). As a result of the above studies we feel that the modelling of synaptic delays, although important within a neuro-physiologically correct model, adds further complexity and reduces our ability to understand the effects of the other factors under study. We have therefore not modelled them in any further detail.

3.1.2 General method

Most of the studies in this thesis involve altering a particular parameter (for example connection strength) over a series of trials. For each trial the parameter is fixed at a set level and the behaviour of the network analysed. To keep the initial conditions the same for each trial, it is necessary to set the state of the nodes in the network to preset but uncorrelated values. To this end, a random starting position (within the normal phase of an oscillation) is selected for each node in the network and recorded as its initial starting position. At the start of each new trial, the node is reset to this initial starting position. This ensures that each node starts at a random position but that this position is the same for each consecutive trial. Using this method, only the variable under study alters from trial to trial.

For each trial of a given parameter value, a number of test runs are performed. Each test run uses different initial starting positions for each oscillator in the network. This approach is taken in order to ensure that the measures we obtain are representative of the average behaviour with the given parameter value and are not chance results. At the end of a set of test runs for a given parameter value, the average of the test runs is taken as the final value. The standard deviation for the test run is also calculated and used to produce a measure of the variation in the result. The standard deviation for a test run is normally displayed in the graph of results as an error bar with the average value as the initial point.

3.1.3 Measuring network coherence

Before considering the effects on coherence of altering network parameters, it is important to consider how we measure the coherence of a network. To measure the coherence in a network, we must first define what we mean by a coherent network. Within the context of the synchronizing oscillator networks discussed in this thesis, a network with high coherence can be roughly defined as one in which nearly all the nodes are at peak activity at or around the same time. We are interested in spike output coherence rather than activation coherence. Whether or not the activation states are similar between peaks should not be significant. Between spikes the activation of a node may move away from the norm due to the effects of noise or phase shift caused by lateral input. What is important is that when the activation of a node peaks, the peak arrives at the same time as the peaks of the other nodes in the network.

One method of examining the coherence of a network is to analyse a spike distribution histogram of the activity of the network. This shows the distribution of spiking in the network across one complete oscillation and therefore indicates where most of the nodes are firing. The absolute position of where a group of nodes fire is not important. All that matters is how tightly they are grouped. If we were to look at a spike distribution histogram over one oscillation period for a group of nodes in an oscillating network, high coherence would be represented by the circular histogram of figure 3.2a. It can be seen that the majority of spikes occur within a relatively small time window. Conversely, low coherence could be represented by the spike distribution histogram of figure 3.2b which shows a dispersed distribution. If we obtained a spike distribution of the form of figure 3.2c where half the nodes fire at one point and the other half fire at a point 180^o out of phase, then this would also be representative of low coherence.

Given the examples of spike distributions shown in figure 3.2, we would like a statistical measure of variance which reflects the degree of coherence present. However, we must be

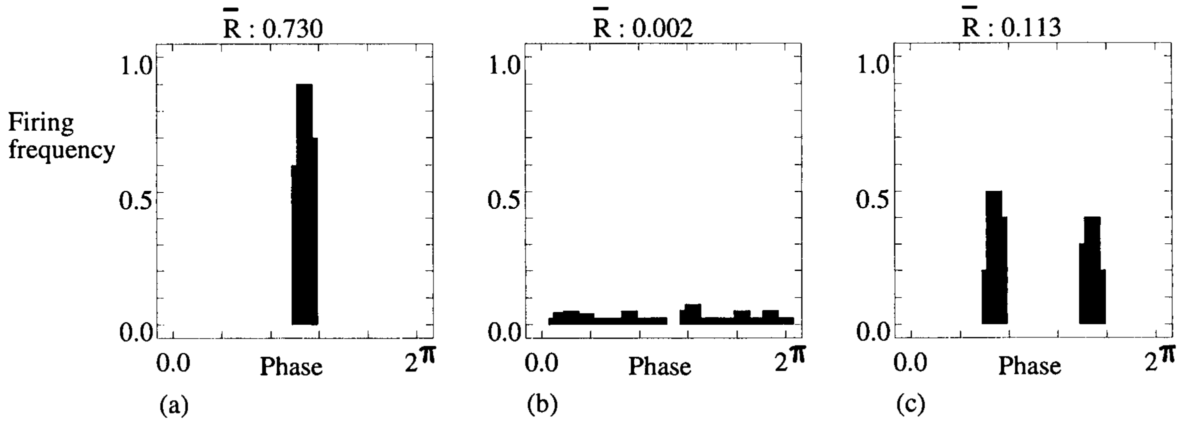

Figure 3.2: **(a)** Circular spike distribution histogram for a network with high coherence. **(b)** Circular spike distribution histogram for a network with low coherence. **(c)** Circular spike distribution histogram for a network with two separate phases.

careful that the measure chosen takes account of the directional property of the data. Models of variance such as the standard deviation do not give a good measure for the coherence represented by the above distributions.

The mean resultant length

A measure which possesses the qualities we require for a measure of coherence is the mean resultant length of a circular distribution (Mardia, 1972). This statistic gives us a measure of the variance in the direction of the data combined with a measure of the magnitude of this variation (equations 3.5...3.7).

$$\bar{C} = \frac{1}{n}\sum_{i=1}^{n} f_i.cos\theta_i \tag{3.5}$$

$$\bar{S} = \frac{1}{n}\sum_{i=1}^{n} f_i.sin\theta_i \tag{3.6}$$

$$\bar{R} = \sqrt{\bar{C}^2 + \bar{S}^2} \tag{3.7}$$

n Number of bins

f_i Fraction of spikes observed vs maximum spikes possible

θ_i Mean direction of bin i

$\bar{R}$ Mean resultant length

To apply this statistic, the spike distribution is allocated to a number of 'bins', the complete set of bins covering one complete oscillation. Each bin represents the mean phase or direction for the spikes that 'fall' into that bin. The number of spikes in a bin denote its significance (in effect length) compared with the complete distribution. Using this data a mean resultant length $\bar{R}$ is calculated. Spikes which are dispersed and are therefore going in different directions produce a mean resultant length near 0. Spikes which are close together are effectively pointing in the same direction and produce a mean resultant length close to 1.

It is an important property of this measure that a distribution requires to be very tightly packed if it is to produce a value approaching 1.0. Unless a network is tightly synchronized, this measure will produce a low value for the level of coherence. For example, the distribution of figure 3.2a which shows a group of nodes firing at or around a similar point, the measure of the mean resultant length $\bar{R}$ is still only 0.730. In order to obtain a value close to 1.0 we require a distribution where almost all of the spikes occur together and the few that do not must fall either side of the main group.

Spike distribution histograms

The spike distribution histograms for the studies in this thesis were obtained by noting the spiking of the network over one complete oscillatory period after some fixed time interval.

The oscillatory period p is divided into a number of bins b. In the majority of these studies b was set to 10. This gave a bin size of 0.6283 radians or $\pm 18^o$ Because the measure of coherence $\bar{R}$ falls off rapidly as the spikes become more dispersed, it is important not to make the number of bins too large; otherwise a false impression may be gained of the level of coherence in the network. For example, with a large number of bins (of the order of 100 or more), even if all the nodes in a network fire within a 15^o window, the value we would obtain for $\bar{R}$ would be relatively small (approximately 0.3 or less). In the studies performed in this thesis, the distribution histograms are normalised by dividing the number of spikes observed in a particular bin by the maximum number possible. This ensures that where networks have different numbers of nodes, the coherence measures obtained are still comparable.

3.2 Connection strength

Since the primary cause of synchronization in laterally coupled networks is the lateral spikes transmitted between nodes, it would be useful to know how varying the strength of the lateral spikes affects network synchrony. Will small inputs achieve synchrony or are quite large lateral inputs necessary before adequate synchrony occurs? Furthermore, do small increases in connection strength cause a correspondingly small increase in network coherency or is the effect more non-linear? The following section details the method and results of a study which investigates this relationship.

Before discussing the study, we must define what we mean by connection strength. In this study, the connection strength σ is a value in the range -1 to 1 where -1 represents negative lateral input (inhibition), 0 represents no lateral input and 1 represents the maximum lateral input. Because the actual value of the maximum allowable lateral input is implementation dependent (due to the different oscillator types), it is easier to represent and manipulate the connection strength in this way. The connection strength can be likened to a filter on the post-synaptic side of a synapse which only lets through a certain degree of the post-synaptic potential generated by a lateral input arriving at the pre-synaptic side (equation 3.8). It is not intended to represent any neurophysiological structure but is only in place as an experimental device. The values we could use for the maximum and minimum lateral inputs are relatively arbitrary. What is important is that we cover a range of inputs that a node is likely to generate or receive. In this simulation we have used the value of the driving input as the maximum value and the negated value of the driving input as the minimum value.

$$l_{ij}(t+1) = \begin{cases} \sigma.\gamma & a_j(t) \geq \alpha \\ 0 & a_j(t) < \alpha \end{cases} \tag{3.8}$$

where σ is the connection strength.

3.2.1 Method

The architecture and method for this study is the same as that discussed in section 3.1 except for the calculation of the lateral input for a given node. Equation 3.1 is now replaced by equation 3.8. The study investigates the effects of altering connection strength on network coherence for the 3 oscillator models discussed in section 3.1.1. For each oscillator model the study is run for 31 trials. Starting with an initial value of -1 on the first trial, the connection strength is increased to a final value of 1 on the last trial. For each trial, 50 separate test runs are performed with different initial starting points for the nodes on each test run. At the end of a test run, a coherence measure for the network is obtained. These are compiled together to produce an average coherence measure for the trial together with a standard deviation for the coherence measures obtained for that trial.

The graphs produced from these experiments are shown in figures 3.3...3.5. The x-axis indicates the level of the connection strength and the y-axis gives the coherence measure obtained. Each point in the graph shows the average coherence measure obtained for the given level of connection strength. The error bars shown for each point indicate one standard deviation in the variance of the test samples taken for that point.

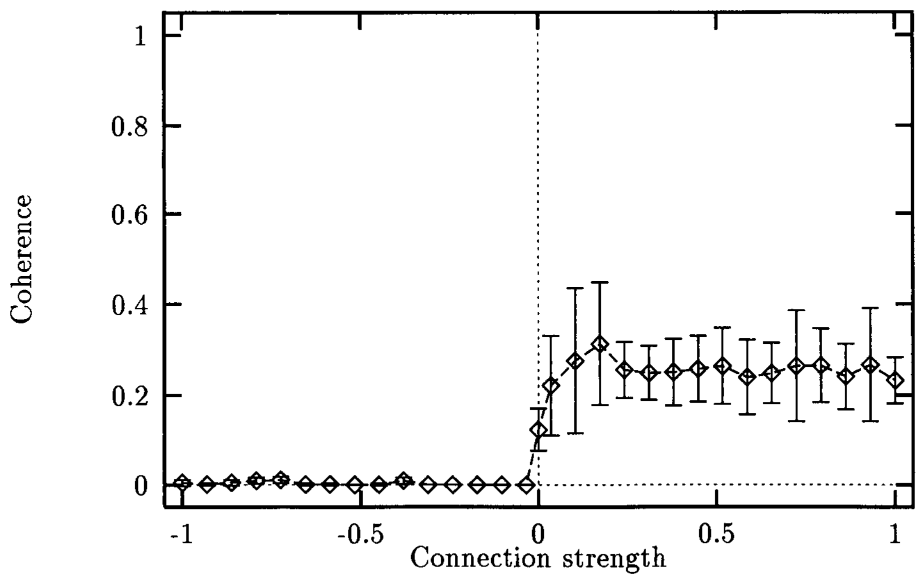

Figure 3.3: Network coherence versus connection strength for the 'leaky integrator' oscillator. Simulation parameters - nodes 20, trials 31, epochs per trial 3000, test runs per trial 50, connection strength -1.0...1.0.

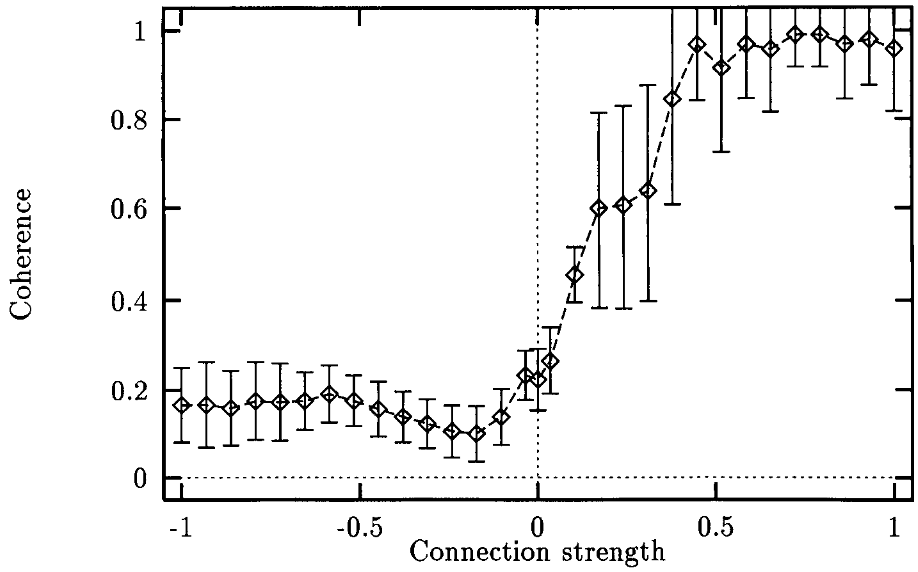

Figure 3.4: Network coherence versus connection strength for the Morris-Lecar oscillator. Simulation parameters - nodes 20, trials 31, epochs per trial 2500, test runs per trial 50, connection strength -0.3...0.3.

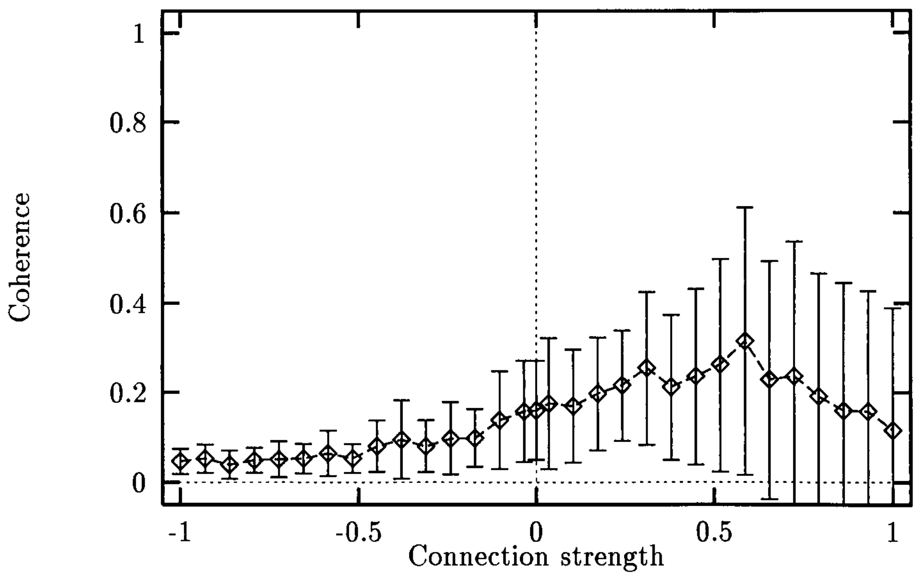

Figure 3.5: Network coherence versus connection strength for the Wang *et al* oscillator. Simulation parameters - nodes 20, trials 31, epochs per trial 2500, test runs per trial 50, connection strength -0.2. . . 0.2.

3.2.2 Discussion

We are more interested in the shape of the graphs in figures 3.3...3.5 than the actual values given. It is better to know how varying the size of the connection strength is likely to affect the behaviour of a network than to have a set of values for which we know the effect of a particular set up. Given that we know which values correspond to the upper and lower bounds in the above graphs, we should be able to gain a useful understanding of how varying the connection strength between them will affect the networks synchronization performance.

The graph for coherence against connection strength for the 'leaky integrator' model (figure 3.3) indicates that it is a relatively poor oscillator for use in synchronizing oscillator networks as the average coherence measures obtained are fairly small. Although the average values obtained for the positive connection strengths are greater than the control measure of 0 connection strength, varying the size of the connection strength seems to have relatively little effect on the synchronization performance apart from the region around -0.1...0.35. An optimal setting for the network connection strength in a network of 'leaky integrator' models would seem to be somewhere in the range of 0.1...0.25 of the driving input. After the connection strength goes beyond 0.3, increases in the connection strength have little effect on synchrony. Inhibition in networks of coupled 'leaky integrator' models would seem to cause a complete breakdown in network synchrony. When considering the size of the variations in the coherence measure (indicated by the error bars) it can be seen that with small lateral inputs, coherence measures of around 0.41 can be obtained. This indicates that the 'leaky integrator' model is capable of some degree of synchrony on occasion but that on average its performance is poor.

The graph of connection strength against coherence for the Morris-Lecar oscillator (figure 3.4) again shows that negative connections cause a drop in network synchrony. Compared with the control measure for network synchrony, negative connections actively desynchronize the network. Small positive connections on the order of about 0.1...0.2 of the driving input cause a marginal increase in network coherence. Increases in network connection strength over and above these levels cause a rapid improvement in coherence although connection strengths of around 0.3 of the driving input produce a wide variation in network coherence. Once the connection strength reaches a value of about 0.5 of the driving input, network synchrony is almost guaranteed. Given this behaviour, we feel that the Morris-Lecar model would be a reasonable candidate for an oscillator in a synchronizing oscillator network. It is robust over a wide range of inputs and is able to guarantee synchrony when strong lateral inputs are used.

Compared with the Morris-Lecar model, the Wang *et al* model does not perform very well. Negative inputs again cause de-synchronization in the network. Positive lateral inputs do not produce a great improvement in network synchrony. The high variation in the results of the mid-range positive connections indicate that the Wang *et al* oscillator is capable of a reasonable degree of synchrony on occasion but like the 'leaky integrator' model, its average performance is relatively poor. Increasing the lateral connection strength does improve the synchronization performance up to about 0.6 of the driving input but after this point, a deterioration in performance occurs. This is possibly due to a saturation effect inherent in the oscillator model which causes a node to remain artificially high when presented with an excessively large input.

The cause of the change in synchronization performance as a function of network connection strength is directly related to the degree of phase shift the lateral input is producing.[1] With small values for the lateral input, the phase shift produced is not sufficient to overcome the basic oscillatory dynamics and little synchronization behaviour takes place. As the size of the lateral input increases, the phase shift produced becomes significantly greater and this enables individual nodes to shift about to a greater extent and therefore increase the chance of synchronization. In the case of the 'leaky integrator' and Wang *et al* model an optimum level is reached after which the phase shifts created by the

[1] The principle of phase shifts occurring as a result of lateral inputs is discussed in more detail in Chapter 5.

stronger lateral inputs cause the relative phases of the nodes to move around too much. The nodes are not able to compensate for this movement and fail to synchronize. The Morris-Lecar model appears more robust and actively benefits from an increase in lateral connection strength. It too reaches a point after which no real improvement occurs but does not appear to deteriorate significantly if given a lateral input larger than normally required to achieve synchrony.

It is worth noting the effect of using negative lateral inputs in these studies. In all three cases, negative lateral input causes an active breakdown in network coherence. Even small negative lateral inputs cause the network coherence measure to fall below the background coherence measure obtained for zero connection strength. Negative lateral inputs cause negative phase shifts to occur at nodes which receive them. This translates to a movement away from the node transmitting the lateral signal and therefore causes it to de-synchronize. This behaviour is reasonably robust and has been shown to operate by other researchers (Nischwitz *et al.*, 1992; Glunder & Nischwitz, 1993).

3.3 Connective scope and coherence

Another principal feature we can alter in a laterally connected network is the degree of lateral inter-connectivity or connective scope. It is not at first obvious what effect this might have on the ability of a network to synchronize. It would seem intuitive that the more nodes a given node is connected to (and therefore the more nodes it is able to influence) the better a network will be able to synchronize. If the connective scope is small (for example nearest neighbour) then we might expect a poor synchronization performance. However, it is not clear how the performance of the network would vary between these limits. Since the level of node inter-connection increases the computational load required to model a network, it would be useful to know if there is an optimal point at which increasing the number of connections produces no further significant gains in synchronization performance. If our study revealed that complete connectivity is a requirement for synchrony then not only would this be a constraint on artificial neural networks but it would also have implications for neuro-physiology. If complete connectivity were a requirement for synchrony then areas of cortex which are not directly linked (but had intermediate connecting sites) should not be able to synchronize. The following section discusses a study which was performed to investigate this property and find some answers to these questions.

3.3.1 Method

There is a problem with performing this type of experiment caused by the fact that as we increase the number of nodes a given node is connected to, we increase the number of inputs it receives. The solution taken here is to normalise the combined lateral input by dividing it by the number of nodes responsible for making up that input. This ensures that as the number of connections in the network is increased, the maximum possible lateral input value to a node remains the same. The problem we have with this solution is that although the sphere of influence of a given node is increasing, the strength of this influence is decreasing. It is realised that this is not perhaps the best solution due to the fact that when the connective scope for a node becomes large, the lateral input it produces will be quite small. In order to reduce this effect, the number of nodes in the network is kept within manageable limits. If the network was to be increased much beyond the size used then a different approach would need to be taken.

The results of the previous section indicate that although the size of the lateral input is significant, it is still possible to improve synchrony with values smaller than the optimum. An alternative to this study would have been to use a sigmoid squashing function on the incoming lateral inputs although it would have been difficult to interpret how the non-linearities introduced by this technique would affect network synchronization.

The architecture used in this study is the same as that for section 3.1.1 except for the connectivity. Instead of each node being connected to all the other nodes, it is only connected to a fixed number of neighbouring nodes either side of the node itself (equation 3.9). A further change to the lateral input values has also been made which takes into account the results of the previous study. The lateral outputs generated for each oscillator are now set to the values which gave the optimal values in the previous study.

The simulation is run over 11 trials. On each trial the number of nodes a given node is connected to is increased by 2 (1 to each side of a node) and this is used to calculate a connectivity measure c (equation 3.10). For each trial value of c, a number of test runs are made (50 for the 'leaky integrator' and Morris-Lecar models and 20 for the Wang *et al* model). At the end of each test run a coherence measure is obtained. This is combined with the results for all the test runs to produce an average measure and standard deviation for the trial. The combined set of trials for each oscillator are used to create the graphs of figures 3.6 to 3.8. Each figure shows the results obtained for one of the oscillator models. In each figure, the x-axis indicates the level of connectivity c and the y-axis the average measure of coherence obtained.

$$L_i(t+1) = \frac{1}{2\nu}\sum_{j=1}^{\nu} l_{i((N+j-i)\mathbf{mod}N)}(t) + l_{i((j+i)\mathbf{mod}N)}(t) \quad (3.9)$$

ν	Number of neighbours on one side of a node
N	Number of nodes in the network
$l_{ij}(t+1)$	Lateral input at synapse from node j to node i
$L_i(t+1)$	Total lateral input for node i at time *t+1*

$$c = \frac{2\nu}{N} \quad (3.10)$$

3.3.2 Discussion

Increasing the connective scope for the 'leaky integrator' network appears to improve the synchronization performance of the network up until around a connectivity level of 0.6. Network coherence then drops off and improves only slightly when complete connectivity (c=1) is used. With mid-range connections, quite high levels of network coherence are obtained on occasion and overall network performance improves compared with the results of the previous section.

