

Department of Computing
Imperial College London

Metastability: An Emergent Phenomenon in Networks of Spiking Neurons

David Bhowmik

Submitted in part fulfilment of the requirements for the degree of

Doctor of Philosophy

of

Imperial College London

June 2014

Statement of Originality

This thesis is an original piece of work produced in entirety by David Bhowmik

Copyright Declaration

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

Abstract

It is widely recognised that different brain areas perform different specialised functions. However, it remains an open question how different brain areas coordinate with each other and give rise to global brain states and high-level cognition. Recent theories suggest that transient periods of synchronisation and desynchronisation provide a mechanism for dynamically integrating and forming coalitions of functionally related neural areas, and that at these times conditions are optimal for information transfer. Empirical evidence from human resting state networks has shown a tendency for multiple brain areas to synchronise for short amounts of time, and for different synchronous groups to appear at different times. In dynamical systems terms, this behaviour resembles metastability — an intrinsically driven movement between transient, attractor-like states. However, it remains an open question what the underlying mechanism is that gives rise to these observed phenomena.

The thesis first establishes that oscillating neural populations display a great amount of spectral complexity, with several rhythms temporally coexisting in the same and different structures. The thesis next explores inter-band frequency modulation between neural oscillators. The results show that oscillations in different neural populations, and in different frequency bands, modulate each other so as to change frequency. Further to this, the interaction of these fluctuating frequencies in the network as a whole is able to drive different neural populations towards episodes of synchrony.

Finally, a symbiotic relationship between metastability and underlying network structure is elucidated, in which the presence of plasticity, responding to the interactions between different neural areas, will naturally form modular small-world networks that in turn further promote metastability. This seemingly inevitable drive towards metastability in simulation suggests that it should also be present in biological brains. The conclusion drawn is that these key network characteristics, and the metastable dynamics they promote, facilitate versatile exploration, integration, and communication between functionally related neural areas, and thereby support sophisticated cognitive processing in the brain.

Acknowledgements

I would like to express my gratitude to Prof. Murray Shanahan for his invaluable guidance throughout the course of this research, and to my family for their constant understanding and support. I would also like to thank everyone in the Computational Neurodynamics Group for making it such an interesting and enjoyable place to work.

Table of Contents

1	Introduction	1
1.1	Motivations and Objectives	1
1.2	Publications	5
1.2.1	Journals	5
1.2.2	Conferences	5
2	Background to the Question of Metastability	7
2.1	Overview	7
2.2	Recording Brain Activity	7
2.3	Oscillations and Behaviour	9
2.4	Correlation between Functional Areas	10
2.5	Anatomical Structure	12
2.6	Dynamical Complexity and Topological Structures	13
2.7	Functional Dynamics and the Brains Anatomy	15
2.8	Discussion	16
3	Modelling Brain Oscillations	17
3.1	Oscillation in the Brain	17
3.1.1	The Function of Inhibition	18
3.1.2	Pyramidal Inter-Neuronal Gamma	19
3.1.3	Neural Coding, Plasticity and Oscillation	20
3.1.4	Oscillation and Neuron Types	21
3.2	Oscillator Model Choice	22

3.2.1	Different Ways of Modelling Neural Populations	22
3.2.2	Justification of Model Choice	23
3.2.3	Model Configuration.....	23
3.3	Modelling Neurons.....	24
3.3.1	Quadratic Integrate-and-Fire Neurons	25
3.3.2	Izhikevich Neurons	26
3.3.3	Hodgkin Huxley Neurons	26
3.4	Modelling Synapses	28
3.4.1	Conductance Synapses	28
3.4.2	Synaptic Reversal Potentials	28
3.4.3	Spike Timing Dependent Plasticity.....	28
3.5	Generating PING Oscillators.....	30
3.6	Hardware Acceleration.....	32
3.7	Analysis of Neuron Models in PING Architecture	33
3.7.1	Oscillating Only in Response to a Learnt Stimulus	34
3.7.2	The Effect of Learning Time on Oscillation	35
3.7.3	The Effect of Stimulus Length on Oscillation	37
3.8	Discussion	40
4	The Collective Behaviour of Neural Oscillator Populations	43
4.1	Synchronisation in Neuroscience	44
4.2	Why Synchrony Occurs	44
4.3	Simple Oscillator Models and Neuroscience Research.....	45

4.4	Experimental Variations.....	46
4.5	Generation of Neural Oscillator Nodes.....	47
4.6	Synchronisation metric.....	50
4.7	The Question of Neural Coupling Strength.....	50
4.8	Critical Coupling Experiment	51
4.8.1	Experimental Setup	51
4.8.2	Results.....	53
4.9	Interaction between Trained and Untrained Nodes.....	61
4.10	Discussion	63
5	Metastability and Inter-Band Frequency Modulation.....	67
5.1	Brain Rhythms and Dynamical Organisation.....	67
5.2	Transient Dynamics.....	68
5.3	Spectral Complexity in Neural Systems.....	69
5.4	Metastability and Frequency Modulation Experiment.....	70
5.4.1	Experimental Setup	70
5.4.2	Extraction of Intermittent Frequency Strands	71
5.4.3	Mean Intermittent Frequency Correlation.....	75
5.4.4	Coalition entropy.....	76
5.4.5	Results.....	77
5.5	Discussion	87
6	Metastability and Plasticity	89
6.1	Transient Synchronisation and Brain Coordination	89

6.2	Local and Global Processing	90
6.3	Plasticity Experiment	90
6.3.1	Experimental Setup	90
6.3.2	Small-world index	91
6.3.3	Modularity	91
6.3.4	Knotty centrality	92
6.3.5	Metastability	93
6.3.6	Results	93
6.4	Discussion	107
7	Conclusion	109
8	Bibliography	113

Table of Figures

Figure 3.1	Pyramidal Inter-Neuronal Gamma architecture used to generate oscillations.....	24
Figure 3.2	The STDP weight modification curve.....	30
Figure 3.3	Raster plot of the excitatory layer of a QIF PING node oscillating at 30 Hz.....	33
Figure 3.4	Trained PING networks response to noise in the learnt stimuli.....	35
Figure 3.5	The effect of learning time on the different neuron networks response to the learnt stimuli.....	36
Figure 3.6	The effect of changing the length of the stimuli.....	38
Figure 3.7	The effect of altering the stimulus length.....	39
Figure 4.1	Mean delays for evolved untrained QIF PING networks.....	48
Figure 4.2	Raster plot of the excitatory layer of an evolved QIF PING oscillator.....	49
Figure 4.3	The effect of synaptic weight and connection ratio on synchrony.....	51
Figure 4.4	Network of oscillators.....	52
Figure 4.5	Global synchrony for 10 PING oscillators.....	53
Figure 4.6	Number of actual coalitions compared to number of intrinsic coalitions.....	55
Figure 4.7	Pairwise synchrony for trained QIF at weight 0.0172.....	57
Figure 4.8	Pairwise synchrony for untrained QIF at weight 0.07.....	58
Figure 4.9	Pairwise synchrony for untrained HH at weight 0.02.....	59
Figure 4.10	Pairwise synchrony at the point of maximum global synchrony reached.....	60
Figure 4.11	The firing behaviour of an excitatory layer of a HH PING node.....	61
Figure 4.12	The mean pairwise synchrony between trained and untrained nodes.....	62
Figure 5.1	Extraction of frequencies from population firings.....	73
Figure 5.2	Synchrony and coalition entropy.....	79
Figure 5.3	The number of coexisting frequencies.....	80
Figure 5.4	The number of correlations found.....	81
Figure 5.5	Mean intermittent frequency correlation.....	83

Figure 5.6	Separation of positive and anti correlations.....	84
Figure 5.7	Average number of correlations, and the peak of modulated exploration.....	86
Figure 6.1	Dynamical metrics prior to learning.....	96
Figure 6.2	Dynamical metrics post-learning.....	98
Figure 6.3	Comparison of metastability metrics.....	99
Figure 6.4	Coalition presence.....	101
Figure 6.5	Metastability and modular small-world structure post-learning.....	104
Figure 6.6	Metastability, modularity, and knotty centredness post-learning.....	106

1 Introduction

1.1 Motivations and Objectives

One of the greatest challenges in science is the question of how cognition arises from the human brain. Much is known about the behaviour of individual neurons, and continual progress is being made in explaining the function of individual brain areas as well as the basic organisational principles of brain structure. It is widely recognised that different brain areas perform different specialised functions. However, it remains an open question how localised neural activity gives rise to complex cognitive functions organised at a global level. In order to efficiently support the myriad changing cognitive demands under which we are placed, a flexible integrative mechanism must exist (Bressler, 1995; Doesburg et al., 2008; Varela et al., 2001).

An influential candidate for the mechanism underlying coordination and communication among brain areas is synchronous oscillation over multiple frequency bands (Varela et al., 2001; Fries, 2005; Fries, 2009; Tognoli and Kelso, 2011). Groups of neurons firing synchronously are hypothesised to underlie many cognitive functions such as attention, associative learning, memory, and sensory selection (Jensen, Kaiser and Lachaux, 2007; Miltner et al., 1999; Siegel, Warden and Miller, 2009; Fries et al., 2001). Recent theories suggest that transient periods of synchronisation and desynchronisation provide a mechanism for dynamically integrating and forming coalitions of functionally related neural areas (Chialvo, 2010; Kelso, 2012; Shanahan, 2012), and that at these times conditions are optimal for information transfer (Buehlmann and Deco, 2010).

The formation of transient coalitions of brain areas has been observed *in vivo* (Betz et al., 2012). Fluctuations have been reported on a time scale of tens to hundreds of milliseconds involving spatially local and remote sites, and resulting in fast reconfigurations of network states (Betz et al., 2012). In dynamical systems terms, the tendency for multiple brain areas to synchronise for short

amounts of time, and for different synchronous groups to appear at different times, resembles *metastability* — an intrinsically driven movement between transient, attractor-like states. However, it remains an open question what the underlying mechanism is that gives rise to these observed phenomena in the brain.

It may be the case that, the tendency for multiple brain areas to synchronise for short amounts of time, and for different synchronous groups to appear at different times is simply an epiphenomenon resulting from different groups of neurons performing different tasks. Alternatively, the phenomenon may be the result of some other mechanism than an intrinsically driven metastable dynamic. Given that synchronous activity between brain areas is hypothesised to facilitate optimal conditions information transfer (Fries, 2005; Fries, 2009; Buehlmann and Deco, 2010), it follows that desynchrony between brain areas would impede information transfer. Therefore, if this phenomenon is driven by an intrinsic metastable mechanism then the question of how this metastable dynamic operates in such a way that human cognition is not disrupted, but instead is stable and effective must be raised. Further to this, questions arise regarding whether such dynamics would form as a basis for cognitive processing, and if so how this would function. Before these questions can be addressed it must first be ascertained whether an intrinsically driven metastable dynamic is at play. The central aim of this thesis is to address this question using a series of computer models of neural dynamics.

Large scale functional brain structure has been identified, for example the in the combined activity of whole cortices and operation of the so-called default network (Buckner, Andrews-Hanna and Schacter, 2008). However, it is unclear what structural and functional characteristics would be required to facilitate metastable dynamics. To help validate the proposal that intrinsic metastable dynamics exist in the brain, it must be shown that such dynamics emerge from known brain structure, as well as how such brain structures themselves emerge. Therefore, this thesis investigates the effect that altering structure through synaptic plasticity has upon the observed dynamics.

In recent years researchers have begun to investigate metastable dynamics in the brain. Abstract oscillator models of interacting neural populations connected according to human connectome data have been used to identify metastable dynamics as a likely cause of such transitions, and the resulting data correlates well with real data observed *in vivo* (Cabral et al., 2011; Cabral et al., 2013; Hellyer et al., 2014). Supporting work using abstract pulse coupled oscillator models, as well as Kuramoto oscillators, placed in a network of communities akin to the types of structure found in the brain, also exhibit metastable dynamics (Wildie and Shanahan, 2011; Shanahan, 2010). However, whether such

simple oscillator models sufficiently capture the behaviour of interacting populations of neurons is unclear. As such, the evidence in favour of metastable dynamics as a cause of the transient formation of coalitions of multiple brain areas is not by any means conclusive. The present investigation takes abstract oscillator models of neural populations as its starting point, but looks into more realistic spiking neuron simulations.

After a background review in chapter 2, chapter 3 discusses the mechanistic causes of oscillation in the brain, leading to the examination of the pyramidal inter-neuronal gamma model of cortical oscillation. The chapter examines in detail the collective oscillatory behaviour of neurons given their individual firing characteristics and internal dynamics. In addition, it explores the effect of plasticity upon oscillating neuronal populations. The chapter closes by proposing an experimental configuration for the studies detailed in later chapters.

A mathematical abstraction that is gaining increasing acceptance for modelling neural information processing, and more recently modelling apparent metastable dynamics in the brain (Cabral et al., 2011; Cabral et al., 2013; Hellyer et al., 2014), is the Kuramoto oscillator model (Kuramoto, 1984). The Kuramoto model is proposed as an elementary unit to represent populations of oscillatory neurons, and to capture fundamental properties of the collective dynamics of interacting communities of oscillatory neurons. Whether abstract oscillator models of this kind are sufficient to model oscillating neural populations is contentious. Whilst the use of simple oscillator models is promising, an exploration of the extent to which they correspond to purportedly equivalent neural systems is required. Of particular concern here is the relevance of models of metastability in the brain that use Kuramoto oscillators. Chapter 4 addresses this issue experimentally by using neural models to replicate the most fundamental of Kuramoto's findings, in which he showed that for any number of oscillators there is a critical coupling value K_C below which the oscillators are fully unsynchronised and another critical coupling value $K_L \geq K_C$ above which all oscillators become fully synchronised (Kuramoto, 1984). In the present study, Kuramoto oscillators are replaced with oscillating populations both of quadratic integrate-and-fire neurons and of Hodgkin-Huxley neurons to establish whether Kuramoto's findings still hold in a more biologically realistic setup. The results establish that oscillating neural populations display more spectral complexity than simple oscillator models, with several rhythms temporally coexisting in different structures.

The thesis next addresses the question of whether there is an intrinsic metastable dynamic present amongst populations of neural oscillators that drives the system towards transient periods of

synchronisation and desynchronisation. Evidence using spiking neuron models of competitive entrainment (Wildie and Shanahan, 2012) suggests that intrinsic dynamical processes between populations of oscillating neurons can selectively pull populations into and out of phase, and thereby are potentially able to support transient coalition formation in the brain as a whole. However, the question remains open of whether a scaled up and more complex version of these mechanics is responsible for the formation of transient coalitions of multiple neural areas in systems of large-scale interacting neural populations, or whether the phenomenon is an artefact of other neural processes, or perhaps even merely an epiphenomenon.

Having already established that oscillating neural populations exhibit high spectral complexity, with several rhythms temporally coexisting in different structures, chapter 5 explores inter-band frequency modulation between neural oscillators using models of quadratic integrate-and-fire neurons, Izhikevich neurons, and Hodgkin-Huxley neurons. The structural connectivity in a network of neural oscillators is varied, the spectral complexity is assessed, and the inter-band frequency modulation is correlated. The correlated frequency modulation is contrasted against measures of coalition entropy and synchrony. The results show that oscillations in different neural populations, and in different frequency bands, modulate each other and bring about changes in frequency. Further to this, the interaction of these fluctuating frequencies in the network as a whole is able to drive different neural populations towards episodes of synchrony. An area in the connectivity space is located in which the system directs itself in this way, facilitating the exploration of a large repertoire of synchronous coalitions, characteristic of metastability.

Chapter 5 highlights anatomical structure as key to facilitating metastability. Others have identified anatomical structure as a major constraint on the functional complexity of neural systems (Sporns, Tononi and Edelman, 2000; Bullmore and Sporns, 2009; Sporns, 2013). Hence, the network structures that are prevalent in the brain are of great interest to those studying its capabilities. Evidence of localised clustering as well as a low characteristic path length between sites has been found in the network structure of large-scale cortical systems (Sporns, Tononi and Edelman, 2000). These findings, characteristic of modular small-world connectivity, have since been substantiated and extended (Bullmore and Sporns, 2009). In addition, Sporns (2013) reports on a growing body of work that draws attention to how the balance between structural segregation and integration is essential for the operation of the brain networks underlying cognitive function. Functional activity is thought to reflect underlying anatomical structure, and correlations have been found to exist between

patterns of structural connectivity and resting-state functional activity (Honey et al., 2009; van den Heuvel et al., 2009).

Chapter 6 explores the relationship between network structure and dynamical complexity, as well as investigating how appropriate connectivity might arise. A computer model wherein synaptic plasticity acting in concert with the modulatory influences between different oscillating populations of neurons is presented. What results is a restructuring of the network topology so that it exhibits modular small-world connectivity. A symbiotic relationship between metastability and underlying network structure is elucidated, in which the presence of plasticity, responding to the interactions between different neural areas, will naturally give rise to modular small-world networks. Modular small-world networks in turn promote metastability, and metastability further enhances these structural features. This seemingly inevitable drive towards metastability in simulation suggests that it should also be present in biological brains.

The conclusion drawn from this study is not only that metastable dynamics naturally arise in systems of oscillating neural populations, but also that the natural formation of modular small-world network characteristics, and the metastable dynamics they promote, facilitate versatile exploration, integration, and communication between functionally related neural areas, and thereby support sophisticated cognitive processing in the brain.

1.2 Publications

The work presented in the thesis has resulted in the following publications.

1.2.1 Journals

Bhowmik D, Shanahan M (2013) Metastability and Inter-Band Frequency Modulation in Networks of Oscillating Spiking Neuron Populations, PLoS One 8 (4), e62234.

1.2.2 Conferences

Bhowmik D, Shanahan M (2012) How Well Do Oscillator Models Capture the Behaviour of Biological Neurons? Proceedings IJCNN 2012, pp. 1-8.

Bhowmik D, Shanahan M (2013) STDP Produces Robust Oscillatory Architectures that Exhibit Precise Collective Synchronization, Proceedings IJCNN 2013, pp. 1-8.

Bhowmik D, Shanahan M (2013) STDP Produces Well Behaved Oscillations and Synchrony. Proceedings 4th International Conference on Cognitive Neurodynamics. In Press.

2 Background to the Question of Metastability

2.1 Overview

The focus of this thesis is the question of whether neuronal oscillation plays an active role in brain function. To set the scene for the study this chapter first discusses the techniques for recording neural activity that have been used to identify brain oscillations. Next an explanation of the characteristics of oscillations that have been observed is presented along with a review of their association to different behaviours. Following this a review of current theory regarding the interaction between oscillations in different neural areas is given. As neuroanatomy places critical constraints on the functional connectivity, the chapter next explores different metrics for analysing network topology. Given such metrics, research work that explores different topological structures and the resulting dynamical complexity they exhibit is presented. The last section focuses on the anatomical structure found in the brain and resulting functional dynamics that have been observed *in vivo*.

2.2 Recording Brain Activity

Neurons are the basic functional units in the brain. They transmit information using electrical signals called action potentials. A neuron has a membrane that serves as a barrier to separate the inside and outside of the cell. The membrane voltage of a neuron is dictated by the difference in electrical potential inside and outside of the cell. Neurons are electrically charged by membrane ion channels that pump ions, which have different electrical charges, across their membranes. Neurons are constantly exchanging ions with the extracellular surroundings in this way. In doing so they can not

only maintain resting potential, but also propagate action potentials by depolarising the membrane beyond a critical threshold (Bear, Connors and Paradiso, 1996).

The recording of neural activity is an important aspect of research for those trying to understand the workings of the brain. There are many techniques available to record neural activity. Single unit recording techniques provide precise recordings from single firing neurons by way of either an invasive “sharp” microelectrode inserted into the soma, or a non-invasive patch clamp that uses a single electrode enclosing a membrane surface area or "patch" that often contains just one or a few ion channel. The latter type of electrode is distinct from the "sharp" microelectrode used to impale cells in traditional intracellular recordings. Other techniques such as electroencephalography (EEG), magnetoencephalography (MEG), and functional magnetic resonance imaging (fMRI) do not facilitate recording at single neuron resolution, but are instead able to record the combined activity of groups of neurons.

Electrophysiological measurements can be obtained using EEG by placing electrodes on the scalp that capture the effect of a process known as volume conduction. When many ions are pushed out of many neurons at the same time they push their neighbours and in doing so cause a wave. When the wave of ions reaches the electrodes on the scalp they push or pull electrons on the metal on the electrodes resulting in a reading. EEG activity reflects the net synchronous activity of thousands or millions of neurons that have a similar spatial location (Niedermeyer and da Silva, 2004).

MEG captures brain activity by using very sensitive magnetometers to record magnetic fields produced due to the ionic currents flowing in the dendrites of tens of thousands of neurons during synaptic transmission. MEG has better spatial resolution than EEG because magnetic fields are less distorted than electric fields by the skull and scalp (Hansen, Kringelbach and Salmelin, 2010). The electric potential generated by an individual neuron is far too small to be identified by either EEG or MEG, and both techniques suffer from not being able to detect activity much below the cortical surface.

fMRI measures brain activity indirectly by detecting changes related to blood flow. As neurons do not have internal reserves of oxygen, firing creates a need for more oxygen to be brought into the neuron. Oxygen is released by the blood to active neurons at a greater rate than to inactive neurons. The difference in magnetic susceptibility between oxygenated or deoxygenated blood, allows an fMRI scanner to detect the magnetic signal variation. Blood flow to a brain region therefore increases

when that area of the brain is in use, and the magnetic effects of this are detected. Unlike EEG and MEG, fMRI can detect deep activity. fMRI also has better spatial resolution than EEG and MEG. The Spatial resolution of fMRI is measured by voxel size. Full brain studies use larger voxels, while those that focus on specific regions of interest typically use smaller sizes. Sizes range from 5 mm to 1 mm. A voxel may contain a few million neurons; however the actual number depends on voxel size and the area of the brain being imaged (Huettel, Song and McCarthy, 2009).

The temporal resolution of fMRI is between 1 or 2 seconds, whereas the temporal resolution of both EEG and MEG is ~ 1 ms. When the brain is at rest, the fluctuations in activity occur at a very slow time scale, hence fMRI has been very successful in studying spatial patterns in such situations. The more precise temporal resolution of MEG and EEG permits the detection of oscillations in the activity of neural populations at frequencies that are much faster than 1 Hz.

The ability to record the fast-acting behaviour of neural populations has allowed the detection of periodic activity. A greater understanding of the biological mechanisms underlying such fast oscillation of neural populations, accompanied with insights gained from advances in the computational modelling of neural systems, has further stimulated interest in brain oscillations (Wang, 2010). It is these oscillations that are the focus of this study and which we shall now discuss.

2.3 Oscillations and Behaviour

The prevalence of periodic behaviour in the brain has long been observed. Oscillatory frequencies in the mammalian brain range from approximately 0.05 Hz to 600 Hz (Buzsáki and Draguhn, 2004). The local nature of most synaptic connections, along with the limits of synaptic and axonal conduction delays (Kopell et al., 2000; Steriade, 2001), restrict the size of the population that can be recruited during an oscillation. As a result, the size of neuronal area involved in an oscillation tends to be small for higher frequencies, whereas larger areas are recruited for slower frequencies (Buzsáki and Draguhn, 2004). The different oscillatory frequencies are categorised into empirically observed bands. The band separation results from the different band classes forming a linear progression on a natural logarithmic scale, such that there is a constant ratio between neighbouring frequencies. Neighbouring frequency bands within the same neuronal structure typically compete with each other (Klimesch, 1999), although several rhythms can temporally coexist in the same structure (Steriade, 2001). It has been suggested that the separation of centre frequencies by a constant non-integer ratio

reduces harmonic frequencies and the effect of phase-modulation interference between bands (Penttonen and Buzsáki, 2003).

It is unlikely that each frequency band is associated with a single cognitive function (Engel and Fries, 2010), and the different rhythms have been noted to have a diverse range of functional associations. Thalamocortical networks display increased delta band (0.1–3.5 Hz) power during deep sleep (McCormick, Sejnowski and Steriade, 1993). Theta (4–7.5 Hz) activity is increased during memory encoding and retrieval (Basar et al., 2000). Alpha band (8–13 Hz) changes are associated with attentional demands (Klimesch, 1999). Beta (14–30 Hz) oscillations have been related to the sensorimotor system (Pfurtscheller, Stancak Jr and Neuper, 1996). Gamma (30–80 Hz) is thought to underlie functions such as attention (Jensen, Kaiser and Lachaux, 2007), associative learning (Miltner et al., 1999), working memory (Siegel, Warden and Miller, 2009), the formation of episodic memory (Lisman, 2005), visual perception (Fries et al., 2001), and sensory selection (Fries et al., 2002) to name just a few. Whilst these different oscillatory rhythms manifest themselves concurrently with the respective functions to which they are associated, it remains an open question as to whether neuronal oscillation plays an active role in these functions and brain operation in general or if it is simply an epiphenomenon resulting from the of activity of neural populations.

As mentioned above, several rhythms can temporally coexist in the same structure (Steriade, 2001). Whilst the nature of the separation of the bands reduces the effect of phase-modulation interference between bands (Penttonen and Buzsáki, 2003), the interaction of frequency bands can occur, and has been observed in both phase (Buzsáki et al., 2003; Palva, Palva and Kaila, 2005; Belluscio et al., 2012) and amplitude (Shirvalkar, Rapp and Shapiro, 2010). For example, the theory of theta-gamma phase-amplitude coupling, as integral to the coding and temporal ordering of items in working memory (Lisman and Idiart, 1995; Tort et al., 2009), was formulated from observations of the phase of a theta oscillations modulating the amplitude of a gamma oscillation in rat hippocampus (Canolty et al., 2006).

2.4 Correlation between Functional Areas

The interaction between oscillations in different brain structures is described by functional connectivity. Functional connectivity between different brain areas is inferred from the statistical correlation of phase and amplitude in EEG or MEG signals or fMRI data during the performance of a task or in the resting state (Park and Friston, 2013). Structural connectivity, based on diffusion MRI,

is able to reveal the cause of much inter-areal functional connectivity. However, diffusion MRI is potentially blind to weak long-range axonal connections, which may serve as weak ties for global integration, and hence cannot tell us the full story. In addition, functional connectivity may exist between anatomically unconnected areas. For example, synchronous activity in two anatomically unconnected areas may be driven by common sources, polysynaptic connections, or perhaps other configurations of bidirectional circuits (Park and Friston, 2013).

When two areas are anatomically connected it is theorised that coherence between them opens up the possibility of communication (Fries, 2005; Fries, 2009). During oscillation in an individual neural population all the neurons are quiescent, or less active, for periods between increased levels of firing. This rhythmic behaviour is often due to local inhibitory neurons that affect the populations' excitatory neurons with a fast and strong synaptic input (Papp et al., 2001). This leaves only a short window for the excitatory neurons to fire after one period of inhibition wears off and before the next one starts (Hasenstaub, Shu and Haider, 2005). Coherence between two neuronal populations aligns the firing of the two populations so that the output window of one falls within the window for which the other is susceptible to input. Gamma band oscillations are sufficiently regular to allow prediction of the next excitability peak. As long as the travelling time from the sending to the receiving group is also reliable, their communication windows for input and output are open at the same times. Conduction delays are typically an order of magnitude shorter than the cycle length of the oscillation allowing sending and receiving to occur within one excitability peak. Packages of spikes can therefore arrive at other neuronal groups in precise synchronisation and enhance their impact. Rhythmic inhibition therefore provides rhythmic modulation of excitatory input gain. Fries calls this hypothesis 'communication through coherence' (Fries, 2005; Fries, 2009).