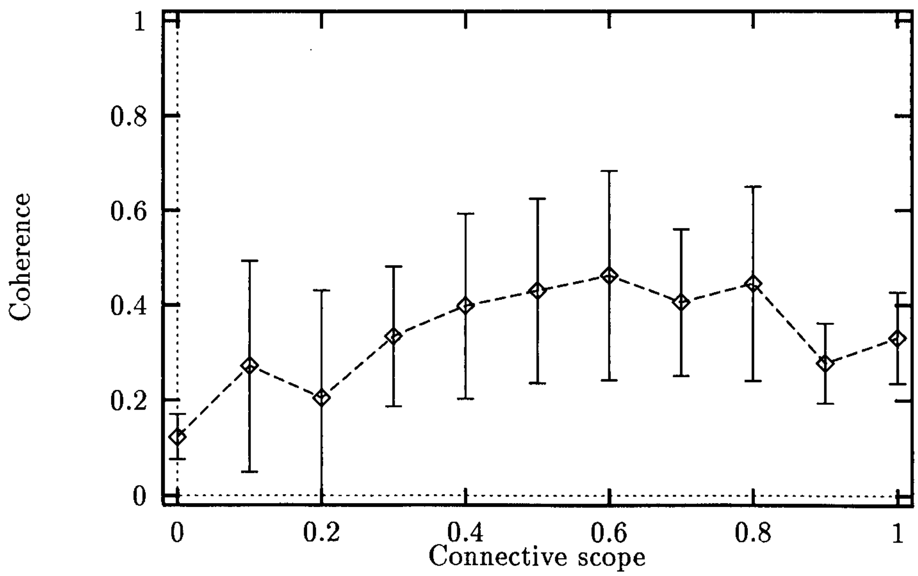

Figure 3.6: Network coherence vs connection scope for the 'leaky integrator' oscillator. Simulation parameters - nodes 20, trials 11, epochs per trial 3000, test runs per trial 50, lateral input 0.2, connective scope 0 (no connection) to 1 (total connectivity).

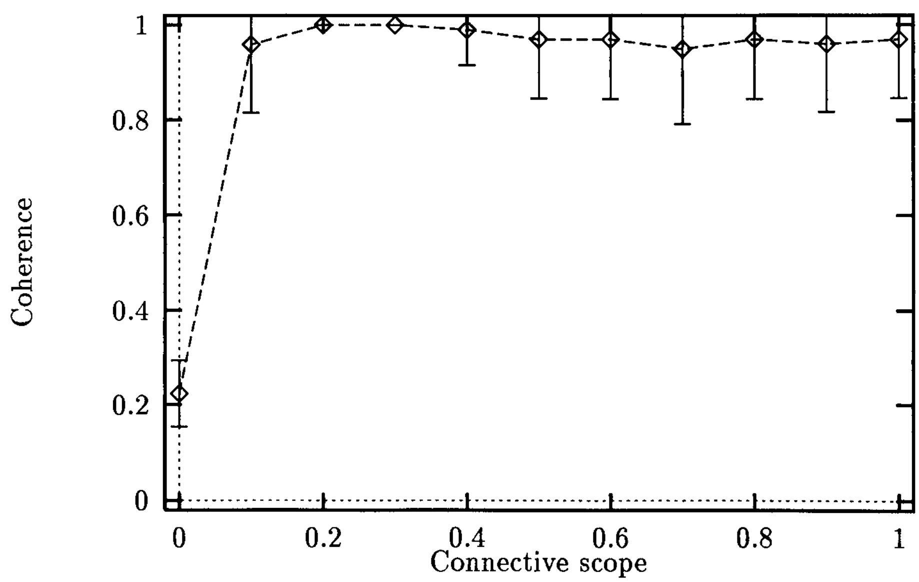

Figure 3.7: Network coherence vs connection scope for the Morris-Lecar oscillator. Simulation parameters -'nodes 20, trials 11, epochs per trial 2500, test runs per trial 50, lateral input 0.3, connective scope 0 (no connection) to 1 (total connectivity).

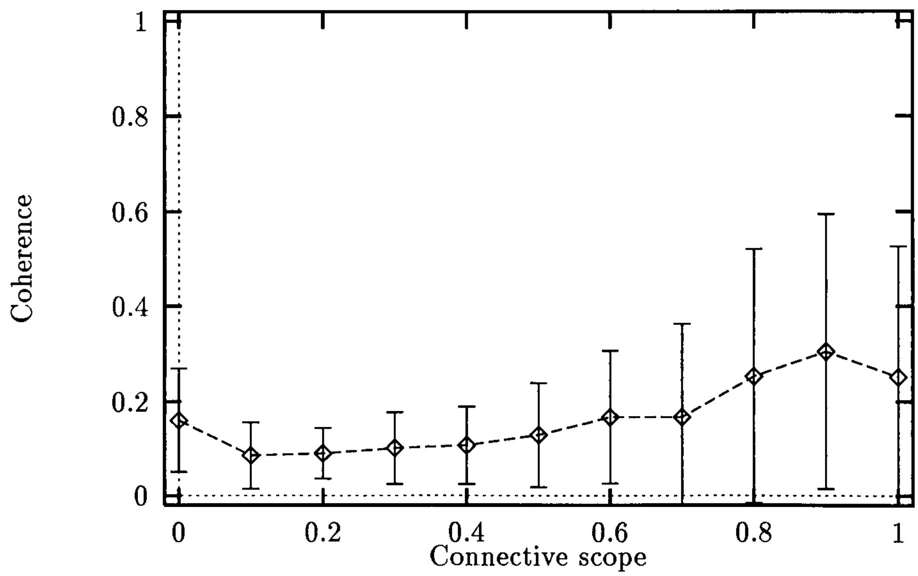

Figure 3.8: Network coherence vs connection scope for the Wang *et al* oscillator. Simulation parameters - nodes 20, trials 11, epochs per trial 2500, test runs per trial 50, lateral input 0.12, connective scope 0 (no connection) to 1 (total connectivity).

Compared with the 'leaky integrator' model, increasing the degree of connective scope in the Morris-Lecar network produces very different results. The high measure of network coherence found when network coherence is increased beyond 0 indicates that even nearest neighbour connectivity is sufficient to almost guarantee synchrony in the Morris-Lecar network. Connection levels in the range 0.2 to 0.5 seem to converge completely on all test runs for each trial. Only when the connective scope increases beyond this point does a larger variation appear in the results. From this behaviour it would seem that although complete connectivity will produce adequate synchronization, the best results for the Morris-Lecar oscillator are obtained with connectivity values of around 0.3.

For the Wang *et al* network, altering the level of connective scope produces only a slight improvement in network coherence. However, even though the average coherence measures are small, a wider variation in the test results is produced with increased connection scope. This indicates that some improvement is being made in the underlying network dynamics but that due to the poor performance of the Wang *et al* oscillator, global synchrony is frequently not achieved. It is as though the network occasionally goes into a de-synchronized state from which it cannot escape.

To understand why short range connectivity allows a network to synchronize completely, it is necessary to consider the process of synchronization in a network with only nearest neighbour connectivity. A network with this structure synchronizes through a process of convergent grouping. Locally connected nodes first form into small synchronized groups. These groups then interact with each other (via the existing connections between nodes in one group and nodes in another group) until a larger synchronized group is formed. This process of convergence towards synchrony continues until the complete network synchronizes.

It is interesting to note such a wide variation in the behaviours of the different model oscillators. Overall, the optimal network connectivity would appear to lie in the range 0.4 to 0.7 but even then this conclusion does not hold very well for the Wang *et al* oscillator. The combined set of results seem to indicate that, compared with the Morris-Lecar oscillator, the 'leaky integrator' and Wang *et al* models are not very good candidates for use in synchronizing oscillator networks. This agrees with the results of the previous section where the 'leaky integrator' and Wang *et al* models performed quite poorly.

It was also noted that with the 'leaky integrator' and Morris-Lecar oscillators, average network synchrony decreased after reaching an optimal point around the 0.6...0.7 range. One possible explanation for this phenomena could be the problems mentioned at the start of this section concerning the decrease in the size of the lateral inputs as the network scope increases. As the connective scope is increased, the lateral influence created by each node decreases. With total network connectivity, the influence of an individual node is at its lowest point. If we were to alter the method used to integrate lateral inputs this might produce some improvement in the synchronization performance for the oscillators when network scope increases. However, although this effect is probably significant, we do not think it is the major contributing factor to the results found here as it does not seem to overly affect the Morris-Lecar oscillator.

For the Morris-Lecar oscillator, the finding that only nearest neighbour connectivity is required for synchrony in laterally connected networks is important within a neurophysiological aspect. It allows for relatively distant areas of cortex which are not directly connected to take part in the same synchronization process (although delays in transmission would eventually overcome this process as the distance increased). Two neurons which are separated by a relatively long distance would not require a direct connection between them in order to synchronize. Provided there were a chain of cells between them which were participating in the synchronization process the neurons could still synchronize. Conversely it is also possible to state that two nodes which are separated by some distance but are observed to be in synchrony could both be participating in the same computational process.

From a computational viewpoint, the finding that nearest-neighbour connectivity is all that is required to achieve synchrony in a Morris-Lecar network is important. An increase in the number of connections in a network means an increase in the computation required.

The fact that we need only use nearest neighbour connections to achieve synchrony means that we are able to minimize the computational load required for this aspect of network implementation.

3.4 Connective scope and speed of synchrony

A further question we can ask about the effects of changing network connection scope is how it affects the change in the *speed* of reaching synchrony in a network. We have already seen how scope affects overall synchrony but we do not have a measure of how well the network performs from a computational viewpoint. Measuring the speed of synchronization would give us a useful performance metric which would compliment the results of the previous section. The following section details a study designed to investigate this query.

3.4.1 Method

This study uses the same architecture and approach as those mentioned in section 3.3 except for the method used to measure the network performance. Instead of letting a network run for a fixed level of time and then measuring the network coherence at this point, a network is left to run until a given level of coherence is exceeded ($\bar{R} = 0.9$). The time taken to reach this level of coherence is then noted against the degree of connectivity used and a graph constructed from the results. If the network fails to reach the set level of coherence by a fixed number of time steps then it is stopped and the value of the limit used as the measure for the given level of connectivity. For each trial of a set level of connection scope, the same test series procedure is carried out as for section 3.3. 50 test runs are performed per trial value for the 'leaky integrator' and Morris-Lecar models and 30 for the Wang *et al* model. The set of time values for each level of connection scope are then used to produce an average time for synchrony and standard deviation for the sample. The results obtained for this study are shown in figures 3.9 to 3.11. The x-axis indicates the level of connectivity in the network and the y-axis the time step at which a coherence

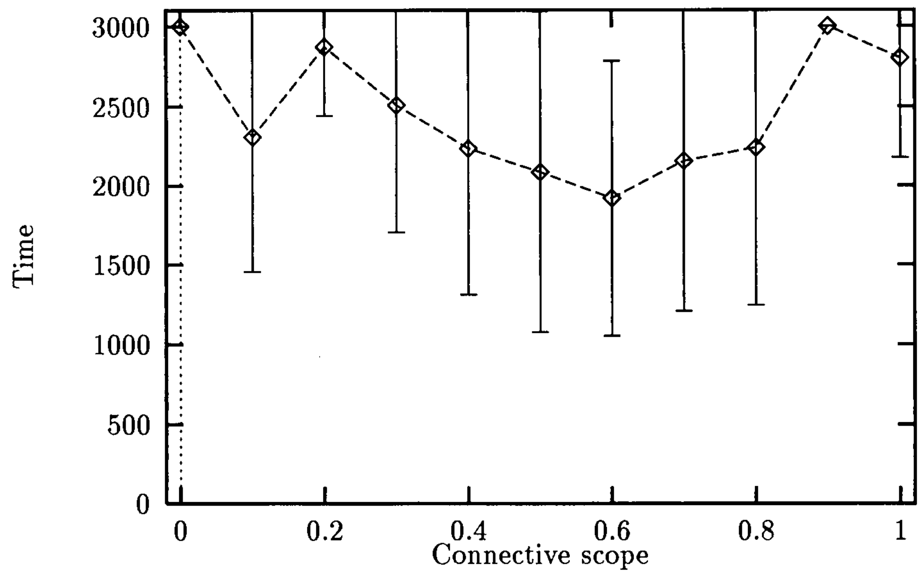

Figure 3.9: Scope of connectivity versus speed of synchronization for the 'leaky integrator' oscillator. Simulation parameters - nodes 20, trials 11, epochs per trial 0...3000, test runs per trial 50, lateral input 0.2, connective scope 0 (no connection) to 1 (total connectivity).

value of 0.9 is exceeded.

3.4.2 Discussion

Increasing the connection scope for the 'leaky integrator' model appears to produce some improvement in network synchronization speed. With only nearest neighbour connectivity (c=0.1), the network converges on average at around 2300 epochs. After some variation around the 0.2...0.4 range performance then improves until an optimum level of 0.6 is reached. After this point, the network coherence again deteriorates. The wide variation in synchronization performance indicated by the error bars shows that although the 'leaky integrator' network is capable of achieving synchrony, its average time to settle is high and that it frequently does not synchronize.

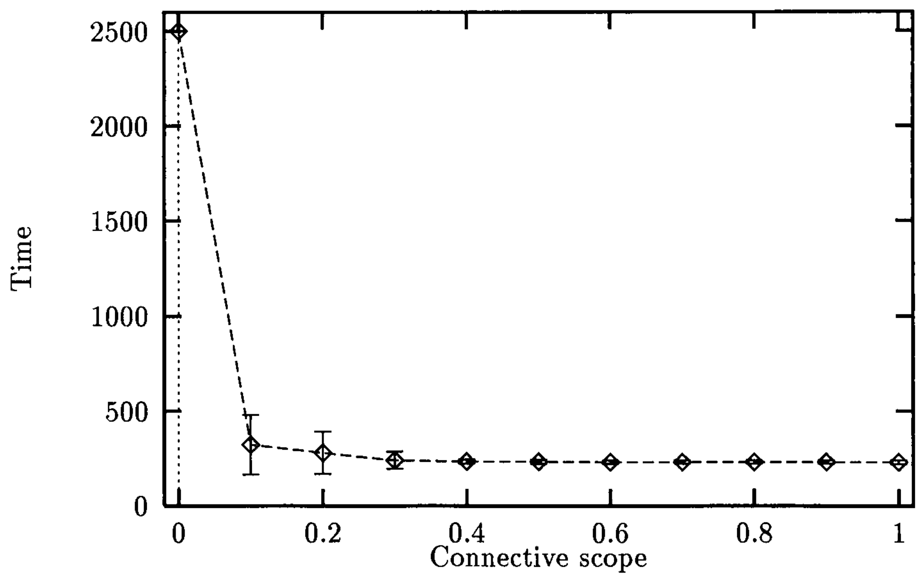

Figure 3.10: Scope of connectivity versus speed of synchronization for the Morris-Lecar oscillator. Simulation parameters - nodes 20, trials 11, epochs per trial 0...2500, test runs per trial 50, lateral input 0.3, connective scope 0 (no connection) to 1 (total connectivity).

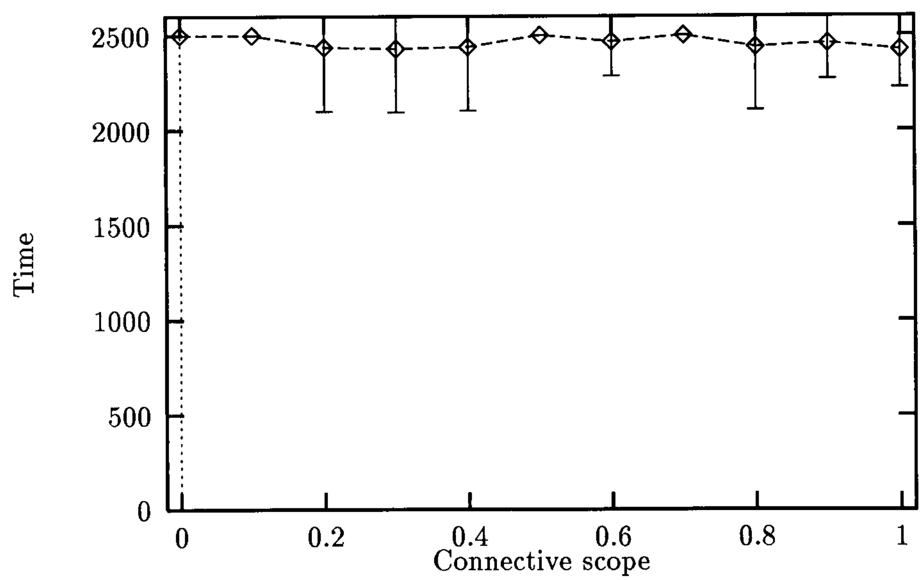

Figure 3.11: Scope of connectivity versus speed of synchronization for the Wang *et al* oscillator. Simulation parameters - nodes 20, trials 11, epochs per trial 0...2500, test runs per trial 30, lateral input 0.08, connective scope 0 (no connection) to 1 (total connectivity).

The results for the Morris-Lecar model show that it rapidly converges to synchrony on all trials where lateral connections are present. Although there is some variation in the time taken to reach network synchrony for the lower levels of connective scope, this disappears after the connective scope exceeds around 0.3. This indicates that if lower levels of connective scope are used in Morris-Lecar networks in order to reduce the computational load, a small trade off will be made in the time the network will take to reach synchrony.

Increasing the connective scope in the Wang *et al* network brings about almost no change in network synchronization speed. Although improvements in network coherence are found when network connectivity is in the region of 0.2 to 0.4 and 0.8 to 1.0, the overall effect is poor. From this result and that of the previous section it would seem that although increasing network connection scope for the Wang *et al* network is capable of enabling synchrony to occur, this effect is largely outweighed by the poor performance of the oscillator.

It is difficult to come to a general conclusion concerning the results found here. It would appear that the effects produced by altering network connectivity are primarily dependent upon the oscillator model used. Altering network connectivity in the leaky-integrator network does appear to have an effect but either this effect is unreliable or more probably the oscillator is not suitable for operation in a synchronizing oscillator network. The results for the Morris-Lecar oscillator are encouraging for the case where a model is suited to operation in a synchronizing oscillator network as it would seem that relatively low levels of connectivity are sufficient. For the Wang *et al* network, the results are very poor. However, it should be noted that some fluctuation in network performance did occur. This indicates that altering the network connectivity was affecting the behaviour of the network.

3.5 Noise in dynamic neural systems

The following section details a study performed on the effects of noise on coherence in a network of laterally coupled oscillators. The purpose of this study is to determine the general robustness of lateral synchronization to the effects of noise using a number of different oscillator models. This will hopefully show, at least from a simulation point of view, that lateral synchronization can cope with the noisy environment of real systems. If this can be proven, then it would also give more support to the possibility that lateral synchronization is present in real neural systems which by their nature are quite noisy. If it is shown that synchronizing oscillator networks do not operate very well under noisy conditions then this would be evidence against the hypothesis that synchronization is used as a computational tool in cortex.

3.5.1 Modelling noise

When considering this type of experiment, it is important to determine what form the noise should take and where it should be applied. Within a neurophysiological framework, it is relatively easy to state what factors might produce noise in a network of coupled cells (for example background chemical activity at the synapse or irregularities in the transmission of signals down an axon). It is not so easy to decide exactly how to model this. A simplified approach has been taken with this study. We have decided to model two types of noise in our network.

1. Noise at the synapse of the lateral connections.

2. Noise in the driving input.

The noise at the synapses of the lateral connections is intended to simulate effects which might realistically be expected to perturb a system which relied on lateral connections to achieve synchronization. By putting noise in at this point, we can determine how resilient the method of synchronization through lateral connectivity is.

The second type of noise, noise in the driving input, is used to investigate how lateral synchronization can help to overcome noise in the input to a system. If there is no noise in the lateral connections but noise in the driving input, then we can determine how effective the lateral connections are at overcoming the perturbations caused by this type of noise. This will let us know to what extent lateral connectivity is capable of causing a system to maintain synchrony despite the fact that the underlying frequencies of the nodes in the network differ.

Noise affecting lateral input

Having decided where to place the noise, we need to decide how to model it. The easiest method of implementing noise would be to perturb a given signal by adding or subtracting a random percentage of its current value from itself. On average this would give an underlying signal which was consistent with the original signal without noise. This is not a very realistic model, particularly in the case of spikes arriving at synapses.

Noise of the above form will not greatly perturb a system which relies on lateral connections to cause a phase shift and therefore achieve synchronization. The noise will only attenuate the signal but not greatly reduce the temporal information contained within it. It is the instant of occurrence of a spike which is the determining factor affecting the node receiving it. With the above method, a large signal arriving at a synapse will still create a large post-synaptic effect. This will be enough to cause the phase of the receiving node to shift in the direction of the phase of the node transmitting the signal. The main effect of this type of noise is to slow down the speed with which a network of laterally coupled oscillators synchronizes. We would like a form of noise which can confuse the receiving node by sending it false information.

To create this effect at the synapses of the lateral connections, a synapse generates noise when there is no pre-synaptic spike present (equation 3.11). If a pre-synaptic spike is present, the post-synaptic spike it generates is presumed to drown out the noise. The noise generated represents background effects such as neighbouring synaptic transmission or residual effects from previous reactions — effects which might cause the synapse to generate a signal when there was little or no pre-synaptic activity. This type of information will cause the receiving node to shift its phase in the direction of the erroneous noise. On average this will be away from the desired phase. It should however be noted that the lateral input signal will still contain some temporal information. On average the background noise will only be 50% of the strength of the lateral input signal. This will usually make the lateral spike distinguishable from the background. A node receiving the signal will therefore be able to engage in some synchronization behaviour even when a high degree of noise is present.

$$l_{ij}(t+1) = \begin{cases} \gamma & a_j(t) \geq \alpha \\ f_{noise}(\gamma . \eta) & a_j(t) < \alpha \end{cases} \tag{3.11}$$

$l_{ij}(t+1)$	Post-synaptic lateral input from node j to node i at time *t+1*
γ	Constant for post-synaptic spike
α	Constant for synaptic threshold
$a_j(t)$	Pre-synaptic activation at time t
$f_{noise}(x)$	Random noise function which returns a value between 0 and x
η	Level of noise $\{\eta, 0 \leq \eta \leq 1.0\}$

Noise affecting driving input

Since the driving input to the oscillating nodes is constant, a simpler form of noise may be applied to it in order to perturb a system. Within certain limits (dependent upon the oscillator), the frequency of the oscillation of a node is proportional to the size of the driving input it receives. Any perturbation in the size of the driving input to a node will cause it to oscillate at a different frequency and therefore attempt to move it out of phase.

It does not make sense to use the form of noise mentioned for the lateral connections as in this case the driving input is always 'on'. We require a form of noise which will cause a controllable degree of perturbation in the driving input but which still retains some relationship with the underlying driving input. In the study discussed below, noise in the driving input is modelled according to equation 3.12.

$$T_i(t) = \tau_i(1 - \eta) + f_{noise}(\tau_i.\eta)_t \tag{3.12}$$

T(t)	Driving input for node i at time t
τ_i	Driving input without noise
$f_{noise}(x)_t$	Random noise function which returns a value between 0 and x for time t
η	Level of noise $\{\eta, 0 \leq \eta \leq 1.0\}$

The amount of noise in the driving input is controlled by the ratio η. The value of η controls the proportion of driving input which is composed of pure noise versus the amount which is 'clean'. With η set to 0.4 for example, 40% of the driving input will be pure noise. The other 60% will be made up with the 'clean' driving input. By varying the value for η we can alter the setting of the driving input and thereby alter the underlying frequency of the node which is been driven by it.

3.5.2 The effects of noise on network coherence

Having decided how to model noise in our networks, we can now investigate its effects. In this study we observe the change in network coherence which occurs when the level of noise in a network is altered. Three studies are performed, one for each of the oscillator models discussed throughout the thesis. For each study, three different experiments are performed. In the first experiment we look at the effects of noise in the *driving* input. In the second experiment we look at the effects of noise in the *lateral* input and in the last experiment we combine the two effects to investigate noise in both driving and lateral input.

Method

The architecture used in this study is the same as is shown in section 3.1.1 except where modifications have been made to simulate the effects of noise. The equations used to calculate the lateral and driving input for the basic model (section 3.1.1, equations 3.1, 3.2) are now replaced by equations 3.11 and 3.12 where appropriate. In the first test case, the driving input T_i in equation 3.3 is replaced with equation 3.12. In the second test case equation 3.1 is replaced by equation 3.11. In the third case both equations are substituted.

For each test, the simulation is performed over 20 trials. For each trial value of a given noise level η, a number of test samples are performed (40 for the leaky-integrator and Morris-Lecar models and 30 for the Wang *et al* model). The combined results for each set of trial samples is used to produce an average response and standard deviation for the given level of noise η. The same method for using a standard set of initial starting points for each trial series is used in this study as was used in the studies of sections 3.2 to 3.4.

On each consecutive trial, the degree of noise η is increased, ranging from 0 in the first trial to 1 in the last trial. At the end of the test series, the results obtained are used to generate a set of scatter graphs showing the relationship between noise and coherence for each of the oscillator models. The results for this study can be seen in figures 3.12...3.17. In each graph, the x-axis indicates the degree of noise in the system and the y-axis indicates the level of coherence found. The simulation parameters used for each of the studies are given in the following table.

Model / *Parameter*	**'leaky-integrator'**	**Morris-Lecar**	**Wang** *et al*
Nodes	20	20	20
Trials	20	20	20
Sim. epochs	3000	2500	2500
Sample test runs	40	40	30
Lateral input	0.2	0.3	0.12

Discussion

The results for noise in the driving input of the leaky-integrator model shown in figure 3.12a, indicate that it is not very resilient to even low levels of this form of noise.

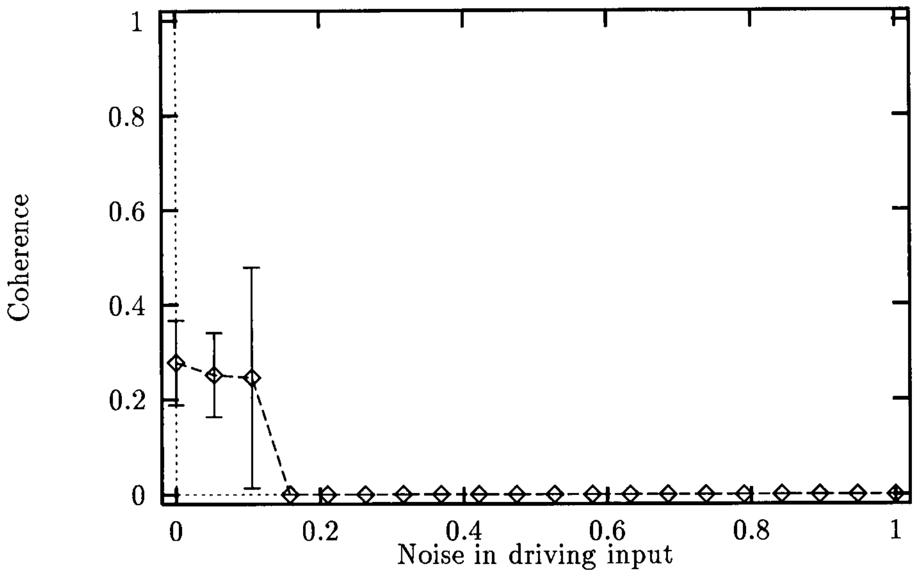

(a)

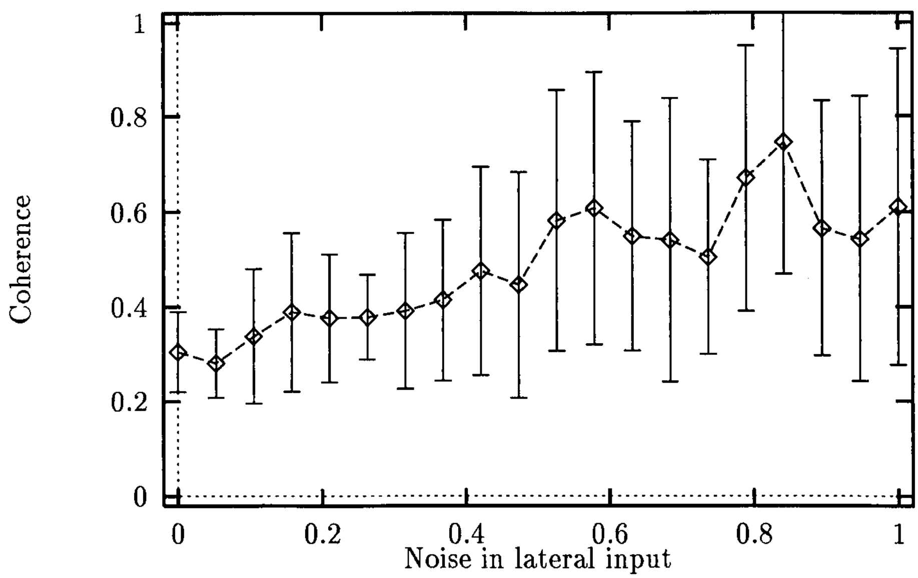

(b)

Figure 3.12: Scatter plot of noise versus coherence for the 'leaky-integrator' oscillator. **(a)** Noise in driving input. **(b)** Noise in lateral input.

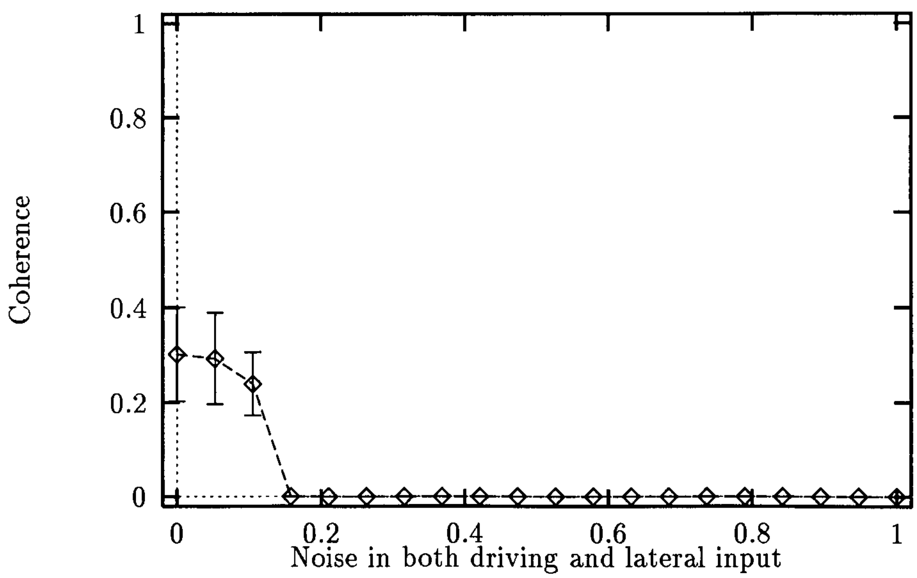

Figure 3.13: Scatter plot of noise versus coherence for the 'leaky-integrator' oscillator with noise in both the driving and lateral input.

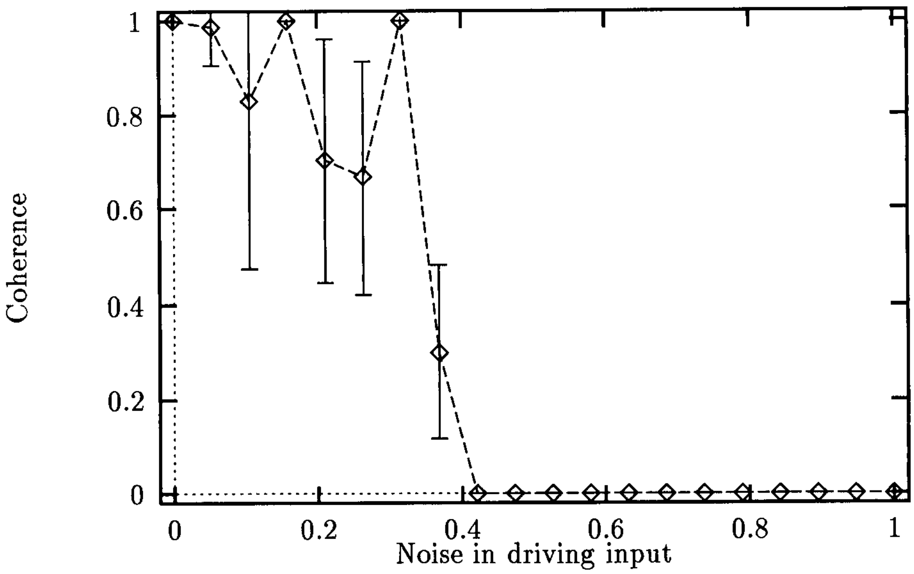

(a)

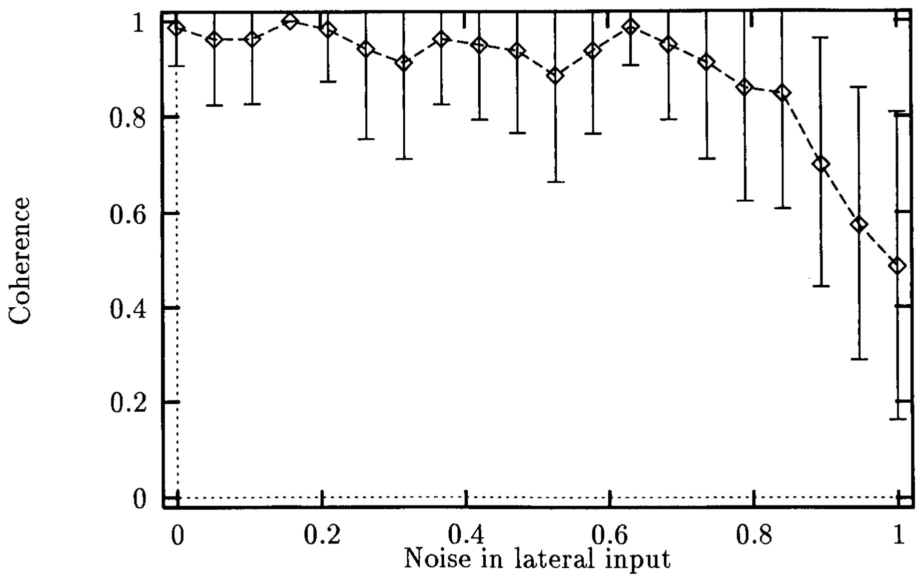

(b)

Figure 3.14: Scatter plot of noise versus coherence for the Morris-Lecar oscillator. (**a**) Noise in driving input. (**b**) Noise in lateral input.

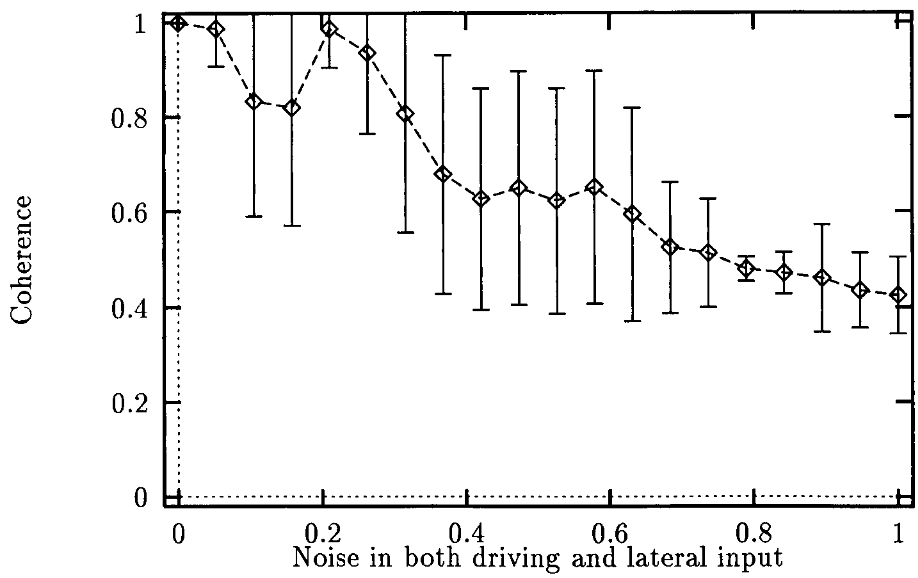

Figure 3.15: Scatter plot of noise versus coherence for the Morris-Lecar oscillator with noise in both the driving and lateral input.

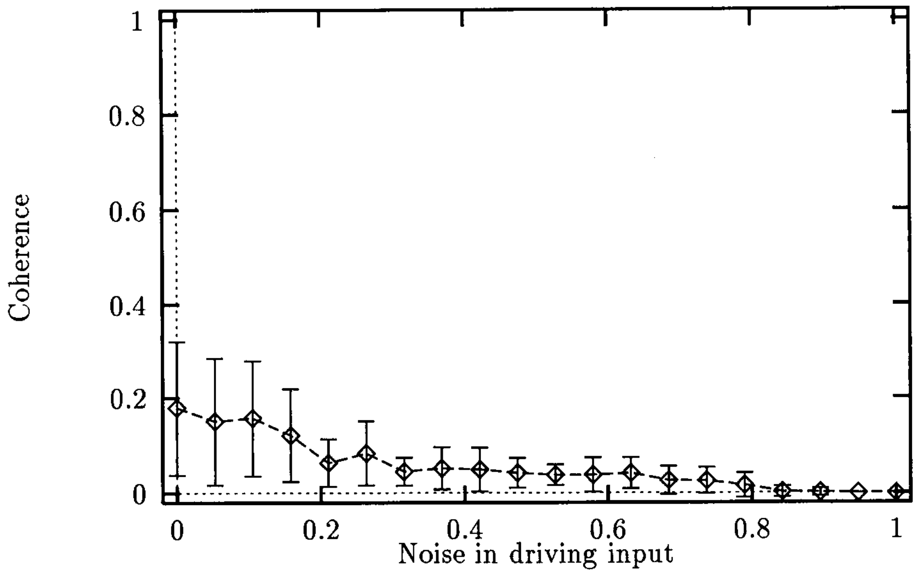

(a)

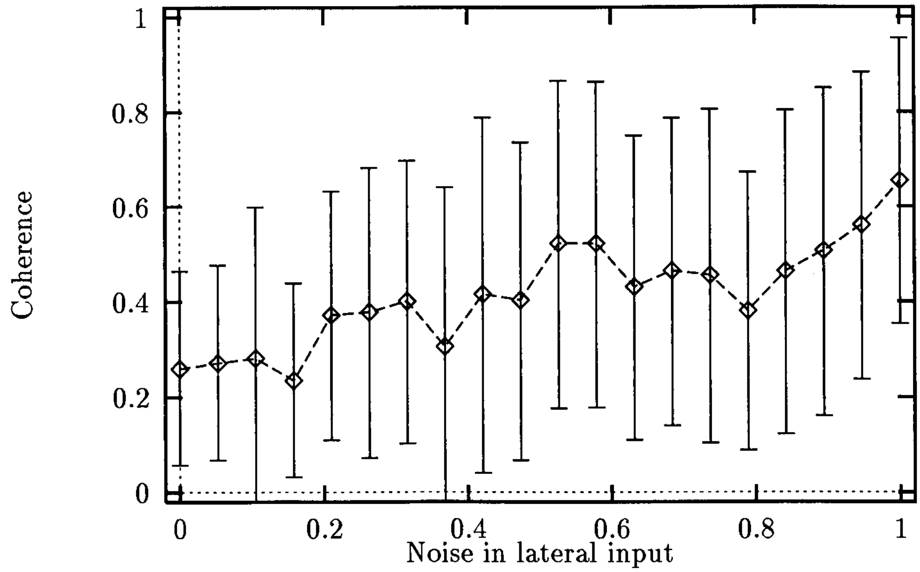

(b)

Figure 3.16: Scatter plot of noise versus coherence for the Wang *et al* oscillator. **(a)** Noise in driving input. **(b)** Noise in lateral input.

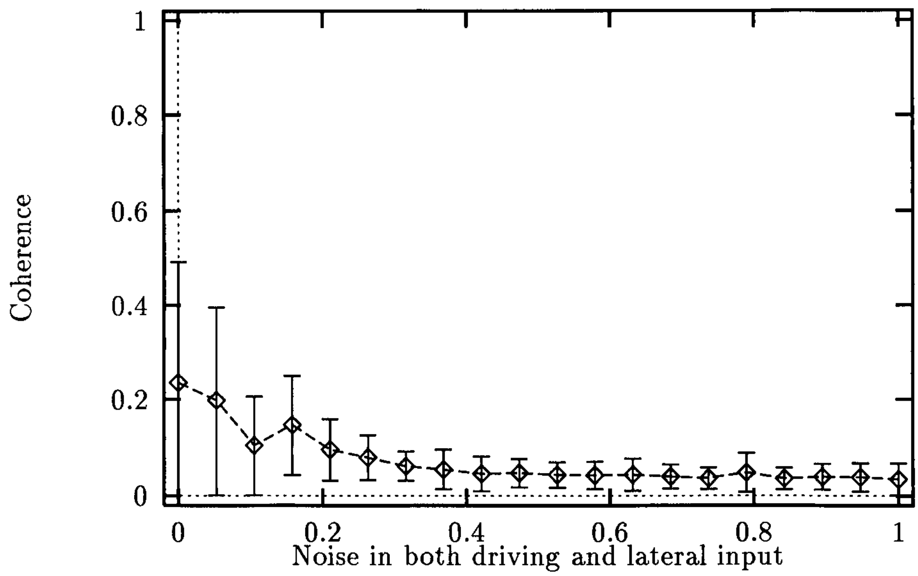

Figure 3.17: Scatter plot of noise versus coherence for the Wang *et al* oscillator with noise in both the driving and lateral input.

Given that the network's coherence measures were already quite poor, added noise causes a total breakdown in network performance. The predominant cause of this deterioration is probably that the nodes in the network cease to fire once the noise levels becomes too high. The leaky-integrator model requires both a consistent and strong driving input in order for it to reach threshold and fire. There is no momentum term in the equations defining the oscillator which will cause the activity of the node to increase if the driving input is removed or decreased for a period of time. When the noise levels rise and the variation in the driving input becomes to great, the nodes in the network therefore fail to fire and a low coherence measure is obtained.

One curious result shown in figure 3.12a is the wider variation in the sample test results when noise in the driving input is at a level of 0.1. This would seem to indicate that for the leaky-integrator model, a small amount of noise in the driving input can be beneficial on occasion. This is made even more evident in the results of figure 3.12b which shows the effects of noise in the lateral input. Noise in the lateral input for this type of network actually improves network coherence. There would seem to be an analogy here with the type of behaviour observed when using simulated annealing in Boltzmann networks (Ackley *et al.*, 1985). High temperature (rapid fluctuation in state values) allows the network to sample a large state space, including areas which it would not enter if the allowable movements were much smaller.

Taking the viewpoint that the desirable position in state space for our network is complete synchrony and that once a leaky-integrator network finds itself in this state, it will remain close to it, the problem would seem to be to get the network near to this state. From the results of sections 3.2...3.4, it seems that the leaky-integrator network is normally quite poor at finding this state, implying that it does not naturally converge toward a state of synchrony. If we add noise to this system then in effect we force the network to enter states it would not normally go into. By increasing the level of noise in the network we are increasing the chance that it will go into the desired state. It also follows that we are probably increasing the chance that the network will then be moved out of this state. However it is possible that it would be easier to get into this state than out of it, particularly if the network is actively trying to keep itself synchronized. From the results and conclusions of above, it would therefore appear that some level of noise in a leaky-integrator network is actually a desirable property. Some of the measures obtained for noise in the lateral inputs of the leaky-integrator network are significantly greater than those found in any of the other studies for this model oscillator.

When noise in the driving and lateral input is combined in the leaky-integrator model, the behaviour is extremely poor (figure 3.13). This is probably due to the effects observed in figure 3.12a. The dependency of the leaky-integrator oscillator on a strong and consistent driving input results in the response of the network to noise in both driving and lateral input being dominated by the response to the driving input.

When noise is placed in the driving input of the Morris-Lecar network, the network remains relatively coherent until the level of noise reaches a level of around 0.35 (figure 3.14a). The network is obviously affected by the noise in the driving input but it is still able to generate a reasonable response. Once the noise level increases beyond this point, network coherence sharply deteriorates. This result is probably due to similar causes to those observed in the leaky-integrator network. In this case however the oscillator dynamics include some degree of momentum and the oscillator is able to withstand more noise before it eventually fails. Another possible cause of the deterioration in coherence is the increasing differences in the underlying frequencies caused by the increase in noise. Each driving input to a node will effectively be a different value and will therefore cause each node to oscillate with a different frequency, thus making the synchronization process more difficult (Cohen, 1987).

In the case where noise is added to the lateral inputs, the network shows very robust behaviour (figure 3.14b). Even with noise at a level of around 0.8, the network is still able to maintain a high degree of synchrony When noise is increased beyond this level the network produces an almost linear decrease in coherence. Considering the fact that the lateral connections are the sole method by which the network achieves synchrony, this result indicates that achieving synchronization through lateral connectivity is a very robust phenomena when using the Morris-Lecar oscillator.

Combining the effects of noise in both the driving and lateral input of the Morris-Lecar network, we observe a deterioration in network synchrony (figure 3.15). It would seem that with noise up to a level of 0.35, the system is able to cope quite well. There is a sufficiently consistent driving input present for the nodes to be able to both oscillate and synchronize. Once the noise level goes beyond this however, the deterioration in the level of driving input begins to dominate.

Noise in the driving inputs of the Wang *et al* network causes a drop in coherence below even the normally low background measure obtained with the Wang *et al* model. This is probably due to a rapid deterioration in the number of nodes which are able to fire. This behaviour is similar to that observed for noise in the driving input of the leaky-integrator model. Neither of these models are capable of maintaining the increase in their activities without a strong and consistent driving input.

The effects of noise in the lateral inputs of the Wang *et al* network seem to be a general improvement in network synchrony. This result is in line with the observations made for the leaky-integrator model which also benefited from the addition of noise. It would seem that where a particular oscillator network does not easily achieve synchrony (as shown by the results of the previous studies), addition of noise can actually improve its performance. In the case of the Morris-Lecar network, a network which synchronizes quite effectively, addition of noise cannot really improve its already adequate performance. This would give a consistent explanation of the differences between the results obtained for each oscillator network.

It is also worth noting that the results for the Wang *et al* network show a high degree of variation. From an analysis of the sample values used to produce the average coherence measures, we determined that the network frequently obtained very high coherence measures ($\bar{R} \rightarrow 1$). However this was then offset by very low measures producing the given mid-range average values. This average value is slightly misleading since most of the coherence measures obtained lie around the upper and lower limits of the error bars. This would seem to indicate that synchrony in this network is something of an all or nothing phenomena.