The dynamics of changing phase coherence between oscillating neuronal populations has been studied computationally by Wildie and Shanahan (2012). They demonstrate a computer model of entrainment using spiking neurons in which a source population oscillating at a particular frequency brings into phase alignment a target population which was initially oscillating at a different frequency. Further to this, they show that competition between two source stimuli, which is driven by differences in coherence of oscillation, aligns the target population to the source population with which it is most coherent. Competition between stimuli of equal coherence results in model output that alternates between the representation of the two stimuli, in a manner strongly resembling the

well known biological phenomenon of binocular rivalry. Transmission of a single selected stimulus is enabled between generating and receiving neurons via ‘communication through coherence’, which they measure through correlation of spikes.

Buehlmann and Deco (2010) explored the hypothesis that synchronisation makes information transport more efficient. They used transfer entropy, an information theoretical measure, to quantify the exchanged information between two neuronal populations. They show that the transferred information depends on the phase relation of the signal, and that the higher the power in the oscillation then the earlier the onset of the information flow and the faster the information transfers. Further to this, the amount of exchanged information also increases as a function of power.

2.5 Anatomical Structure

Neuroanatomy places critical constraints on the functional connectivity of the cerebral cortex (Sporns, Tononi and Edelman, 2000). Therefore the relationship between structural and functional connectivity has been of growing interest to neuroscientists. The balance between segregation and integration is believed by many to be essential for the operation of the distributed networks underlying cognitive function (Tononi, Sporns and Edelman, 1994; Fox and Friston, 2012). Communities and hubs in brain networks promote such integration and segregation. Communities are sets of regions that are densely interconnected internally, while connections between members of different communities are more sparse. Network hubs link communities to one another and ensure efficient communication and information integration (Sporns, 2013).

Small-world structure described by Watts and Strogatz (Watts and Strogatz, 1998; Watts, 1999) has been the focus of recent research that assesses its relationship to the topologically of brain networks (Sporns, Tononi and Edelman, 2000; Bullmore and Sporns, 2009). Small-world networks have a few long-range connections and are highly clustered locally. As a result they exhibit a short average path length between any two nodes. The small-world property is formally characterised by the relationship between two values, the average path length between any two nodes and the fraction of all possible edges between neighbours that are actual edges (Watts and Strogatz, 1998). Measures have been defined for quantifying the degree of small-world connectivity within existing networks (Humphries, Gurney and Prescottt, 2006). Sporns *et al* (2000) investigated the structure of large-scale cortical systems and found evidence of localised clustering as well as a low characteristic path length

between sites, the attributes of modular small-world connectivity, and these findings have since been substantiated and extended (Bullmore and Sporns, 2009).

The metric known as modularity measures the degree to which a graph can be divided into highly coupled sub-graphs with only a few connections between those sub-graphs. Within the brain such community structure is assumed to reflect functional segregation, with highly coupled nodes believed to share similar functional association. Local modularity and clustering surrounding a node are measured by the clustering coefficient (Watts and Strogatz, 1998), which is defined as the fraction of a nodes neighbours that are also neighbours of each other.

There are many different ways of identifying a node as significant within the structure of a network. Centrality is a key concept for doing this and has also been considered in the analysis of brain connectivity (Sporns, Honey and Kötter, 2007; Joyce et al., 2010; Zuo et al., 2012). The degree is the sum of edges on a node, and measures the direct interaction of a node with others in the network. Variations of this, such as the average shortest path length between a node and all others, which is termed closeness centrality have also been proposed (Freeman, 1979). Another measure in the analysis of brain networks is betweenness centrality, which is the fraction of all shortest paths in a network that pass through a given node (Freeman, 1977; Cheng et al., 2012; Kuhnert et al., 2012). An area of high betweenness centrality is commonly interpreted as lying on a significant informational pathway that connects different brain regions. Hagmann *et al* (2008) used this measure to identify central hub nodes within the structure of the brain. Such hub nodes have proved significant in the centrality of functional networks (He et al., 2009; Lohmann et al., 2010).

2.6 Dynamical Complexity and Topological Structures

The extent to which complex dynamics are dependent upon network structure has been the focus of much research (Bullmore and Sporns, 2009; Sporns, 2013). Sporns and Kötter (2004) gained insights into the rules governing the structure of complex networks by investigating their composition from smaller network building blocks, called motifs. A motif may consist of a set of brain areas and pathways that can potentially engage in different patterns of interactions depending on their degree of activation, the surrounding neural context or the behavioural state of the organism. Sporns and Kötter distinguish between structural motifs that quantify anatomical building blocks, and functional motifs that represent elementary processing modes of a network. Their assumption is that a particular structural motif can support a repertoire of functional motifs that may, or may not, be called upon for

neuronal computations. They conclude that a large repertoire of functional or effective circuits are obtained because the connection patterns of real brain networks maximise functional motif number and diversity, while at the same time they minimise the number and diversity of structural motifs, thus promoting efficient assembly and encoding. By rewiring random networks and imposing a cost function that maximises functional motif number, they generated network topologies that exhibit the small-world attributes of real brain networks.

Sporns (2013) reports on a growing body of work that draws attention to how the balance between structural segregation and integration is essential for the operation of the brain networks underlying cognitive function. A study by Sporns and Tononi (Sporns and Tononi, 2001) explored the simulated generation of different networks with differing dynamical properties of integration, segregation and complexity. Complexity captures the extent to which a system is both functionally integrated such that large subsets tend to behave coherently, and functionally segregated such that small subsets of the system tend to behave independently. They found that networks optimised for complexity separated into densely connected clusters with fewer connections between the clusters. They compared the networks optimised in simulation for functional complexity and cortical connection matrices derived from real datasets, and found a significant overlap in their structural and functional characteristics (Sporns, Tononi and Edelman, 2000).

Modular community structured networks akin to those found in the brain have also been used to study complex dynamics in systems of phase lagged, delayed and pulse coupled oscillators (Shanahan, 2010; Wildie and Shanahan, 2011). These systems exhibit interesting phenomena such as: metastability, chimera-like states and coalition entropy. Metastability is quantified by the variance of synchrony within an individual oscillator cluster over time, averaged for all clusters in the system, and so characterises the tendency of a system to continuously migrate between a variety of synchronous states. Fixing time and calculating the variance across clusters gives an index of how chimera-like the system is, that is to say the level of spontaneous partitioning into synchronised and desynchronised subsets (Abrams and Strogatz, 2004). Coalition entropy measures the variety of metastable states entered by a system of oscillators and is calculated from the number of distinct states the system can generate and the probability of each state occurring. As a collection, these measures capture the ability and tendency of a system to fully explore the space of dynamic synchronous coalitions. In the afore-mentioned work by Shanahan, a key area within the oscillator network parameter space was identified where the combination of these measures is optimal.

2.7 Functional Dynamics and the Brains Anatomy

Anatomical structure is believed to have a major effect on functional dynamics. Brain regions that are connected by a large number of cortical projections are likely to be functionally related. Functional activity is thought to reflect underlying anatomical structure, and correlations have been found to exist between the degree of structural connectivity and resting-state functional activity (Honey et al., 2009; van den Heuvel et al., 2009). The relationship between the two is, however, very complicated, and the extent to which the dynamical properties of the brain are constrained by structure still remains unclear (Honey, Thivierge and Sporns, 2010).

Cabral *et al* (2011) used a structural connectivity matrix between 66 regions of the human brain, downsampled from the high resolution connection matrix of 998 regions of interest obtained and described in Hagmann *et al* (2008), to study functional dynamics between brain regions. They performed simulations in which simple Kuramoto oscillators modelled the intrinsic oscillatory dynamics of functional neural areas set within a human connectome. They report the presence of complex dynamics and describe metastability, the formation of transient synchronous coalitions of brain areas. They further demonstrated that the resulting phenomena reproduced the fluctuations in spontaneous blood oxygen level dependence (BOLD) connectivity from human fMRI resting state data obtained *in vivo*. Cabral *et al* (2013) have extended this work to model the faster oscillatory dimension revealed by MEG recordings, and explore metastable partial synchronisation in functional connectivity in the 8-30 Hz range. They hypothesise that fluctuations in the synchrony degree may also modulate the oscillators' frequency, leading to frequency-specific amplitude fluctuations, making the link between fMRI and MEG expressions of resting-state activity. During periods of lower synchrony, brain areas oscillate in the gamma-band (30-80 Hz), implying an acceleration of the temporal dynamics which have been shown to cause BOLD signal increases at the local level. During periods of synchronisation, the temporal dynamics is slowed down (10.5-21.5 Hz) and consequently the BOLD signal is decreased.

Hellyer *et al* (2014) investigated the relationship between brain network activity, metastability, and cognitive state in humans. Their results from a choice reaction time task suggest that the balance of activity in the frontoparietal control/dorsal attention network and default mode network might control global metastability, providing a mechanistic explanation of how attentional state is shifted between an unfocused/exploratory mode characterised by high metastability, and a focused/constrained mode characterised by low metastability. They further simulated the neural dynamics arising in these

distinct cognitive states using a computational model of brain function consisting of a network of Kuramoto oscillators.

2.8 Discussion

The work of Cabral and others exemplifies a trend to use simple oscillator models such as the Kuramoto model (Kuramoto, 1984) as abstraction representations of oscillating neural populations (Breakspear, Heitmann and Daffertshofer, 2010). However, it is unclear how accurate and appropriate such models are for reproducing the behaviour of connected populations of oscillating neurons. Clearly a study using lower level neural models of oscillating neural populations would provide stronger evidence in support of metastability. Such a study should assess any intrinsically driven movements between transient attractor-like states that result in episodes of synchrony and desynchrony between functional areas. Further to this, a study of the structural requirements that promote metastability would provide greater insight into the phenomenon. These issues will be addressed in the forthcoming chapters, beginning with the next chapter, which presents an investigation of oscillation in the brain that are caused by inhibitory inter-neurons.

3 Modelling Brain Oscillations

This chapter reviews the mechanistic causes of oscillation in the brain, leading to the proposal and examination of a simulation model for oscillation to be used in later experiments.

Oscillations are formed by the repetitive synchronous firing behaviour of many different neurons within a local population. The neurons in a population collectively fire producing a spatio-temporal population code, which is often triggered by a stimulus they have learnt to respond to. There are many different neuron types in the brain, and the firing characteristics of individual neurons that have different types depend on their internal dynamics. Given that neurons can have different firing behaviour in response to learnt stimuli, and that the collective firing of different neurons is the cause of oscillation, the relationship between individual neuron behaviour and collective oscillatory dynamics requires examination.

After an initially exploring of the causes of oscillation in the brain, and proposing an architecture to model this, a study of the collective oscillatory behaviour of neurons given their firing characteristics and internal dynamics is carried out in the context of the proposed architecture. The chapter closes with a review of the results and proposes an experimental configuration for the studies detailed in later chapters given the findings presented.

3.1 Oscillation in the Brain

Studies have shown that groups of neurons firing together rhythmically can occur because of common input from a pacemaker neuron or due to the intrinsic firing patterns of excitatory principle cells. The latter is more common both in the cortex and the hippocampus and is an emergent property of interactions between excitatory principal cells and inhibitory inter-neurons (Whittington et al., 2000). Hasenstaub *et al* (2005) demonstrate both *in vivo* and *in vitro* that intracellular injection of

synaptic current, which have the same characteristics as recorded natural excitatory and inhibitory pre-synaptic potentials, reveal that inhibitory pre-synaptic potentials play an important part in modulating the timing and probability of action potential generation in pyramidal cells. Given that inhibition plays a significant role in oscillation, a review of the function of inhibition in oscillation follows.

3.1.1 The Function of Inhibition

An inhibitory input to a neuron may not always come from an inhibitory source neuron. This is because incoming action potentials to a neuron are subject to reversal. When neurotransmitters reach the postsynaptic neuron of a synapse, these molecules can bind to either ionotropic receptors or metabotropic receptors that are clustered in a protein-rich portion of the postsynaptic cytoskeleton called the postsynaptic density. Each of these two channels has a specific reversal potential. Reversal potentials dictate whether the incoming signal has its normal function, or whether the effect is reversed so that excitatory signals become inhibitory and vice versa. Gamma-aminobutyric acid (GABA) is one of the two neurotransmitters that facilitate inhibition in mammalian brains. The synaptic reversal potential of GABA_A receptor currents varies widely among cell types, and probably even between compartments within same cell. When the synaptic reversal potential is below the resting potential of the neuron, inhibition will be hyperpolarising. GABAergic synapses can be excitatory if the synaptic reversal potential is above the action potential threshold of the neuron (Andersen et al., 1980). Clearly, if the configuration of the synaptic reversal potentials plays an important role in the control of inhibition, and inhibition plays a vital role in oscillation, then the configuration of reversal potentials must be considered in the proposed model of oscillation.

It is interesting to note that if the synaptic reversal potential lies between the resting potential and the action potential threshold of the neuron, then GABAergic synapses will be *shunting* (Alger, 1979; Andersen et al., 1980). The term shunting inhibition derives from the fact that synaptic conductance short-circuits currents that are generated at adjacent excitatory synapses by reducing locally input resistance causing subsequent excitatory postsynaptic potentials to be reduced. Shunting synaptic currents has a selective inhibitory influence on locally synapsed afferents while aiding more remotely placed excitatory synapses. Anderson *et al* (1980) propose the term ‘discriminative inhibition’ for this behaviour.

Fries (2009) states that, inhibitory inter-neurons inhibit each other and excitatory neurons resulting in synchrony. For efficient synchronisation, inhibition is generally assumed to be hyperpolarising (Bartos *et al.*, 2002). Whilst the inhibitory model used in this work uses synaptic reversal potentials, in line with Bartos *et al.*, reversal potentials are configured such that inhibition is generally hyperpolarising. This mechanism by which inhibition causes synchronous firing is called *pyramidal inter-neuronal gamma*, which shall now be discussed.

3.1.2 Pyramidal Inter-Neuronal Gamma

Variations of the *pyramidal inter-neuronal gamma* (PING) mechanism can give rise to both fast gamma oscillations, as well as slower oscillations such as theta in the cortex and the hippocampus (Nyhus and Curran, 2010). Excitatory neurons drive the entire local network, including inhibitory inter-neurons. The most strongly driven inhibitory neurons will fire first and provide inhibition to numerous other inhibitory neurons. The inhibitory effect on all these neurons will wear off at approximately the same time. Affected inhibitory neurons will then fire roughly together, causing large numbers of inhibitory neurons to be entrained to a rhythm within just a few oscillatory cycles (Vida, Bartos and Jonas, 2006).

This rhythmically synchronised inhibition also influences the network's excitatory neurons with a fast and strong synaptic input (Papp *et al.*, 2001). Whittington *et al.* (2000) explain that inhibition must wear off at a rate that creates a window for when excitatory cells are ready to fire again, and excitatory cells must be able to fire fast enough so that they fire every time inhibition wears off and the window is open.

Rhythmic firing can happen at different temporal scales depending on the duration of inhibition, and different spatial scales depending on the length of the connection from inhibitory inter-neuron to excitatory cell. Bartos *et al.* (2002) report that if the synaptic decay time constant is fast, then high coherence is mainly reached with high excitatory drive resulting in a frequency in the upper gamma band (37–79 Hz). Conversely, if the decay time constant is slow, high coherence is mainly reached with low excitatory drive resulting in a frequency in the lower gamma band (18–55 Hz). Decreased conduction delay also increases the frequency, whereas increased conduction delays reduce the frequency. However, both decreased and increased conduction delays can reduce the average coherence. The particular effects of inhibition are exemplified by the difference in generation of theta and gamma oscillations:

“The neural mechanism responsible for gamma and theta rhythmic firing in the cortex and the hippocampus is similar, but primarily differs in the speed of inhibitory neurons. Whereas gamma rhythmic firing occurs through the interaction of excitatory principle cells and fast basket cell inhibitory inter-neurons acting on fast Gamma-aminobutyric acid (GABA_A) receptors, theta rhythmic firing occurs through the interaction of excitatory principle cells and slow stellate cell inhibitory inter-neurons acting on slow GABA_A receptors.” (Nyhus and Curran, 2010)

3.1.3 Neural Coding, Plasticity and Oscillation

Neural coding is a subject of immense debate. Traditionally, neuron firing rates have been predominant in coding theory. However, modern recording techniques have enabled the detection of both synchronous and correlated firing with high temporal resolution. Experimental results suggest that important functions that are beyond the scope of rate coding are manifest in spatio-temporal spike coding, which in turn is based on individual spike timing and synchrony (Masuda and Aihara, 2004). Groups of synchronous firing neurons are also believed to encode object information (Masuda and Aihara, 2004). Neurons as well as groups of synchronous firing neurons come to encode object information by adapting the weights on incoming synapses. A brief review of the relationship of synaptic plasticity to oscillation follows. We shall study plasticity in more detail later in the chapter.

Spike timing dependent plasticity (STDP) is an empirically derived learning rule for spiking neural networks (Abbott and Nelson, 2000). Long-term plasticity depends on the exact timing relation of the spikes from pre- and post-synaptic neurons. In a pairwise system, a pre-synaptic spike followed by a post-synaptic spike induces long term potentiation (LTP) of the synapse, whereas a post-synaptic spike followed by a pre-synaptic spike causes long term depression (LTD) of the synapse. STDP reduces the latency of a neuron's response to a given input (Masquelier, Guyonneau and Thorpe, 2008). Due to this, the phase of network oscillation upon which a given neuron fires may be determined by the strength of the synaptic input to that neuron. Increasing synchrony between neurons may induce LTP of the synapse connecting them. As a result, the stimulus properties that have activated these neurons are associated with each other, and synchrony is also further increased (Axmacher et al., 2006).

Many observations support the view that gamma synchronisation (30-100 Hz) binds the neurons that represent the attended features of a stimulus (Fries et al., 2001; Engel and Singer, 2001; Fell et al., 2003). Slower oscillations do not have a small enough time difference between the firings of the pre-

synaptic and post-synaptic neurons to fall within the STDP time window within which plasticity is applied. Conversely, faster oscillations have more than one cycle during the STDP window. Hence, the post-synaptic neuron may receive inputs both before and after having generated a spike, and thus cause both LTP and LTD (Axmacher et al., 2006).

Given the significance of the relationship between oscillation and STDP, the analysis of the chosen oscillator model presented later in this chapter considers oscillators that have learnt to respond to a given stimulus through STDP.

3.1.4 Oscillation and Neuron Types

The fundamental element in oscillation is the spike elicited by a neuron. As neurons have different discharge patterns, and the spikes discharged by many neurons are the basis of oscillation, we must now consider effect of neuron discharge patterns on oscillation.

Neuron discharge patterns can be classified according to their electrophysiological characteristics. Some neurons are typically tonically active. Neurons that fire in bursts are called phasic. Some neurons are notable for their high firing rates. Hodgkin (1948) was the first to classify neurons by their firing characteristics. He identified them as follow:

Type I: Action potentials can be generated with arbitrary low frequency depending on the strength of input current.

Type II: Action potentials are generated at a certain frequency that is relatively insensitive to the strength of incoming current.

Type III: A single spike is generated in response to a pulse of current. Tonic spiking can only be generated for extremely strong currents.

During cortical gamma rhythms, single neurons recorded *in vivo* display irregular firing which is at a slower rate than the frequency of the oscillation recorded in the local field potential (Colgin et al., 2009; Geisler, 2005; Kondgen et al., 2007; Pesaran et al., 2002). This effect has been observed for both pyramidal and inhibitory inter-neurons (Csicsvari et al., 1999). Coherent network oscillation and intermittent firing of individual neurons have both been observed as a result of high levels of input noise to the network acting in concert with strong recurrent inhibition (Brunel, 2000; Brunel and Wang, 2003). The latter highlights that a learnt response is not necessary in order to generate

oscillation but that recurrent inhibition is required. Further to this, intermittent firing of individual neurons is a key characteristic of oscillation.

Given the irregular and slower than the gamma frequency rate of firing of individual neurons during population oscillation, the different classes of neuron firing identified by Hodgkin may affect the coherence of oscillation in different ways. It is therefore pertinent to consider the effect of different neuron models during population oscillation, and we shall do so in section 3.7.

3.2 Oscillator Model Choice

3.2.1 Different Ways of Modelling Neural Populations

There are several ways to model neural oscillations within brain regions and the interaction of oscillations between brain regions. The simplest and most abstract method is that of a system of weakly coupled phase oscillators. The Kuramoto model (Kuramoto, 1984) is one such example. It captures the activity of a population of neurons by its circular phase alone and ignores the amplitude of oscillations. It modifies the phase over time of each oscillator by changing it according to the oscillators' own intrinsic frequency and the effect of phase difference and coupling strength between it and the other oscillators that connect to it. Simulations using the Kuramoto model are very popular, and have for example been used with realistic long-range cortical connectivity and time-delayed interactions to model the emergence of slow patterned fluctuations that are comparable to resting-state BOLD functional maps, which have been measured using fMRI (Cabral et al., 2011).

Neural field models are a differently abstracted model important in studying neural oscillations. For example, the Wilson-Cowan model (1972) approximates the mean firing rate of a group of excitatory and inhibitory neurons and their average interactions. It models the activity of large numbers of neurons as a spatially continuous neural network. The model has been used to provide mathematical descriptions of neural oscillations and EEG rhythms (Barlow, 1993; Jirsa et al., 2002).

A spiking neural network models a population of physically interconnected neurons each of which has internal neuron dynamics that produce spiking behaviour comparable to biological neurons. The dynamics in the network as a whole arise from interactions between individual neurons in the circuitry. Oscillations arise due to the synchronisation of spiking activity that results from these local interactions. Models of interacting pyramidal cells and inhibitory inter-neurons have been shown to generate brain rhythms such as gamma activity (Whittington et al., 2000).

3.2.2 Justification of Model Choice

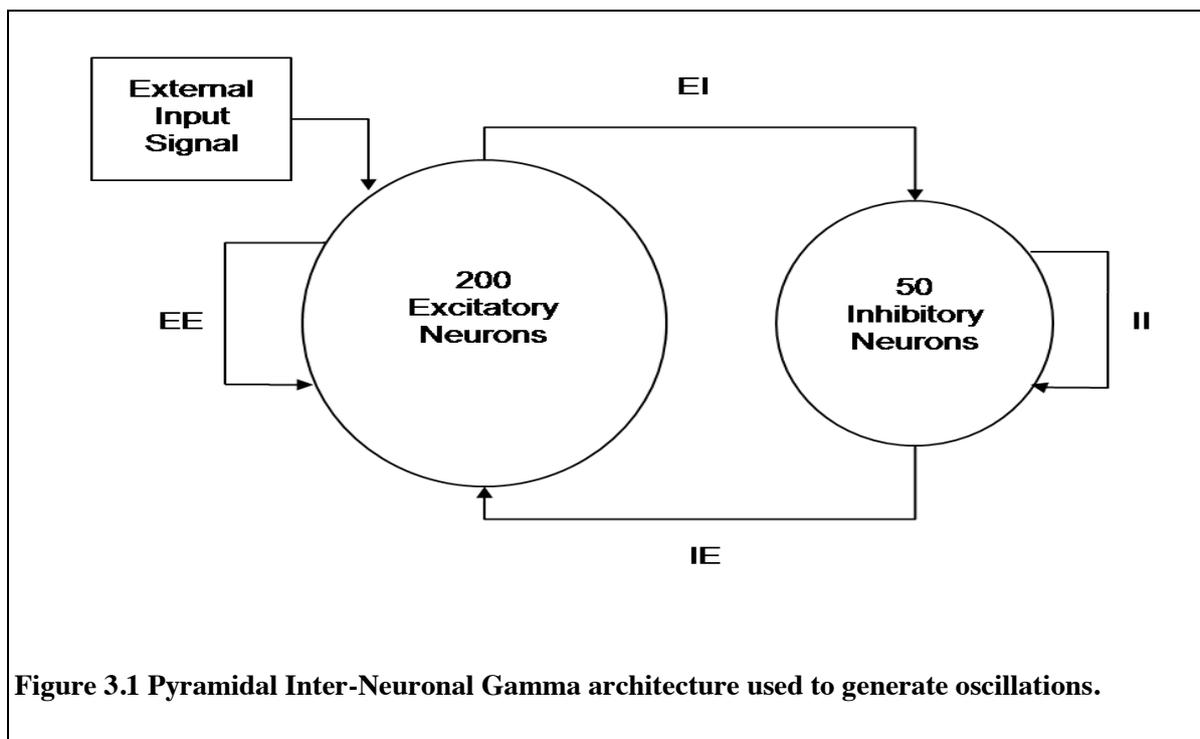
The central aim of this thesis is to understand the relationship between neural dynamics and structural topology. The highly abstract system of weakly coupled oscillators captures the topological structure of the network between oscillators, but does not represent the distinction between inhibitory and excitatory neurons within a population and the connectivity between them. The latter as we shall see in chapters 4 and 5 plays an important role in controlling the dynamics between neural oscillators in a large scale network. Whilst neural field models capture the activity of both inhibitory and excitatory neurons, they sacrifice the details of connectivity such as the strength of synapses. The latter is a key consideration when attempting to understand the relationship between topology and dynamics, and will be expanded upon in chapters 4 and 6.

Given the above mentioned requirements, the work presented uses the more detailed spiking neuron models. In making this choice, a compromise must be made regarding the number of neurons in each oscillator given the computational burden of simulating at such a high level of detail. The chosen number of neurons per oscillator is 250. The largest neural network simulations run in this work required 64 neural oscillator nodes each of which consisted of 250 neurons. This resulted in 16,000 neurons and $\approx 36,256,000$ synapses, and entailed an immense computational burden across the entire parameter space sweep in the experiments. 250 neurons is a very small number, especially when the number of neurons in a neural population in the brain that we wish to model may number into the millions. However, the belief is that the dynamical properties that arise from the network structure extrapolate to larger scales.

3.2.3 Model Configuration

The neural oscillators used in this work conform to the PING architecture. Whilst the general PING architecture is well understood, the specific details required for particular oscillatory frequencies vary and involve a large space of parameter values within the general PING framework. In order to provide a wide range of different intrinsic oscillatory frequencies for the neural oscillator nodes used in the experiments presented, it was decided to obtain these parameter values by use of a genetic algorithm (described below). The genetic algorithm evolved, within biologically plausible bounds, the oscillatory frequency of 30 Hz for the each of the neuron models used. The evolutionary mechanisms were constrained so that each neural network was evolved in accordance with the general PING architecture mentioned above. All neural populations for the PING oscillators used an excitatory layer of 200 neurons and an inhibitory layer of 50 neurons. The excitatory layer drives the

entire network and so is the only one to receive external input. The input is generated from a Poisson process with parameter $\lambda = 4.375$. The inputs were scaled using S in order to provide sufficient stimulus to induce firing. The value of S differed for each neuron model. For the QIF neuron $S = 45$, for the Izhikevich neuron $S = 50$, and for the HH neuron $S = 2.5$. The networks were wired up with connections between inhibitory neurons (II), from excitatory to inhibitory neurons (EI), from inhibitory to excitatory neurons (IE), and from excitatory to excitatory (EE). The PING architecture used is illustrated in figure 3.1. In addition to the synaptic weight, a scaling factor of 7 was used on all synaptic current in the oscillatory populations to simulate networks of a larger size than could feasibly be modelled, given the number of individuals in a population and the number of generations in an evolutionary run, as well as the large number of simulation runs for the networks of neural PING oscillators in the later experiments.



3.3 Modelling Neurons

Neurons can have different internal dynamics which in turn dictate different firing characteristics. Therefore the choice of neuron model used for the experiments in this work may affect the results obtained. In order to provide a balanced perspective, in this chapter different neurons with different behavioural characteristics are assessed in the context of PING oscillation.