The results for noise in both the driving input and the lateral input of the Wang *et al* network show that network behaviour is predominantly affected by the consistency of the driving input. As the level of noise increases, the consistency of the driving input decreases. This causes the network to mimic the behaviour found for noise in the driving input alone. Without a sufficient and consistent driving input, the nodes in the network cease to oscillate and low coherence measures are obtained.

The main conclusion to draw from this study is that, where an oscillator is suitable for operation in a synchronizing oscillator network, it should be robust to a relatively high degree of noise. Where it is not suitable, or at least does not synchronize effectively, addition of noise to the lateral inputs can cause an improvement in network behaviour. If we allow ourselves to draw analogies between the behaviour of neuro-physiologically based oscillators such as the Morris-Lecar oscillator and real neurophysiological systems, then we can also conclude that these systems should be sufficiently robust to withstand the noise inherent within them. This is a tentative step to make but if it were the case, it would certainly give more support to the possibility of synchronized oscillations being present in cortex. A further supposition we can make from this point is that, given these systems are robust to noise, it could be part of their function to overcome noise in cortex. If signals transmitted simultaneously from one area in cortex arrive at their destination slightly out of synchrony due to transmission delays, then the dynamics of the synchronizing systems would ensure that the signal were re-synchronized. This would be a very useful role for synchronization in neural systems.

3.6 Summary

In this chapter we have looked at a number of principal factors which affect the behaviour of laterally coupled networks of synchronizing oscillators. The first section investigated the relationship between lateral connection strength and network coherence. The main findings from this section were that increasing connection strength produces an improvement in network coherence although the exact behaviour is dependent upon the oscillator model used. For the leaky-integrator and Morris-Lecar oscillator, increasing the lateral connection strength generates an improvement up to a point after which network coherence reaches an asymptote and no further significant change occurs. The asymptote for the leaky-integrator model is at quite a low value, reflecting its poor overall synchronization performance. Increasing network connection strength for the Wang *et al* model causes some improvement but there would seem to be an optimal level for connection strength in the Wang *et al* network. Further increases in connection strength caused a deterioration in network coherence.

From these results we concluded that the main cause of the fluctuation in network coherence is the phase shift produced by the lateral inputs. When the lateral inputs are small, the phase shifts are correspondingly small and therefore the degree of general movement in the state of the network is small. This allows the network to slowly search for the synchronized state without overshooting or moving about too rapidly. As the network connection strength is increased, the degree of movement in the network increases. This can lead to an improvement in network synchrony as in the case of the leaky integrator network or Morris-Lecar network. However, it can lead to a deterioration in network synchrony as the network moves about too much and is unable to stabilize as in the case of the Wang *et al* network. The effect produced is dependent upon the dynamics of the oscillator model used, however we think this principle underlies the behaviour observed for all the oscillator models studied.

The second section in this chapter investigated the relationship between connective scope and network coherence. We found that total connectivity does not necessarily produce the most optimal coherence levels. For the leaky-integrator model, network interconnectivity at the level of around 0.6 gave the most optimal performance. With the Morris-Lecar network it was shown that total interconnectivity is neither necessary nor the most optimal architecture to use. Increasing network connectivity beyond the level of 0.5 actually brings about a slight deterioration in performance. Increasing connective scope with the Wang *et al* network produced a small improvement in network coherence indicating that, at least for this model, increasing scope is beneficial. The results for this model also showed a greater variation in the average coherence measure obtained as the connective scope was increased. From this we concluded that increasing connective scope was causing the network to reach higher levels of synchrony than before but that frequently the network did not synchronize very well, therefore bringing down the average coherence measure obtained.

A particularly useful conclusion we drew from the Morris-Lecar result is that, when an oscillator is well suited for operation in a synchronizing oscillator network, it should be possible to use relatively low levels of interconnectivity. This cuts down the computational overhead required to use the network. From a neurophysiological viewpoint, we can conclude that since only local interconnectivity is required, cortex would be capable of using synchronization as a computational tool. Indeed it would seem advantageous to do so as it would allow for binding across areas which are not directly connected yet need some method of relating global information.

The third section in this chapter investigated the effect of altering connection scope in terms of synchronization speed. The results obtained in this section were directly related to those of the previous section but gave a different perspective on synchronization performance. In terms of synchronization speed, we concluded that increasing network inter-connectivity in the leaky-integrator network produced some improvement in performance. We found that with the optimal level of connective scope deduced from the results of the previous section (around 0.6), we also obtained the fastest synchronization times.

The results for this network did however vary widely, indicating that obtaining synchronization in this network was an all or nothing phenomena. The conclusions drawn for the Morris-Lecar network were in line with those of the previous study. With even just local connectivity, the network synchronized rapidly. The small variation observed in network behaviour when using localized connections was again found. From this, we concluded that to guarantee synchrony, a connective scope of around 0.3...0.4 would be optimal. Increasing the connective scope for the Wang *et al* oscillator produced very little change in its synchronization performance. From this we again concluded that, although altering connective scope can have an effect on the synchronization performance of a network, the result produced has some dependence on the model used.

The fourth section in this chapter concerned the robustness of synchronizing oscillator networks to noise. For the case of noise in the driving input of the three networks, it was found that all three model networks were vulnerable to this type of noise. The leaky-integrator and Wang *et al* models showed a rapid deterioration in network coherence in the presence of noise levels of even 0.2. The Morris-Lecar model was shown to be more robust but it too rapidly deteriorated when the noise level exceeded 0.4.

For noise in the lateral inputs, the responses of the different oscillator models were found to be dependent upon the suitability of the oscillator model to synchronization. For the leaky-integrator and Wang *et al* models, synchronization performance was found to improve with the addition of noise. The Morris-Lecar model maintained a robust resistance to relatively high levels of noise and then deteriorated. It was concluded that when noise was added to systems which were poor at synchronizing, it enabled them to synchronize more effectively. If the oscillator model was already effective then no real improvement could be made.

Combining noise in the driving and lateral inputs of the nodes in the network caused a general breakdown in network coherence. It was concluded that this effect was similar to that seen in the other two networks and that it was primarily due to insufficient driving input resulting in a failure of the nodes to oscillate.

Chapter 4

Computation using synchronizing oscillators

Given that we have a framework for the requirements necessary to achieve synchrony in laterally coupled networks of oscillators, we now investigate some computational processes which use these networks. This chapter covers a number of studies which investigate some of these computational processes using the Morris-Lecar oscillator. It begins with a discussion of a study intended to produce a simplified replication of the results of Gray *et al* and Eckhorn *et al.* This is followed by a discussion and illustration of the potential use of learning in synchronizing oscillator networks. The last section in the chapter investigates the hypothesis of using multiple phases to simultaneously bind a number of objects.

4.1 The split line experiment

The following section discusses a simplified simulation intended to investigate the results of Gray *et al* and Eckhorn *et al* (Chapter 2, section 2.2.1). The network used in this simulation is only one dimensional but it serves to illustrate the principles which are suggested to be behind the results of Gray *et al* and Eckorn *et al.*

4.1.1 Architecture

The architecture used in this study is very similar to that used in the previous chapter except for the degree of connectivity and the number of nodes used. Connection scope in this study is fixed at 3 units either side of each unit with wrap-around at the edges. This connectivity structure is sufficient to allow synchronization between connected groups to occur yet short enough to allow for different, unconnected groups to develop within the network.

The structure of the network is fixed as a one dimensional array of 20 nodes where the activity of each node is modelled using the Morris-Lecar oscillator (Appendix A.2.2). It is accepted that the architecture used here is a one dimensional simulation of what is physiologically a two (or more probably three) dimensional system. However adding a second dimension to the network produces no real gain as the effect of the input, a one dimensional line, remains the same. It would only stimulate nodes in a one dimensional subspace of the network.[1]

4.1.2 Method

A 'line' of length 8 units with a gap in it is presented in each of four different experiments. With each successive experiment the size of the gap in the line is increased, starting with no gap and going up to a gap of 3 units (the maximum distance an output of a single node will reach with the defined level of connectivity). A gap of three units between the two lines results in there being no direct connection between a node receiving input from one of the line segments and a node receiving input from the other line segment. Any interactions between these nodes must occur through an intermediary node which is receiving no external input.

The only difference between each experiment is the input used. For each of the experiments, each node is initialised to a random position in phase space but this initial position

[1] We have performed a similar experiment with a two dimensional network and the results were found to be effectively the same.

for each node is kept constant over the different experiments. This ensures that each node is initially out of phase with its neighbours but that the same initial conditions prevail from experiment to experiment. Each experiment is run for 2000 time steps and the activation levels for each of the nodes recorded. The external input to each node is kept constant over the entire period of the experiment.

4.1.3 Results

The results of the experiments are used to generate an activation contour plot. The graphs show the synchronization behaviour produced as a result of the external input presented to the networks. An activation contour plot is used to display the results as it was found to be the most informative method of illustrating the data. A three dimensional perspective of the results of the first experiment is also included to show the relationship between the contour plots and the three dimensional graphs from which they are derived (figure 4.1). In the three dimensional plot, the x-axis represents time, the y-axis the the number of a node and z-axis the activity of a node. The contour plots maintain the same format except that the z-axis is represented by brightness. Highly coherent peaks of activity occur as white bands and low or negligible activity is represented as a broad area of gray.

It can be seen from the graph of figure 4.1 that the number of plot points for each node over time is relatively low compared to the number of time steps over which the experiment was run. This does not represent a low degree of accuracy in the data produced by the experiment. The data used to produce the graphs was sampled from the output of the experiment at a relatively low rate to allow manageable conversion of the data into the graphical output shown.

The graphs of the results produced in this study are shown in figures 4.1..4.5. Each figure shows the input used for the experiment which produced the contour plot together with the contour plot itself.

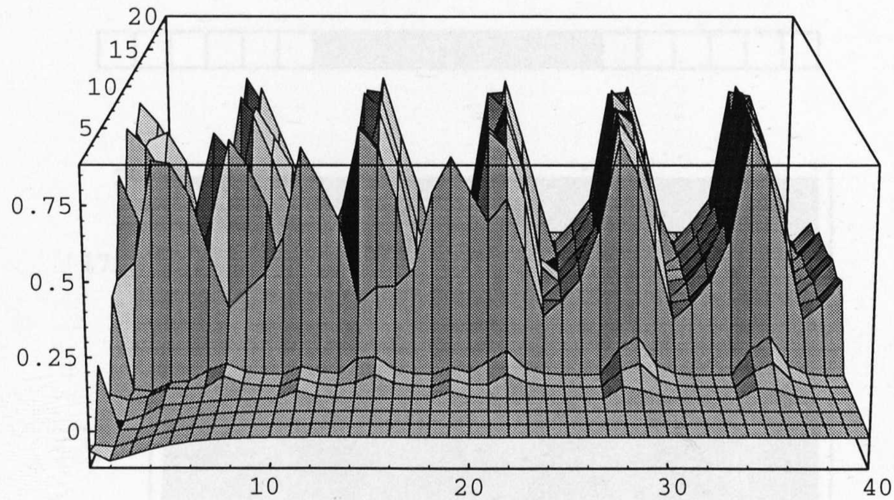

Figure 4.1: A three dimensional plot of the network activity over time for experiment one — the single line. The x-axis indicates time, the y-axis indicates the number of an individual node and the z-axis shows node activity.

4.1.4 Discussion

It can be seen in figures 4.2...4.5 that as the two bars are steadily moved apart the correlation between the two groups drop. This is a reflection of the reduced effect each group is having upon the other as the distance increases. Where there is no direct connection between the group representing one bar and the group representing the other, the two groups remain out of phase with each other (figure 4.5). Within each group however, the nodes go into and maintain synchrony. Where two groups are split but still have lateral connections maintained between them, each group first synchronizes locally and then begins to

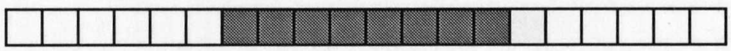

(a)

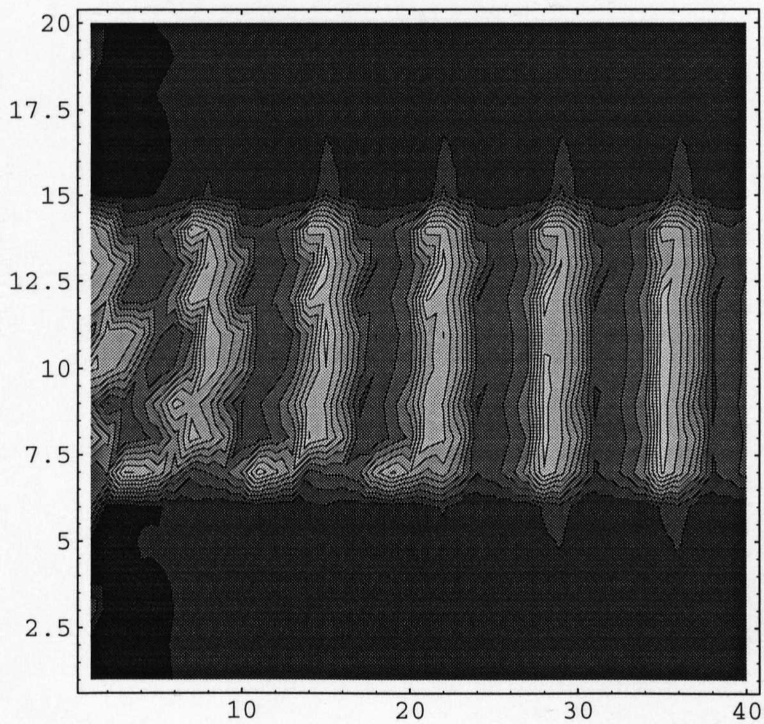

(b)

Figure 4.2: Experiment one : A single line. **(a)** Input **(b)** Contour plot of activity over time.

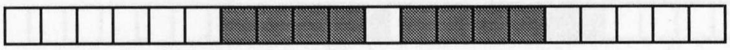

(a)

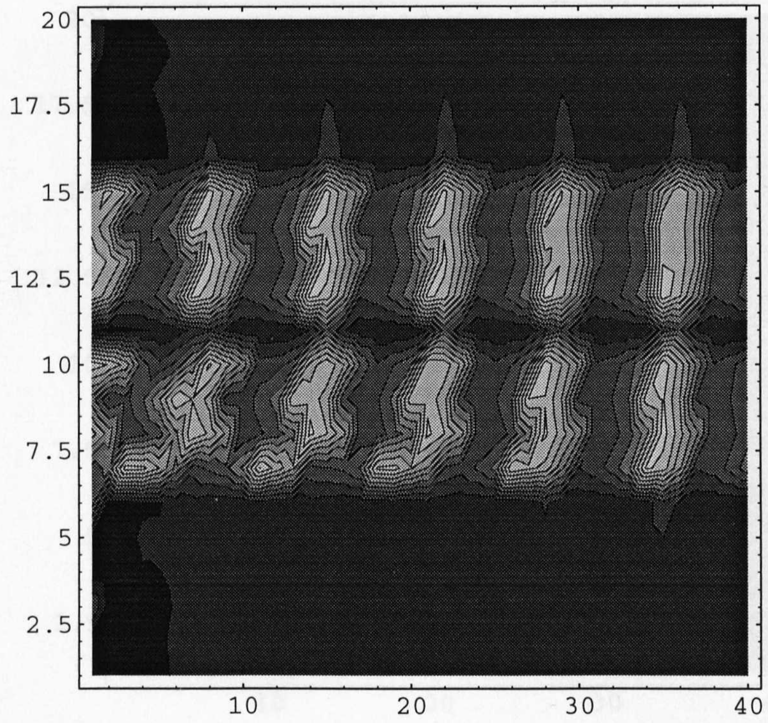

(b)

Figure 4.3: Experiment two : Split line with a gap of 1 unit. **(a)** Input **(b)** Contour plot of activity over time.

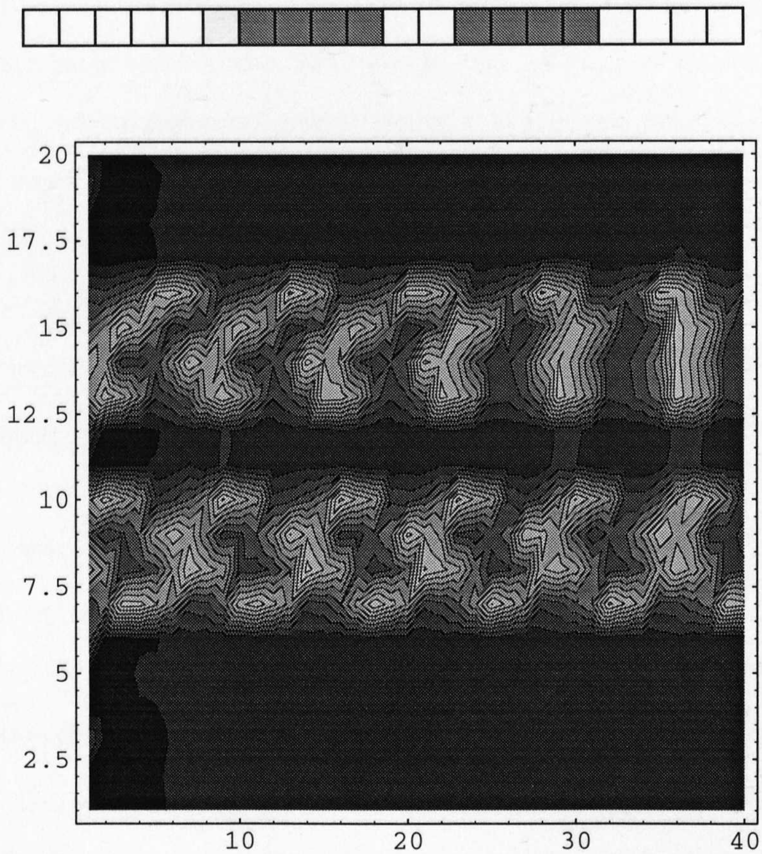

(a)

(b)

Figure 4.4: Experiment three : Split line with a gap of 2 units. **(a)** Input **(b)** Contour plot of activity over time.

(a)

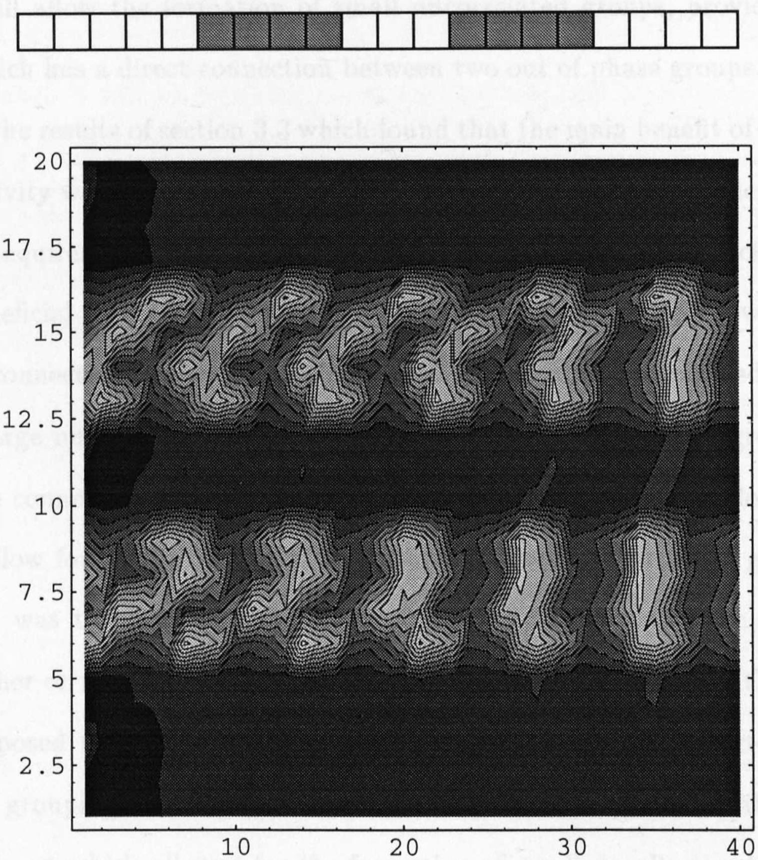

(b)

Figure 4.5: Experiment four : Split line with a gap of 3 units. **(a)** Input **(b)** Contour plot of activity over time.

synchronize with the other group.

One of the important conclusions from this study are its implications for the connectivity of a synchronizing neural network. It shows that the degree of connectivity is related to the size of synchronized groups obtained with the network. A net with extended connectivity will tend to stabilise into one synchronized group whereas one with only local connectivity will allow the formation of small uncorrelated groups, provided there is no active node which has a direct connection between two out of phase groups. These conclusions relate to the results of section 3.3 which found that the main benefit of using extended lateral connectivity was an increase in the speed with which a group of nodes synchronized. Unless it is a requirement that a large number of nodes are kept in synchrony, it would seem to be beneficial to keep the connectivity of the network relatively low.

With low connection scope, it is still possible to obtain synchronized groups which encompass a large number of nodes, provided there are no gaps in the groups which are larger than the connection scope. The main disadvantage of low connection scope is that it would not allow for the formation of large but dispersed synchronized groups of nodes where the aim was to bind a number of unconnected fragments in the input into one 'clump'. Whether or not this would be a disadvantage would be dependent upon what the network is supposed to be doing. For example, if we require a net to perform some form of coarse scale grouping, then the above behaviour would be quite desirable. If however we would like a net which allowed for the formation of small, locally synchronized groups where each group was out of phase with the other, then high connection scope would not be an advantage.

The results of this study illustrate how the principle of synchronization through lateral connections is affected by the degree of connectivity between synchronizing groups. It shows that connection scope has a direct bearing on the type of grouping that a synchronizing oscillator network will perform. Furthermore, it concurs (at least in principle) with the findings of Gray *et al* and Eckhorn *et al* as regards the observed loss of synchrony which is found when the two bars discussed in their experiments are moved apart and offers at least one explanation as to why this is so.

4.2 Adaptive synchronizing networks

Having shown how synchronizing oscillator neural networks can bind groups of nodes together through common phase, we now ask whether it is possible to extend the abilities of these nets to learn the patterns which are to be bound. Can we configure a network in such a way that it will learn which nodes to bind and which to ignore without prior hard-wiring? The following section investigates these possibilities from a practical point of view and discusses where previous theories on learning in static neural networks might be applied to synchronizing neural networks. It also discusses how learning rule behaviour and interpretation differs from that of static neural networks.

4.2.1 Learning rules for synchronizing oscillator networks

One method of applying learning to synchronizing oscillator networks is to use weighted connections to encode binding information. Because activation transfer through lateral connections between nodes is the principle cause of synchronization, weights on the lateral connections can be used to encode whether two nodes should synchronize. A lateral connection between two nodes which has a large weight (near 1), will allow most of the lateral activation to pass from one node to another and will therefore make synchrony between the nodes more likely. A low weight will not permit sufficient lateral activation to pass through, leaving the states of the two nodes relatively unaffected.

If a particular pattern of activity is established in a group of nodes and it is desirable that nodes which are nearly in phase should synchronize then adaptation can be applied to the lateral weights in between them in such a way as to cause an increase in the weight. If there is no decay affecting the lateral weights then the excitatory connection established, will encourage the nodes to synchronize if a similar temporal pattern re-occurs on a later presentation.