Neurons can be described by their bifurcation properties and the period of their super-critical limit cycle. Type I neurons have a saddle node bifurcation and have a zero frequency super-critical limit cycle. Type II neurons can have a saddle node or an Andronov-Hopf bifurcation, but have a fixed frequency super-critical limit cycle. For saddle node bifurcation neurons, the resting state of the neuron is at a stable equilibrium point. Incoming spikes are integrated and move the neuron voltage to a saddle point at which it enters a super-critical limit cycle and produces a spike. For Andronov-Hopf bifurcation neurons, there is a small sub-critical limit cycle around a stable fixed point. The position the neuron is at in this limit cycle will dictate the effect that incoming spikes have upon the limit cycle, and in turn the effect of future incoming spikes. The internal neuron dynamics can therefore resonate to the incoming signal. When the sub-critical limit cycle approaches a large amplitude spiking limit the neuron enters a super-critical limit cycle which elicits a spike (Izhikevich, 2007).

In order to validate the work presented, comparative experiments were performed using different neural models. By using three different neuron models, both Type I and Type II spiking properties, as well as saddle node and Andronov-Hopf bifurcations were able to be assessed.

3.3.1 Quadratic Integrate-and-Fire Neurons

The quadratic-integrate-and-fire (QIF) model (Latham et al., 2000) displays Type I neuron dynamics (Ermentrout, 1996). The time evolution of the neuron membrane potential is given by:

$$\frac{dV}{dt} = \frac{1}{\tau}(V - V_r)(V - V_t) + \frac{I}{C}$$

where V is the membrane potential, with V_r and V_t being the resting and threshold values respectively. C is the capacitance of the cell membrane. τ is the membrane time constant such that $\tau = RC$ with R being the resistance. I represents a depolarising input current to the neuron.

An action potential occurs when V reaches a value V_{peak} at which point it is reset to value V_{reset} . The QIF model is equivalent to the theta neuron model described by Ermentrout and Kopell (1986) if one sets the reset condition $V_{peak} = \infty$ and $V_{reset} = -\infty$. Like Börgers and Kopell (2005) values used are $V_r = V_{reset} = 0$ and $V_t = V_{peak} = 1$, which reduces the above equation to:

$$\frac{dV}{dt} = aV(V - 1) + \frac{I}{C}$$

Here $a=1/\tau$ and is set to the value 2 for all experiments carried out in the paper. When working with the QIF model it is assumed that the membrane potential is between the biologically plausible range $V_r = -65$ mV and $V_t = -45$ mV.

3.3.2 Izhikevich Neurons

The Izhikevich (IZ) neuron model (Izhikevich, 2003) is a two variable system that can model both Type I and Type II neurons depending upon how it is parameterised. The model simulates a refractory period, which is an advance on the QIF model when a Type I neuron is simulated. The time evolution of the model is defined as follows:

$$\frac{dV}{dt} = 0.04V^2 + 5V + 140 - U + I$$

$$\frac{dU}{dt} = a(bV - U)$$

$$\text{if } V > 30, \text{ then } \{V \leftarrow c, U \leftarrow U + d$$

$$\text{if } U > 15, \text{ then } \{U \leftarrow 15$$

I is the input to the neuron. V and U are the voltage and recovery variable respectively, and a , b , c and d are dimensionless parameters. For excitatory neurons, the values chosen for these are as follows: $a=0.02$, $b=0.2$, $c=-65+15 \times r^2$, and $d = 8-6 \times r^2$, where r is a value between 0 and 1 chosen from a uniform distribution. For inhibitory neurons, the values chosen for these are: $a=0.02+0.08 \times r$, $b=0.25-0.05 \times r$, $c=-65$, and $d = 2$, where again r is a value between 0 and 1 chosen from a uniform distribution. The chosen parameter values dictate that the Izhikevich neurons used are Type II neurons with a saddle node bifurcation. The extra term limiting U from going above 15 prevents over saturation of the recovery variable caused by high levels of input.

3.3.3 Hodgkin Huxley Neurons

The Hodgkin-Huxley (HH) model (Hodgkin and Huxley, 1952) is widely considered as the benchmark standard for neural models. It is based upon experiments on the giant axon of the squid. Hodgkin and Huxley found three different types of ion current: sodium (Na^+), potassium (K^+), and a leak current that consists mainly of chloride (Cl) ions. Different voltage-dependent ion channels

control the flow of ions through the cell membrane. From their experiments, Hodgkin and Huxley formulated the following equation defining the time evolution of the model:

$$C \frac{dV}{dt} = g_K n^4 (V - E_K) - g_{Na} m^3 h (V - E_{Na}) - g_L (V - E_L)$$

$$\frac{dn}{dt} = \alpha_n(V)(1-n) - \beta_n(V)n$$

$$\frac{dm}{dt} = \alpha_m(V)(1-m) - \beta_m(V)m$$

$$\frac{dh}{dt} = \alpha_h(V)(1-h) - \beta_h(V)h$$

C is the capacitance and n , m and h describe the voltage dependent opening and closing dynamics of the ion channels. The maximum conductances of each channel are: $g_K=120$, $g_{Na}=36$ and $g_L=0.3$. The reversal potentials are set so that that $E_K=-12$, $E_{Na}=115$ and $E_L=10.6$. The rate functions for each channel are:

$$\alpha_n(V) = \frac{(0.1 - 0.01V)}{e^{1.0 - 0.1V} - 1.0}$$

$$\beta_n(V) = 0.125 e^{\frac{-V}{80.0}}$$

$$\alpha_m(V) = \frac{2.5 - 0.1V}{e^{2.5 - 0.1V} - 1.0}$$

$$\beta_m(V) = 4.0 e^{\frac{-V}{18.0}}$$

$$\alpha_h(V) = 0.07 e^{\frac{-V}{20.0}}$$

$$\beta_h(V) = \frac{1.0}{e^{3.0 - 0.1V} + 1.0}$$

All work in this paper using the HH model adjusts the neuron resting potential from 0 mV of the standard HH implementation to the more accepted value of -65 mV (Gerstner and Kistler, 2002).

3.4 Modelling Synapses

3.4.1 Conductance Synapses

The synaptic model for simulations using the QIF neuron model and the IZ neuron model is a current synapse that simply multiplies the incoming spike by a synaptic weight:

$$I_j(t) = \sum_i w_{ij} \sum_k^n \delta(t - d_{ij} - t_{i,k})$$

where $I_j(t)$ is the input to neuron j and time t . w_{ij} is the synaptic weight from neuron i to neuron j , and d_{ij} is the synaptic delay from neuron i to neuron j . A list of all n spikes produced from neuron i during a simulation are denoted by their spike times $t_{i,k}$, where $k=1,2,\dots,n$. δ is a delta function applied to $t - d_{ij} - t_{i,k}$, such that adjusting the current time t by the synaptic delay d_{ij} identifies the spike production time at neuron i for which a spike is due to arrive at neuron j at time t . If $t_{i,k}$ matches this spike time then the delta function produces an output value 1.

3.4.2 Synaptic Reversal Potentials

The synaptic model for simulations using HH neuron model uses conductance in the synapses, and also uses reversal potentials to further scale incoming spikes. The latter model is as follows:

$$I_j(t) = \sum_i (V_{rev} - V_j) w_{ij} \sum_k^n \delta(t - d_{ij} - t_{i,k})$$

The additions to the previous conductance synaptic model are, V_{rev} which is the reversal potential, and V_j , which is the voltage of the target neuron. The reversal potentials for the model are set to the same values in all experiments. For excitatory inputs the reversal potential is set to 0 mV, and for inhibitory inputs the reversal potential is -70 mV. Not using a synaptic reversal model for the QIF and IZ models is equivalent to using a synaptic reversal model with reversal potentials set to $+\infty$ mV for excitatory neurons and $-\infty$ mV for inhibitory neurons.

3.4.3 Spike Timing Dependent Plasticity

STDP is an empirically derived refinement of the Hebbian learning principle for spiking neural networks. STDP displays strengthening of correlated groups of synapses, the basic feature of Hebbian learning, as well as other desirable features such as firing-rate independence and stability (Song, Miller and Abbott, 2000). Long-term plasticity depends on the exact timing relation, on the

time scale of milliseconds, of the spikes from the pre-synaptic neuron and the spikes from the post-synaptic neuron.

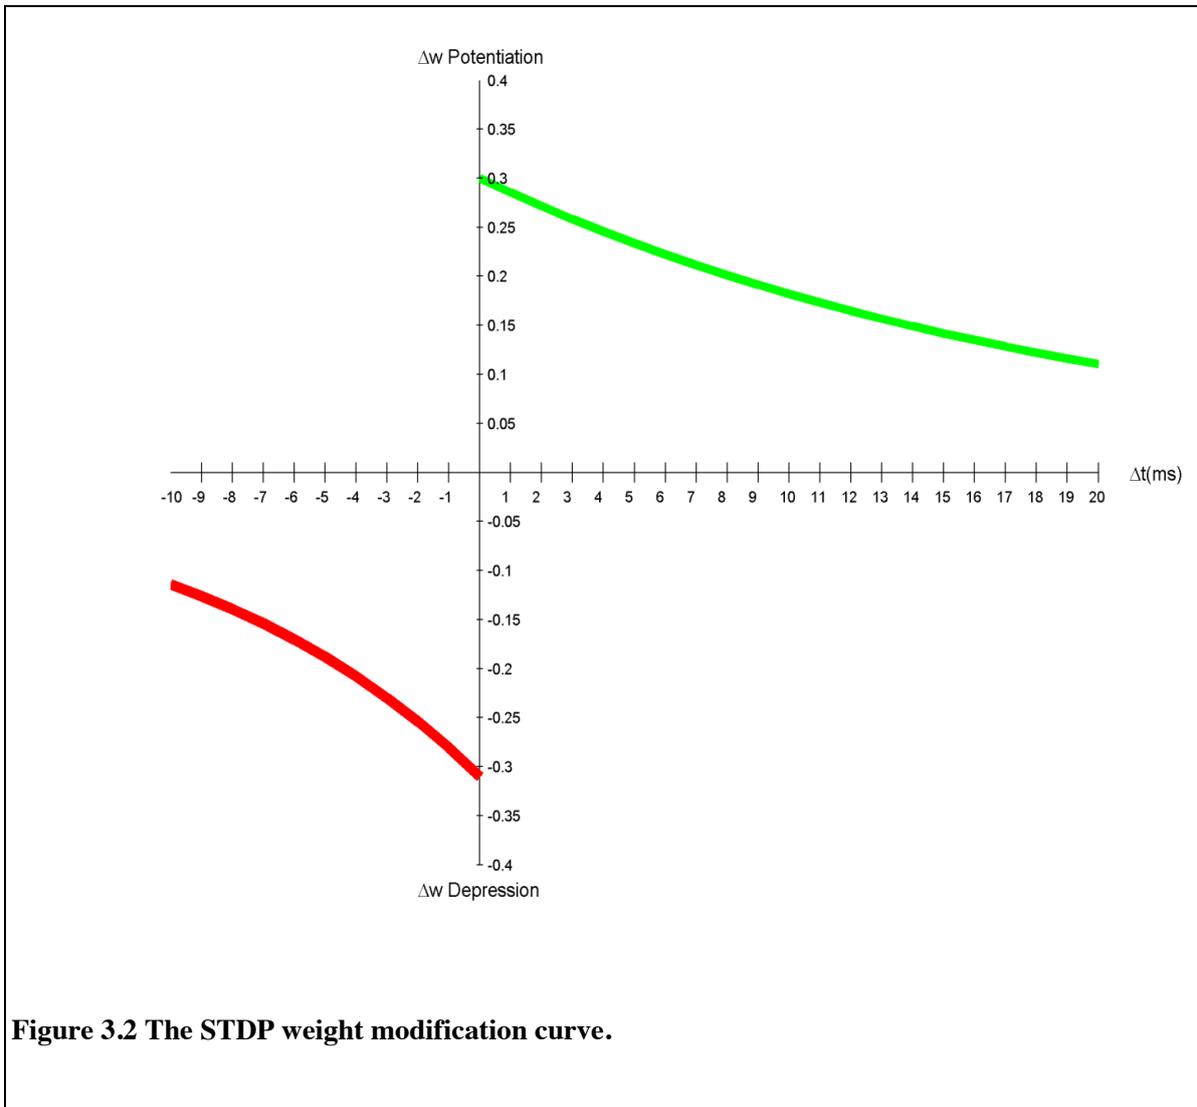
When the post-synaptic neuron fires at time t it initiates the synaptic weight update rule. The update rule considers pre-synaptic spike times ($t-\Delta t$) within a given window (τ). The update method used in this work is an ‘additive nearest neighbour’ scheme, in which only the spike temporally nearest the time of the post-synaptic spike is considered, and the weight change is not dependent upon the current weight value (Morrison, Diesmann and Gerstner, 2008). A pre-synaptic spike followed by a post-synaptic spike potentiates the synaptic weight, where as a post-synaptic spike followed by a pre-synaptic spike depresses the synaptic weight. The change in weight (Δw) is affected by the exponential of the time difference (Δt) and the learning rate constant (η):

$$\Delta w = \eta e^{\frac{-\Delta t}{\tau}}$$

In this work, for potentiation the learning rate value η is set to 0.3, and the window τ is set to 20 ms. For depression, the learning rate value η is set to -0.3105 and the window τ is set to 10 ms.

A post synaptic neuron firing rate is initially fairly insensitive to the timing of the pre-synaptic firing but instead responds to total input. As a result, there are roughly equal numbers of pre-synaptic action potentials before and after each post-synaptic spike. The STDP sampling of pre-synaptic spikes is based on the trigger of post-synaptic firing, and therefore it is averaged over the entire potentiating and depressing spiking influence of the pre-synaptic neuron around that time. Figure 3.2 shows a graph displaying the potentiating and depressing influences on STDP. The asymmetry in areas under the positive and negative portions of the STDP modification curve means that depression will be preferred. As the synapses are weakened, the post-synaptic neuron generates a more irregular sequence of action potentials that are more tightly correlated with the pre-synaptic spikes that evoke them. Therefore, post-synaptic action potentials are triggered through statistical fluctuation, at which time there tend to be more excitatory pre-synaptic spikes before than after a post-synaptic response (Song, Miller and Abbott, 2000). Song *et al* explain that:

“STDP thus modifies excitatory synaptic strengths until there is a sufficiently, but not excessively, high probability of a pre-synaptic action potential occurring before a post-synaptic spike. This causes the neuronal response to be sensitive to the timing of input fluctuations.” (Song, Miller and Abbott, 2000)



3.5 Generating PING Oscillators

In order to provide a wide range of different intrinsic oscillatory frequencies of neural PING oscillators, it was decided to obtain the PING architecture parameter values by use of a genetic algorithm. A genetic algorithm is a blind search and optimisation technique based upon the theory of natural selection (Holland, 1975). Parameters are encoded in a pseudo genome, and are used to instantiate an individual, in this case a neural network. A population of individuals are tested and scored for their fitness at performing the test. Based upon their fitness ranking, pairs of individuals are chosen to produce offspring for the next generation via crossover of their genomes. Mutation is

then applied to some parameters in the new offspring genome. As this process continues over generations, individuals in the population become optimised at performing the target task.

The parameters that were evolved in the generation of PING oscillators were the synaptic weights and delays. Both of these were generated during genome expression of each individual in each generation using an approximately Gaussian distribution, with the means and variances for the weights and the delays being the parameters in the genome evolved. The distribution is only approximately Gaussian as the weights were bound to evolve values between 0 and 1 for excitatory connections and 0 and -1 for inhibitory connections. Delays were similarly bound. Long delays are quite unrealistic for a cluster of neurons in which all neurons are anatomically close together. In the cortex synaptic latency ranges from 0.2 ms to 6 ms (Markram et al., 1997). In order to produce realistic results, excitatory delays were bounded between 1 ms and 10 ms. The inhibitory delays were allowed to have a maximum value of 50 ms to simulate the effect of the speed of slow acting inhibitory inter-neurons, the behaviour of which was otherwise not modelled.

PING architectures were evolved that were able to use STDP to learn to respond synchronously to a given stimulus but no other. To achieve this, after genome expression a network was first trained for a given amount of time by using STDP and given the stimulus it was supposed to respond to. The stimulus was generated using a Poisson process with parameter λ , and was scaled by S in order to provide sufficient input to induce firing (the values of λ and S are detailed in section 3.2.3). After training individuals were tested for 5000 ms of simulated time in which they again received the same stimulus. The fitness function for this test consisted first of taking the spike firing times of the excitatory population and converting it to a continuous time-varying signal. This was achieved by binning the spikes over time, and then passing a Gaussian smoothing filter over the binned data. Next a Fourier transform was performed on the mean centred signal to produce the frequency spectrum s of the signal. The main fitness term was calculated by creating a scaled Gaussian centred around the desired frequency f in the spectrum of the form:

$$\text{clip} = 20G\left(f, \frac{1}{1000}\right)$$

The frequency spectrum s was subtracted from this and normalised:

$$\text{fitness} = \frac{-|\text{clip} - s|}{\sum \text{clip}}$$

An extra penalty term was introduced to discourage frequencies outside the desired range. This was achieved by multiplying the frequency spectrum by -0.002 in the areas further away from the desired frequency whilst ignoring the area at and immediately around the desired frequency. The result was then normalised and added to $fitness_1$. The stimulus to be learnt was changed on each test so as to ensure that the networks developed the ability to learn to respond to any stimulus that they were trained upon.

A second test was performed in which an alternative random stimulus was applied for which it was desired that the network should not respond. A different fitness term was used to evaluate this second test. It consisted of ensuring the amplitude of the peak frequency response for that test was below 0.5 so as to discourage any firing. This was achieved by obtaining a frequency response spectrum as described above, and then locating the frequency with the highest amplitude A . Given this the second fitness term was calculated as follows:

$$fitness_2 = -\frac{20 - (A - 0.5)}{20}$$

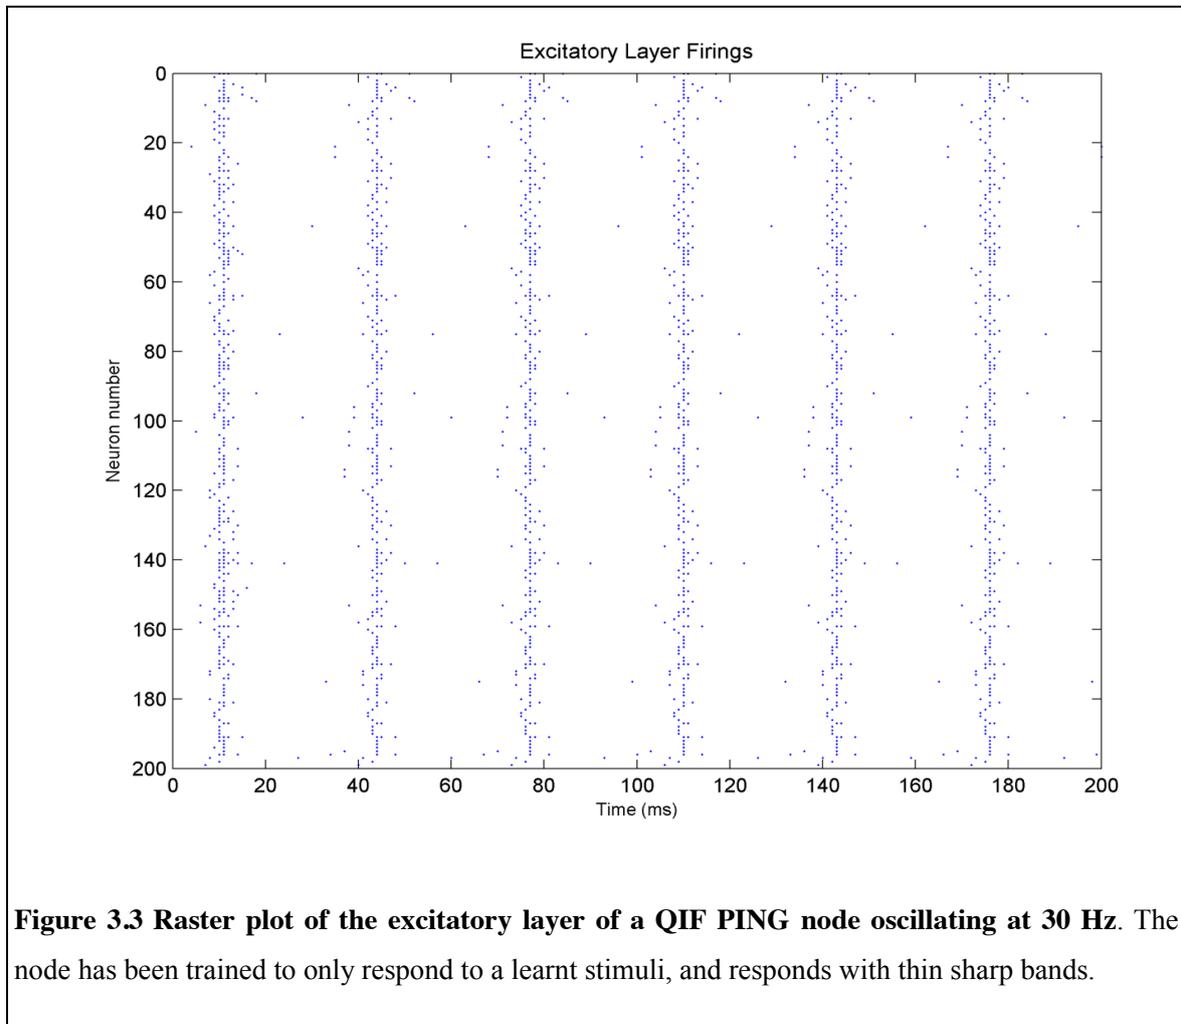
The two fitness terms were summed to give the overall fitness of the individual. The evolutionary population consisted of 20 individual genomes. After each individual was tested and rated for fitness individuals were probabilistically selected for the next generations' parents based upon its fitness ranking. Crossover was performed on parent genomes after which mutation was applied to the offspring with a probability of 0.1.

3.6 Hardware Acceleration

The largest neural network simulations run in this thesis required 64 neural PING nodes each of which consisted of 250 neurons. This resulted in 16,000 neurons and more than 32,256,000 synapses, and entailed an immense computational burden across the entire parameter space sweep in the experiments. To cope with this, the NeMo neural network simulator was used, which processes neurons concurrently on general purpose graphics processing units (GPUs) (Fidjeland and Shanahan, 2010). The NeMo software permits the addition of user plugins for neural models, which allowed the implementation of the QIF, IZ and HH models facilitating the work presented here.

3.7 Analysis of Neuron Models in PING Architecture

To analyse the effect that different neuron models have on the ability of a PING architecture to oscillate, the question of how the choice of neuron model affects the learning of oscillation through STDP is posed. STDP has previously been studied in relation to oscillations. Hosaka *et al* (2004) demonstrate oscillatory dynamics in a network of excitatory and inhibitory Nagumo-Sato neurons that has been trained using STDP with an external spatio-temporal stimulus that was repeatedly applied. In this work a synchronous response gradually emerges, and the synchrony becomes sharp as learning proceeds. The authors state that the generation of synchrony itself does not depend on the length of the cycle of external input, however they found that synchrony emerges once per cycle of the length of the external stimulus trained upon. This work is here extended by assessing the effect of

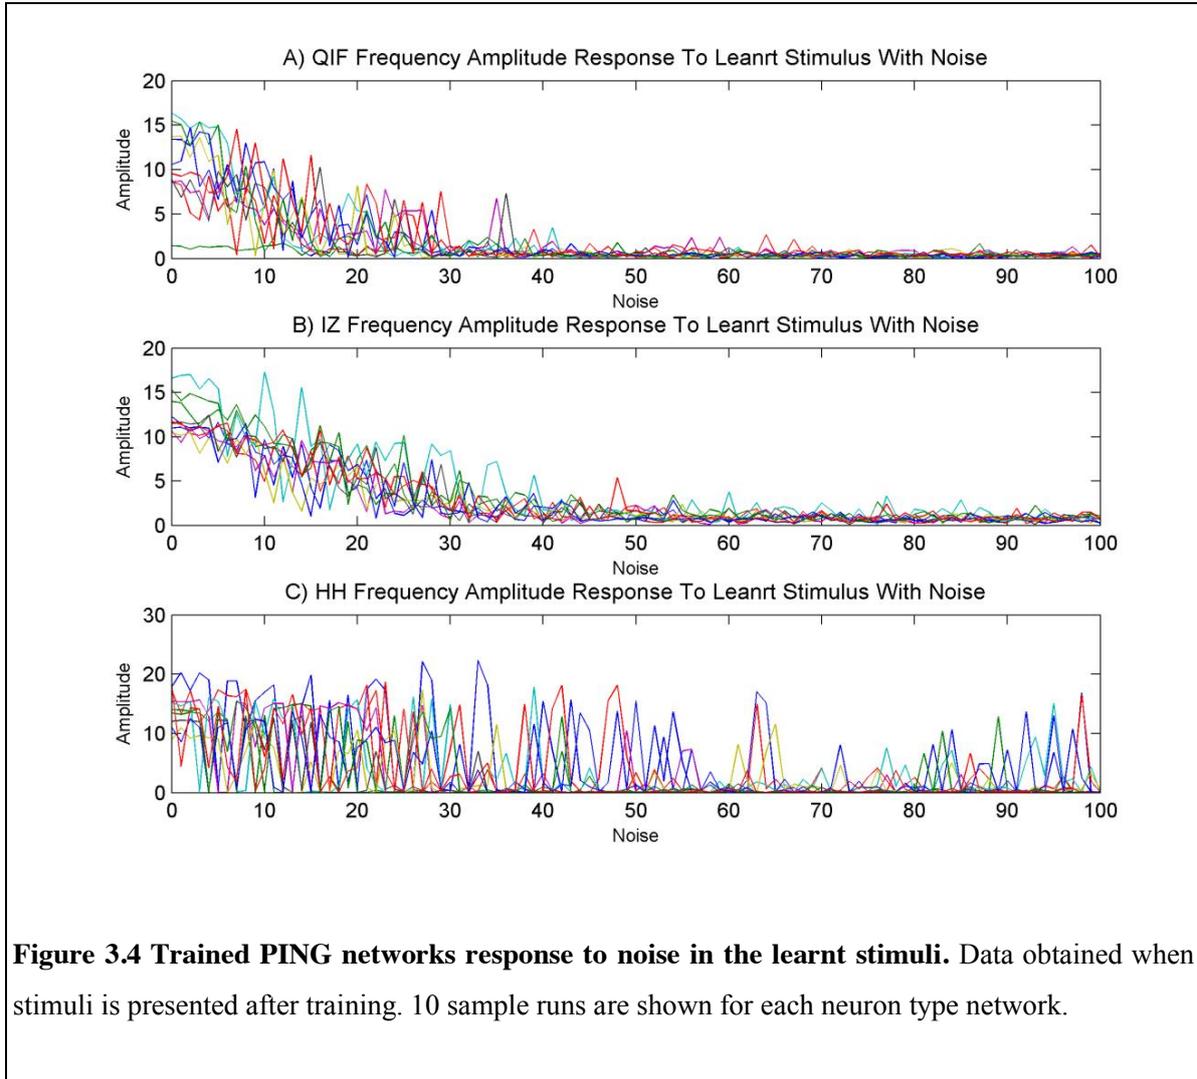


different neuron models in the same experimental context.

In order to explore the question, neural PING oscillators were evolved to oscillate in response to a learnt stimulus at 30 Hz for QIF, IZ and HH neuron models using the evolutionary process described above. The evolutionary process optimises the networks both for their ability to oscillate at the desired frequency in response to a learnt stimulus, as well as their ability to not respond to a non-learnt stimulus. Figure 3.3 shows a raster plot of the firings of the excitatory layer from the evolved QIF solution when it has been presented with a learnt stimulus after training. In accord with the finding of Hosaka *et al* (2004) the network fires regularly at the stimulus presentation, and has narrow and pronounced periodic bands. These thin bands appear approximately every 33 milliseconds giving the 30 Hz oscillation desired.

3.7.1 Oscillating Only in Response to a Learnt Stimulus

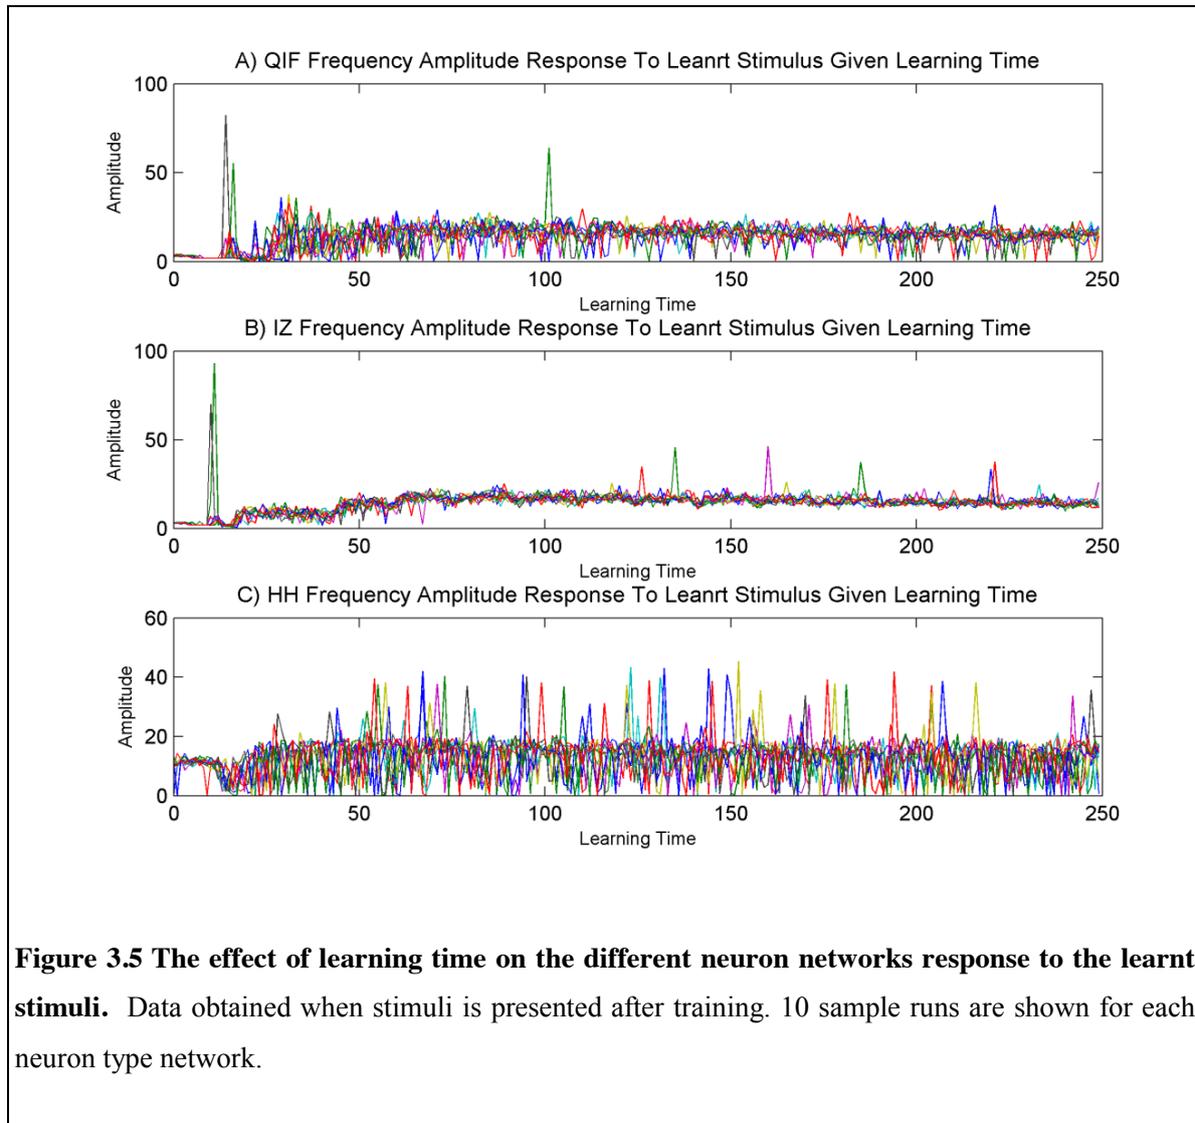
This section presents a study that ascertains the networks' ability to only respond by oscillating to the learnt stimulus and no other. Figure 3.4 shows how the evolved networks respond to noise in the learnt stimulus. To obtain these graphs, after evolving the PING networks, they were trained with a stimulus. After training, the stimulus trained upon was replaced with a percentage of random data drawn from the same distribution. The network was then tested with the new noisy stimulus for 5000 ms and the amplitude of the desired frequency response was measured. This was done for every noise percentage from 0% to 100%. The test was performed 10 times for the QIF, IZ and HH models. Each test is displayed as a different plotted line in the respective model graphs. A 100% noise level represents a completely different pattern from the training stimulus, for which it is expected that the network not to respond at all. The QIF network performs the best with regard to not responding to noisier stimuli, showing a gradual decline in the amplitude of the frequency response until it reaches a minimal response at 42% noise. Less than 0.5 amplitude implies that only a few neurons are firing, and hence no response is really being produced. With such low response to noise threshold the QIF model is the most highly selective to only its learnt stimulus. However, there is a wide variation in responses in the region of decline over the 10 test runs. This may be because the network is highly sensitive to particular afferents which may or may not appear in any of the 1000 runs performed at their particular noise level on a sample run. The IZ model has similar but not so much variation in the region of decline, and performs almost as well as the QIF model in regard to not responding to noisy stimuli with a minimal response achieved at around 50%. The HH model performs poorest with a much less pronounced frequency amplitude decline as noise rises, and also a very varied response



across the 10 tests and over all noise levels. Whilst this may mean that the network is very sensitive to particular parts of the learnt stimulus which may or may not appear in any of the sample runs due to noise, the difference to the QIF and IZ models behaviour which could also suffer from the same problem, suggest that the network is just inherently more volatile.

3.7.2 The Effect of Learning Time on Oscillation

This section presents a study that explores the effect that learning time has upon the networks. Each network type was trained on a stimulus for a given time t . The network was then tested on the stimulus for 5000 ms and the desired frequency amplitude response was measured. This was done for



every learning time t between 1 ms and 250 ms. The test was performed 10 times for the QIF, IZ and HH models. Each test is displayed as a different plotted line in the respective model graphs shown in Figure 3.5. Some unpredictability can be seen below 20 ms learning time for all models. This is expected relative to a 33 ms stimulus. After this point the amplitude rises for QIF and IZ models until it stabilises around 100 ms. The QIF model shows some variance over the 10 sample runs highlighting the volatility of the network, which in contrast does not appear for the IZ model. Beyond this time point there is some dip with not much variation. The learning time is very quick with only three stimulus presentations required to learn a maximal frequency amplitude response. The HH

model performs poorest with a less stable response throughout indicated by the high variance in the frequency amplitude over the 10 sample runs. Variance of amplitude in this case is not due to particular afferents being present on particular runs as all are present on all runs, and so is due to the network being more volatile. This confirms that the variance in amplitude in the noise experiments is also due to volatility.

3.7.3 The Effect of Stimulus Length on Oscillation

Figure 3.6 shows the effect of changing the length in milliseconds of the stimulus. To achieve this, a stimulus of length t was used to train the networks, after which they were tested for 5000 ms on the same stimulus. This was done for every stimulus length t between 1 ms and 100 ms. The experiment was performed 10 times for each neuron model. All learning stages for all stimulus lengths t had the same learning time. The frequency with the highest amplitude was then located. From the figure we can see that none of the models respond significantly to stimuli less than 10 ms long. Beyond this, the figure shows that for both QIF and IZ models, the length of the stimulus is roughly proportional to the frequency (f), with $f=1000/t$. The variations over the 10 sample runs are due to harmonic frequencies having higher amplitudes than the desired frequency response. The proportionality between the stimulus length and frequency cannot be seen for the HH model which is unable to use the same network architecture to learn to oscillate at different frequencies, given only a change in the stimulus length.

Having found a dependency on stimulus length, the inhibitory layer was removed from the networks. It was found that it made no noticeable difference to the performance of QIF, IZ and HH models. It can therefore be concluded that regular repetition of a stimulus to a network that has been trained using STDP will cause oscillation at the frequency of presentation. For the HH model this further means that whilst stimulus length is important in achieving the result, the tuning of other variables is necessary to achieve the desired oscillation.

The fact that oscillatory frequency is dependent upon the length of the presentation can be elucidated by the work of Masquelier *et al* (2008). They report that during learning with STDP, uncorrelated firings are depressed, whilst the synaptic connections with the afferents that took part in the firing of a neuron are potentiated. Further to this they report:

“Each time the neuron discharges in the pattern, it reinforces the connections with the presynaptic neurons that fired slightly before in the pattern. As a result next time the pattern is presented the

neuron is not only more likely to discharge to it, but it will also tend to discharge earlier.” (Masquelier, Guyonneau and Thorpe, 2008)

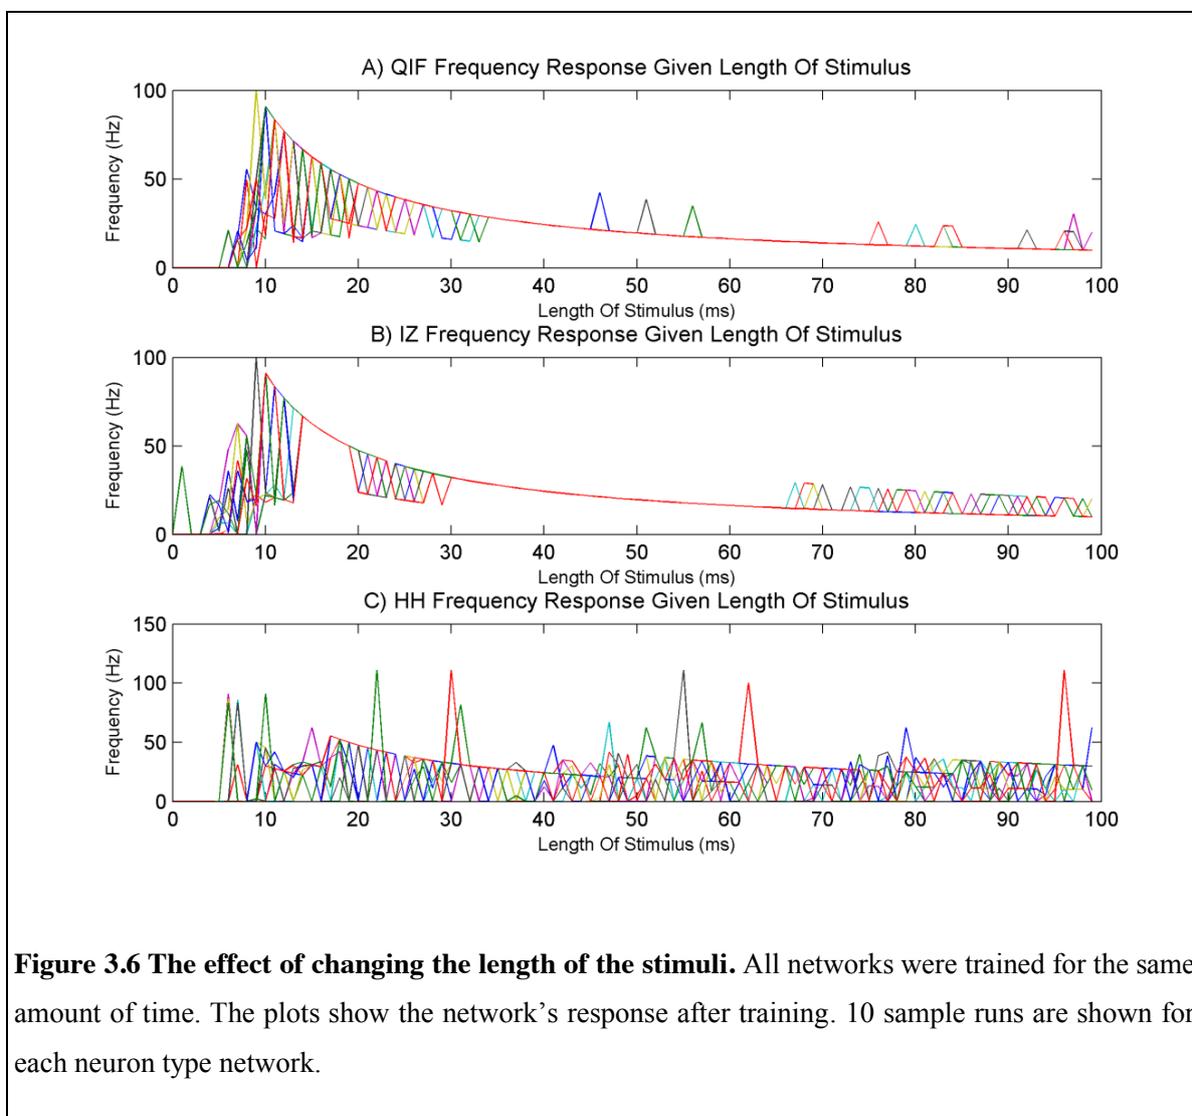
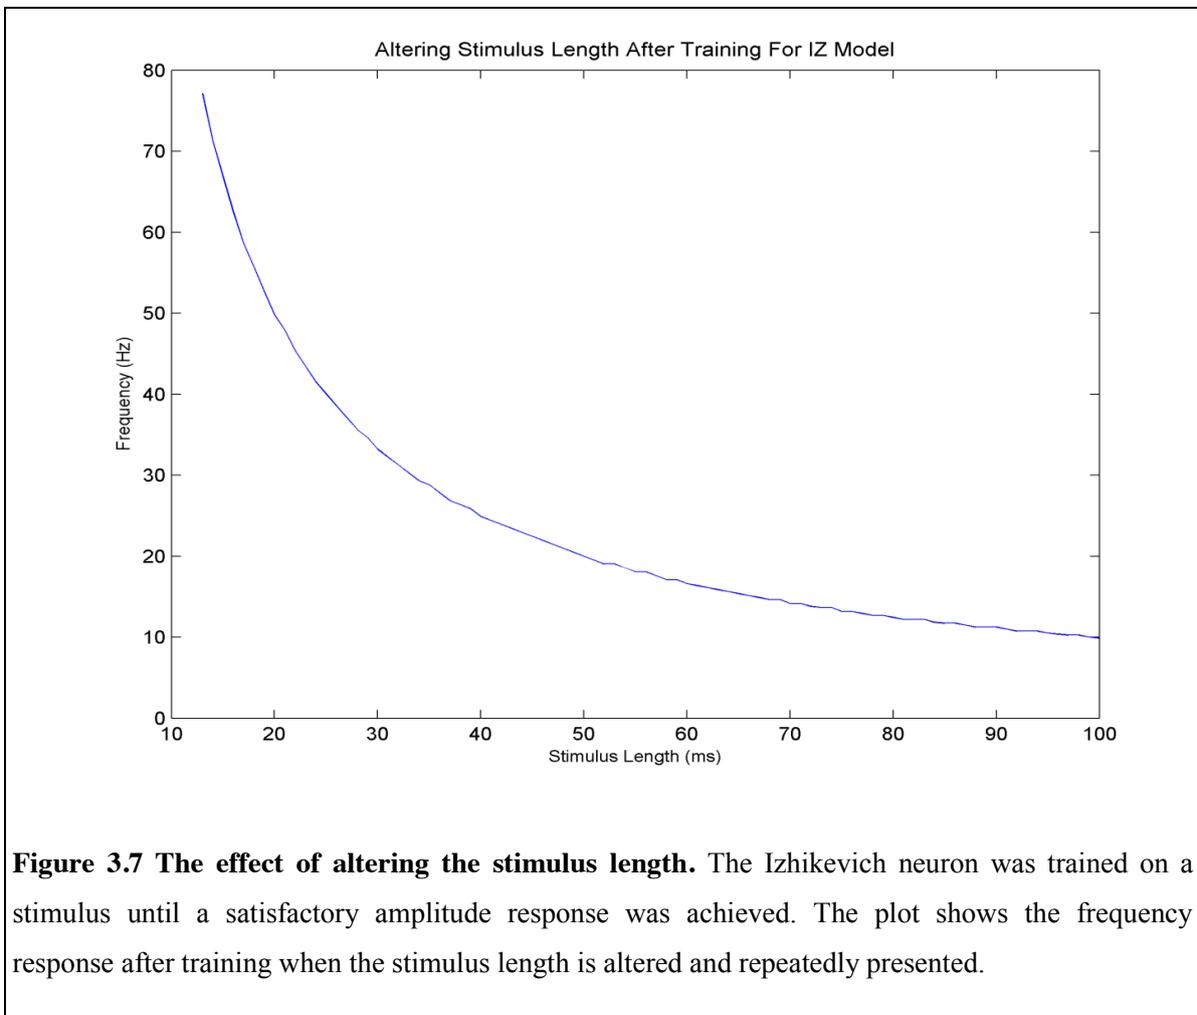


Figure 3.6 The effect of changing the length of the stimuli. All networks were trained for the same amount of time. The plots show the network’s response after training. 10 sample runs are shown for each neuron type network.

The fact that neurons learn to always respond to a particular stimulus implies that the regular repetition of a stimulus would cause the network to fire regularly at the stimulus presentation, and that this firing would become earlier and sharper, in the sense of producing narrower and more pronounced periodic bands, as learning proceeds. Hence, the resulting synchrony.

It follows from this that after an appropriate period of learning the frequency of the oscillation can be adjusted by simply altering the length of the stimulus, as it is only the beginning of the stimulus that

is required to induce firing. To test this hypothesis a stimulus of 100 ms was generated and the IZ network was trained on it repeatedly until a satisfactory amplitude response was attained. After this, the network was tested for 5000 ms with the stimulus but only using the first t milliseconds repeatedly. This was done for every value of t between 13 and 100 ms. As can be seen by the results shown in Figure 3.7, the hypothesis is correct. Hosaka *et al* (2004) state that in a network of excitatory and inhibitory neurons, STDP transforms a spatiotemporal pattern to a temporal pattern. However, from the evidence above it may be concluded that the resultant temporality obtained using STDP is not due to the network dynamics that result from the PING architecture, but is an artefact of repeated *periodic presentation* of a learnt stimulus. The network will respond ‘synchronously’ whenever the stimulus is presented.



3.8 Discussion

The analysis presented demonstrates that STDP generates robust synchronous responses. After learning, the networks are highly selective for their learnt stimulus and do not respond to other stimuli. Effective learning is possible within only three stimulus presentations. Given that the resultant oscillatory frequency is dependent upon the length of the presentation, the hypothesis that the frequency of the neural oscillator can be adjusted by simply altering the length of the stimulus was experimentally proven.

The frequency of the oscillations in PING architectures is caused by the feedback time over the EI/IE loop. From the experiments it is concluded that the resultant temporality obtained through learning is not due to the network dynamics that result from the PING architecture, but is an artefact of repeated periodic presentation of a learnt stimulus. The network will respond ‘synchronously’ whenever the stimulus is presented. Hence, it can be concluded that repeated post-learning presentation of a stimulus can override or interfere with the oscillations that would otherwise be caused by the PING architecture. A fast EI/IE loop will feed back and subside before the next learnt stimulus response. In this case oscillations from the periodic stimulus will take precedence over PING oscillations. However, a longer EI/IE loop could provide inhibition at the time at which the next learnt stimulus response is due and therefore cause interference. This suggests that over-training on a single stimulus may not be desirable. The motivation of this work is to understand the relationship between network structure and dynamics. The PING architecture is an important aspect of the network structure. Therefore, it is advisable that the assessment of PING oscillators in the larger experimental settings presented in the forthcoming chapters do not make use of pre-training oscillators using STDP, as this may cause interference.

Type I and Type II neuron behaviour, as exhibited by QIF and IZ models respectively, does not make any difference in learning to respond to the temporality of stimuli, nor to the robustness thereof. However, the HH model does not perform in the same manner. The difference in the HH neuron, which is classified as a Type II model, is the Andronov-Hopf bifurcation and the neuron’s synaptic reversal potential. The result is a less robust network that is also unable to use the same architecture to learn to respond to stimuli that have a variety of presentation times, but requires specific tuning of parameters to achieve desired oscillatory frequencies. It is interesting to note that the more biologically realistic model is less robust and requires specific parameter tuning, leaving open the

question of how the brain facilitates this in order to achieve a broad variety of oscillatory frequencies in response to different stimuli.

Regarding the application to future experiments presented in this work, it is clear that experiments must be performed with neuron models that exhibit both saddle node and Andronov-Hopf bifurcations. This will ensure that the work offers a balanced assessment of the use of PING architectures for oscillations in a larger experimental context, since these differences in internal neuron dynamics may impact on the results obtained.

4 The Collective Behaviour of Neural Oscillator Populations

This chapter discusses the phenomenon of synchrony in nature and neuroscience. Abstract models of oscillators are presented and their collective synchronous behaviour compared to neural oscillators by means of computer simulation.

Synchrony is a ubiquitous phenomenon in nature that is manifest when groups of oscillators are connected. When two distinct objects oscillate at different frequencies and phases they are desynchronised, but when they both oscillate at the same frequency and phase they are said to be synchronous. Synchrony can be likened two people hitting a drum at the same tempo and phase. The tempo is the frequency of the beat. The phase refers to when the beats happen. Imagine two people drumming at the same tempo. Even though each is at the same tempo, one person may hit the drum when the other person is quiet, and vice versa. If this is so they are said to be completely out of phase. If they hit the drum at the same time and therefore quiet moments also happen at the same time they are said to be in phase.

It is shown through the experiments presented in this chapter that the collective behaviour of neural oscillators broadly conforms to the same synchronous behaviour exhibited by mathematically abstract oscillator models. Further to this, the chapter will demonstrate that dynamical systems that have external influences that themselves do not receive feedback from the system display internal dynamics that differ dramatically from dynamical systems that are completely self contained, whilst at the same time exhibiting comparable external behaviour.

4.1 Synchronisation in Neuroscience

There has been growing interest in brain dynamics and oscillatory behaviour within neuroscience communities due to the realisation that different perceptual and behavioural states are associated with different brain rhythms. It has been proposed that disparate groups of neurons firing synchronously provide a mechanism that underlies many cognitive functions, such as attention (Jensen, Kaiser and Lachaux, 2007), associative learning (Miltner et al., 1999), working memory (Siegel, Warden and Miller, 2009), and the formation of episodic memory (Nyhus and Curran, 2010). A role for synchronisation has been proposed in opening up communication channels between neuron groups (Fries, 2005). As Buszaki and Draguhn claim:

“The synchronous activity of oscillating networks is now viewed as the critical ‘middle ground’ linking single-neuron activity to behaviour” (Buzsaki and Draguhn, 2004).

4.2 Why Synchrony Occurs

Synchrony may be contrasted with the concept of resonance. With resonance, one object is initially oscillating whilst a second is not, but then the second starts to oscillate at the same frequency as the first. Resonance therefore describes an object that is oscillating in sympathy with another. The reason an object oscillates in sympathy is due to some connection between the two. For example, one oscillating object (such as a violin string) may be causally linked to another (such as a nearby guitar string) via the surrounding atmosphere, causing vibrations in the air that affect the second object so that it too starts oscillating.

Synchrony is similarly caused by a connection between two objects. But unlike resonance, where one object is originally oscillating and the other is not, with synchrony both objects are originally oscillating. However, the frequencies at which each is originally oscillating are different. Eventually they both come to be oscillating at the same frequency. When they synchronise, they may synchronise to a frequency that is different from either of the original frequencies. So for example, consider two pendulums hanging from the same beam. Suppose one is swinging at 2 Hz and the other at 4 Hz. The beam connecting the two creates a causal interaction. After a while and after much interaction both may end up oscillating, for example at 3 Hz. This phenomena was first reported by Huygens in 1665. Both are synchronised to the same frequency, but at a different frequency different from both initial frequencies. The reason for the different final frequency is that the causal interaction

between the pendulums is bidirectional. The oscillation from one is affecting the oscillation of the other, and vice versa. With resonance, the causal effect is one way, so the second object oscillates in sympathy at the frequency of the first.

In the brain there may for instance be one population of neurons oscillating at one frequency and another population oscillating at another frequency. The neurons in one population are causally connected to the other by synapses, and vice versa. Over time the oscillation in each synchronises so that the bursts of firing in each population are at the same frequency. Not only that but they may also have same phase.

4.3 Simple Oscillator Models and Neuroscience Research

An increasingly common level of abstraction for modelling neural information processing is one in which simple phase oscillators are used as elementary units representing populations of oscillatory neurons (Acebrón et al., 2005; Breakspear, Heitmann and Daffertshofer, 2010). A popular phase oscillator model used to capture the collective dynamics of such interacting communities is the Kuramoto model (Kuramoto, 1984; Acebrón et al., 2005; Strogatz, 2000), which is defined as follows:

$$\frac{d\theta_i}{dt} = \omega_i + \frac{1}{N} \sum_{j=1}^N K_{ji} \sin(\theta_j(t-d) - \theta_i(t) - a)$$

ω_i is the intrinsic frequency of oscillator i and $\theta_i(t)$ is the phase of oscillator i at time t . K_{ji} is the strength of connection from oscillator j to oscillator i . N is the number of oscillators in the system. a is an optional phase lag. The phase of oscillator j (θ_j) that feeds into oscillator i (θ_i) may also be subject to a time delay d , such that the phase of oscillator j used to update at time t is the phase of oscillator j at time $t-d$. A basic Kuramoto model sets a and d to 0. As can be seen the strength of connection and differences in phase between oscillators affects the current phase of an oscillator. Amplitude of oscillation plays no role in this model. The model therefore describes the most basic property of oscillators and the interaction between them.

The Kuramoto model has proved valuable in modelling neural systems, from replicating real data to producing behavioural architectures. Cabral *et al* (2011) use the Kuramoto model in simulation to explore the relationship between the slow modulation of gamma-band activity in the functional connectivity of a resting state network as seen in fMRI results, and its relationship to anatomical

connectivity and oscillations in localised neural areas. Investigating metastable chimera states in small-world topologies akin to those that appear in the brain, Shanahan (2010) built a community-structured network of Kuramoto oscillators and identified, as did Cabral *et al* (2011), the rich dynamics that results from the interplay between long-range connectivity of a large-scale network and interactions at a local level. In order to engage neuroscience with insights into the situated and embodied nature of cognitive dynamics Santos *et al* (2011) explored metastable dynamical regimes in a network of Kuramoto oscillators that was embodied in an agents sensorimotor loop.

Unlike the Kuramoto model, pulse coupled oscillators interact with each other at discrete times that are dictated by a phase response curve. Whilst there is a greater perceived affinity to neural systems when moving from phase lagged, to delayed, to pulse coupled oscillator system, it is not obvious that these models are a good representation of neural systems, which are far more complex. Given that the simplest model, the Kuramoto model, is widely used to capture fundamental properties of the collective dynamics of interacting communities of oscillatory neurons, this chapter poses the question: How well or accurately does the Kuramoto model represent oscillating population of neurons? The question is addressed by assessing if neural models behave in the same manner as simple oscillator models when replicating the most fundamental of Kuramoto's findings regarding oscillator interaction.

In a much cited monograph, Kuramoto (1984) showed that for an infinite number of oscillators with different intrinsic frequencies that are all uniformly connected with one another, there is a critical coupling value K_c below which the oscillators remain fully unsynchronised. Further to this there is another critical coupling value $K_L \geq K_c$ above which all oscillators become fully synchronised (Jadbabaie, Motee and Barahona, 2004). In this chapter this most fundamental property of oscillators is evaluated. The critical coupling experiment is emulated, but using populations of oscillating neurons in place of the simpler Kuramoto oscillators.

4.4 Experimental Variations

In the experiment presented in this chapter, numerous simulations are performed in which many neural PING oscillators are wired together to form a network. In each simulation the oscillators are wired together with a particular connection strength, and the connection strength is different on each simulation run. The network's overall synchrony can therefore be assessed according to the connection strength on each simulation run.