Static versus dynamic learning behaviour

Given the dynamic behaviour of synchronizing oscillator networks, we need to approach the problem of learning in a slightly different way. In static auto-associative networks the activity of a node is usually rate coded, the activation value reflecting the average firing frequency of a node. Learning with a static network is applied according to whether or not a node is responding to meaningful data and a weight change derived from the differences between the activation states of the nodes concerned. A frequent format for weight change is to cause nodes which are highly co-active to have their weights increased and nodes which differ in their activity levels to have their weights decreased or left unchanged (Rumelhart & McClelland, 1987a; Rumelhart & McClelland, 1987b; Hertz *et al.*, 1991). We would like to maintain the principle of this behaviour in order to encode the information that two oscillating nodes are near in phase and should therefore be encouraged to synchronize. Since we are no longer dealing with an approximation to the firing rate of a node but the actual process of firing itself, applying the static auto-associative method of learning in its simple form to a synchronizing oscillator network would not be sensible.

The aim of our learning rule should be to strengthen connections between nodes which are close in phase, no matter what their frequency. A system which penalised low firing frequencies as being in-active would therefore not be appropriate. In effect we require a learning rule which detects close (but not absolute) coincidence between signals and strengthens the connections between them. It would also be useful to have some flexibility which allowed small movements in the phase of a node in between spikes to have no adverse affect on the strength of the connections.

One solution to this problem is to use an activation threshold to determine whether a node is 'on' or spiking. The threshold is set to detect the spiking region of an oscillation. Adaptation of the lateral connection going from one node to another then follows a standard Hebbian type learning rule. The weight change rule used in these studies is the Stent-Singer rule shown in figure 4.6 (Stent, 1973; Singer, 1985). If a pre-synaptic node and post-synaptic node are both 'on', the weight on the lateral connection going from the pre-synaptic to the post-synaptic node is increased by a fixed amount. If the pre-synaptic node is 'off' but the post-synaptic node is 'on' then the weight is decreased by a fixed amount. Weight values are limited to the interval 0 to 1.

	Pre	
Post	Off	On
Off		
On	-	+

Figure 4.6: Post-synaptic learning rule.

This form of learning rule causes near or in phase nodes to have their lateral weights increased and out of phase nodes to have their weights decreased. If both nodes are 'off' then it is indeterminate as to whether they are in phase or not and the weights are left unchanged. Using this learning rule a decrement occurs only when the post-synaptic node is active. This occurs for a short and relatively fixed period of time. Because the period of over which a node is determined to be 'off' (not spiking) is significantly longer than when it is 'on', it would require a different decrement to be used in the case where the post-synaptic node is 'off' and the pre-synaptic node is 'on'. For this reason, a post-synaptic only rule is used in these studies.[2] It would be possible to include the pre-synaptic case mentioned above but this would make the implementation more complicated without great benefit as it is shown in section 4.2.2 that the post-synaptic only learning rule works reasonably well.

Unlike static learning rules where the learning rate is usually set according to various statistical properties of the data being presented to a net, the learning rate for the dynamic learning rule discussed here has to be tuned to the number of time steps used to simulate one oscillation. Because there are often many time steps to even just one oscillation,

[2] A purely pre-synaptic rule would also work just as effectively, problems only occur when we attempt to combine elements of both types of rule.

(on the order of 600 or more for the simulation of the Morris-Lecar oscillator used in this study), it is important to scale the weight change accordingly. For example, if you wanted a pair of nodes which were firing in synchrony to develop a connection strength of c within s oscillations and where the period of one spike (the period over which the activation of the node is above threshold) is p time steps, then the learning rate l should be set to $l = \frac{c}{s.p}$.

Choosing an activation threshold

A balance also has to be made between the value chosen for the activation threshold and the value chosen for the learning rate. If the activation threshold for a node is relatively low then weight adaptation for that node's weights will occur frequently and consequently a small learning rate should be used. If however the activation threshold is high then the weight adaptation will behave in a manner more similar to that of weight adaptation in static auto-associative neural networks and a larger increment can be used.

A further point to be considered when choosing a suitable threshold is the type of grouping that will occur when a low or high threshold is chosen. If the threshold is low then a significant proportion of the net will be considered to be in the 'on' state at the same time. This will cause the lateral connections between a large number of nodes which are not necessarily close in phase to be strengthened causing them to become synchronized. If however the threshold is high (near what might be considered as the firing point of the node), then only a few nodes will be considered to be 'on'. This will result in a large number of the connection strengths being reduced and only a select few being increased. This will produce a low tolerance to temporal variation in the signals. Admittedly we do not want the case where all nodes have their lateral weights increased regardless of their phase, however we also do not want the case where the lateral weights of nodes are only increased (and the node subsequently synchronized) when the nodes are already tightly synchronized. Apart from learning a temporal pattern which had been established without synchronization, this would not assist in the synchronization process itself.

The value chosen for the threshold depends upon the particular oscillator model being studied but can be related to a common principle across each model. The level of the threshold effectively defines a period in the phase of a node over which it is considered to be 'on'. The size of the threshold directly affects the width of the period. The wider this period is, the fewer the number of independent phases the adaptive network will be able to support. Given that all the nodes oscillate with roughly the same frequency, a wide response period will produce a small number of independent phases where a node can become synchronized without being affected by nodes in another phase. If p is the period of a node and o is the period over which a node is considered to be on then there will be $v = [\frac{p}{o}]$ independent phases where a node could be linked to only one group of nodes in the same phase. It can be readily seen that if we decrease the size of o by increasing the threshold then we will be able to maintain more phases but at the cost of reduced tolerance to noise. It is also worth noting that we achieve independent phases in an adaptive network by keeping nodes which are to be in a different phase unconnected. It is shown in the following chapter that it is very difficult to maintain independent phases in a synchronizing network when nodes which are supposed to be in a different phase are connected together.

When altering the activation threshold we must be careful to note that the relationship between a change in the value of the activation threshold and that of o is dependent upon the model used and is usually of a non-linear form. An example of this can be seen in figure 4.7 where two examples of shifting the activation threshold are shown. In each case, the amount by which the threshold is shifted is the same. In the second case however, the change in the size of 'o' is significantly greater than the change caused by the first case.

4.2.2 An example of learning in a synchronizing oscillator network

The following section discusses a study performed to illustrate the points covered in the previous sections. The grouping behaviour of an adaptive network is shown with three different activation thresholds and the effects of the thresholds on the types of patterns

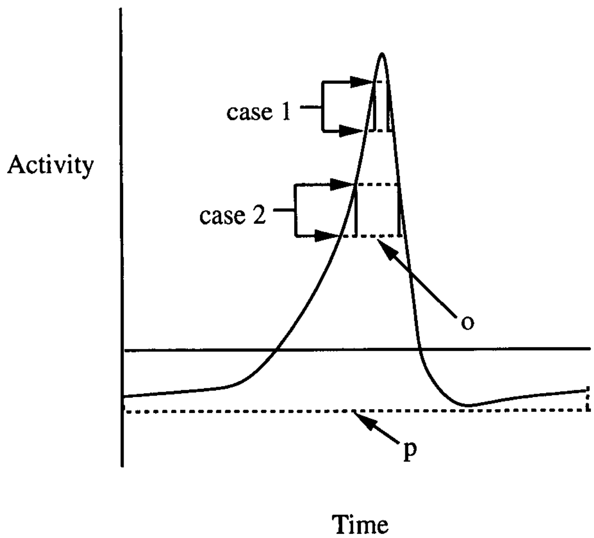

Figure 4.7: An example of the non-linear relationship between changing the activation threshold and the change caused in 'o', the period over which a node is considered to be on.

formed are discussed. The study aims to show that it is both possible and potentially useful to apply learning to synchronizing oscillator networks. A further aim is to illustrate how altering the activation threshold discussed above affects the size of groups formed by a particular learning rule.

An adaptive architecture

A general architecture is used for each of the experiments in this section. The architecture follows the same structure as that for section 4.1 on the split line experiment except for the configuration of the lateral connections. Connectivity is still maintained at three nearest neighbours on either side of a node with wrap-around as before. To facilitate learning, all lateral connections are now weighted with an initial starting value of 0. For each time step, the weights are altered according to the learning rule discussed in the previous section.

Weights on connections transmitting the external input remain fixed at 1 and are non-adaptive throughout the study. Lateral input to a node is now a normalised weighted sum rather than the normalised sum in the previous studies (equation 4.1).

$$L_i(t+1) = \frac{1}{(N-1)} \sum_{j=1, j \neq i}^{N} w_{ij} l_{ij}(t+1) \tag{4.1}$$

N	Number of nodes in the network ($N = 20$)
l_{ij}	Lateral input at synapse from node j to node i
w_{ij}	Lateral weight at synapse from node j to node i
$L_i(t+1)$	Total weighted lateral input for node i at time $t+1$

Method

For this experiment, three trials are performed, the value of the activation threshold being decreased for each successive trial. The three trials are intended to illustrate the behaviour of the net with a high threshold, a medium threshold and a low threshold. All the nodes in the network receive the same level of external input but are started off at different points in their limit cycle. The same initial starting points are used for each trial. The nodes numbered 1..5 and 16..20 have random initial starting points but nodes 6..15 all start relatively close together on the limit cycle. For each trial, only the threshold is altered. The value of the threshold for each trial is 0.2, 0.1 and 0 respectively and each trial is run for 1600 time steps.[3] The output spike threshold α is set to 0.21 with a constant spike size $\gamma = 0.3$.

A learning rate of 0.01 is used for each trial. This produces relatively rapid weight adaptation sufficient to allow tight synchronization to occur within a small number of oscillations. It is perhaps more sensible to use a smaller learning rate (on the order of 0.005 or less) to give less rapid weight change. This would require a number of coherent

[3]Note: An activation threshold value of 0 does not represent the lowest available threshold value. The activation (in this case voltage) of the Morris-Lecar model can go as low as -0.3 in the trough of a recovery period. The value 0 is used in this case because at this limit most connections will eventually be strengthened.

oscillations to take place before a strong connection is established. A higher rate is used in these studies in order to be able to illustrate the net behaviour over a relatively small number of oscillations as a result of constraints on displaying the data. The primary behaviour is however the same for a lower learning rate and this has been verified through experiment.

Results

At the end of each trial, the activation values for the nodes in the network are collated to produce a contour plot of the network activity over time. The following results show the activity contour plot obtained for each trial together with a weight matrix. The activity contour plot shows the activation of each node in the network over time where time goes across the page and the nodes in the network are numbered down the page. The weight matrix shows the size of the weight connections obtained at the end of each trial. A blank space in the weight matrix indicates no connection.

Pre / Post	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1		0.0	0.0	0.0														0.0	0.0	0.0
2	0.1		0.0	1.0	0.0														0.0	0.0
3	0.0	0.0		0.0	0.0	1.0														1.0
4	0.1	1.0	0.0		0.0	0.0	0.0													
5		0.0	0.0	0.0		0.0	0.0	0.0												
6			1.0	0.0	0.0		1.0	1.0	1.0											
7				0.0	0.0	1.0		1.0	1.0	1.0										
8					0.0	1.0	1.0		1.0	1.0	1.0									
9						1.0	1.0	1.0		1.0	1.0	1.0								
10							1.0	1.0	1.0		1.0	1.0	1.0							
11								1.0	1.0	1.0		1.0	1.0	1.0						
12									1.0	1.0	1.0		1.0	1.0	1.0					
13										1.0	1.0	1.0		1.0	1.0	0.0				
14											1.0	1.0	1.0		1.0	0.0	0.0			
15												1.0	1.0	1.0		0.0	0.0	1.0		
16													0.0	0.0	0.0		0.0	0.0	0.0	
17														0.0	0.0	0.0		0.0	0.9	0.0
18	0.0														1.0	0.0	0.0		0.0	1.0
19	0.0	0.0														0.0	1.0	0.0		0.0
20	0.0	0.0	1.0														0.0	1.0	0.0	

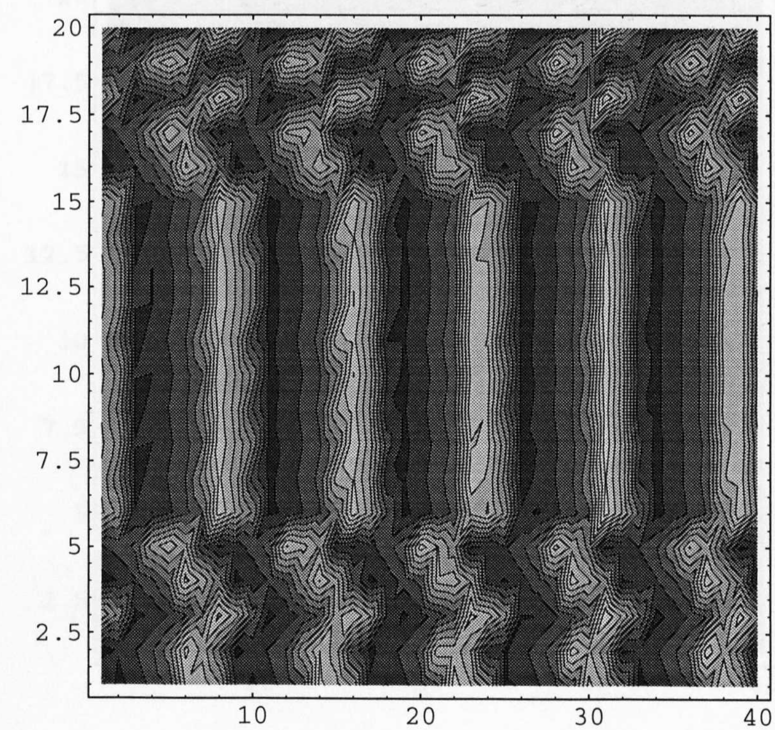

Figure 4.8: Weight matrix and activity contour plot for an activation threshold of 0.20.

Pre / Post	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1		0.3	1.0	0.9														1.0	0.0	1.0
2	0.8		0.2	1.0	0.0														0.0	0.2
3	0.8	0.0		0.1	0.0	0.9														1.0
4	1.0	0.8	0.4		0.0	0.5	0.5													
5		0.4	0.0	0.2		0.0	0.0	0.0												
6			1.0	0.2	0.0		1.0	0.9	0.9											
7				0.4	0.0	1.0		1.0	1.0	1.0										
8					0.0	1.0	1.0		1.0	1.0	1.0									
9						1.0	1.0	1.0		1.0	1.0	1.0								
10							1.0	1.0	1.0		1.0	1.0	1.0							
11								1.0	1.0	1.0		1.0	1.0	1.0						
12									1.0	1.0	1.0		1.0	1.0	1.0					
13										1.0	1.0	1.0		1.0	1.0	0.7				
14											1.0	1.0	1.0		1.0	0.8	0.0			
15												0.9	0.9	1.0		0.5	0.0	1.0		
16													0.9	0.9	0.7		0.0	0.6	0.0	
17														0.0	0.0	0.0		0.0	1.0	0.0
18	0.9														0.9	0.4	0.0		0.0	1.0
19	0.0	0.3														0.1	1.0	0.0		0.0
20	0.8	0.0	1.0														0.0	0.9	0.0	

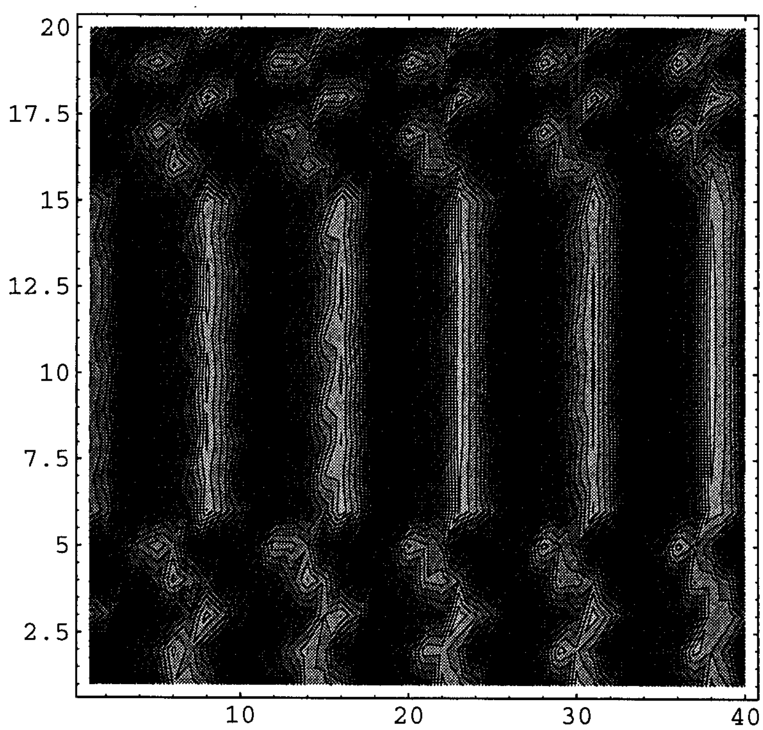

Figure 4.9: Weight matrix and activity contour plot for an activation threshold of 0.10.

Pre/Post	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1		0.9	1.0	1.0														1.0	0.6	1.0
2	1.0		1.0	1.0	1.0														0.9	0.9
3	0.9	0.8		0.9	0.4	1.0														1.0
4	1.0	0.9	1.0		0.7	1.0	1.0													
5		1.0	0.7	0.8		0.6	0.7	0.7												
6			1.0	0.9	0.4		1.0	1.0	0.9											
7				0.9	0.4	1.0		1.0	1.0	1.0										
8					0.4	1.0	1.0		1.0	1.0	0.9									
9						1.0	1.0	1.0		1.0	1.0	1.0								
10							1.0	1.0	1.0		1.0	1.0	1.0							
11								1.0	1.0	1.0		1.0	1.0	1.0						
12									1.0	1.0	1.0		1.0	1.0	1.0					
13										1.0	1.0	1.0		1.0	1.0	1.0				
14											1.0	1.0	1.0		1.0	1.0	0.0			
15												0.9	0.9	0.9		0.9	0.0	1.0		
16													1.0	1.0	1.0		0.0	1.0	0.4	
17														0.2	0.2	0.3		0.2	1.0	0.2
18	0.8														1.0	0.9	0.0		0.2	1.0
19	0.8	1.0														0.7	0.6	0.6		0.6
20	0.8	0.6	0.9														0.0	0.9	0.1	

Figure 4.10: Weight matrix and activity contour plot for an activation threshold of 0.

Discussion

It can be seen from the results that the inner group of nodes (6..15) which start of relatively close in phase rapidly synchronize in all three cases. The main difference between each trial concerns the outer groups of nodes (1..5 and 16..20). With the inner group of nodes the weights on the lateral connections increase as they are all determined to be more or less 'on' at the same time - whatever the threshold is set at. Once a lateral link is established between the inner group of nodes, the synchronization process is able to take place. With tighter synchrony the learning process proceeds more rapidly and the cycle repeats itself until the weights reach the limits set by the learning rule.

With the outer groups of nodes, the determining feature as to whether or not the nodes synchronize with the main group is the activation threshold. With the threshold set at 0.2, the outer groups of nodes are effectively excluded from the inner group (except where they happen to be in phase by chance). Weights between the inner and outer groups are observed to be fairly small or 0. As the threshold is reduced, more of the outer groups of nodes are now classified as being 'on' at the same time as the inner group and are thus brought into closer synchrony through the increase in the lateral weights. Once this convergence process begins, the nodes move closer and closer into phase until eventually they are completely synchronized with the main group and the lateral weights are fixed at 1.

In the case where the threshold is set at 0 all the nodes eventually become synchronized as they are all considered to be 'on' at the same time, despite the wide variation in phase. Even if there is a case where a node is still out of phase with the main body, there is still sufficient room for an intermediate node to have a phase in between the main group and the out of phase node such that as the intermediate node is pulled into synchrony with the main group, it will also pull the out of phase node with it. As can be seen, a threshold set as low as this is of little practical use in an adaptive network unless it is actually desirable that all nodes in the network should eventually synchronize.

In a separate study on the potential for learning in synchronizing oscillator networks, Konig *et al* illustrated how learning could produce behaviour similar in principle to that shown in the results of Gray *et al* and Eckhorn *et al* (Konig & Schillen, 1991). In their network, they used a more advanced two threshold learning model to encode the necessary binding information (Artola *et al.*, 1990). Their results indicated that it was possible to use learning to correctly encode binding information and that their more advanced ABS rule produced an improvement in performance. Their network was able to segregate different oscillatory assemblies more effectively than the Hebb like rule used in this study.

We feel that it is reasonably evident from these results that it is possible to perform learning with synchronizing oscillator networks. The form and use of learning in a synchronizing oscillator network is different to that of static networks. The key difference shown here is the role learning can play in enhancing or defining the regions of synchronization. Learning can be used to define and determine which nodes are to be encouraged to synchronize in future examples and which are not.

4.3 Multiple phase networks

One of the possible explanations for synchronization phenomena put forward by Gray *et al* (Gray *et al.*, 1989a) was that the brain could be using a number of synchronous phases to maintain objects simultaneously. Different objects could be bound through synchronization to a different phase. It would be interesting to see if this principle could be applied to oscillatory neural networks. If it could be shown to work in simulations then it could be used to multiplex information in neural nets, a feature which would significantly increase their power. It would also serve to validate the hypothesis put forward by Gray *et al*, giving further insight into how the brain manipulates a number of different concepts at the same time. The following section covers a number of experiments which were performed to determine the feasibility of this idea within the context of the oscillatory neural nets discussed in this thesis. The experiments investigate the behaviour of a synchronizing oscillator network which attempts to maintain a number of different phases simultaneously using different types of oscillator models.

4.3.1 Architecture

The network architecture used in this study is similar to that of section 4.1. An 18 node network connected in a ring formation is used. Each node in the network is connected to its 3 nearest neighbours on either side. Each node receives a constant input from an external source together with the lateral input from its neighbouring nodes. Three different oscillator models are used in the study, the Morris-Lecar, Wang *et al* and leaky integrator models first discussed in chapter 2.

4.3.2 Method

Three experiments were performed in this study, one for each oscillator model mentioned above. Since the aim of this study was to determine how effectively a network of laterally connected oscillators could maintain different phases, it was decided that for each experiment half the nodes in the network would be started exactly 180° out of phase with the other half. As a common reference point half the nodes were initialised at the point of peak activity in their oscillation and the other half placed at the point which corresponded to 180° away from this. The choice of the peak as one of the starting points is arbitrary, it was merely an identifiable location common to all the oscillator models.

Each node in the network received the same external driving input (tailored for the particular oscillator model under study) and was initialised at one of the two starting points. The first 9 nodes in the network were started at one of the points and the second 9 at the other. The driving inputs and output thresholds for each model are given in figure 4.11.

For each experiment, the network is left to run for a period of time and the activation behaviour over time collected. The period of time each network is left to run for is selected in relation to the oscillatory behaviour of the model used and has no correlation between each experiment. When the network behaviour settles, the experiment is stopped.

Model	**External input**	**Output threshold**
Morris-Lecar	0.3	0.3
Wang *et al*	0.25	0.26
'leaky integrator'	1.0	19.8

Figure 4.11: External input and output threshold parameters used for each model oscillator.

4.3.3 Results

The results obtained for this study are shown in figures 4.12...4.14. For each experiment, they show the network activity over time represented as a contour plot. The method used to collect this data is the same as used for the previous studies in this chapter although the sample rate has been altered to take into account the different oscillatory periods of the Wang *et al* and 'leaky integrator' oscillators.

4.3.4 Discussion

It can be seen from the above results that none of the networks were able to maintain the initial phase states that they were placed in. In the case of the Morris-Lecar and Wang *et al* experiments, both networks eventually converge toward a stable, single phase state. The leaky integrator model did not stabilise at all, small groups of neighbouring nodes synchronized with each other but did not settle.