In the previous chapter a difference was found in the behaviour of oscillatory networks whose neurons had different bifurcation properties. In order to assess the critical coupling experiment fully, the classic experiment is performed first using quadratic integrate-and-fire (QIF) neurons and then using the Hodgkin-Huxley neuron model (HH). The two models capture the properties of Type I neurons with a saddle node bifurcation, and Type II neurons with an Androv-Hopf bifurcation respectively.

The frequency of the oscillations in PING architectures is caused by the feedback time over the EI/IE loop. The previous chapter highlighted how training using STDP may interfere with the behaviour of this architecture. In this case the resultant temporality obtained through learning was not due to the network dynamics that result from the PING architecture, but is an artefact of repeated periodic presentation of a learnt stimulus. After training, a population of neurons will respond ‘synchronously’ whenever the stimulus is presented. Given that there are differences in the behaviour of trained and untrained architectures both shall be assessed in the context of the critical coupling experiment.

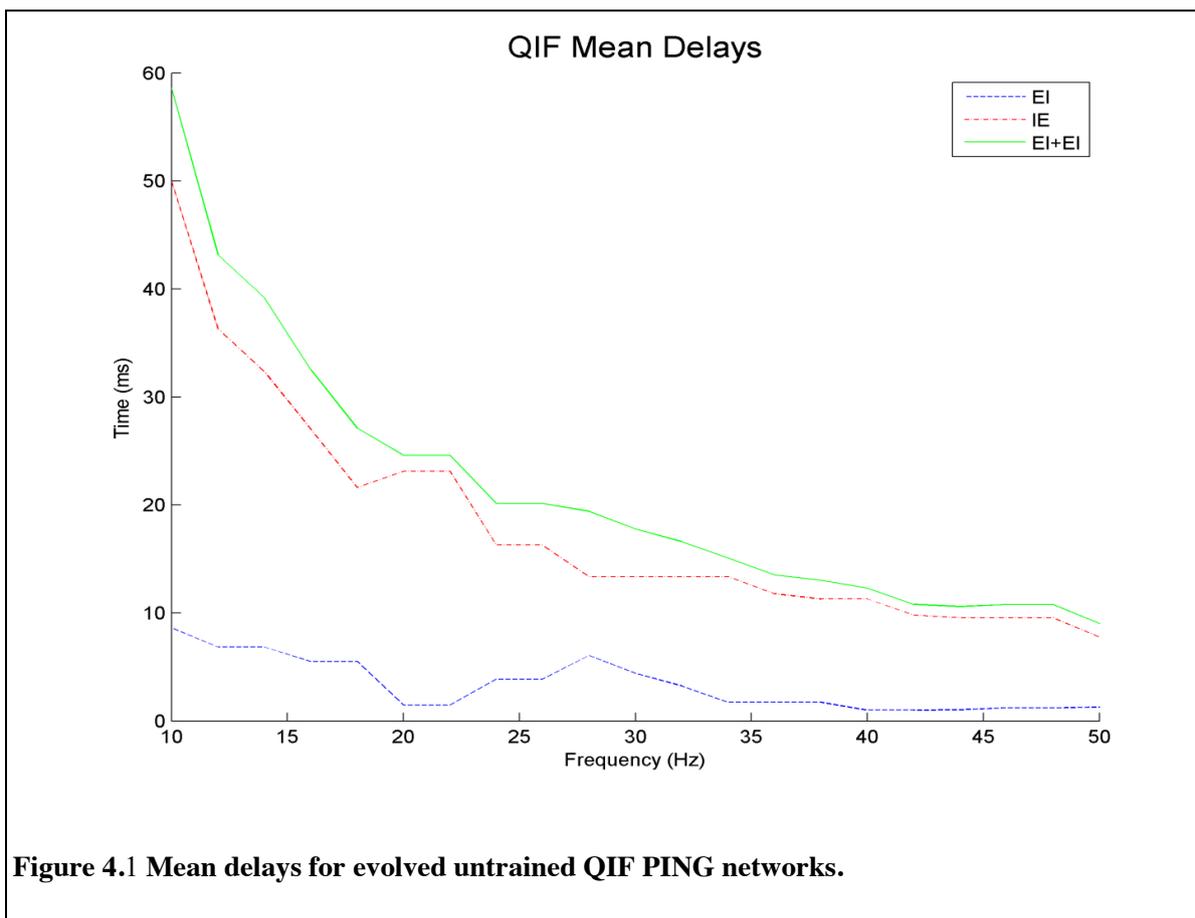
4.5 Generation of Neural Oscillator Nodes

The neural oscillators used in this work conform to a PING architecture. Whilst the general PING architecture is well understood, the specific details required for particular oscillatory frequencies and particular neuron models vary and involve a large space of parameter values within the general PING framework. In order to provide a wide range of different intrinsic oscillatory frequencies for the neural oscillator nodes used in the experiment, it was decided to obtain these parameter values by use of a genetic algorithm (described in section 3.5). The genetic algorithm evolved, within biologically plausible bounds, every oscillatory frequency between 10 Hz and 50 Hz for the QIF and HH models. QIF and HH neural oscillators were evolved that responded to any stimulus, whereas neural oscillators that were trained to only respond to a learn stimulus were only evolved for the QIF model.

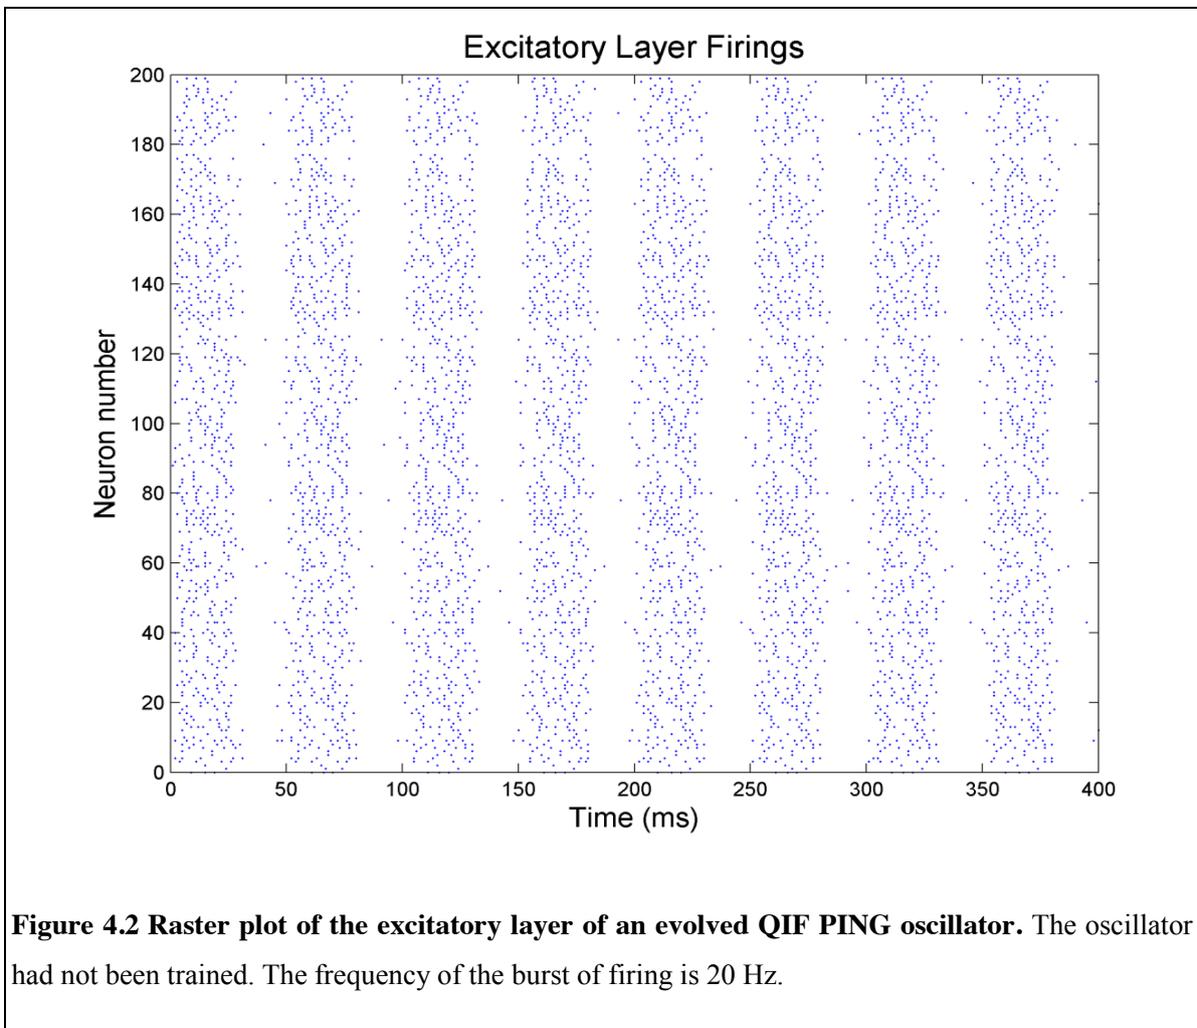
The oscillator networks that were not trained to respond to a single given stimulus, but instead responded by oscillating to all stimuli, had only one test phase in which they were presented with a random stimulus as described in the previous chapter. After this they were scored for fitness using only the term $fitness_1$ from chapter 3. The input is generated from a Poisson process with parameter $\lambda = 4.375$. In order to provide sufficient stimulus to induce firing, the untrained QIF neuron networks scaled inputs by a factor $S = 8$, and the untrained HH neurons networks scaled inputs by $S = 1$. The

QIF oscillators that were trained to respond only to a single learnt stimulus were evolved exactly as described in section 3.5 and the model parameters were as described in section 3.2.3.

The networks were wired up with connections between inhibitory neurons (II), from excitatory to inhibitory neurons (EI) and from inhibitory to excitatory neurons (IE). Excitatory to excitatory (EE) connections were included for networks that were trained using STDP, but excluded for networks that were not trained. Excluding the EE connection limited saturation effects in the HH model (meaning that all neurons were firing all of the time). Saturation effects tended to arise in the simulations in which many neural PING nodes were wired together. The possibility of saturation was not otherwise catered for in the evolutionary process due to the PING networks being evolved in isolation. Network that used untrained oscillator used a scaling factor of 5 on all synaptic connections in these pathways, similarly the trained oscillators used a scaling factor of 7.



The evolved QIF solutions for networks that did not use STDP but relied on the PING architecture to produce different frequencies of oscillation all had weights with very high means and small variances. The respective HH solutions showed greater variation in the weight means across evolved solutions for different frequencies, indicating greater sensitivity in the model and solution such that they required a very specific balance of the parameters for each particular frequency solution. The means for the delays evolved for both QIF and HH solutions had a similar form, from which it can be concluded that the EI mean delay + IE mean delay $\approx 1000/2f$. This is illustrated for the QIF model in Figure 4.1. Figure 4.2 shows a raster plot of the firings of the excitatory layer of a QIF node that had *not* been trained using STDP and therefore has no EE connections. The oscillator is producing regular wide bursts of firing at a frequency of 20 Hz.



4.6 Synchronisation metric

A measure of the networks global synchrony was required for each simulation run in the experiment. In this work each neural PING oscillator in a network consisted of an excitatory layer and an inhibitory layer. The synchrony was only calculated for the excitatory neuron layers in the PING oscillators in a network. To achieve this, the spikes of each neuron in each excitatory layer were first binned over time, and then a Gaussian smoothing filter was passed over the binned data to produce a continuous time varying signal. Following this, a Hilbert transform was performed on the mean-centred filtered signal in order to identify its phase. No band-pass filtering was performed during this process. The synchrony at time t was then calculated as follows:

$$\varphi(t) = \left| \frac{1}{N} \sum_j^N e^{i\theta_j(t)} \right|$$

$$\varphi = \frac{1}{t_{max}} \sum_t \varphi(t)$$

where $\theta_j(t)$ is the phase at time t of oscillatory population j . i is the square root of -1, N is the number of oscillators, and t_{max} is the length of time of the simulation. $\varphi(t)$ is the instantaneous synchrony and φ is the mean frequency.

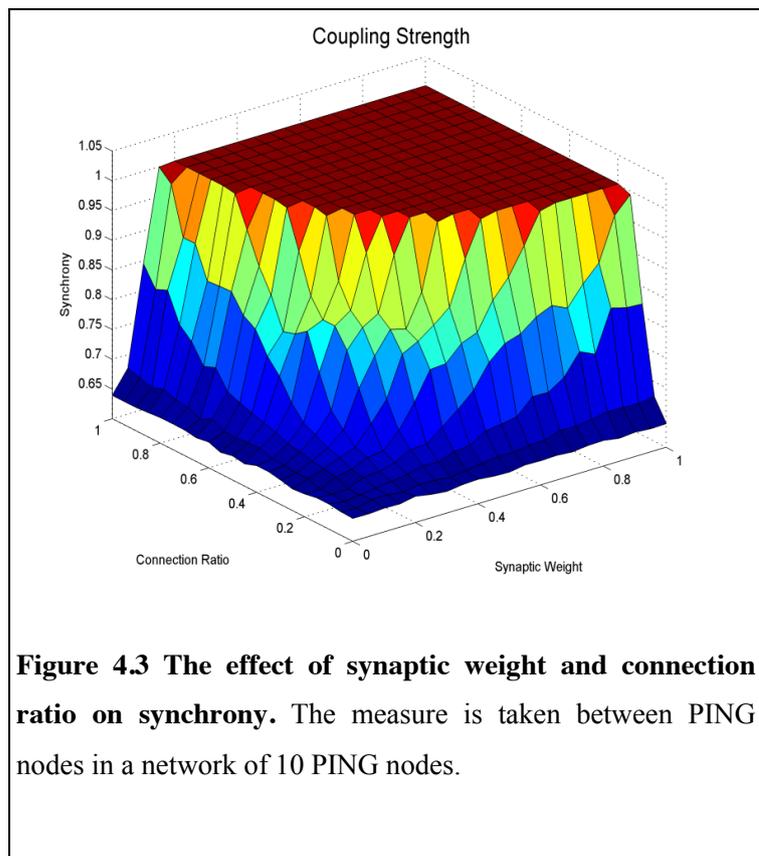
4.7 The Question of Neural Coupling Strength

In replicating the critical coupling experiment with neural systems we are faced with an immediate question: How does the notion of coupling strength in Kuramoto's oscillator model, in which each node is a simple phase oscillator, relate to a neural model in which each oscillatory node is made up of many neurons? We have two options: the coupling between two nodes may refer to the number of synaptic connections between the neurons, or it may refer to the strength of the synaptic weights.

To address this issue the results from a parameter sweep of connection ratios and synaptic weights is first presented. For each set of parameter values in the sweep, the experiment takes 10 QIF oscillator nodes (not trained) that have each been evolved to oscillate at a different frequency using a PING architecture as described above. Each node is chosen by frequency from a uniform distribution ranging between 10 Hz and 50 Hz. The excitatory layers in each node are connected to one another with a given connection ratio to form a network of nodes. A connection ratio of 1 represents all-to-all

connectivity so that each neuron in one node's excitatory layer is connected to each neuron in another node's excitatory layer. Once the connections have been established they are given a weight as defined by the point in the sweep. The sweep is two dimensional, ranging from a connection ratio of 0 to 1 and a synaptic weight from 0 to 1. An increment of 0.05 is used for both the connection ratio and the synaptic weight in the sweep. At each point in the sweep the overall synchrony of the network is measured.

The result of the sweep is shown in figure 4.3. It can be seen that connection ratio and weight have a similar effect with neither showing a marked importance over the other and with the graph appearing symmetrical. From this result it is safe to assume that the use of both connection ratio and synaptic weight in the main Kuramoto experiment will not be biased by using a prescribed value for the connection ratio and sweeping across only the synaptic weight in order to explore Kuramoto notion of coupling strength. As such all further experiments will use a connection ratio of 0.2 and simply sweep the synaptic weight parameter.



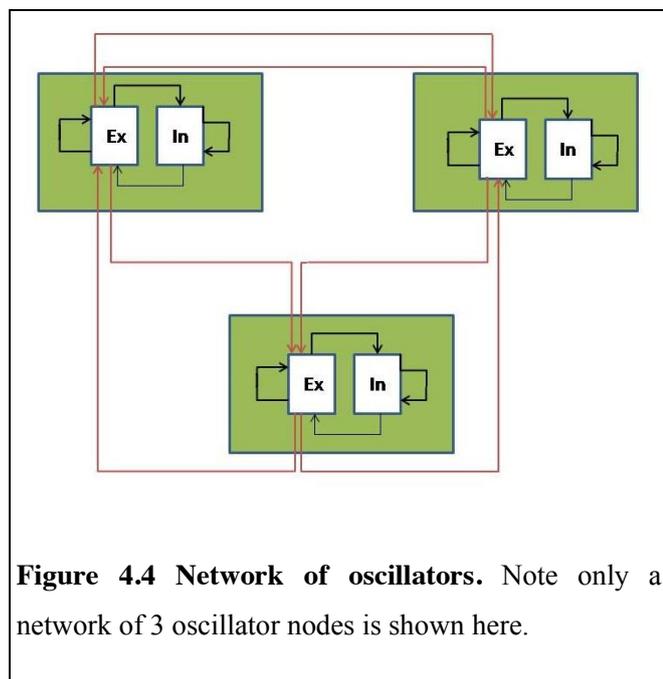
4.8 Critical Coupling Experiment

4.8.1 Experimental Setup

PING architectures were evolved for every frequency between 10hz and 50hz for both QIF and HH neuron models, as well as the QIF models that have been trained using STDP. A series of simulations were performed for each type of model.

Kuramoto calculated analytically the critical values in the case of an infinite number of oscillators connected all-to-all for a few well-known distributions of intrinsic oscillator frequencies (Jadbabaie, Motee and Barahona, 2004). For simplicity, Kuramoto assumed that the distribution of oscillator intrinsic frequencies was unimodal and symmetric about its mean frequency, as in a Gaussian distribution for example (Strogatz, 2000). In line with Kuramoto's specification the evolved oscillators were selected using a Gaussian distribution with a mean of 30 Hz. The variance that was chosen in order to ensure a good spread of different oscillator frequencies was 10 Hz.

In all experiments 64 neural oscillator nodes were used to form a network of nodes. Each node receives external input to its excitatory layer as with the evolutionary setup described above. Figure 4.4 illustrates the network setup (but for a network of 3 nodes). For architectures that were not trained, the external input along with the evolved PING architectures induces the intrinsic oscillatory frequency. For those that were trained, the frequency of presentation of the learn stimulus induced the intrinsic oscillatory frequency of the node.



The phase of each oscillator was determined by the time at which external input to the oscillator was started, which varied from 0 ms to 100 ms. The slowest oscillator was 10 Hz and therefore a random start point ranging from 0 ms to 100 ms allowed for 10 Hz oscillators (as well as all oscillators of higher frequency) to be completely out of phase with each other. The neurons in the excitatory layers of each node were synaptically connected to the neurons in the excitatory layers of each other node with a connection ratio of 0.2. The experiments involved a sweep of synaptic weights for all inter-node connections. These weights were all set to the same value on the same iteration in the experimental sweep. On each iteration this synaptic weight value was changed. At each point in the sweep the overall synchrony of the network was measured. The network was simulated for 2000 ms for each iteration of the sweep. Each network comprised 16,000 neurons and $\approx 32,256,000$ synapses.

4.8.2 Results

100 simulations were performed for each model type, with a different sample weight for each simulation. After a particular high coupling value, all networks models exhibited ‘saturation’, meaning that all excitatory neurons in all nodes were firing continuously. The results shown here display data up to the respective point of saturation for each model type as data beyond this point is

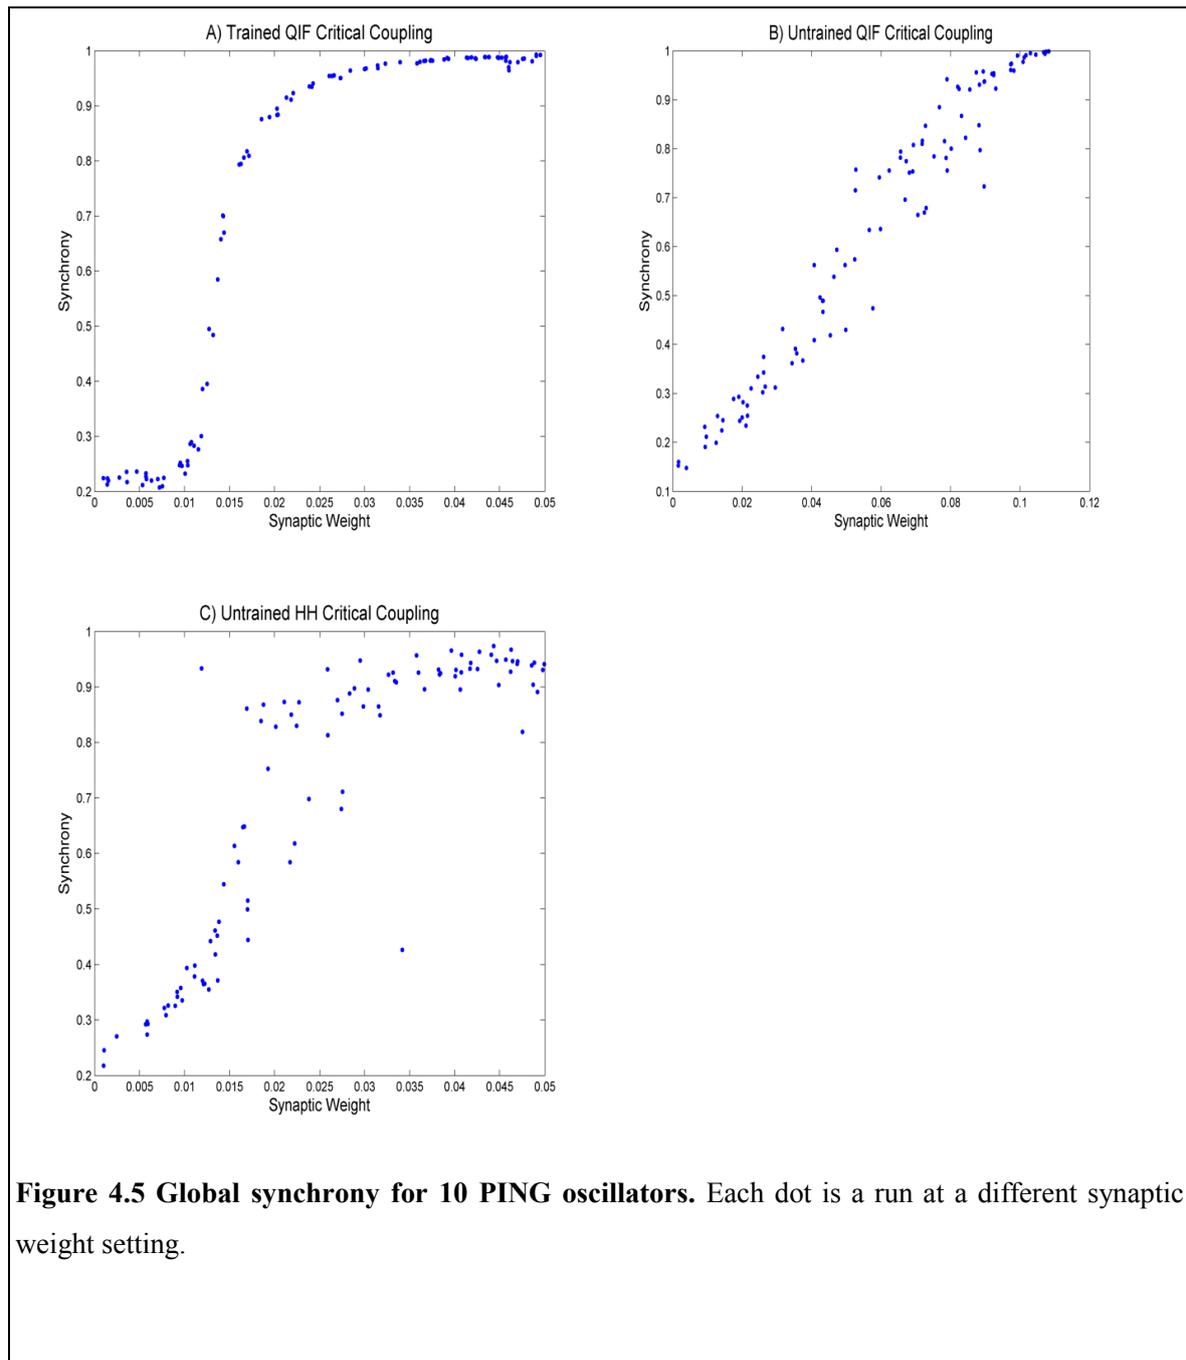


Figure 4.5 Global synchrony for 10 PING oscillators. Each dot is a run at a different synaptic weight setting.

not noteworthy.

All network types have an initial synchrony value of ≈ 0.2 which corresponds to no synchrony apart from coincidental phase alignment. The weight sweeps for all models can be seen in figure 4.5. All models show a steady increase in synchrony until saturation is reached, which is at weight 0.11 for QIF networks and 0.05 for HH networks and trained QIF. For untrained QIF and HH models synchrony rises with connection strength but so too does the spread of the dots, indicating some variation in behaviour with these systems. For trained QIF models synchrony can be seen to increase smoothly with a critical region of connection strength in which synchrony rapidly increases, in accord with Kuramoto's findings. The connection strength is effective at different levels from the untrained PING models due to different sensitivities in the different architectures, Poisson process parameters, and scaling factors. However the behaviour is the key difference to note. There is a very sharp sigmoidal increase for trained QIF models in figure 4.5A, indicating little variation in behaviour with these learnt systems, unlike those in figure 4.5B and 4.5C. There are also no discontinuities.

A coalition is a collection of nodes that are oscillating at the same unique frequency. In each simulation there may be several coalitions. The synchrony between nodes is best understood by comparing the number of coalitions that are present on each simulation. As we are using discrete intrinsic oscillatory frequencies selected from a Gaussian distribution, in one run for one weight setting there may be several nodes with the same intrinsic frequency. Therefore comparing the number of coalitions of oscillators sharing the same frequency during the experiment is only meaningful in relation to the number of coalitions that can be formed from the intrinsic frequencies on each trial. Figures 4.6A, 4.6B and 4.6C shows for trained QIF, untrained QIF and HH models respectively the number of coalitions of nodes sharing the same frequency. Green dots show intrinsic coalitions prior to running the simulations and red crosses show the number that results from the interaction between nodes when the simulations are run. Weight 0.013 for trained QIF networks, weight 0.05 for untrained QIF networks, and weight 0.015 for HH networks marks the critical point K_c beyond which the number of coalitions becomes less than the number of intrinsic frequency coalition. This indicates that nodes are affecting one another, pulling each other away from their natural frequencies towards a common frequency. It is interesting to note that prior to the point K_c both untrained QIF and HH network models display the ability to pull apart from their intrinsic frequency groups into more frequencies than prescribed by the initial Gaussian selection. The trained

QIF model reaches a point of full synchrony K_L , at ≈ 0.02 , beyond which one frequency exist for all nodes and is maintained until saturation at 0.05. The untrained QIF model reaches the same point at ≈ 0.04 . The HH model displays full synchrony at ≈ 0.026 showing only one frequency group. It is

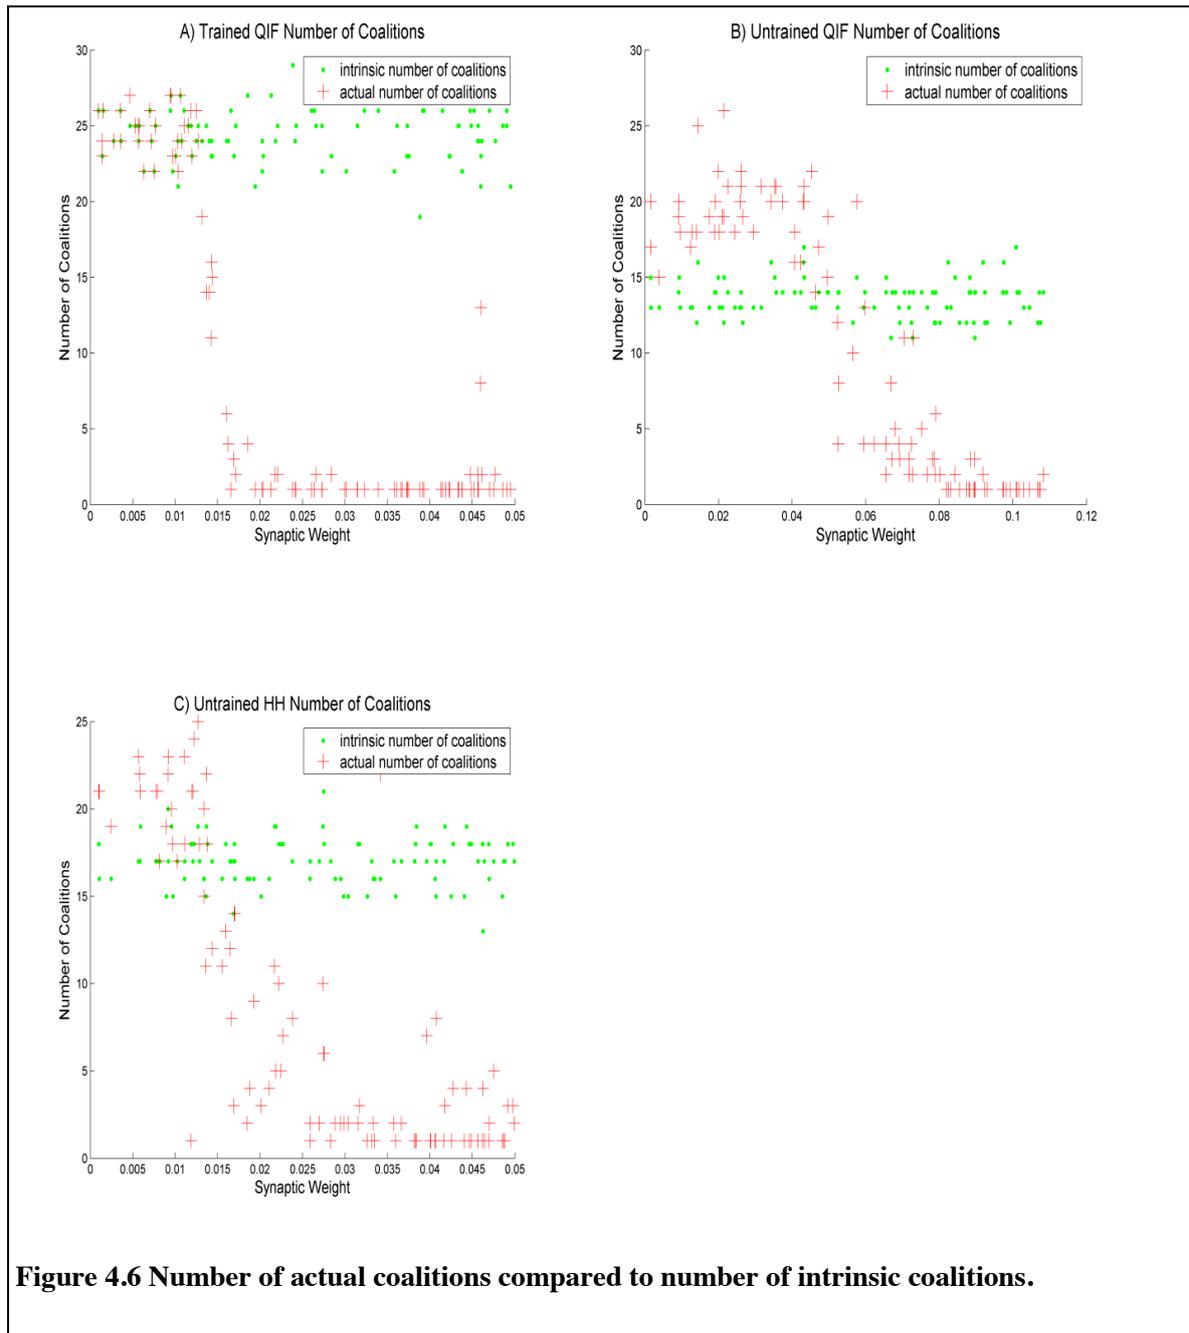


Figure 4.6 Number of actual coalitions compared to number of intrinsic coalitions.

clear from figure 4.6C that K_L for the HH model is much less well defined than for the QIF model in 4.6B.

Taking a closer look, plots of pairwise synchrony between a selected, representative node (node 1) and all other nodes in the network at particular weight points in the sweep are next discussed. Figure 4.7 shows pairwise synchrony for the trained QIF network at weight 0.0172, a point after K_c . Each sub-plot shows the nodes in a coalition sharing the same frequency (note that for figures 4.7, 4.8 and 4.9 plots for coalitions containing only one node are not displayed). For the trained QIF model, oscillators sharing the same frequency show major variations in phase.

Figure 4.8 shows pairwise synchrony for the untrained QIF network at weight 0.07, a point after K_c . Each sub-plot shows the nodes in a coalition sharing the same frequency. As can be seen the behaviour is very precise with the nodes in each group moving closely together both in the same pattern and with little offset indicating matching phases as well. However they do not maintain a constant offset from the phase of node 1 indicating that whilst the frequency and phase of a coalition remains the same within itself, a more complex synchronised behaviour is at play in the network of nodes as a whole. A few desynchronous moments appear for the group with frequency 25.88 Hz indicating that the nodes themselves are not fully stable at this coupling strength. The HH model behaves a lot more noisily. Figure 4.9 displays the pairwise synchrony at weight 0.02, a point beyond K_c . Nodes sharing frequency 25.88 Hz show distinct phase offset, whereas nodes sharing frequency 46.88 Hz do not appear very synchronous at all.

Figure 4.10 shows the pairwise synchrony for the networks at their highest respective values for synchrony before saturation. The untrained oscillator networks show deviations from full synchrony in which the network separates into sub-groups, which although they diverge, show similar phase movements indicating mutual influence between the groups. The trained architectures show little group separation behaviour but instead single oscillators seem to separate into their own phases away from full synchrony. The same behaviours are manifest at lower global synchrony levels albeit that the deviations are greater. The behaviour may be intuitively explained by the fact that in the untrained architectures the individual intrinsic oscillatory frequencies of the nodes as well as the interaction between them are generated by the network architecture which forms a complete system, whereas in the trained architectures the individual intrinsic oscillatory frequencies are created by an external stimulus that is separate from the network system, does not receive dynamic feedback, and therefore facilitates more individual rather than group behaviour.

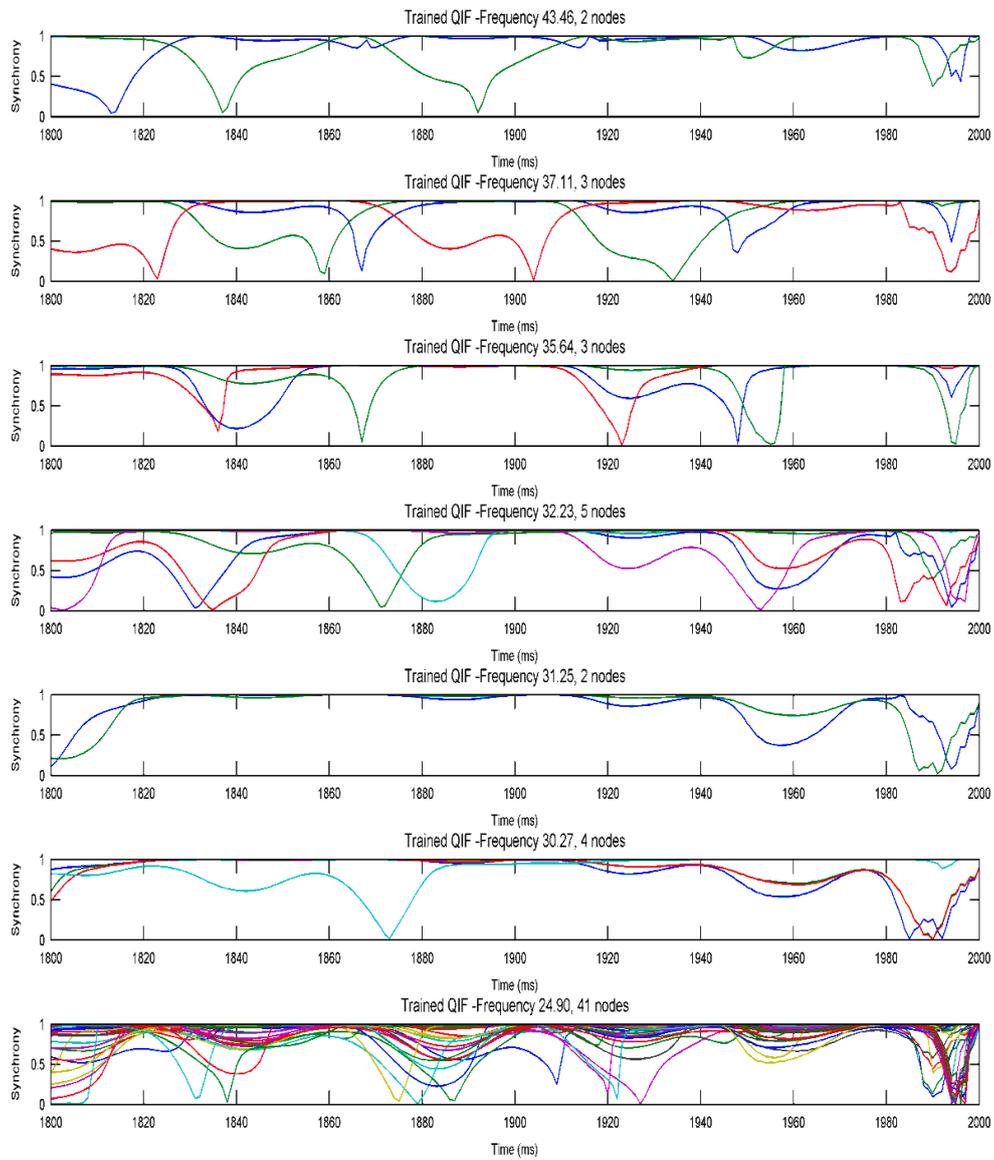


Figure 4.7 Pairwise synchrony for trained QIF at weight 0.0172. This is a point after K_c . Coalitions of 1 node are not shown.

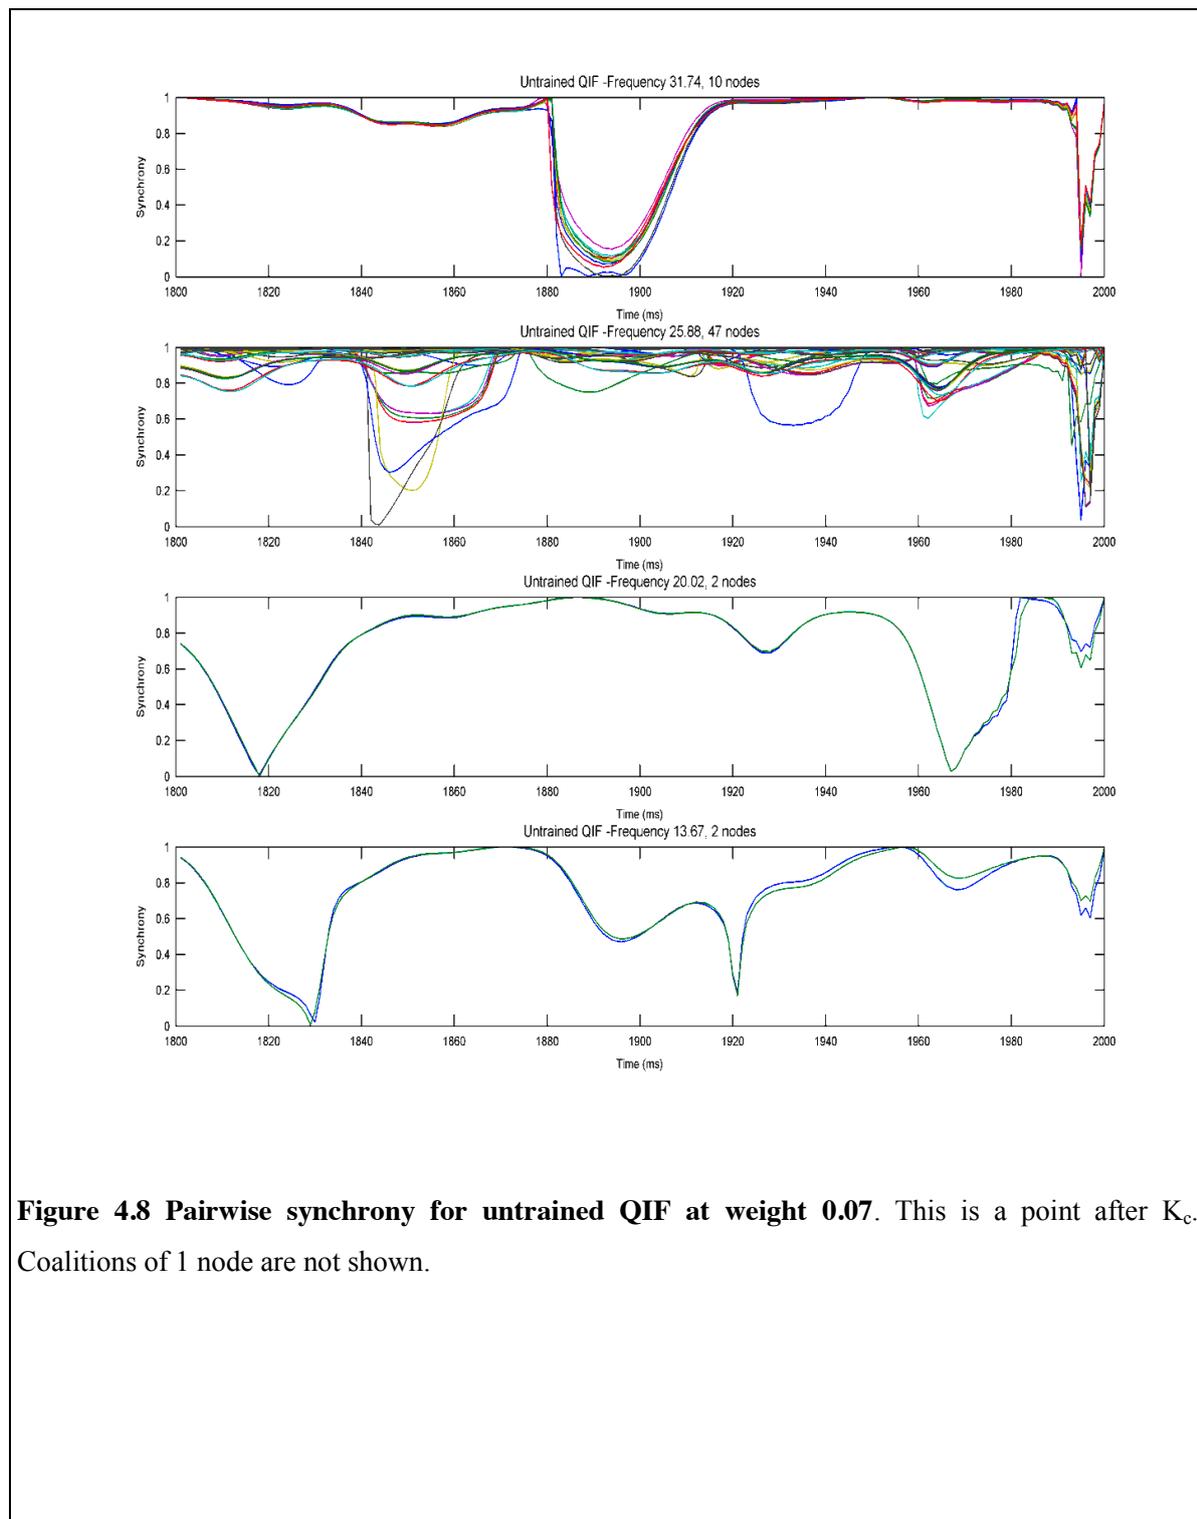
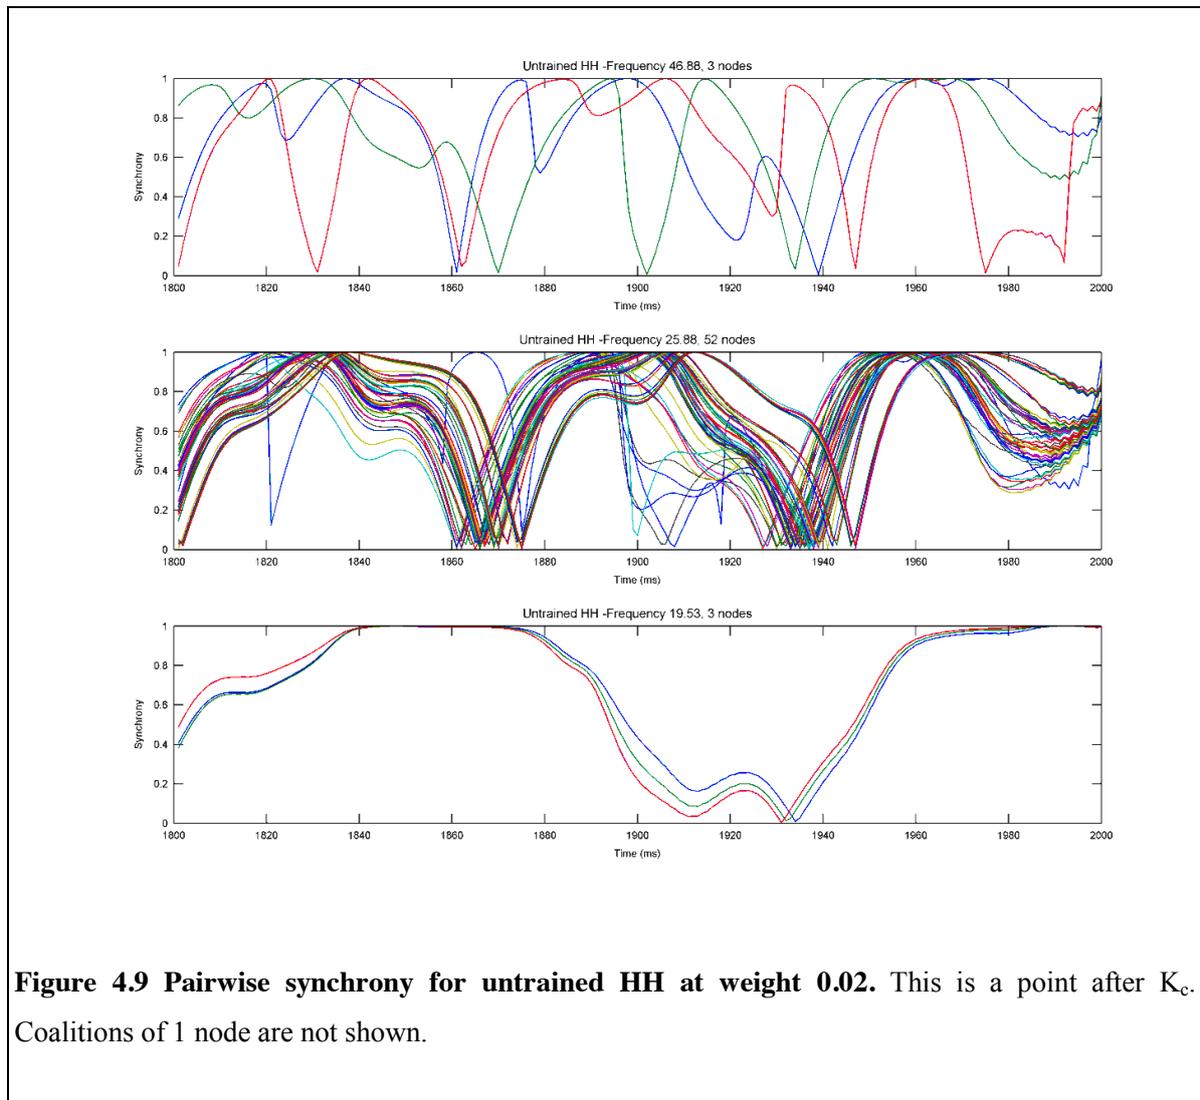
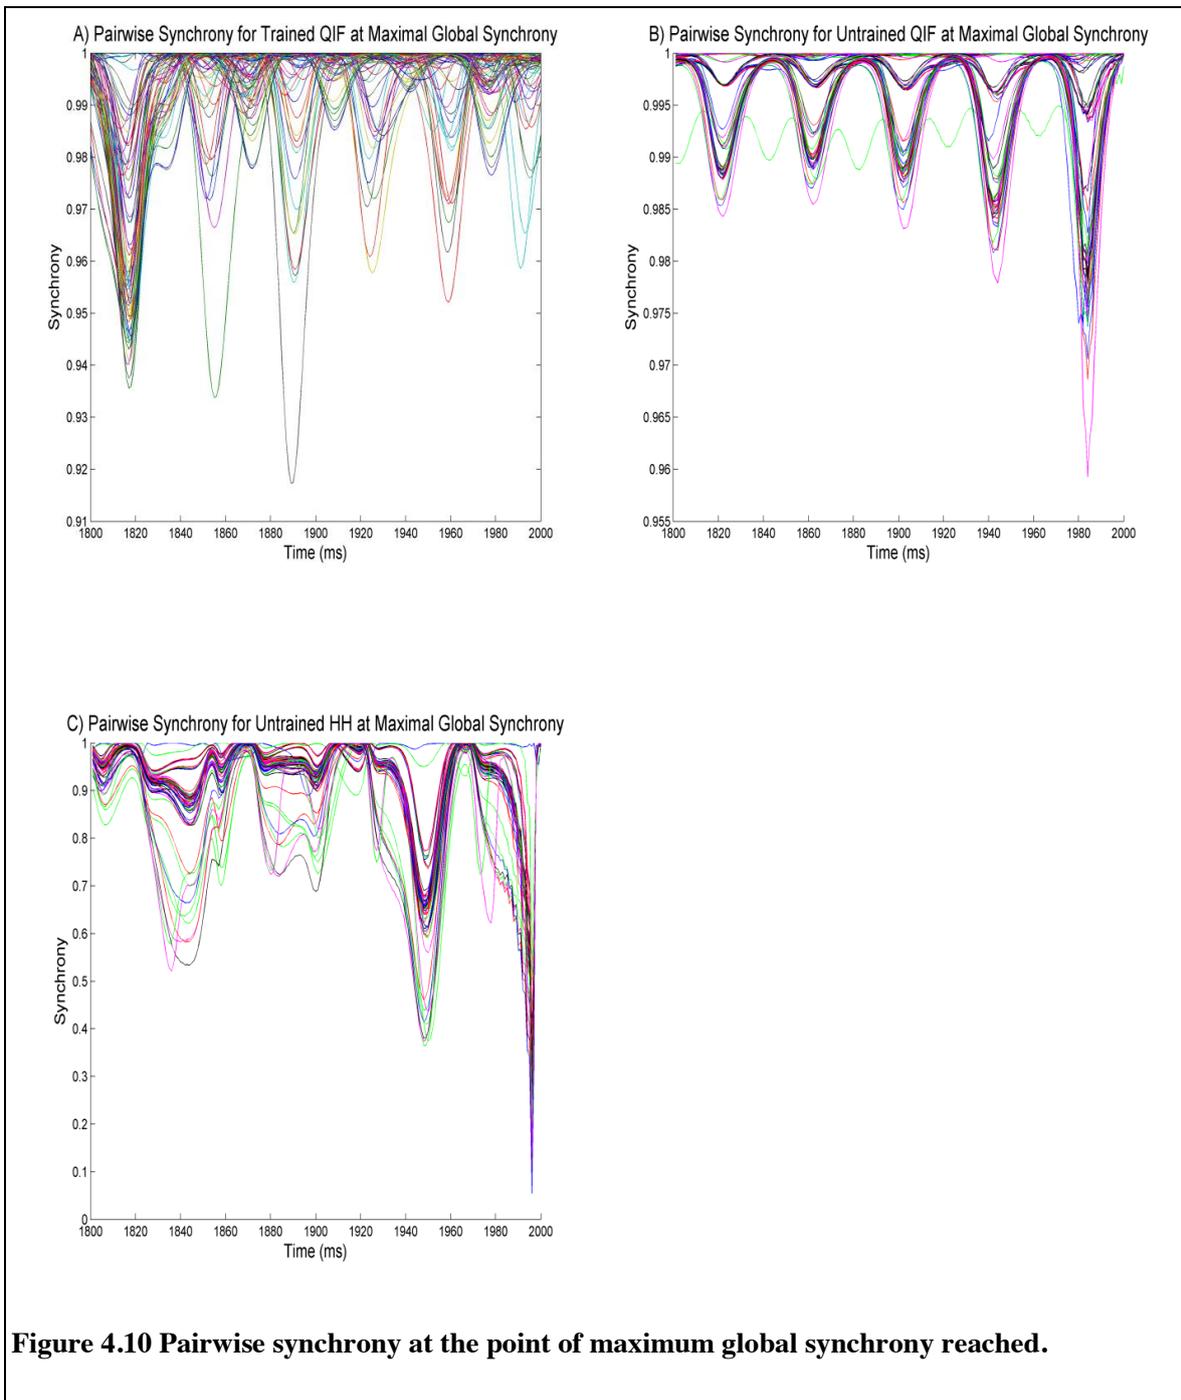


Figure 4.8 Pairwise synchrony for untrained QIF at weight 0.07. This is a point after K_c . Coalitions of 1 node are not shown.



The pairwise synchrony for the untrained QIF network is shown in figure 4.10B. All nodes are synchronous with minor periodic deviations of up to 0.05 from maximum synchrony, resulting in general synchrony over 0.955 which is very high. A closer inspection of the spiking behaviour for this group shows that whilst they share a common main frequency their firing patterns are much noisier with other less well defined frequencies present. Similar behaviour is seen in figure 4.10C which shows the HH network at weight 0.04, a point at which all nodes share the same frequency. However, the scaling of the y-axis is larger than figure 4.10B showing greater deviations from full synchrony. Again on inspection of the firing behaviour shown in the raster plot in figure 4.11, a much noisier behaviour is seen in which other frequencies underlie the main frequency of 20.51 Hz.