These results are not conclusive proof that it is impossible to set up a network of laterally connected oscillators which will support more than one phase. They do however indicate that supporting multiple phases is not a robust phenomenon. It may be possible to maintain more phases if certain network parameters are adjusted. Increasing the threshold at which an output spike is generated would reduce the interference between two oscillating nodes unless they were both already close in phase. This in turn would allow nodes which

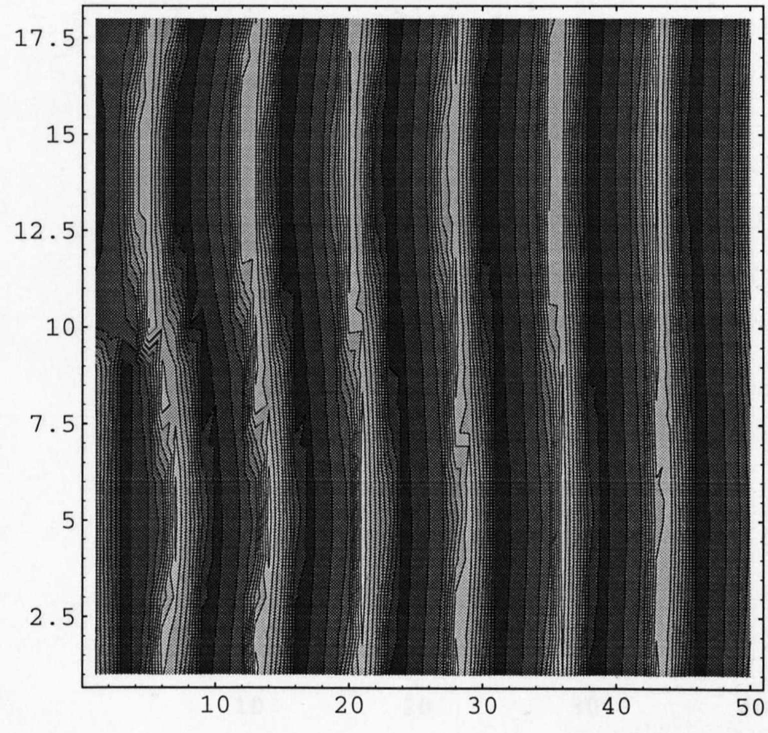

Figure 4.12: Contour plot of activity over time for the Morris-Lecar model.

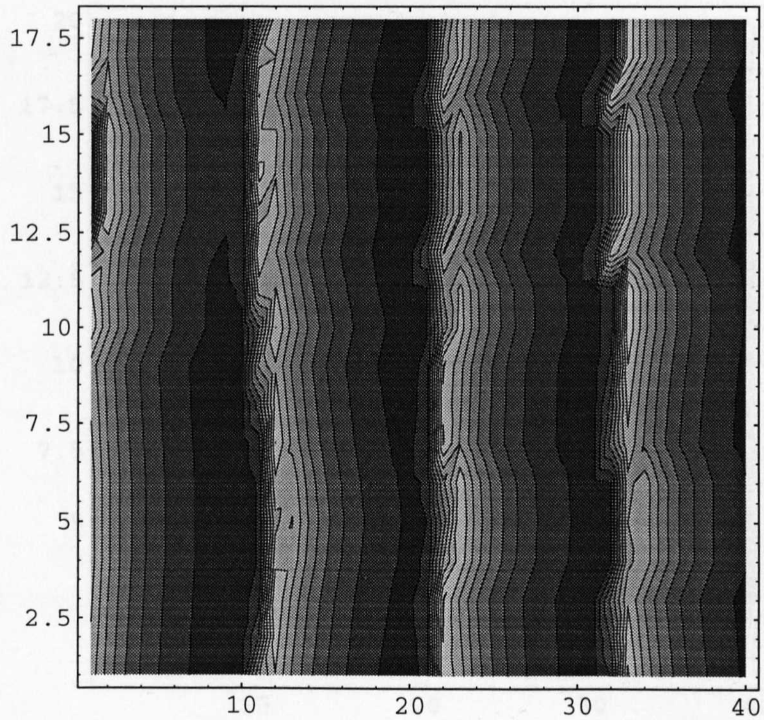

Figure 4.13: Contour plot of activity over time for the Wang *et al* model.

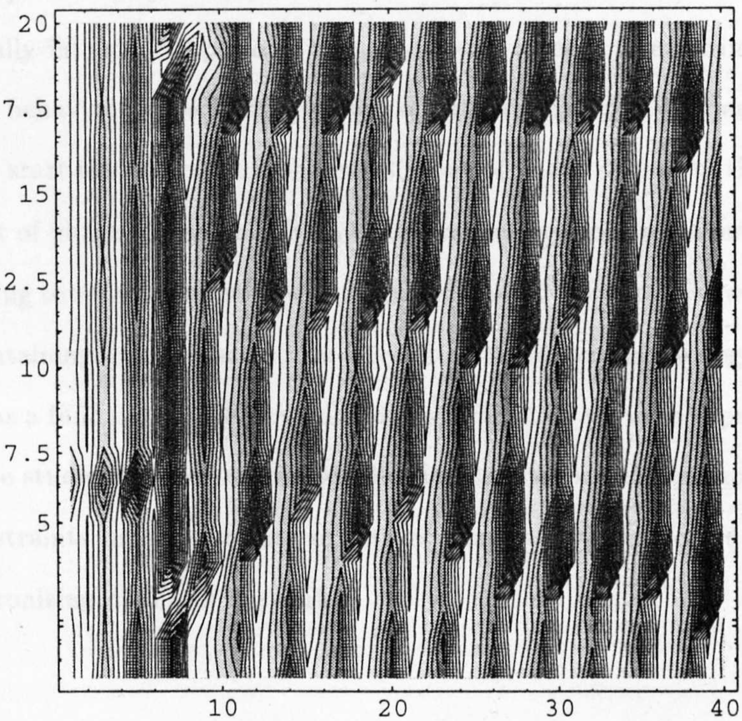

Figure 4.14: Contour plot of activity over time for the 'leaky integrator' model.

were neighbouring but further out of phase to remain unsynchronized, thus allowing for the possibility of more phases.[4]

We attempted adjusting a number of network parameters to improve the performance of the networks but met with little success. It appeared that even when 2 phases were established the network was not very stable. If there were to be any noise in the system, it would be unlikely that the 2 phase state would be maintained. The natural behaviour of the networks was to either converge toward global synchrony or 'wander' about in a continuous attempt to find synchrony.

It was originally intended to perform this study with a series of experiments aimed at investigating the behaviour of oscillating networks with a number of different phases. The first series would start the nets with 2 phases 180^o out of phase as shown, the second with 3 phases 120^o out of phase and so on. Given the above results however, there seemed little point in attempting to set up a network with 3 phases when it appeared we would find great difficulty in maintaining even 2 phases. As a result of these studies we began to question whether there was a fundamental reason behind the difficulties we were experiencing. This led directly to the studies carried out in the following chapter which attempt to determine if there are constraints on the number of phases that can be supported in a laterally connected synchronizing oscillator network.

---

[4]The trade-off in this situation would be a reduction in the synchronization performance of the individual nodes as they would have less effect on each other.

4.4 Summary

This chapter has looked at a number of ways of using synchronization to perform computational tasks. In the first section we showed how temporal synchrony can be used to bind objects and that as the distance between objects increases, a breakdown in coherence between groups will occur. We attributed this behaviour to the connectivity structure of the network and concluded that we can use connective scope to define the size and type of regions formed in a synchronizing oscillator network. Local connectivity will produce small regions of locally synchronized nodes. A more extensive connectivity will produce larger regions of synchronized nodes. However, local connectivity will still allow for large synchronized groups to form, provided there is a continuous link between different areas. The most flexible solution would therefore seem to be the use of local connectivity unless we explicitly need a guarantee of global synchrony.

The second section in this chapter investigated the potential for using learning in synchronizing oscillator networks. We illustrated that it is possible to use learning to encode the binding information for a set of nodes and that this information was capable of improving the synchrony between nodes in a network. We also illustrated how varying the learning threshold in our learning rule affects the size of the regions which are synchronized. A high threshold permits only nodes which are close to synchrony to have their lateral weights increased and therefore cause improved synchronization. A lower threshold causes greater numbers of nodes to have their lateral weights increased. This allows for nodes which are relatively out of phase to become synchronized producing greater global synchrony. Depending on our requirements, we can use a high threshold to produce a network which has a number of synchronized regions but where each region is effectively unconnected and therefore out of phase. Alternatively we can use a lower threshold to produce global synchronization of those nodes which are oscillating against a background of overall noise.

In the last chapter, we discussed a study aimed at investigating the potential for using multiple phases in networks of coupled oscillators to allow simultaneous binding of multiple objects. This study indicated that, at least in principle, systems of coupled oscillators with excitatory lateral connections find difficulty in maintaining distinct phases. Despite attempting to achieve a state where two groups of nodes are exactly 180^o out phase with three different oscillator models, we discovered that none of the models maintained the required behaviour. Although this was not conclusive proof that multiple phases could not be achieved, it leads us to question whether or not there is some limitation affecting the formation of multiple phases in networks of laterally coupled synchronizing oscillators. This question is addressed in the following chapter.

Chapter 5

Constraints on multiple phases

Although the principle of using multiple phases in synchronizing oscillator networks is not unreasonable, the results from the last chapter would seem to indicate that it is not necessarily a stable behaviour for the networks we have investigated. From these results and also from those of other groups, it would appear that there are problems with the number of stable phases a synchronizing oscillator network can sustain (Horn & Usher, 1991). In order to try and determine whether or not there is a constraint on the number of co-existing phases a network can maintain, we decided to study the behaviour of the model oscillators discussed in this thesis through the analysis of phase response graphs. The phase response graph of an oscillator indicates the dependency of the phase shift produced in an oscillator on the arrival time of the entraining signal that causes it. Furthermore, it also enables us to determine the effect of signal amplitude on phase shift, indicating how varying the amplitude of the signal alters the phase shift produced. This information gives an insight into the synchronization behaviour of the oscillator. By studying this behaviour we should gain an insight into whether there is a limit or whether it is possible to maintain multiple phases in our networks if we were to use the correct initial conditions.

5.1 Phase response dynamics

The technique for producing phase response graphs is taken from a study by Rinzel and Ermentrout (Rinzel & Ermentrout, 1989) on the dynamics of theoretical oscillators. However it has also been used to study the dynamics of real neuro-physiological oscillators (Chance *et al.*, 1973; Perkel & Mulloney, 1974). To produce a phase response graph we must determine the relationship between the arrival time of an entraining signal and the subsequent phase shift produced. For example, in figure 5.1a we have a node which is oscillating and producing a regular output spike. At some point an entraining signal arrives from another node. This is integrated into the activation of the oscillator and will therefore affect the timing of the next spike. To calculate the phase shift we must measure the difference between the expected arrival time of the next output spike (as measured by the original frequency) against the actual arrival time (Figure 5.1b). The direction and magnitude of this shift will be dependent upon the particular dynamics of the oscillator under study. In order to obtain a complete phase response graph of this behaviour we can calculate the phase shifts produced by entraining inputs arriving at a range of points in the phase of an oscillation.

Before examining the results produced from our study, it might prove useful to ask what kind of phase response graph we would expect for an idealized oscillator. We are interested in two possible behaviours for our oscillator: how it should operate in a single phase system and how it should operate in a multiple phase system. Furthermore, are these two behaviours complimentary or would they be exclusive. For the case of the single phase system, we would like the phase of the oscillator to converge toward the phase of the originator of the entraining signal. The direction of the phase shift would be dependent upon how far the phase of the receiving node is displaced from the sender. If the signal arrives in the first half of phase space then it would be better to have a negative phase shift (the receiving node will fire later than it otherwise would have done). On the other hand, if the signal arrives in the second half of phase space then the phase shift should

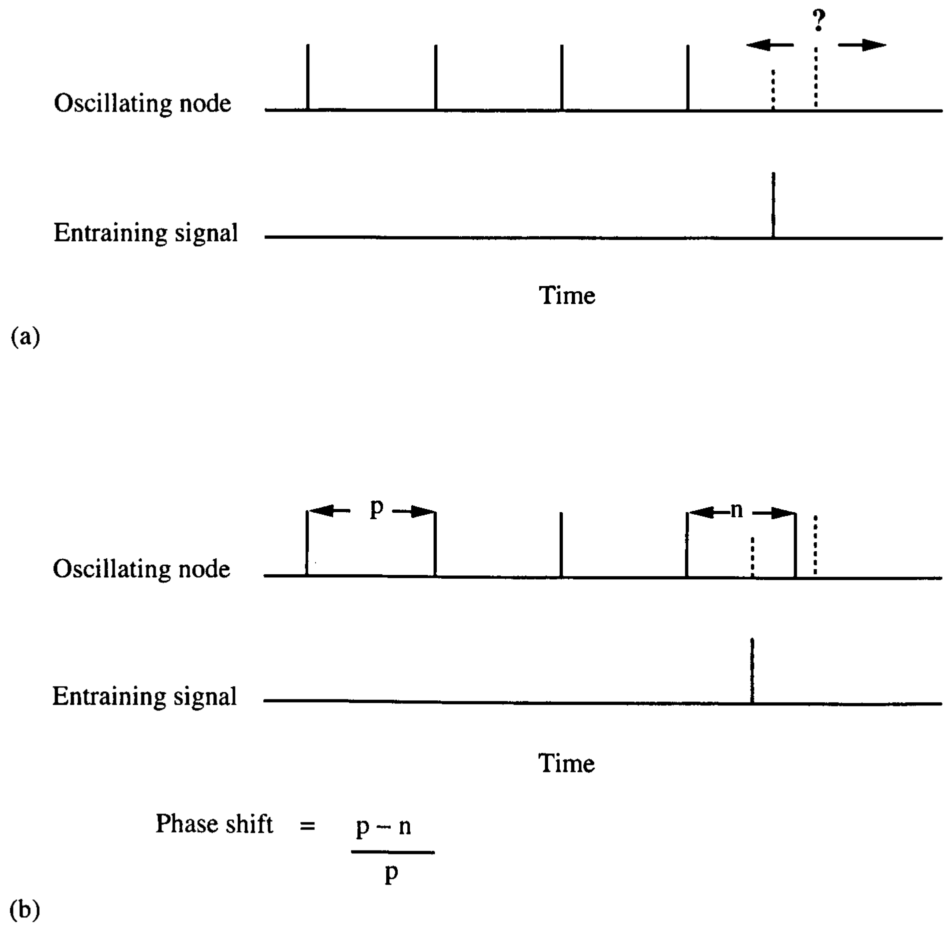

Figure 5.1: **(a)** Phase shift **(b)** Calculation of phase shift.

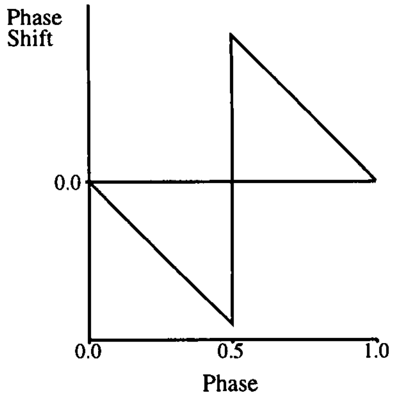

Figure 5.2: Idealized phase response for an oscillator in a single phase system.

be positive, causing the node to fire sooner than it would have done and therefore moving nearer to the firing point of the the originator of the signal. The magnitude of the phase shift should either be sufficient to bring the receiving node exactly into phase with the sending node or bring it closer (less than or equal to the phase difference between the two nodes). If the phase shift produced is greater than the phase difference between the two nodes then the receiver will overshoot. An illustration of one possible solution to these requirements is given in figure 5.2.

We should now consider the case for multiple phase systems. Oscillators operating in a multiple phase system should still possess synchronization behaviour within a small window around their peak activation. This is to allow them to synchronize with nodes oscillating around the same phase.[1] In a phase response graph, this corresponds to two areas of convergence around the peak activation of the node. In between these two parts should be an area of insensitivity, in effect a 'dead' zone where entraining inputs arriving from other nodes have no effect on the phase state of the receiving node. This behaviour would allow the node to remain unperturbed by inputs arriving from nodes which were

[1] The size of this window will be inversely proportional to the number of phases the system can maintain.

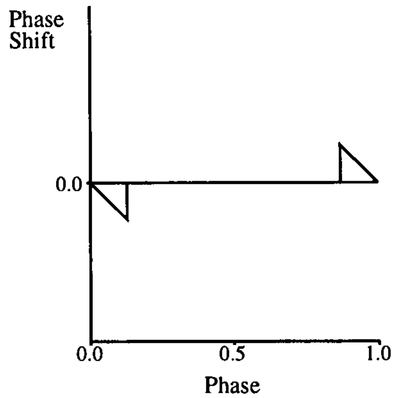

Figure 5.3: Idealized phase response for an oscillator in a multiple phase system.

not nearly in phase with it and therefore allow it to remain unaffected by other groups which are in different phases. Combining these different requirements we can arrive at an idealized phase response graph for a multiple phase system as shown in figure 5.3.

5.2 A study on phase response graphs

Having investigated the foundations necessary for maintaining single or multiple phases in networks of synchronizing oscillators, we can now investigate the behaviour of the oscillators used throughout this thesis. The following section discusses a study performed to obtain a phase response graph for each of these oscillators and is followed by a discussion on the results.

5.2.1 Method

To obtain the phase response graphs, each oscillator is driven by a constant input until the period of the oscillation has stabilized. This gives a base period λ_b. Driving the oscillator by the constant input and starting from a point just after peak activation, trials are made for points across the phase of the oscillation. For each successive trial, the point of delivery of an entraining signal to the oscillator is increased by $\frac{2\pi}{1000}$. Each entraining input is of constant size and is delivered for a constant proportion (0.025) of the period of the oscillator. The entraining input causes a change in the period of the oscillator resulting in a new period for the oscillation λ_n. A measure of the relative phase shift caused by the entraining input is then calculated using equation 5.1. The results of performing a phase response analysis for each of the oscillators is shown in figure 5.4. The parameters used to obtain these graphs are the same as those given in appendix A.2.

$$\Delta\theta = \frac{\lambda_b - \lambda_n}{\lambda_b} \tag{5.1}$$

$\Delta\theta$ Phase shift λ_b Normal period λ_n New period

5.2.2 Discussion

The phase response graph for the 'leaky integrator' reflects the capacitance charging effect inherent within it. External inputs arriving within the first third of the phase of the oscillator have relatively little effect upon it. This is because they are quickly absorbed into the activation of the node. External inputs arriving later on, near the point where the activation has reached its asymptote, cause the activation to suddenly 'trip' over the threshold (rather than slowly approach it) and cause the node to fire earlier than it otherwise would have done. For this reason, large phase shifts can be produced but always in a positive direction.

Unlike the phase response graph for the 'leaky integrator', the phase response graph of the Morris-Lecar oscillator contains regions of both positive and negative phase shift. Inputs arriving in the first third of the phase space cause small negative shifts. This occurs because the inputs cause the activation of the oscillator to remain high longer than normal thus causing an increase in the period of the oscillation. After this point there is a small area of insensitivity corresponding to the refractory period of the model oscillator. There is a further area of negative phase shift after this point, again corresponding to an area

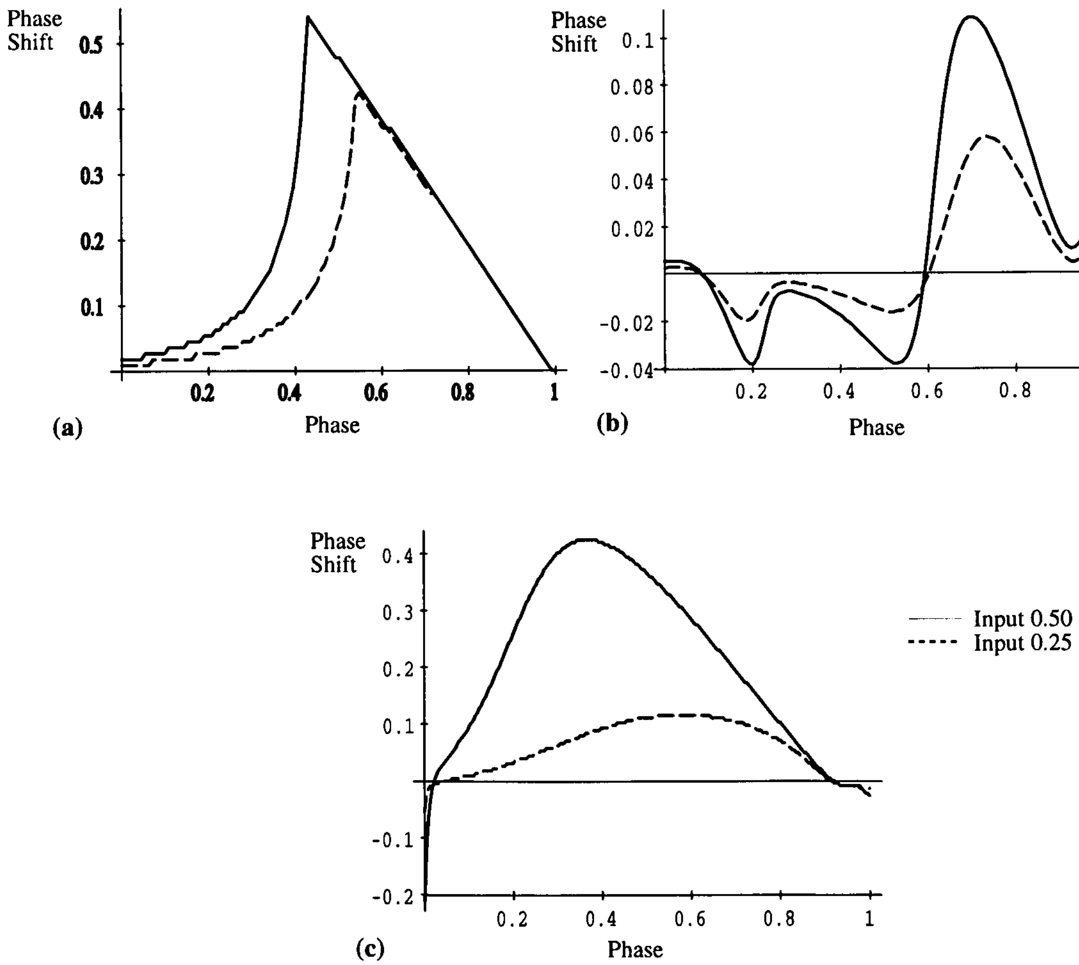

Figure 5.4: Phase response graphs.

(a) Leaky integrator model. (b) Morris-Lecar cell membrane model. (c) Wang *et al* cell cluster model. Each graph shows the change in phase which occurs when an oscillator is perturbed at a given point in its phase. The x-axis gives the phase of the oscillator at the point it is perturbed and the y-axis the degree of perturbation in terms of phase shift. The amount of entraining input I by which the oscillator is stimulated by is given as a fraction of the driving input of the oscillator.

where the activation is raised by the input but succeeds only in increasing the period of the oscillation. At some point the entraining input moves over a threshold and instead of prolonging the period of the oscillation, causes it to fire. This corresponds to the sudden shift in the phase response graph from a negative to a positive phase shift. Once the dynamics go beyond the point where a positive phase shift occurs, it merely becomes a question of the degree of phase shift which is induced.