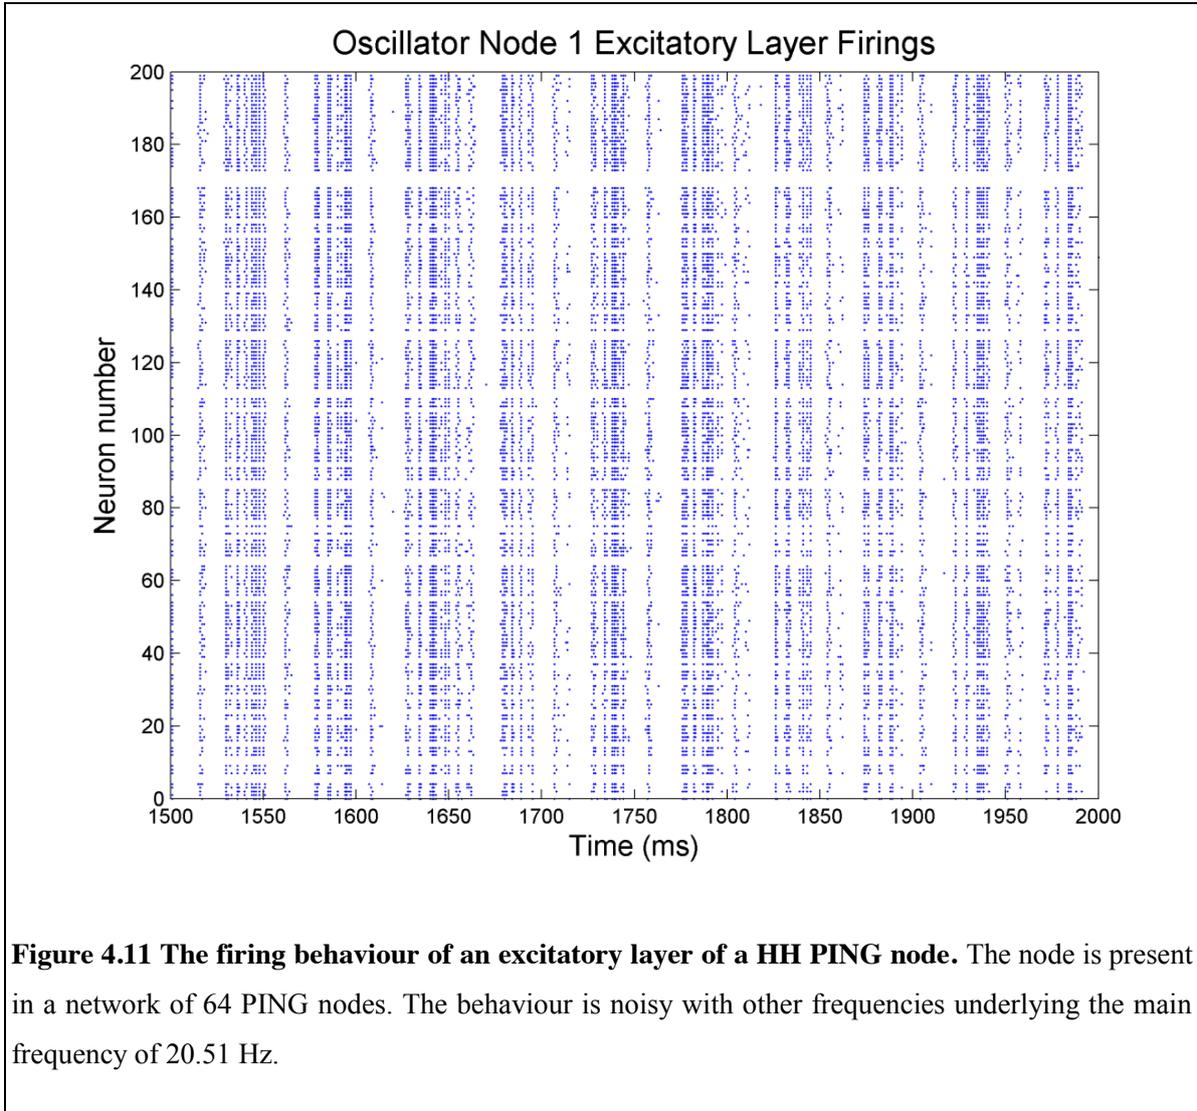


Figure 4.11 The firing behaviour of an excitatory layer of a HH PING node. The node is present in a network of 64 PING nodes. The behaviour is noisy with other frequencies underlying the main frequency of 20.51 Hz.

4.9 Interaction between Trained and Untrained Nodes

Given the marked difference in internal behaviour during the different critical coupling experiments for nodes that have been trained to oscillate to a learnt stimulus and those that have not, a further experiment is presented in this section that explores the interaction between the two types of node. The critical coupling experiment was again performed for QIF PING oscillators but this time on each simulation run 63 nodes were chosen from the set that had been evolved to respond to any stimulus and one node was chosen from the set that had been trained to only respond to a learnt stimulus. The

intrinsic frequency for each node was chosen as before. The experiment presents a situation in which the entire network forms a complete dynamical system except one node which operates more independently.

After each simulation at a given coupling strength the mean pairwise synchrony was calculated between the trained node and all other nodes. The pairwise synchrony was calculated using the equation in section 4.6 but setting $N=2$ and only using the phases for the trained node and one other node on each calculation. This was performed for all combinations of the trained node and all other nodes and then averaged. The mean pairwise synchrony was also calculated for each untrained node and all other untrained nodes. These were then averaged to give the mean pairwise synchrony between all untrained nodes.

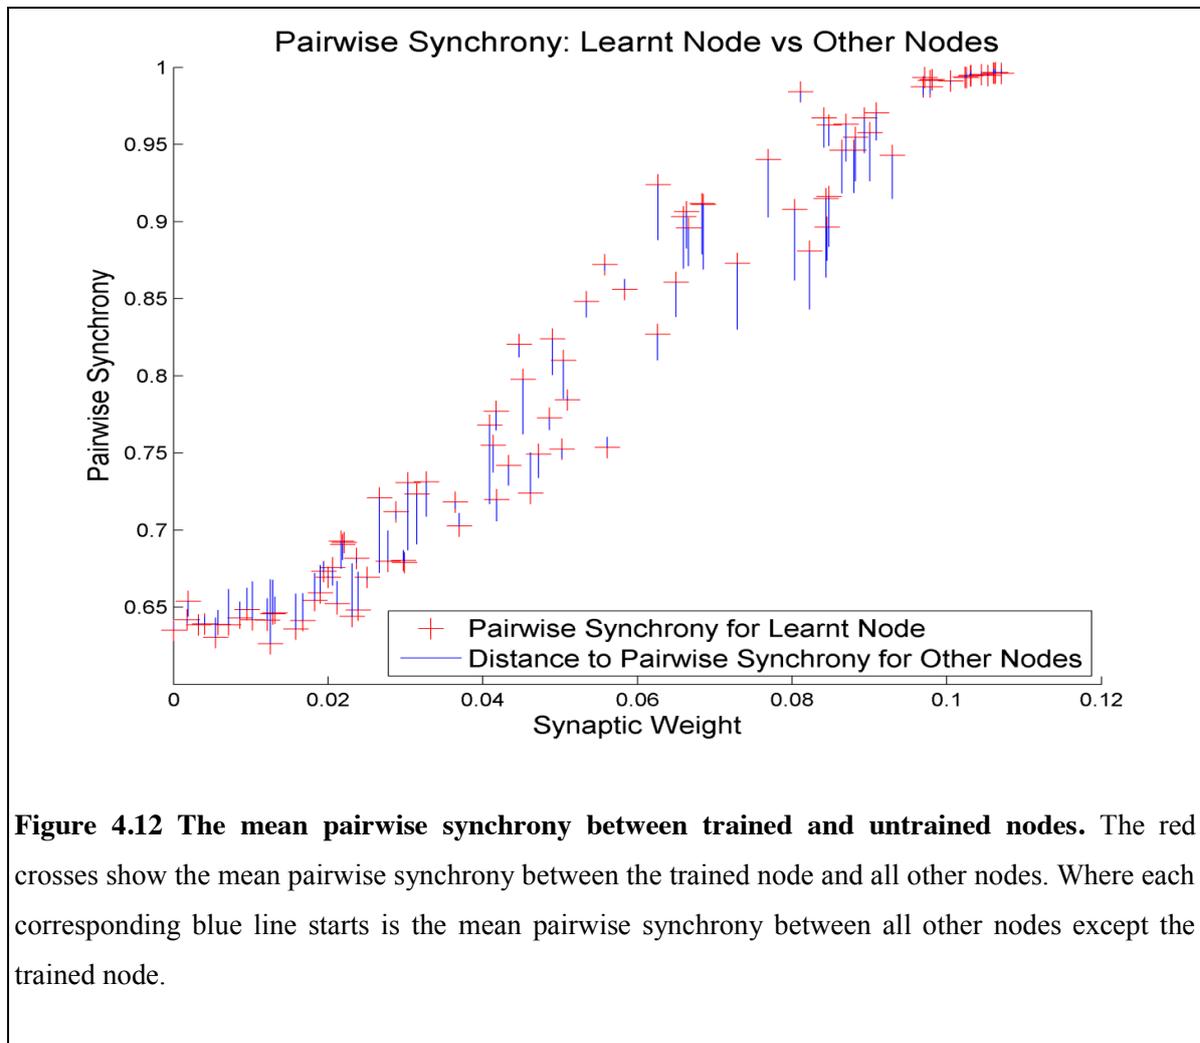


Figure 4.12 shows the results. The red crosses indicate the mean pairwise synchrony between the trained node and all other nodes on each run. Where each corresponding blue line starts is the mean pairwise synchrony between all other nodes except the trained node on the same run. As can be seen, at low weight levels before coupling starts to take effect the trained nodes are less synchronised to the untrained nodes than the untrained nodes are with each other. As coupling takes effect the trained nodes display more synchrony with the untrained node than the untrained nodes do with each other.

From these results it can be concluded that the trained node is initially more isolated from the rest of the system, but as coupling strength increases the trained node entrains more untrained nodes to be synchronous with it than untrained nodes do with each other.

4.10 Discussion

The experimental results presented in this chapter show that suitably connected oscillatory populations of Type I neurons with saddle node bifurcations and Type II neurons with Andronov-Hopf bifurcations broadly conform to critical coupling findings found with Kuramoto oscillators. However both models display much greater spectral complexity than that which the simple Kuramoto oscillator can capture. In fact, the more detailed and biologically plausible the model is then the greater the spectral complexity. This greater complexity may have implications for neural simulation using the Kuramoto abstraction, as this simple model only partially captures the range of temporal phenomena which we find with the more biologically plausible spiking models.

The critical coupling experiment for trained oscillators demonstrates that the collective behaviour is far better defined and more precise, compared to those that were not been trained. However, the internal dynamical behaviour differs dramatically between the trained and the untrained architectures. The untrained oscillator networks show deviations from full synchrony in which the network separates into sub-groups or coalitions, which although they diverge, show similar phase movements indicating mutual influence between the groups. The trained architectures show little group separation behaviour and influence, but instead single oscillators seem to separate into their own phases away from full synchrony. The behaviour may be intuitively explained by the fact that in the untrained architectures the individual intrinsic oscillatory frequencies of the nodes as well as the interaction between them are generated by the network architecture which forms a complete system. In contrast the individual intrinsic oscillatory frequencies of the trained architectures are created by an external stimulus that is separate from the network system, and hence is unable to receive dynamic

feedback so as to form a closed system. The latter facilitates more individualistic rather than group behaviour.

Chapter 3 detailed the behaviour of trained oscillators, which ignored the dynamics inherent in the PING architecture but instead oscillated at the same frequency as the external stimulus. This may be regarded as a similar situation to photosensitive epilepsy. Photosensitive epilepsy is a form of epilepsy in which seizures are triggered by visual stimuli that form patterns in time or space. These may be visual effects such as flashing lights, bold, regular patterns, or regular moving patterns. With photosensitive epilepsy an external intermittent photic stimulation (IPS) *entrains* the neuronal populations at gamma frequencies. This is known as a photic following response (Parra et al., 2003). The trained oscillators presented in Chapter 3 exhibit similar behaviour in that an external stimulus controls the frequency of oscillation of the neural population.

In a review article, Michel Le Van Quyen (2005) discusses the collective behaviour of neural assemblies when an epileptic focus is present. He points out that, the pathologically discharging neuronal population of the epileptic focus exhibits a state of decreased synchronisation with the rest of the system. This state may isolate the pathologically discharging neuronal population from the influence of activity in wider brain areas. The simulations presented in this chapter that used only trained oscillators presents an unrealistic situation in which each oscillator is like an individual epileptic focus. The results support the view that an epileptic focus is more disassociated from normal dynamics which may exhibit properties such as coalition formation. In this case each oscillator operates separately from the global dynamics. The critical coupling experiment which used one trained node and all other untrained nodes presents another situation in which there is a single epileptic focus amongst an otherwise normal system. Before critical coupling influence takes effect the epileptic focus is more separated from the behaviour of the rest of the system in line with Le Van Quyen's findings.

Le Van Quyen further states that under normal conditions, control mechanisms exist in the brain to counteract the initial increase in epileptic synchrony. In epileptics these do not function. Measurements from pairs of electrodes show that populations near the epileptic focus may decrease their synchronisation before seizures. This decrease is the exact converse of phase-locking, and is best described as phase-scattering. As the pathologically discharging neuronal population is isolated from the influence of activity in wider brain areas, it facilitates the development of pathological recruitments with these local and disassociating populations. As global coupling influence is

increased in the experiment with a single epileptic focus (the trained node), the results presented in section 4.9 showed that normal untrained nodes in the system become more synchronous with the epileptic focus than they are averagely with each other. This indicates that more nodes are being entrained by the epileptic focus, in line with Le Van Quyen's findings.

At the macroscopic level a prevalent feature of the brain activity is the presence of apparently metastable spatio-temporal patterns of neural activity distributed over many distinct brain regions (Le Van Quyen, 2005). With regard to epilepsy, Parra *et al* hypothesise:

"the enhancement of synchrony in the gamma band of photosensitive patients may represent a loss of control of the brain over a high-frequency oscillatory process that normally operates to transiently connect neural assemblies involved in the cerebral cortex." (Parra et al., 2003)

There has been growing interest in neuroscience communities in seemingly metastable dynamics that serve to transiently connect neural assemblies as well as the possible relationship between these dynamics and brain function. This chapter has shown how a normally behaving dynamical system of oscillatory neural nodes will, given enough coupling influence between the nodes, form sub-groups or coalitions of connected neural assemblies. The next chapter shall investigate *transiently* connecting neural assemblies and the metastable dynamical mechanisms that form them.

5 Metastability and Inter-Band Frequency Modulation

An important question in neuroscience is how localised neural activity gives rise to complex cognitive functions organised at a global level (Sporns et al., 2004; Tognoli and Kelso, 2011; Shanahan, 2012). Somehow the brain must orchestrate neural processes in different cortical areas so that they form coordinated coalitions. An influential candidate for the mechanism underlying coordination and communication among brain areas is synchronous oscillation over multiple frequency bands (Fries, 2005; Varela et al., 2001). However, the question of how coordinated coalitions of brain processes form and break apart remains open.

This chapter, builds upon a growing body of modelling work supporting the view that the brain, regarded as a dynamical system, is inherently metastable, and suggests that metastable dynamics are responsible for producing transient episodes of synchrony between different neural areas (Cabral et al., 2011; Shanahan, 2010; Wildie and Shanahan, 2011; Cabral et al., 2013; Hellyer et al., 2014).

5.1 Brain Rhythms and Dynamical Organisation

As discussed in chapter 2, periodic activity is found in the brain when taking measurements with EEG and MEG equipment. The various different rhythms that have been identified have diverse associations. Thalamocortical networks display increased delta band (0.1–3.5 Hz) power during deep sleep (McCormick, Sejnowski and Steriade, 1993). Theta (4–7.5 Hz) activity is increased during memory encoding and retrieval (Basar et al., 2000). Alpha band (8–13 Hz) changes are associated with attentional demands (Klimesch, 1999). Beta (14–30 Hz) oscillations have been related to the sensorimotor system (Pfurtscheller, Stancak Jr and Neuper, 1996). Of all the frequency bands the role

of gamma (30-80 Hz) is thought to be most extensive and is hypothesised to provide a mechanism that underlies many cognitive functions such as: attention (Jensen, Kaiser and Lachaux, 2007), associative learning (Miltner et al., 1999), working memory (Siegel, Warden and Miller, 2009), the formation of episodic memory (Lisman, 2005), visual perception (Fries et al., 2001), and sensory selection (Fries et al., 2002).

The evidence suggests that basic modes of dynamical organisation are reflected in brain rhythms (Steriade, Jones and Llinás, 1990). In addition the ‘communication through coherence’ hypothesis proposes that such synchronisation opens up communication channels between distant neuronal groups (Fries, 2005), providing optimal conditions for information transfer (Buehlmann and Deco, 2010). With these insights in mind it has also been suggested that transient periods of synchronisation and desynchronisation provide a mechanism for dynamically integrating and forming coalitions of functionally related neural areas (Shanahan, 2010).

5.2 Transient Dynamics

One way of studying neural dynamics is through modelling and simulation. The previous chapter demonstrated a close similarity between populations of neurons that collectively oscillate and simple oscillator models. Transient dynamics have been demonstrated in systems of phase-lagged, delayed and pulse-coupled oscillators that have been organised into a modular community-structured small-world networks akin to those found in the brain (Shanahan, 2010; Wildie and Shanahan, 2012). These systems exhibit interesting phenomena such as metastability, chimera-like states and high coalition entropy. Metastability is quantified by the variance of synchrony within an individual oscillator module over time, averaged for all modules in the system, and so characterises the tendency of a system to continuously migrate between a variety of synchronous states. Fixing time and calculating the variance across modules gives an index of how chimera-like the system is, indicating the level of spontaneous partitioning into synchronised and desynchronised subsets. Coalition entropy measures the variety of metastable states entered by a system of oscillators and is calculated from the number of distinct states the system can generate and the probability of each state occurring. As a collection these measures capture the ability and tendency of a system to best explore the space of dynamic synchronous coalitions. In the aforementioned work in which these transient dynamics were demonstrated, a key area within the oscillator network parameter space was identified where the combination of these measures is optimal. An embodied neural oscillator system tuned to

such a sweet spot would facilitate versatile exploration, integration and communication of functionally related areas throughout the behavioural problem solving process. However, the aforementioned work made use of simple abstract oscillator models rather than neural oscillator populations. The previous chapter demonstrated that neural systems display greater spectral complexity than simple abstract oscillator models, and we shall now discuss this.

5.3 Spectral Complexity in Neural Systems

It is increasingly common for simple oscillator models to be used as abstractions of oscillating neural populations in brain modelling (Breakspear, Heitmann and Daffertshofer, 2010). Whilst there is a greater perceived affinity to neural systems when moving from phase-lagged, to delayed, to pulse-coupled oscillator systems, the previous chapter experimentally demonstrated that such oscillator models display close behavioural similarities to networks of oscillating neural populations. However, these simulations illustrate how neural models display greater spectral complexity during synchronisation than more abstract oscillator models, with several oscillatory frequencies coexisting within an individual neural oscillator population. The chapter explored the relationship between simple oscillator models and their neural population cousins by emulating neurally the Kuramoto critical coupling experiment (Kuramoto, 1984) which showed an increase in synchrony as connection strength is increased in a uniformly connected network of simple oscillators. It was demonstrated that at the point of maximum synchrony the neural systems not only displayed several coexisting frequencies within an individual oscillator population but that the system also showed deviations from a measure of full synchrony likely caused by these additional fluctuating influences.

The spectral complexity of neural systems has been observed *in vivo* (Steriade, 2001). It has been hypothesised that slower oscillations provide a framework for other faster oscillations to operate such that fast oscillations communicate content while slow oscillations mediate transient connectivity (Nyhus and Curran, 2010). Very large networks are recruited during slow oscillations whereas higher frequency oscillations are confined to a small neuronal space (Buzsáki and Draguhn, 2004). Widespread slow oscillations modulate faster local events. Some such interactions have received particular attention, for example the nesting of gamma in theta during memory formation (Axmacher et al., 2006; Roopun et al., 2008). However, the phenomenon as a whole is not well understood. Within the same neuronal structure, neighbouring frequency bands, which are typically associated with different brain states, coexist but compete with each other. However, several rhythms

temporally coexist not only in the same structure but also in different structures and interact with each other (Buzsáki and Draguhn, 2004). How these different frequencies affect each other across populations is a question demanding further exploration, which is the aim of this chapter.

5.4 Metastability and Frequency Modulation Experiment

5.4.1 Experimental Setup

Much research has focused on measuring the effect when different populations of neurons synchronise to the same frequency (Buehlmann and Deco, 2010; Wildie and Shanahan, 2012; Womelsdorf et al., 2007), with further interest in correlations across frequency bands, as for example assessed by the mean local time-frequency energy correlation (Lachaux, Chavez and Lutz, 2003). It has been shown that, within a single neural population, coexisting oscillatory frequencies in different bands start, stop and restart. The simulations in this chapter demonstrate that the frequency of an oscillating population does not remain at a constant but may speed up and slow down over time. In other words the frequency fluctuates.

The aim of this chapter is to understand how the fluctuation in the frequency of one neural population's oscillation over time affects the other neural populations it is connected to. The approach taken is to build simulations of many neural oscillator populations, with each oscillating at a different intrinsic frequency. Fluctuations in the oscillatory frequency of each population become manifest when the populations are connected to each other. The results in this chapter demonstrate that the fluctuation in frequency in one neural population modulates the fluctuation in frequency in other neural populations, and that this influence increases with greater structural connectivity between the populations. It is shown that, this interaction of fluctuating frequencies in the network as a whole is able to drive different populations towards episodes of synchrony.

A number of different measures are then used to assess the network's overall dynamics. The frequency modulation between neural oscillators, averaged in the network as a whole, is measured using the *mean intermittent frequency correlation*. This is obtained by capturing in detail the intermittent fluctuating frequencies in each oscillator as fragments of times series (time series *strands*), and correlating these strands against other such stands, across bands, and across neural populations. This correlation measure is averaged for the network as a whole in order to give the global measure of modulatory influences in the network. Coalition entropy measures the variety of

metastable states entered by a system of oscillators and is calculated from the number of distinct states the system can generate and the probability of each state occurring. By identifying high levels of mean intermittent frequency correlation and coalition entropy it is possible to discern whether the neural populations in a network are modulating each other's oscillatory behaviour so as to explore many different coalitions. If the network also displays levels of synchrony (section 4.6) that promote information transfer during coalition formation, it is claimed that the network exhibits metastability. This is characterised by the tendency of a system to continuously migrate between a variety of synchronous states.

5.4.2 Extraction of Intermittent Frequency Strands

The work presented in this chapter aims at assessing the correlation between the fluctuating frequencies in different neural oscillators that are connected together in a network. In order to achieve this it is first necessary to extract the instantaneous frequency responses for each neural oscillator at each moment in time during a simulation. The standard techniques for doing this are to either use a short-time Fourier transform or a wavelet transform. To perform either first requires converting the firings of an oscillatory neural population into a continuous time signal upon which one of these transforms can be performed. Only the excitatory layer in a neural oscillator is used when producing this signal. The signal is obtained by first binning the number of spikes at each moment in time for the excitatory layer, and then passing a Gaussian smoothing filter over the data. Finally the signal is centred around its mean to obtain the continuous time signal upon which the transform can be performed.

Both Fourier and wavelet based approaches for extracting the time-frequency information from a signal suffer from shortcomings due to the time-frequency uncertainty principle (Yu et al., 2005). A Gabor wavelet transform has been chosen for use in this work, because the responses of Gabor wavelets have optimal properties with respect to the time-frequency uncertainty principle (Yu et al., 2005). The Gabor wavelet used had a centre frequency of 0.6 Hz and was applied with a continuous wavelet transform using scales from 1 to 100 in increments of 0.1, and a delta of 0.001. Figure 5.1A shows the scalogram of a wavelet transform taken from the excitatory layer of a neural oscillator that was placed in a network of oscillators, with each oscillator oscillating at a different intrinsic frequency. In this setup there is a given probability of connecting each oscillator to another in the network, and a given synaptic connection weight for the connections formed between oscillators. Figure 5.1B shows a raster plot of the firing behaviour of the same excitatory layer between time

points 1000 ms and 1500 ms in the simulation. It can be seen that the spacing of the bursts of firing between 1050 ms and 1150 ms is wider and thus at a different, slower frequency to the spacing of the bursts of firing between 1200 ms and 1275 ms. The wavelet response to the slow and then faster bursting can be seen as a change from low to high frequency on the scalogram in the same temporal area and around the frequency range from 30 Hz-50 Hz. These responses are deviations from the regular 33 Hz bursting that the PING oscillator was evolved to fire at and are due to the interaction with the other oscillator nodes.

The Gabor wavelet produces a blurred impulse response around given frequency responses at each point in time. The blurring from a Gabor wavelet is in the form of a Gaussian (Yu et al., 2005), as illustrated by a time slice at time point 1460 ms shown in figure 5.1C taken from the scalogram in figure 5.1A. Further techniques have to be applied to the transformed data in order to extract the instantaneous frequency information. Standard ridge and skeleton methods do not perform well when there are many components, some of which remain very close for a while and separate again, or when they can die out, or when new ones can appear from nowhere (Daubechies and Maes, 1996). As can be seen by the scalogram in figure 5A the data in the work presented here is of this type. Drawing upon the Gaussian nature of the impulse response from the Gabor wavelet, a technique of fitting a sum-of-Gaussians model to the transformed data at each point in time is applied (Yu et al., 2005). Figure 5.1C shows such a fitting. Identifying the means and magnitudes of the means of the fitted Gaussians gives the instantaneous frequencies and their amplitudes respectively.

The next stage in preprocessing the data requires forming a time series of the instantaneous frequencies as they fluctuate over time, what I call a *strand*. These fluctuating frequency responses may also be intermittent due to the frequency response dropping out and starting again. Zero values are substituted into the time series strands during the drop out moments to indicate the absence of a frequency response at those times. There may be many coexisting frequencies for each neural oscillator at each time point in a simulation, and therefore many coexisting strands. To obtain these strands, after the instantaneous frequencies at each point in time for an oscillator are calculated, the movement of each frequency is tracked over time so as to link them together into a single time series fragment.