Despite the complexity of the Wang *et al* model, the phase response graph yields rather poor results. Its route to synchrony is not dissimilar to that of the 'leaky integrator' model. Rather than lengthening its period when receiving entraining inputs in the first half of its phase, it continues to move forward. A consequence of this behaviour (which also holds for the 'leaky integrator' model) is that if an oscillating node of this form is converging in phase with another node and suddenly overshoots (due to noise) then it will take longer to synchronize than if negative phase shifts occurred on overshoot.

5.2.3 Single or multiple phase systems?

In the view of these results, the following section discusses the suitability of our model oscillators to a single or multiple phase system. In the case of a single phase system, it can be seen that the model oscillators studied all show at least the basic requirements for operation in a single phase environment. Although the phase response graphs produced by our study do not closely match the idealized phase response graph, they possess sufficient similarities to indicate that they will exhibit synchronization behaviour. All three models have an almost linearly decreasing phase shift from around phase 0.5 indicating that, whilst in this region, they will converge and synchronize. The major discrepancy between each of the oscillators is the degree of phase shift produced in the first half of their phase. Only the Morris-Lecar model produces a negative phase shift. The positive phase shift produced by the other two oscillators results in less efficient behaviour with the oscillators taking longer to synchronize and being more susceptible to noise and overshoot. This phenomena perhaps offers some explanation for the results observed in Chapter 3. It was observed in this chapter that both the 'leaky integrator' and Wang *et al* models did not synchronize very reliably. They exhibited weak synchronization behaviour and did not cope very well with noise. The Morris-Lecar oscillator on the other hand synchronized very effectively, even under quite high levels of noise. We think that one of the possible causes of this difference is the negative phase shift present in the Morris-Lecar phase response dynamics. This allows it to quickly correct any overshoot that might occur due to effects such as noise or excessively large lateral inputs. Lacking in this feature, the 'leaky integrator' and Wang *et al* oscillator are unable to move efficiently back into phase when they overshoot and therefore take longer to synchronize and become more susceptible to noise.

In the case of a multiple phase system, it can be seen that none of the oscillators show the type of phase response required. The only oscillator that even approaches the type of behaviour required is the Morris-Lecar model which has a small 'dead' zone around phase 0.3.[2] The other two oscillators are clearly totally unsuited for operation in a multiple phase system. It becomes clear from this analysis that an oscillator which is suitable for operation in a single phase system will be unsuitable for operation in a multiple phase system due to the conflicting requirements of each system.

Given the above results, we are forced to conclude that systems constructed using the given architecture (direct lateral coupling) and using similar oscillators will not be able to support very many phases. If we allow ourselves to draw an analogy between this conclusion and theories on neural computation then it would seem that these theories will be bounded by a very small number of phases. We are therefore left with two possibilities. Either we allow ourselves to make changes to our system which enable multiple phases to be maintained or we forego the principle of computation using multiple phases and examine where that leaves us.

[2]This has made itself evident in simulations where 3 phases have been maintained with difficulty.

5.3 Input modulation

As we have seen, it would seem unlikely that the current system of oscillators which uses direct lateral coupling, will be capable of sustaining multiple phases. One method by which it is possible to produce the desired phase response behaviour of an oscillator for a multiple phase system is to use input modulation. The phase response graph of an oscillator for a multiple phase system is essentially that for a single phase system but with the middle section removed. It is possible to create this effect in the dynamics of an oscillator normally used in a single phase system by making a distinction between the way in which the driving input and entraining input are integrated into the activation of an oscillator. If the entraining input is filtered by modulating the amount that gets through by the current activation level, such that activation levels corresponding to the middle phase of an oscillation cause little entraining input to get through, then a multiple phase response behaviour is produced. A modulator can be placed in between a node and the lateral inputs arriving at the node. If the node is highly active then the modulator lets most of the lateral input through. If on the other hand the node is only weakly active then little or no lateral input gets through (figure 5.5).

For this system to work well however, the underlying dynamics of the oscillator used should give a phase response graph similar to that of our ideal single phase oscillator 5.2. Positive phase shifts should occur before spiking and negative phase shifts just after spiking. If the oscillator lacks the negative phase shift then it will find difficulty in synchronizing even to neighbours that are close in phase as it will frequently overshoot and be unable to quickly move back. Given these conditions, it is clear that neither the 'leaky integrator' nor the Wang *et al* oscillator are suited to this task.

However, we can illustrate this behaviour with the Morris-Lecar oscillator as it possesses most of the phase response requirements we need. To implement this feature, we require a function which lets through a proportion of the incoming lateral signal depending upon the current activity of the receiving node. As the activity of the node increases, an increasingly

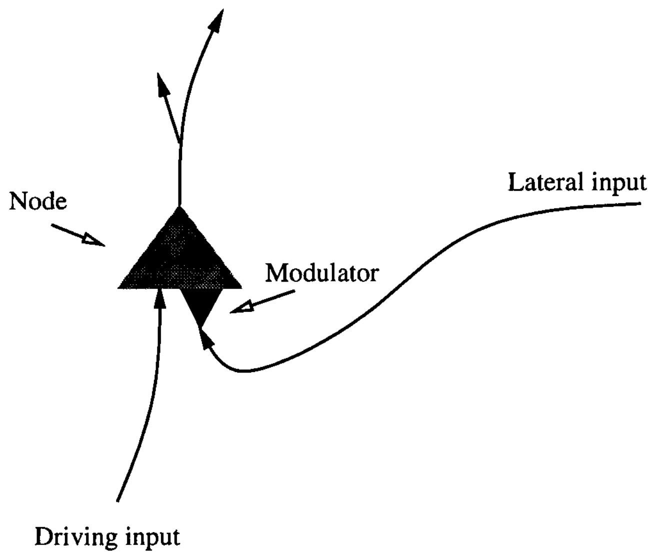

Figure 5.5: Input modulation : The modulator filters the amount of lateral input arriving at the node dependent upon the current activity of the node.

larger amount of the lateral input should get through. We can use a standard sigmoid function to give us this behaviour, tailored to respond correctly over the desired range of activation values. An example which gives the correct behaviour for the Morris-Lecar oscillator is given in equation 5.2. A graph of its behaviour is shown in figure 5.6. Once we have calculated the lateral effect, we can then compute the total lateral input using equation 5.3.

$$f_{mod}(l_{ij}, a_i) = \frac{l_{ij}}{\omega + e^{\mu(\iota a_i - \nu)}} \tag{5.2}$$

$$L_i(t+1) = \frac{1}{N-1} \sum_{j=1}^{N} f_{mod}(l_{ij}(t+1), a_i(t)) \tag{5.3}$$

where ω=0.3, μ=-20, ι=4 and ν=0.5. If we modulate the lateral input coming into a node using equations 5.2 and 5.3 then we obtain a new phase response graph which possesses some of the desired properties for our multi-phase system. Repeating the study

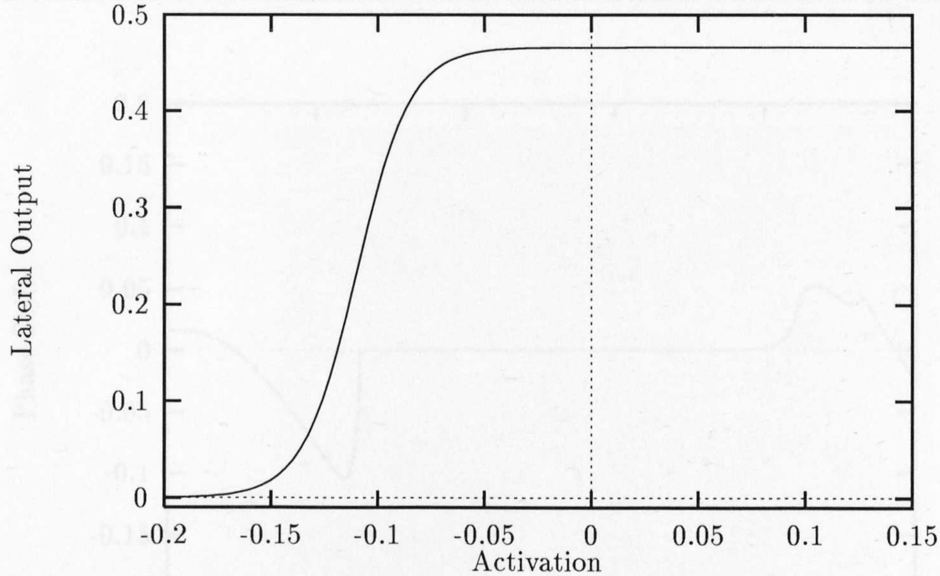

(a)

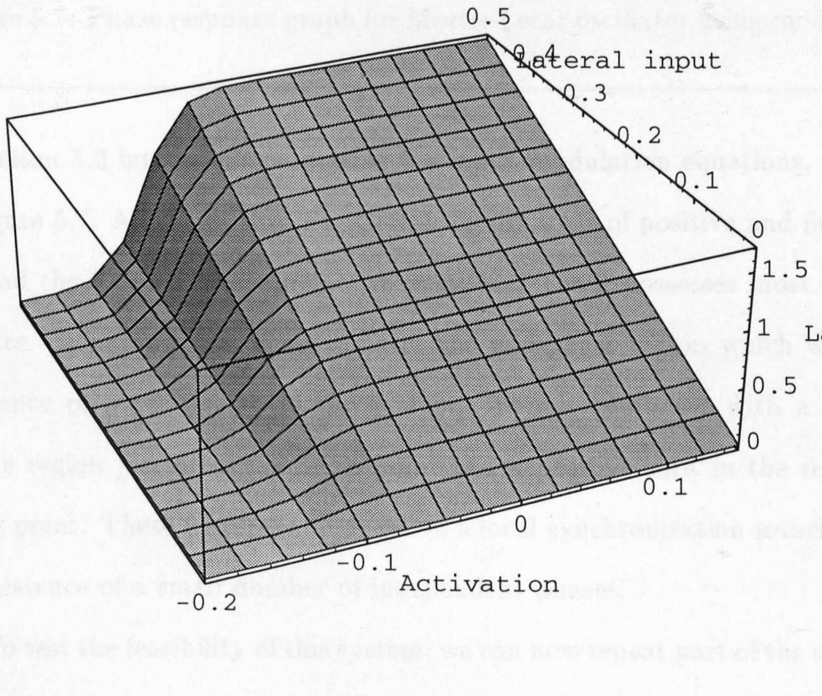

(b)

Figure 5.6: Modulation of lateral input : **(a)** Modulation of a constant lateral input (l_{ij}=0.14). **(b)** Modulation effect over a range of lateral inputs. The x-axis shows the activation of a node, the y-axis the unmodulated lateral input arriving at a node and the z-axis, the modulated lateral output produced.

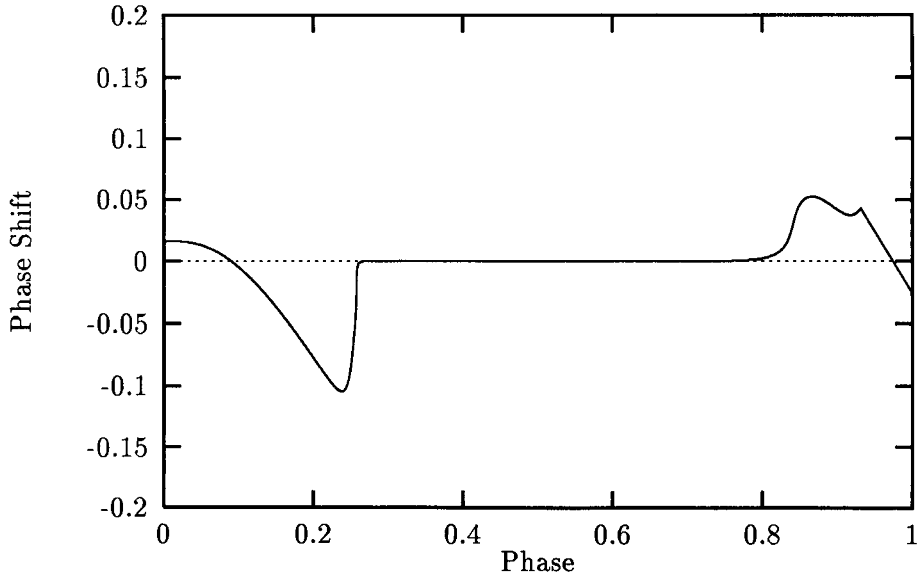

Figure 5.7: Phase response graph for Morris-Lecar oscillator using modulated lateral input.

of section 5.2 but in this case using the input modulation equations, we obtain the graph of figure 5.7. Although this is not ideal due to areas of positive and negative shift directly around the 0 and 1 phase points respectively, it still possesses most of the properties we require. It has a large 'dead' zone in the mid-phase region which will allow for the co-existence of other phases in the system. This is combined with a small negative shift in the region just after the firing point and a positive shift in the region just before the firing point. These features will preserve a local synchronization principle yet allow for the co-existence of a small number of independent phases.

To test the feasibility of this system, we can now repeat part of the study on maintaining multiple phases in oscillator networks (Chapter 4, Section 4.3) using the Morris-Lecar oscillator and the new equations. The method and architecture for this study are identical to that of section 4.3 except for the new equations. Half the nodes in the network are started exactly 180^o out of phase with the other half and the network is left to run for a

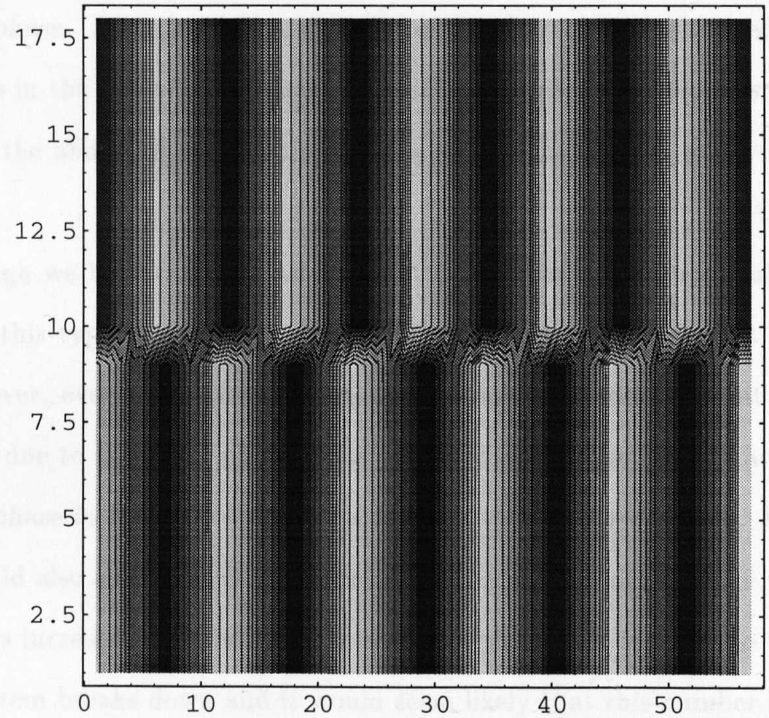

Figure 5.8: Contour plot of activity over time for the Morris-Lecar oscillator. The x-axis indicates time and the y-axis indicates the number of a node.

period of time. The results of running the study are shown in figure 5.8. It can be seen that this time the network easily maintains the two phases in which it is initially placed. Local synchronization behaviour keeps the two groups locked to their given local phase and no interference occurs between the two groups. The input modulation effectively disconnects the two groups provided they remain sufficiently out of phase with each other. This study was also repeated using three distinct phase locked groups but the network was unable to maintain this state. Although it began in 3 distinct phases, the middle group soon split up with half the nodes moving toward one of the synchronized groups and the other half moving toward the second group. The two groups then maintain a stable 2 phase state. This behaviour is primarily due to the width of the 'dead' zone being insufficient to allow for an extra phase. In order to maintain an extra phase, we would need to reduce the response areas in the modulation equations. Although this would allow us more phases, it would reduce the ability of synchronized nodes to remain coherent under the influence of noise.

Even though we have demonstrated that it is possible to produce multiple phase behaviour with this type of system, it would be ambitious to suggest that it is present in cortex. However, even if this is were true, such a system would exhibit relatively slow phase locking due to the small phase shifts produced and nodes would have to be reasonably close in phase in order for them to phase lock to the correct group. From the above results, it would also appear that a further restriction will be found as the requirement for extra phases is increased. There are only so many phases we can fit into the 'dead' zone before the system breaks down and it would seem likely that this number is fairly low (in the order of 3 or 4).

5.4 Synchronized activation

If we take the position that it is unlikely that a computational process based upon using multiple phase labels is present in cortex, where does this leave us? We still have the neurophysiological results of Gray *et al* and Eckhorn *et al* which support the case for some form of synchronization behaviour occurring in the brain even if the exact interpretation of these results is under debate (Young *et al.*, 1992). A possible alternative interpretation which would still be in line with the results found here is computation using synchronization as a process itself (Crick & Koch, 1990; Sompolinsky *et al.*, 1990; Koch & Schuster, 1992).

So far in this thesis the lateral connections have been treated as a means of achieving coherent firing. Usually in neural networks these connections have another role, that of capturing a model of previous experience. This can either be in terms of storing a limited number of memories as in an associative network, or alternatively in terms of a model to perform statistical inference. Both Hinton (Hinton & Sejnowski, 1986) and Mackay (MacKay, 1991) show that with a particular choice of lateral weights (based on previous experience) the network will find a solution of maximum likelihood to the input or at least a local maximum of this function. This provides a different framework for understanding the role of synchronous firing.

Coherent activity can be treated not as an all or nothing phenomena but as a continuous one. The stronger the connections between units, the more tightly phase locked they will be. If the weights between units were chosen so that the system of spiking neurons was performing statistical inference then the degree of synchronization for a particular state is a simple and robust function of the likelihood of the state given the model inherent in the weights (Baddeley, 1992). This gives a simple and alternative interpretation of the coherence phenomena. Synchronization is used to convey the likelihood of a solution to other networks in a system. Relating back to the studies by Gray *et al* for example, a single straight line will be a likely solution based on previous experience and will thus be labeled via a highly synchronous activity. A collection of randomly oriented lines will be unlikely and will be labeled by incoherent activity.

5.5 Summary

It can be seen from our results that the oscillators in the study are best suited for operation in a single phase system. Indeed, given the requirements of an oscillator for a single phase system, it is reasonably true to say that most systems which contain coupled oscillators will exhibit some synchronization behaviour. It is clear however that the oscillators studied are not well suited for operation in a multiple phase system. The requirements for an oscillator to operate in a multiple phase environment are almost contradictory to one which is to operate in a single phase system. These conclusions imply that, given a phase limit in the region of only 2-3 phases, models of neural computation which use synchronized oscillations and multiple phases to explain binding and local memory are likely to be impaired by this limit.

It is possible to circumvent some of these problems by imposing architectural constraints such as input modulation. These however are not necessarily neuro-physiologically feasible. There are other interpretations we can make of the evidence for synchronization. It is not necessary for a system to be able to support multiple phases before it can perform useful computation. Synchronization alone can enable a system to perform figure/ground segregation or other forms of attention definition. Alternatively we can use synchronization in a different role as a statistical measure of the likelihood of a solution. To conclude, coherency in neural systems is not necessarily an epiphenomenon. Even though it is unlikely that multiple phases are present, coherent oscillations on their own can still have a useful computational function.

Chapter 6

Conclusion

This thesis has been concerned with the function and application of synchronizing oscillator networks. The role of lateral connectivity as a means to achieving synchrony has been studied and a number of examples of its use have been illustrated. The following chapter reviews the conclusions which have been drawn from the studies within this thesis and discusses where this work might be taken further.

6.1 Summary of Results

6.1.1 The evidence for synchronization

In Chapter 2, the neurophysiological evidence which led to the work in this thesis was reviewed. As a result of the studies of Gray *et al* and Eckhorn *et al*, it was concluded that there was sufficient evidence to show that synchronization was present in cat visual cortex. Furthermore, it was also concluded that a possible mechanism for this phenomena was the extensive lateral connections present in cortex. A network of laterally connected oscillators was defined on this basis and its basic synchronization phenomena illustrated.

6.1.2 Factors affecting synchronization

In Chapter 3, we investigated a number of factors which affect the performance of synchronization in laterally coupled networks of oscillators. These studies concentrated primarily on the effects of altering factors affecting the lateral connections since they are the principal means of synchronization in the networks used in this thesis. The first study of this chapter investigated the role of lateral connection strength on synchronization performance for the three oscillator models used throughout this thesis. The main finding from the study was that increasing connection strength has a tendency to improve synchronization performance. The exact effect observed was however dependent on the oscillator model used. For the leaky-integrator and Morris-Lecar models, synchronization performance improved until it reached a point where increases in connection strength produced no further improvement. For the Wang *et al* model there appeared to be an optimal level of connection strength beyond which further increases in lateral connection strength actually cause a deterioration in performance.

From the results of this study, it was concluded that the main effect of altering the size of the lateral connection strength was to alter the phase shift produced by the lateral inputs. With small lateral inputs, small phase shifts are induced between coupled oscillators causing only limited movement in network state. When network connection strength is increased, a greater phase shift is induced between nodes. This causes an increase in the fluctuation of the network state and although it becomes more subject to overshoot, it also enhances the probability of the network reaching synchrony. Eventually a level of phase shift is reached where it is not possible for the network to reach synchrony any faster. At this point further increases in network connection strength will have one of two possible effects. A network will either remain in synchrony, absorbing the increased input, or it will become desynchronized. The increased connection strength will cause it to move around too rapidly for a node to compensate and will therefore cause overshoot to occur.

In the second study of Chapter 3, the effect of altering connection scope on network coherence was investigated. It was concluded from this study that total connectivity is neither required nor optimal for achieving synchrony in networks of laterally connected oscillators. The exact results produced were again dependent upon the oscillator model used. For the leaky-integrator model, a lateral interconnectivity of around 0.6 gave the best results. With the Morris-Lecar model, a connective scope of 0.1 (nearest neighbour) gave very good performance with an optimal result around 0.2 to 0.3. The Wang *et al* model gave best results with a connective scope of 0.9 but showed a deterioration after this point.

One of the main conclusions drawn from this study concerned the behaviour observed for the Morris-Lecar model. For this oscillator model, it was observed that low levels of connective scope were sufficient to enable high network coherence to occur. Where an oscillator model is suitable for operation in a synchronizing oscillator network, it should be capable of operating with low levels of connectivity. This feature is important in two respects. From a simulation point of view it significantly reduces the processing load required to model a system. From a neurophysiological viewpoint it is critical that this behaviour is possible. Given the three dimensional nature of cortex it would be impossible for distant cortical areas to be completely connected. If synchronization is to take place, it would have to rely upon indirect connections between these areas. The results of this study showed that, at least in principle, this is possible.

The third study of Chapter 3 concerned the relationship between connection scope and network synchronization speed. The results of this study were related to and in agreement with the results of above. The optimal connection levels observed in the previous study also tended to produce the fastest synchronization times. It was however difficult to come to a general conclusion from these results as the effect produced was again dependent upon the oscillator model used. For the leaky-integrator model, the network gave the best performance with a connectivity level of 0.6. For the Morris-Lecar model, it was again observed that even only with nearest neighbour connectivity, the network synchronized rapidly. Increasing connective scope for the Wang *et al* network produced only small fluctuations in the network synchronization performance but on average the network failed to synchronize.

The final study of Chapter 3 dealt with the effects of noise on synchronization performance in networks of laterally coupled oscillators. Two forms of noise were studied: the first concerning noise in the driving inputs and the second with noise in the lateral inputs. Noise in the driving inputs to the oscillators caused a general breakdown in synchronization performance. This resulted from a failure of the nodes in the network to oscillate when the noise caused the driving input to become too variable. Both the leaky-integrator and Wang *et al* models require a consistent and strong input in order to oscillate. The Morris-Lecar network fared slightly better, managing to withstand a greater degree of noise but its performance also eventually deteriorated.

Noise in the lateral inputs of the different networks caused one of two effects. For the leaky-integrator and Wang *et al* models which have consistently shown poor synchronization performance, noise in the lateral inputs actually improved network coherence. It was concluded that this effect was due to the noise causing greater movement in the network state (in phase space) than would normally occur. This caused the networks to move through areas of state space that they would not normally enter and increased the likelihood that they would placed in a state closer to synchrony. For the Morris-Lecar model, there was little or no room for improvement. The only possible effect this type of noise could cause was a deterioration in network coherence. Given this, the Morris-Lecar network was still able to withstand a high degree of noise before deterioration set in. It was concluded from this observation that if an analogy were drawn between the Morris-Lecar model and cortex, then it should be possible for neurophysiological oscillators to withstand a considerable amount of noise within the systems in which they operate. Furthermore, it was concluded that given this robustness, one of the functions of these systems could be to counteract the effects of noise in signal transmission between different cortical areas.

When noise in the driving input and lateral input was combined, it was observed that the effects of noise in the driving input tended to dominate the results. Noise in the driving input caused both a drop in the ability of an individual node to oscillate and a difference in the frequency between individual nodes. Without consistent underlying oscillatory behaviour network coherence consequently deteriorated.

6.1.3 Computation using synchronizing networks

In Chapter 4 we considered ways in which networks of synchronizing oscillators can be used to perform computational tasks. In the first study of this chapter, it was shown how temporal synchrony can be used to encode binding information. It was also illustrated how the type of synchronization which occurred was dependent upon the distance between the objects being presented to the network. We concluded from these results that the type of binding which occurs is dependent upon the network inter-connectivity used. Local inter-connectivity (as used in this study) will produce small regions of synchronized nodes with synchrony breaking down where there is a gap between nodes greater than the connective scope. Longer range inter-connectivity will produce large synchronized groups, even where small gaps are present. The most optimal solution would appear to be local inter-connectivity. It would still be possible to create large synchronized groups provided the individual groups making up the input were not too dispersed but it would also allow for smaller groups to form when required.

The second study in Chapter 4 concerned the potential for using learning in synchronizing oscillator networks. By placing adaptive synapses on the lateral connections, we showed that it was possible to use learning to encode binding information. Nodes which were nearly in phase had their lateral weights increased, inducing them to move further into synchrony. Nodes which were out of synchrony had their lateral weights decreased toward a limit of 0, discouraging further synchronization behaviour.

An activation threshold was also investigated in this study. Its function was to determine whether a node should be considered to be 'on' or not and therefore participate in the learning process. It was shown that it is possible to use this threshold to determine the type of binding which occurs. If two nodes are out of synchrony and a high threshold is used then no learning takes place. If a lower threshold is used, then both nodes are more likely to be considered as being 'on' together and the weights between them will be increased. Depending on requirements, we can use the threshold to set the type of binding information encoded. If we want a number of different, unsynchronized regions to form then we will need quite a high threshold whereas if we wish to detect a group of nodes oscillating against a background of overall noise then we can use a lower threshold.

The final study in Chapter 4 concerned the issue of using multiple phases in networks of synchronizing oscillators. It was the intention of this study to attempt to initialise the model networks in such a way that they would maintain independent and stable phases simultaneously. We found that under the conditions used, the networks all failed to maintain their set pattern. The Morris-Lecar and Wang *et al* models both moved into synchrony. The leaky-integrator model moved around in state space attempting to achieve synchrony but failing to form a stable behaviour. We concluded from these results that, while the results obtained were not conclusive proof against synchronizing oscillator networks being able to maintain multiple phases, they were at least indicative that it was not a standard behaviour. This led us to question the principle of using multiple phases in networks of synchronizing oscillators and resulted in the studies of Chapter 5.

6.1.4 Constraints on the number of stable phases

From the results of Chapter 4, it was concluded that the laterally connected architecture used throughout this thesis did not easily support multiple phases. This result was not however certain and in Chapter 5 we investigated this principle further by performing a phase response analysis of the oscillators used in this thesis. Phase response graphs for each of the oscillator models used in the thesis were generated and their behaviour analysed. It was concluded from the phase response analysis that none of the oscillators studied were suitable for operation in a multiple phase synchronizing oscillator network. Furthermore, where an oscillator is likely to be suitable for operation in a single phase system, it would be unsuitable for operation in a multiple phase system. An alternative architecture was suggested which would enable multiple phases to be maintained. The new architecture relied on an input modulator to distinguish between the lateral and driving input. Depending upon the current activation state of a node, the input modulator would allow a degree of the lateral input through to a node. A phase response graph for the Morris-Lecar oscillator combined with an input modulator was generated and shown to have the desired response properties for operation in a multiple phase network. Using this configuration, the last study of Chapter 4 was repeated for the Morris-Lecar network and it was shown that the network was now able to maintain two stable phases. We concluded from these results that although it was possible to generate multiple phases using input modulation, the neurophysiological feasibility of this method was not certain. Furthermore, even if it were possible, it would only allow for a small number of phases and would cause a reduction in synchronization capability. From this we concluded that even though multiple phases in networks of synchronizing oscillators are difficult to maintain, synchronization as a process in itself has a useful role to play in real or artificial neural systems.

6.2 Further work

There are a number of areas where the work performed in this thesis could be extended. This section reviews some of these possible extensions and indicates where they might enhance network performance. In Chapter 2 we compared two possible architectures which were capable of exhibiting synchronization behaviour. The use of the laterally connected architecture was selected over the phase comparator architecture primarily on the grounds that it was a more flexible system and was closer to the neurophysiology of cortex. However in their comparison of the phase comparator versus laterally connectivity , Kammen *et al* concluded that the phase comparator architecture was more effective at maintaining and achieving synchrony (Kammen *et al.*, 1990). Since there is evidence for both feedback and lateral connections in the neurophysiology of cortex, it might prove useful to attempt to combine both of these systems. The combined system might be able to exhibit improved performance by using the feedback connections yet be capable of encoding binding information in the synapses of the lateral connections of the system.

It has been previously noted that signal transmission delays have not been implemented in the architectures used in this thesis. This was mainly due to issues of added complexity and simulation overhead. It would be useful to investigate this area. In a separate study not covered in this thesis, we have shown that combining signal transmission delays with inhibition can have rather counter-intuitive results (Smith *et al.*, 1993). Depending upon the exact transmission delay used, synchronization or desynchronization will take place. This was only a preliminary investigation and there are many issues that could be studied within this area which are likely to show useful results.

The investigation into learning in Chapter 4 of this thesis concentrated mainly on illustrating the function and feasibility of this principle. As a consequence, the learning rule used was of a fairly simple form. It would be useful to extend this study to consider other more flexible rules. One step that could be taken would be to replace the standard increment/decrement value with a continuous variable which was dependent upon the difference between the activation threshold and the current activation values for the nodes concerned. This would give a more sensitive response. Where the spiking regions of two nodes only overlapped by a small fraction, a relatively small weight increment would take place as compared with the current case. It might also prove useful to investigate the effects of other types of learning rules. It is possible that different rules will produce different types of grouping effects, producing separate types of adaptive synchronizing networks for different tasks.

In Chapter 5, we investigated the possibility of maintaining multiple phases using input modulation. However the system was shown maintaining only two phases. We feel that the example shown in Chapter 5 could be extended to at least 4 phases although there would be a subsequent trade off in synchronization performance. It would be interesting to investigate the upper limits of this system.

Appendix A

Model Oscillators

A.1 Terminology

An attempt has been made at keeping a standard terminology throughout this thesis. Unless otherwise stated, the following symbols have the given meaning:

a_j	Pre-synaptic activation for node j
A_i	State of node i
α	Constant for synaptic threshold
c	Measure of connective scope (equation 3.10
η	Level of noise $\{\eta, 0 \leq \eta \leq 1.0\}$
$f_{noise}(x)$	Random noise function which returns a value between 0 and x
γ	Constant for post-synaptic spike
$I_i(t+1)$	Total input for node i at time $t+1$
l_{ij}	Lateral input at synapse from node j to node i
L_i	Total lateral input for node i
ν	Number of neighbours on one side of a node
N	Number of nodes in a network
T_i	Driving input for node i
τ_i	Driving input for node i without noise
τ_n	Driving input with noise

A.2 Oscillator models

A.2.1 Leaky integrator

$$x(t+1) = \begin{cases} T + E + x(t)k & x(t) < \theta \\ 0 & x(t) \geq \theta \end{cases} \tag{A.1}$$

T	Driving input (1.0)	$k = 0.95$
E	Entraining input (0.5,0.25)	$\theta = 19.93$

A.2.2 Morris-Lecar

The Morris-Lecar oscillator, a simplified derivative of the original Hodgkin-Huxley model of cell membrane activity (Hodgkin & Huxley, 1952; Morris & Lecar, 1989; Rinzel & Ermentrout, 1989).

$$\frac{dv}{dt} = -i_{ion}(v, w) + T + E \tag{A.2}$$

$$\frac{dw}{dt} = \phi \frac{[w_{\infty}(v) - w]}{\tau_w(v)} \tag{A.3}$$

$$i_{ion}(v, w) = \bar{g}_{Ca} m_{\infty}(v)(v - 1) + \bar{g}_K w(v - v_K) + \bar{g}_L(v - v_L) \tag{A.4}$$

$$m_{\infty}(v) = 0.5 * [1 + tanh\{(v - v_1)/v_2\}] \tag{A.5}$$

$$w_{\infty}(v) = 0.5 * [1 + tanh\{(v - v_3)/v_4\}] \tag{A.6}$$

$$\tau_w(v) = 1/cosh\{(v - v_3)/(2 * v_4)\} \tag{A.7}$$

v	Voltage	$\bar{g}_{Ca} = 1.1$	$v_1 = -0.01$	$\phi = 0.2$
w	Fraction of K^+ channels open	$\bar{g}_K = 2.0$	$v_2 = 0.15$	$v_K = -0.7$
T	Driving input (0.28)	$\bar{g}_L = 0.5$	$v_3 = 0.0$	$v_L = -0.5$
E	Entraining input (0.14,0.07)		$v_4 = 0.3$	

The equations are solved using the Euler method with a time step of 0.05.

A.2.3 Wang et al oscillator

The Wang et al model of interacting excitatory / inhibitory neural systems which exhibits oscillatory behaviour (Wang *et al.*, 1990).

$$\frac{dx_i}{dt} = -\frac{x_i}{\tau_x} + G_x[T_{xx}\frac{x_i}{\bar{x}} - T_{xy}F(\frac{y_i}{\bar{y}}) + S_i + I_i - H_i] \tag{A.8}$$

$$\frac{dy_i}{dt} = -\frac{y_i}{\tau_y} + G_y(-T_{yy}\frac{y_i}{\bar{y}} + T_{yx}\frac{x_i}{\bar{x}}) \tag{A.9}$$

$$H_i = \alpha \int_0^t x_i(\tau) exp[-\beta(t-\tau)]d\tau \tag{A.10}$$

$$G_r(v) = \frac{1}{1 + exp[-(v-\theta_r)/\lambda_r]} \qquad r \in \{x, y\} \tag{A.11}$$

$$F(x) = (1-\eta)x + \eta x^2 \qquad (0 \leq \eta \leq 1) \tag{A.12}$$

$\tau_x = 0.9$ $\theta_x = 0.4$ $T_{xx} = 1.0$ I_i : Driving input (0.3)

$\tau_y = 1.0$ $\theta_y = 0.6$ $T_{xy} = 1.9$ S_i : Entraining input (0.15,0.075)

$\bar{x} = 0.2$ $\lambda_x = 0.05$ $T_{yx} = 1.3$

$\bar{y} = 0.2$ $\lambda_y = 0.05$ $T_{yy} = 1.2$

$\eta = 0.4$ $\alpha = 0.2$ $\beta = 0.14$

The equations are solved using the Euler method with a time step of 0.05.

References

Abbott, L.F., & Kepler, T.B. 1990. Model Neurons: From Hodgkin-Huxley to Hopfield. *Statistical Mechanics of Neural Networks,Proceedings (Barcelona).*

Ackley, D.H., Hinton, G.E., & Sejnowski, T.J. 1985. A Learning Algorithm for Boltzmann Machines. *Cognitive Science, 9, pp147-169.*

Arbib, A.A., & Amari, S. (eds). 1989. *Dynamic Interactions in Neural Networks : Models and Data.* Springer-Verlag.

Artola, A., Brocher, S., & Singer, W. 1990. Different voltage-dependant thresholds for inducing long-term depression and long term potentiation in slices of rat visual cortex. *Nature (London), 347, pp69-72.*

Atiya, A., & Baldi, P. 1989. Oscillations and Synchronisations in Neural Networks : An Exploration of the Labeling Hypothesis. *International Journal of Neural Systems, 1(2).*

Baddeley, R.J. 1992. Phase coherence for solution likelihood estimation in neural networks. *Technical Report CCCN-11, Centre for Cognitive and Computaional Neuroscience, University of Stirling.*

Baldi, P., & Meir, R. 1990. Computing with Arrays of Coupled Oscillators : An Application to Preattentive Texture Discrimination. *Neural Computation, 2(4), pp458-471.*

Chagnac-Amitai, Y., & Connors, B. 1989. Horizontal Spread of Synchronised Activity in Neocortex and its Control by GABA-Mediated Inhibition. *Journal of Neurophysiology, 61(4).*

Chance, B., Pye, E.K., Ghosh, A.K., & Hess, B. (eds). 1973. *Biological and biochemical oscillators.* Academic Press.

Cohen, A.H. 1987. Effects of oscillator frequency on phase-locking in the lamprey central pattern generator. *Journal of Neurosciene Methods, 21, pp113-125.*

Connors, B.W. 1984. Initiatation of synchronised bursting in neocortex. *Nature, 310, pp685-687.*

Crick, F., & Koch, C. 1990. Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2, pp263-275.*

Damasio, A.R. 1989. The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation, 1, pp123-132.*

Dumenko, V.N. 1988. Spatial synchronization of cortical potentials and high frequency components of neocortical electrical activity during learning. *Neuroscience and Behavioural Physiology, 18(3), pp179-187.*

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H.J. 1988. Coherent Oscillations : A Mechanism of Feature Linking in the Visual Cortex? *Biological Cybernetics, 60, pp121-130.*

Eckhorn, R., Reitboeck, H.J., Arndt, M., & Dicke, P. 1989. Feature Linking via Stimulus-Evoked Oscillations: Experimental Results from Cat Visual Cortex and Functional Implications from a Network Model. *Proc. of the International Joint Conference on Neural Networks (Washington), pp723-730.*

Eckhorn, R., Reitboeck, H.J., Arndt, M., & Dicke, P. 1990. Feature Linking via Synchronization among Distributed Assemblies : Simulation of Results from Cat Visual Cortex. *Neural Computation, 2(3), pp283-292.*

Gilbert, C.D., & Weisel, T.N. 1983. Clustered intrinsic connections in cat visual cortex. *Journal of Neuroscience, 3, pp1116-1133.*

Gilbert, C.D., & Wiesel, T.N. 1979. Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature, 280, pp120-125.*

Glass, L., & Mackey, M.C. 1988. *From Clocks to Chaos : The Rhythms of Life.* Princeton University Press.

Glunder, H., & Nischwitz, A. 1993. On spike synchronization. *Brain Theory, Aertsen, A.(ed).*

Gray, C., & Singer, W. 1989. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *PNAS USA, 86, pp1698-1702.*

Gray, C.M., Konig, P., Engel, A.K., & Singer, W. 1989a. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature, 338, pp334-337.*

Gray, C.M., Konig, P., Engel, A.K., & Singer, W. 1989b. Synchronisation of oscillatory responses in visual cortex: A plausible mechanism for scene segmentation. *Proceedings of the International Symposium on Synergetics of Cognition.*

Grossberg, S., & Somers, D. 1991. Synchronised oscillations during co-operative feature linking in visual cortex. *Artificial Neural Networks.*

Hertz, J., Krogh, A., & Palmer, R.G. 1991. *Introduction to the Theory of Neural Computation.* Addison-Wesley.

Hinton, G.E., & Sejnowski, T.J. 1986. Learning and relearning in Boltzmann machines. *In Parallel Distributed Processing, 1, pp282-317, Rumelhart, D.E. and McClelland, J.L. (eds).*

Hodgkin, A.L., & Huxley, A.F. 1952. A quantitive description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology (London), 117, pp500-544.*

Horn, D., & Usher, M. 1991. Parallel Activation of Memories in an Oscillatory Neural Network. *Neural Computation, 3(1), pp31-43.*

Horn, D., Sagi, D., & Usher, M. 1991. *Segmentation, Binding and Illusory Conjunctions.* Technical Report. Weizmann Institute of Science.

Kammen, D.M., Holmes, P.J., & C., Koch. 1990. Cortical Architecture and Oscillations in Neuronal Networks: Feedback Versus Local Coupling. *In Models of Brain Function, R.M.J. Cotteril, ed., pp273-284, Cambridge University Press.*

Kisvarday, Z.F. et al. 1986. Synaptic targets of HRP-filled layer III pyramidal cells in the cat striate cortex. *Experimental Brain Research, 64, pp541-552.*

Koch, C., & Schuster, H. 1992. A simple network showing burst synchronization without frequency locking. *Neural Computation, 4(2).*

Konig, P., & Schillen, T.B. 1991. Stimulus-Dependant Assembly Formation of Oscillatory Responses : I Desynchronisation. *Neural Computation, 3, pp167-178.*

Lumer, E.P., & Huberman, B.A. 1992. Binding Hierarchies: A Basis for Dynamic Perceptual Grouping. *Neural Computation, 4(3), pp341-355.*

MacKay, D.J.C. 1991. Maximum entropy connections: neural networks. *In Maximum Entropy and Bayesian Methods, pp237-244, Grandy, W.T. and Schick, L.H. (eds), Kluwer Academic Publishers.*

Malsburg, C. von der, & Buhmann, J. 1992. Sensory Segmentation with Coupled Neural Oscillators. *Biological Cynernetics, 67, pp233-242.*

Malsburg, C. von der, & Schneider, W. 1986. A neural cocktail-party processor. *Biological Cybernetics, 54, pp29-40.*

Mardia, K.V. 1972. *Statistics of directional data.* Academic Press.

Mitchison, G., & Crick, F. 1982. Long axons within the striate cortex : Their distribution, orientation and patterns of connection. *Proc. of National Academy of Science (U.S.A.), 79, pp3661-3665.*

Morris, C., & Lecar, H. 1989. Voltage oscillations in the barnacle giant muscle fibre. *Biophys. J., 35, pp193-213.*

Mozer, M., Zemel, R., & Behrmann, M. 1991. *Learning to segment images using dynamic feature binding.* Technical Report, CU-CS-540-91. University of Colorado - Boulder.

Nischwitz, A., Glunder, H, & Klausner, P. 1991. Synchronization of Spikes in Populations of Laterally Coupled Model Neurons. *Artificial Neural Networks, Kohonen T.,Makisara K., Simula O., Kangas J. (eds), pp1771-1774.*

Nischwitz, A., Glunder, H., von Oertzen, A., & Klausner, P. 1992. Synchronization and Label-switching in Networks of Laterally Coupled Model Neurons. *Pages pp851–854 of:* Aleksander, I., & Taylor, J. (eds), *Artificial Neural Networks 2.* North-Holland.

Othmer, H.G. 1985. Synchronization, Phase Locking and Other Phenomena in Coupled Cells. *In Temporal Order, Rensing, L. and Jaeger, N. (Eds), pp130-143.*

Pavlidis, T. 1973. *Biological oscillators : Their mathematical analysis.* Academic Press.

Perkel, D., & Mulloney, B. 1974. Motor Pattern Production in Reciprocally Inhibitory Neurons Exhibiting Postinhibitory Rebound. *Science, 185.*

Poppel, E, & Logothetis, N. 1986. Neuronal oscillations in the human brain. *Naturwissenschaften, 73, pp267-268.*

Rinzel, J., & Ermentrout, G.B. 1989. Analysis of Neural Excitability and Oscillations. *In Methods in Neuronal Modeling - From Synapses To Networks, Koch, C. and Segev, I. (eds).*

Rockland, K.S., & Lund, J. 1982. Widespread Periodic Intrinsic Connections in the Tree Shrew Visual Cortex. *Science, 215, pp1532-1534.*

Rumelhart, D., & McClelland, J.L. 1987a. *Parallel Distributed Processing, Vol. 1: Foundations.* MIT Press.

Rumelhart, D., & McClelland, J.L. 1987b. *Parallel Distributed Processing, Vol. 2: Psychological and Biological Models.* MIT Press.

Schillen, T.B., & Konig, P. 1989. *Coherency Detection by Coupled Oscillatory Responses - Synchronizing Connections in Neural Oscillator Layers.* Technical Report. Max-Planck Institute for Brain Research.

Shastri, L., & Ajjanagadde. 1993. From simple associations to systematic reasoning : A connectionist representation of rules,variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences 16, pp417-494.*

Shepherd, G.M. (ed). 1990. *The Synaptic Organization of the Brain.* Oxford University Press.

Singer, W. 1985. Activity-dependent self-organization of the mammalian visual cortex. *In:* Rose, D., & Dobson, V.G. (eds), *Models of the Visual Cortex.* Wiley (New York).

Singer, W. 1989. Search for coherence: A basic principle of cortical self-organization. *Concepts in Neuroscience, 1(1).*

Skarda, C.A., & Freeman, W.J. 1987. How brains make chaos in order to make sense of the world. *Behavioural and Brain Sciences, 10, pp161-195.*

Smith, L.S., Cairns, D.E., & Nischwitz, A. 1993. Spike synchronization with time-delayed inihibitory lateral connections. *Submitted to Neural Computation.*

Sompolinsky, H., Golomb, D., & Kleinfeld, D. 1990. Global processing of visual stimuli in a neural network of coupled oscillators. *Proc. Natl. Acad. Sci. USA, 87, pp7200-7204.*

Sporns, O., Gally, J.A., Reeke(Jnr), G.N., & Edelman, G.M. 1989. Reentrant signaling among simulated neuronal groups leads to coherency in their oscillatory activity. *Neurobiology, 86, pp7265-7269.*

Stent, G.S. 1973. A physiological mechanism for Hebb's postulate of learning. *Proceedings of the National Academy of Science USA 70, pp. 997-1001.*

Steriade, M., Jones, E.G., & Llinas, R.R. (eds). 1990. *Thalamic oscillations and signaling.* John Wiley and Sons.

Vidal, J.J., & Haggerty, J. 1988. Synchronization In Neural Nets. *Neural Information Processing Systems.*

Wang, D., Buhmann, J., & Malsburg, C. von der. 1990. Pattern Segmentation in Associative Memory. *Neural Computation, 2, pp94-106.*

Wilson, H.P., & Cowan, J.D. 1972. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal.*

Wilson, M., & Bower, J. 1990. Computer Simulation of Oscillatory Behaviour in Cerebral Cortical Networks. *Pages 84–91 of:* Touretzky, D.S. (ed), *Advances in Neural Information Processing Systems*, vol. 2. Morgan Kaufmann Publishers Inc.

Young, M.P., Tanaka, K., & Yamane, S. 1992. On oscillating neuronal responses in the visual cortex of the monkey. *Journal of Neurophysiology, 67, pp1464-1474.*