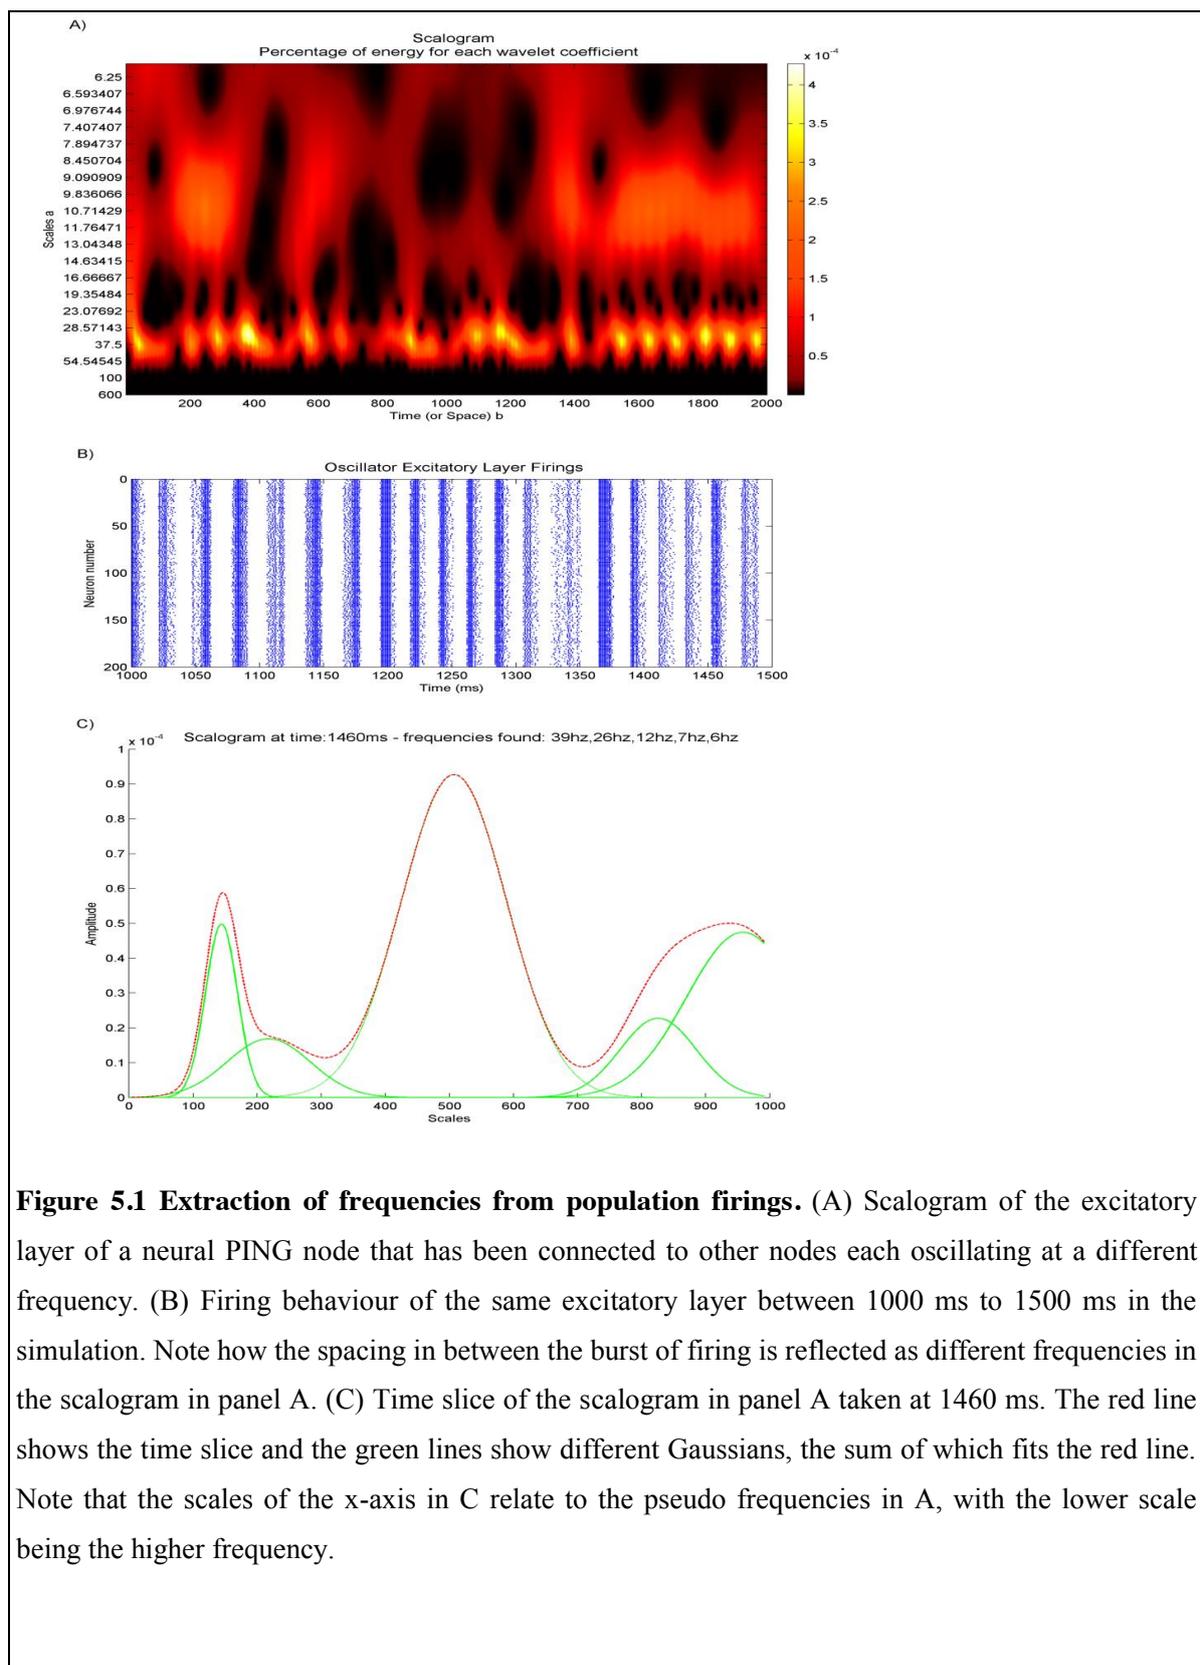


Figure 5.1 Extraction of frequencies from population firings. (A) Scalogram of the excitatory layer of a neural PING node that has been connected to other nodes each oscillating at a different frequency. (B) Firing behaviour of the same excitatory layer between 1000 ms to 1500 ms in the simulation. Note how the spacing in between the burst of firing is reflected as different frequencies in the scalogram in panel A. (C) Time slice of the scalogram in panel A taken at 1460 ms. The red line shows the time slice and the green lines show different Gaussians, the sum of which fits the red line. Note that the scales of the x-axis in C relate to the pseudo frequencies in A, with the lower scale being the higher frequency.

The algorithm for forming the frequency time series strands has three parts. The first part is simply to sequence the nearest frequencies in time into a strand as follows:

```

T=start_time.

while T is not equal to end_time.

    for each unassigned instantaneous frequency at time point T

        Create a new strand containing that frequency.

    end

    T=T+1.

    while there are strands and frequencies within the distance limit L.

        Find the strand at time point T-1 with the closest frequency to one
        of the instantaneous frequencies at time point T.

        if the frequency is within limit L

            Add the frequency to the strand and remove the strand from further
            consideration until the next iteration.

        end

    end

end

```

An oscillator may be in a different state at different times. A state is described as several coexisting frequency responses in different frequency bands present at one time. The state may evolve over time such that the response in each frequency band fluctuates. Further to this, the state may dramatically change such that a frequency response in any particular band dramatically changes such that there is no smooth transition from one moment in time to the next. Such dramatic changes are called bifurcations. It is necessary to cope with bifurcations in the oscillator behaviour when an oscillator in a particular state A1, flips to another state B1, and then flips again to a state A2 such that states A1 and A2 have the same number of coexisting frequencies and these frequencies have approximately

the same values. Hence the system returns to its original state (A1) after the middle state B1. In each state there may be several coexisting frequency strands. It is desirable that the strands in the original state A1 and its return state A2 after the middle state, are stitched together so as to maximise the strand length and as a result the correlation. The distance between the frequencies in the strands in the state A1 and state A2 may be near enough within a limit L to make a direct match as in the previous algorithm, due to the fact that there is a close continuation between the two. However, there are situations in which the values of the frequencies in A2 are not near enough for direct matching, but instead have values similar to how state A1 would have been at that time if the bifurcations had not occurred and state A1 had instead continued developing. That is to say, whilst being the same state as A1, state A2 is in a later stage of development. In such a situation regression is used to project where the frequencies of strands in the original state A1 would have progressed to, and match these projections to the frequencies of the strands in A2. A maximum frequency distance limit L is applied as before on this projected matching.

In order to stitch states in this way, all the strands which share the same start time are first grouped together so as to identify them as being in the same state. The state matching algorithm then preferentially matches states nearest to each other whose strands have the closest frequencies or projected frequencies. There is a maximum time limit between states for which such stitching is allowed to occur. In order to get the best matching between states, the algorithm is first performed with the constraint that stitched states must contain the same number of strands, and then performed again without this constraint.

After the state stitching has been carried out the individual strands that are contained within the states are extracted, as we only consider pairs of individual strands during correlation. The strands are time series of fluctuating frequencies sequenced by closeness. Each strand will have a start and end and may contain zero values in its time series where the frequency dropped out due to a bifurcation.

5.4.3 Mean Intermittent Frequency Correlation

For each pair of oscillators, m and n , there is a collection of fluctuating frequency strands scattered over the frequency domain and stretching over time. Each strand in oscillator m is correlated with each strand in oscillator n , for all oscillator combinations in the network. This is done by passing a 100 millisecond window across time in incremental steps of 1 millisecond. In each of these windows the time series data is taken for all pairs of strands i and j , where i and j are from different oscillators

m and n respectively. For each pair of windowed strands, time points are removed from each strand where both strands do not have a response at that time point in the window, or when both the frequencies in the strands at one time point are the same as at the previous time point. This results in two time series strands for the window at time t , $w_{m,i}(t)$ and $w_{n,j}(t)$. Both are the same length and are potentially shorter than the window size. Each time series only contains data where both original strands have a frequency response and they are both fluctuating. These two series are then correlated. Correlations are only selected where the coefficient is greater than or equal to 0.5 or the coefficient of the anti-correlation is less than or equal to -0.5, and the p-value for either is less than 0.05. By randomising the order of one of the series and performing the same correlation and selection process a phantom correlation is obtained. Phantom correlations are used to confirm the importance of the measure of real correlation found. For both types of correlation the *mean intermittent frequency correlation* is calculated as follows:

$$\text{MIFC} = \frac{1}{t_{\max} - 100} \sum_t^{t_{\max} - 100} \sum_n \sum_{m \neq n} \sum_i^{I(t)} \sum_j^{J(t)} \frac{W(t)}{100} \left| \text{coef}(w_{m,i}(t), w_{n,j}(t)) \right|$$

where n and m are oscillators, $I(t)$ and $J(t)$ are the total number of strands in the window at time t for each oscillator respectively, and i and j are particular strands within each oscillator. *coef* defines the value of a significant correlation coefficient as previously described. $W(t)$ is the length of the two series $w_{m,i}(t)$ and $w_{n,j}(t)$, that only contain time points that have a fluctuating frequency response in both original windowed strands. 100 is the length of the window. Thus the significant correlation coefficient is normalised according to the length of the two series in that window. t is the time of the particular window and t_{\max} is the length of the simulation time. The metric calculates all pairwise significant frequency correlations between all oscillators, normalises them by their length, and averages them over time.

5.4.4 Coalition entropy

Coalition entropy measures the variety of metastable states entered by a system of oscillators (Shanahan, 2010). Coalition entropy is only calculated using the excitatory neuron layers in the oscillators. As with the synchrony metric of section 4.6, the phase of each oscillator is calculated at each time point t using a Hilbert transform. For coalition entropy, clustering was then performed at each time point by picking the two most synchronous oscillators/coalitions using the first equation defined for the synchrony metric. Once a pair was identified they were joined to form a new coalition,

and the new coalition's mean complex exponential phase was calculated for use in the future most synchronous pair selection process. A threshold of 0.05 from full synchrony was used to limit the cluster merging. The process was repeated until no oscillators/coalitions fell within the threshold to allow merging into a new coalition.

Having identified the synchronous coalitions at each point in time, the probability $p(s)$ of each coalition s occurring was calculated from the number of times it appeared throughout the simulation. The coalition entropy H_c was then calculated as follows:

$$H_c = -\frac{1}{\log_2 |S|} \sum_{s \in S} p(s) \log_2(p(s))$$

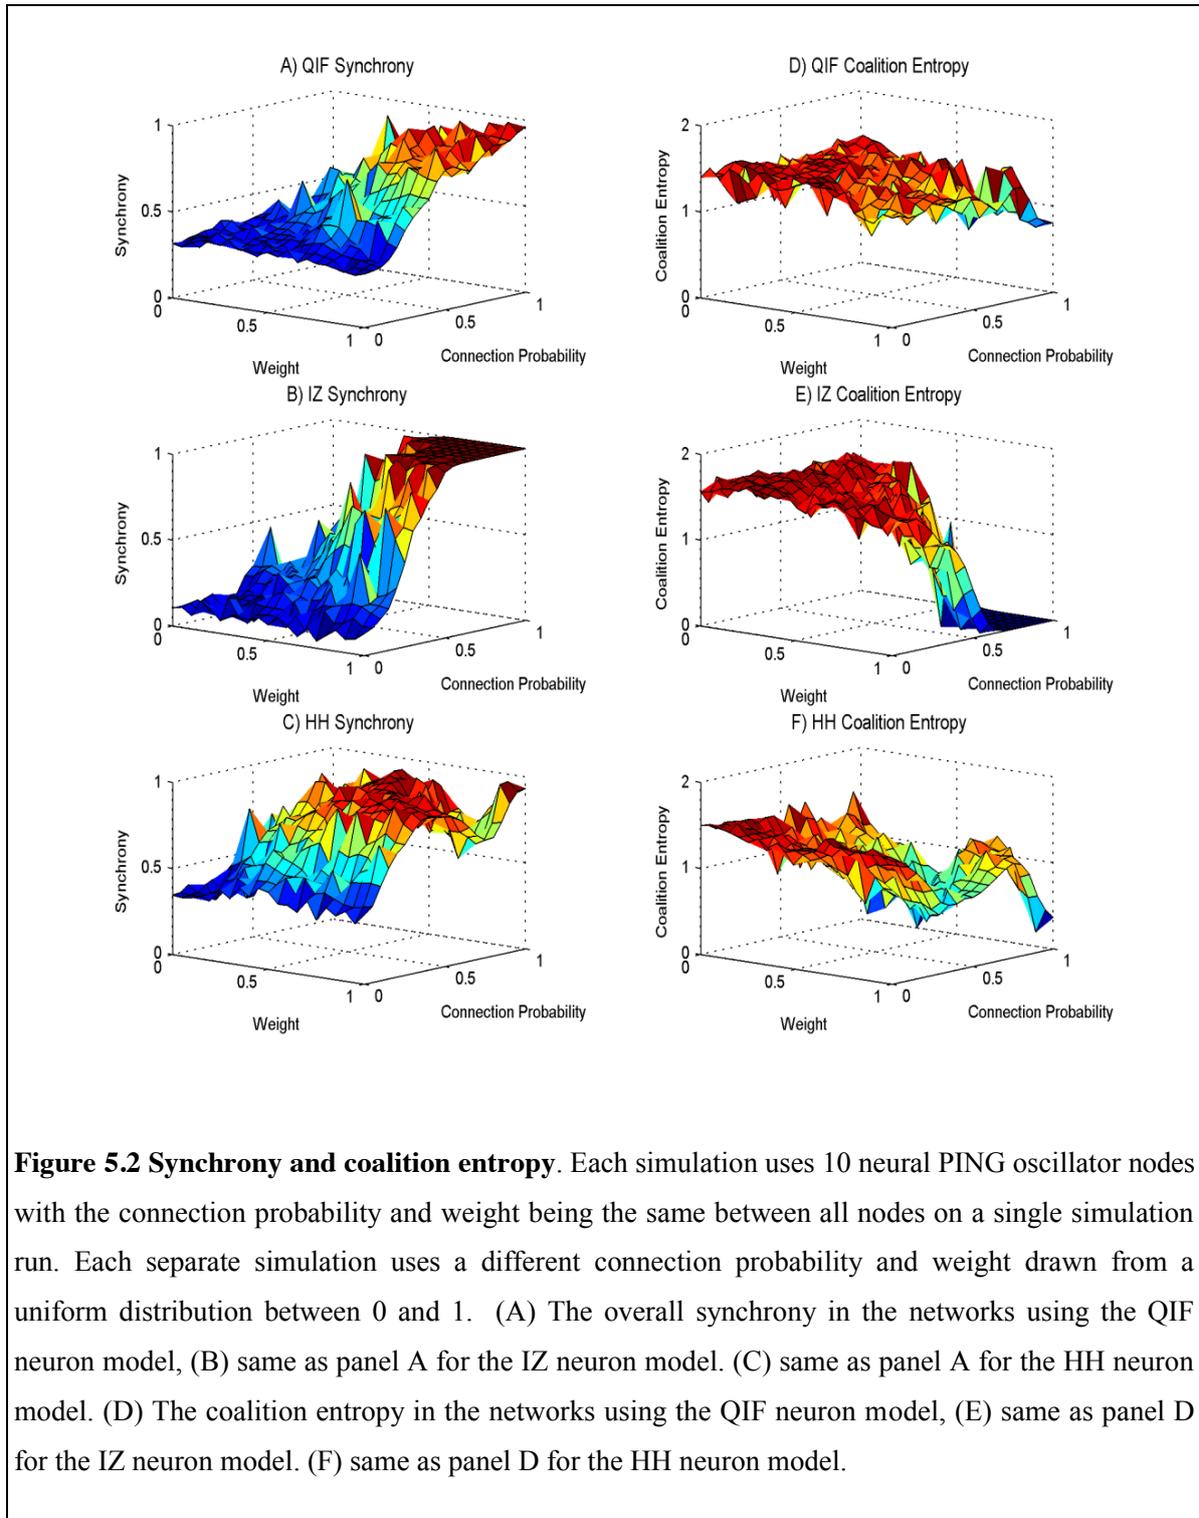
where $|S|$ is the number of possible coalitions given the number of oscillators in the system given the set of all possible coalitions S .

5.4.5 Results

A series of experimental simulations were performed in each of which 10 neural PING oscillators were chosen from the set that had been evolved with intrinsic frequencies ranging from 30Hz to 50Hz. The probability of one oscillator providing neural input to another was determined with a given probability C . The probability C was the same for all oscillator to oscillator connections in the same experimental simulation. Given that a connection was established from oscillator n to oscillator m the excitatory neurons in oscillator n would form synaptic connections to the excitatory neurons in oscillator m . The number of synaptic connections formed was 20 percent of the 40000 possible synaptic connections from the 200 excitatory neurons in oscillator n to the 200 excitatory neurons in oscillator m . These were selected at random. For all synaptic connections formed the weight of the synapse was set to W . The value for W and C were randomly chosen at the beginning of each experimental simulation from a uniform distribution between 0 and 1. 250 simulations were performed for the QIF neural model, 250 simulations for the IZ neural model, and 250 simulations for the HH neural model. As the weight and connection probability for each simulation were chosen at random these data points are scattered throughout the parameter space. The 250 simulations thus constitute a scattered sweep of weight and the inter-oscillator network connection scarcity. Figures 5.2-5.7 show various measures taken from these 250 simulations of QIF, IZ and HH neuron models. These are analysed and discussed in detail below. In each of these figures a surface has been fitted to the underlying trend of the 250 data point for each measure depicted.

Throughout each simulation, external stimulus input was provided to each neural oscillator from a Poisson process with parameter $\lambda = 4.375$. For each model, the inputs were scaled in order to provide sufficient stimulus to induce firing. Each experiment was run for 2000 ms of simulated time. After each experiment, the firing activity of the excitatory layers in each oscillator was used to calculate synchrony, coalition entropy and the mean intermittent frequency correlation as described above. The first 500 ms of each simulation were discarded in the calculation of these metrics to eliminate initial transients.

Figures 5.2A, 5.2B and 5.2C show the synchrony through the parameter sweep for QIF, IZ and HH models respectively. Unsurprisingly, and in accord with the findings of the previous chapter, synchrony increases as connectivity increases. At their maximum value for synchrony the QIF and HH neural systems exhibit deviations from full synchrony. The value of 0.25 synchrony in the area of low weight and low connection probability represents no synchrony at all except coincidental alignments in phase. It is interesting to note that the IZ model shows a critical area in which synchrony increases dramatically as connection strength increases, until it eventually completely flattens out at full synchrony. Figures 5.2D, 5.2E and 5.2F show the coalition entropy through the parameter sweep for QIF, IZ and HH models respectively. The trend for coalition entropy takes the reverse form to synchrony, decreasing as synchrony increases, and the oscillators become more aligned in phase for more of the time. The measure of coalition entropy used cannot distinguish coincidentally synchronous coalitions from those that are genuinely coupled. However, when coalition entropy is contrasted with the graphs for synchrony we can get an idea of what is happening. Regions of the parameter space with low weight and low connection probability exhibit high coalition entropy, but the same regions present low values for synchrony. This suggests that the many coalitions that appear are constituted by very short coincidental alignments in phase that are not capable of significant information transfer (Buehlmann and Deco, 2010). The mid parameter space area shows fairly high values for synchrony, indicating the capacity for substantial information transfer, as well as high coalition entropy indicating transfer between many different groups at different times. The region of the parameter space in which the weight and connection ratio are high facilitates more information transfer but less variation in coalitions. It is interesting to note that the IZ model shows a much more pronounced slope decreasing coalition entropy as connection strength increases until it levels off with no coalition entropy. This levelled off area appears as an artefact on all IZ measures.



Figures 5.3A, 5.3B and 5.3C show the average number of coexisting frequencies in a single oscillator at any one time in each simulation for QIF, IZ and HH models. Interestingly, for the QIF and IZ models, as the causal influences increase through stronger weights and greater connectivity, the number of coexisting frequencies rises. This indicates that stronger causal interactions between neural populations, that otherwise oscillate at a single intrinsic frequency, are a source of increased spectral complexity. In the high connectivity area in which synchrony is maximal the IZ model shows low numbers of coexisting frequencies, indicating that the system is converging on a single synchronous frequency. The HH model shows a dip in the mid area of the parameter space, after which the number of coexisting frequencies rises. Whilst this latter area also demonstrates that causal interactions increase spectral complexity in areas where there are stronger inter-oscillator influences, it is interesting to note that in the area of weaker

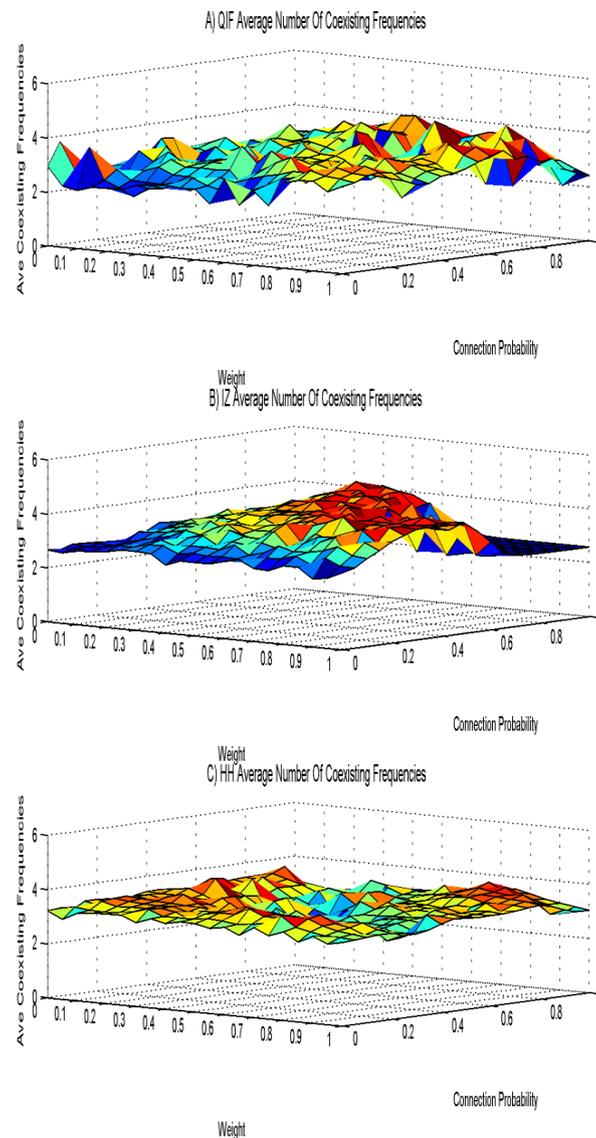


Figure 5.3 The number of coexisting frequencies. (A) The average number of coexisting frequencies per oscillator at each time point in the networks using the QIF neuron model, (B) the same as panel A for the IZ neuron model, (C) the same as panel A for the HH neuron model.

influences this model also generates a large number of coexisting frequencies. This latter phenomenon will be elucidated later.

Taking a look at the number of significant correlations found through the parameter sweep, figures 5.4A, 5.4B and 5.4C show that QIF, IZ and HH models display an increased number of correlations as the synaptic weight and connectivity increases, although the HH model has a less pronounced incline. The increase has a similar trend to that of synchrony. The data indicate that correlated

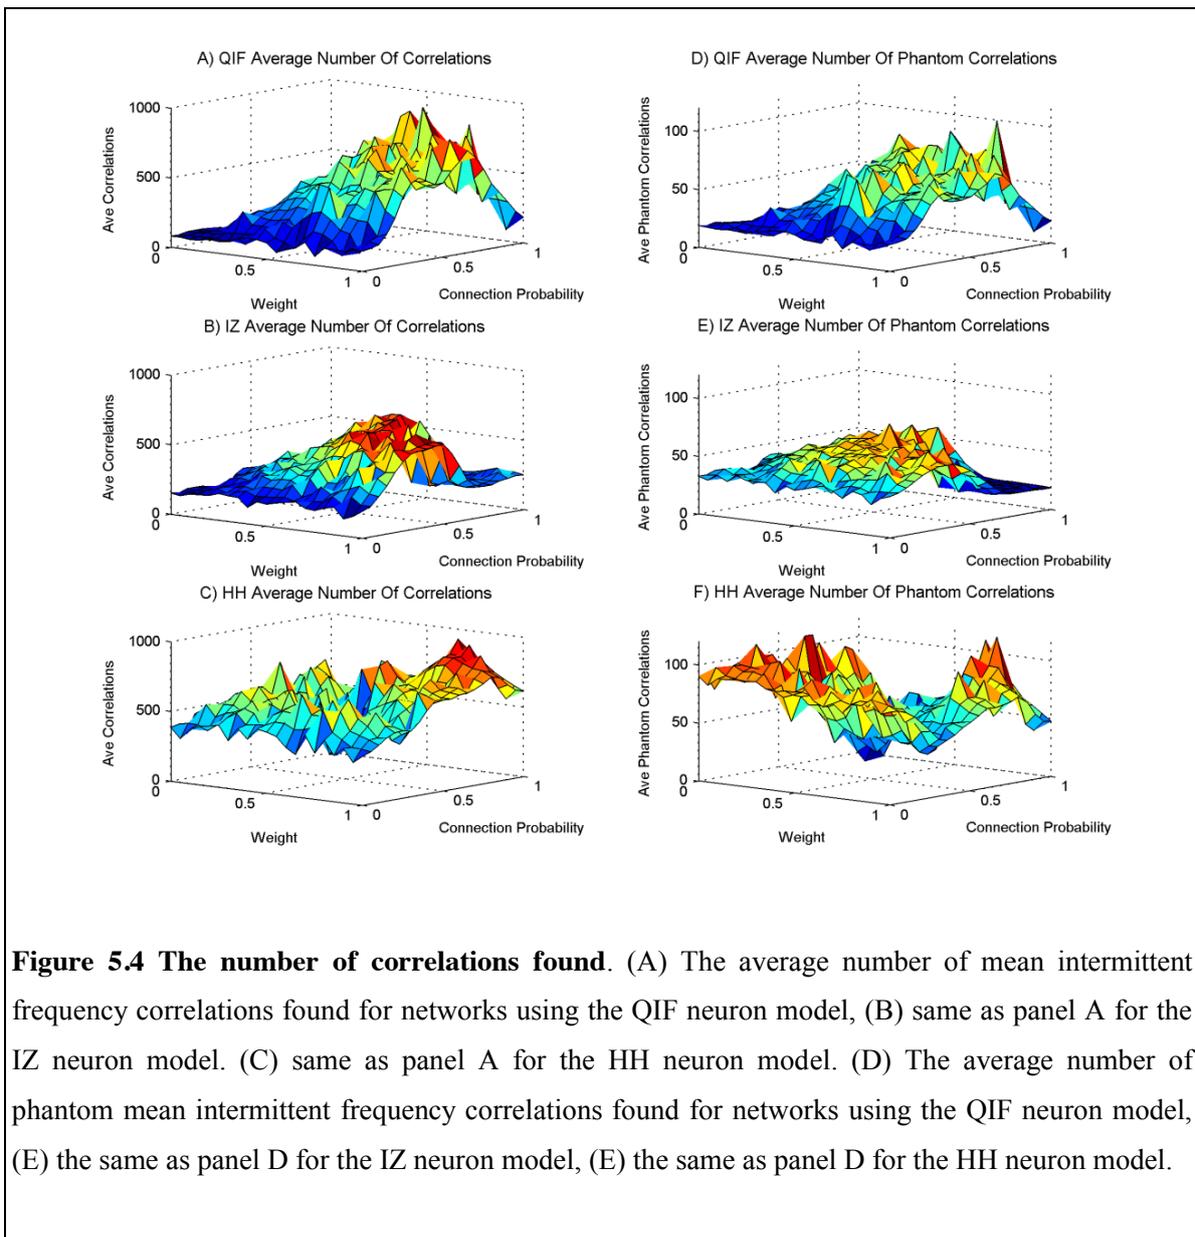
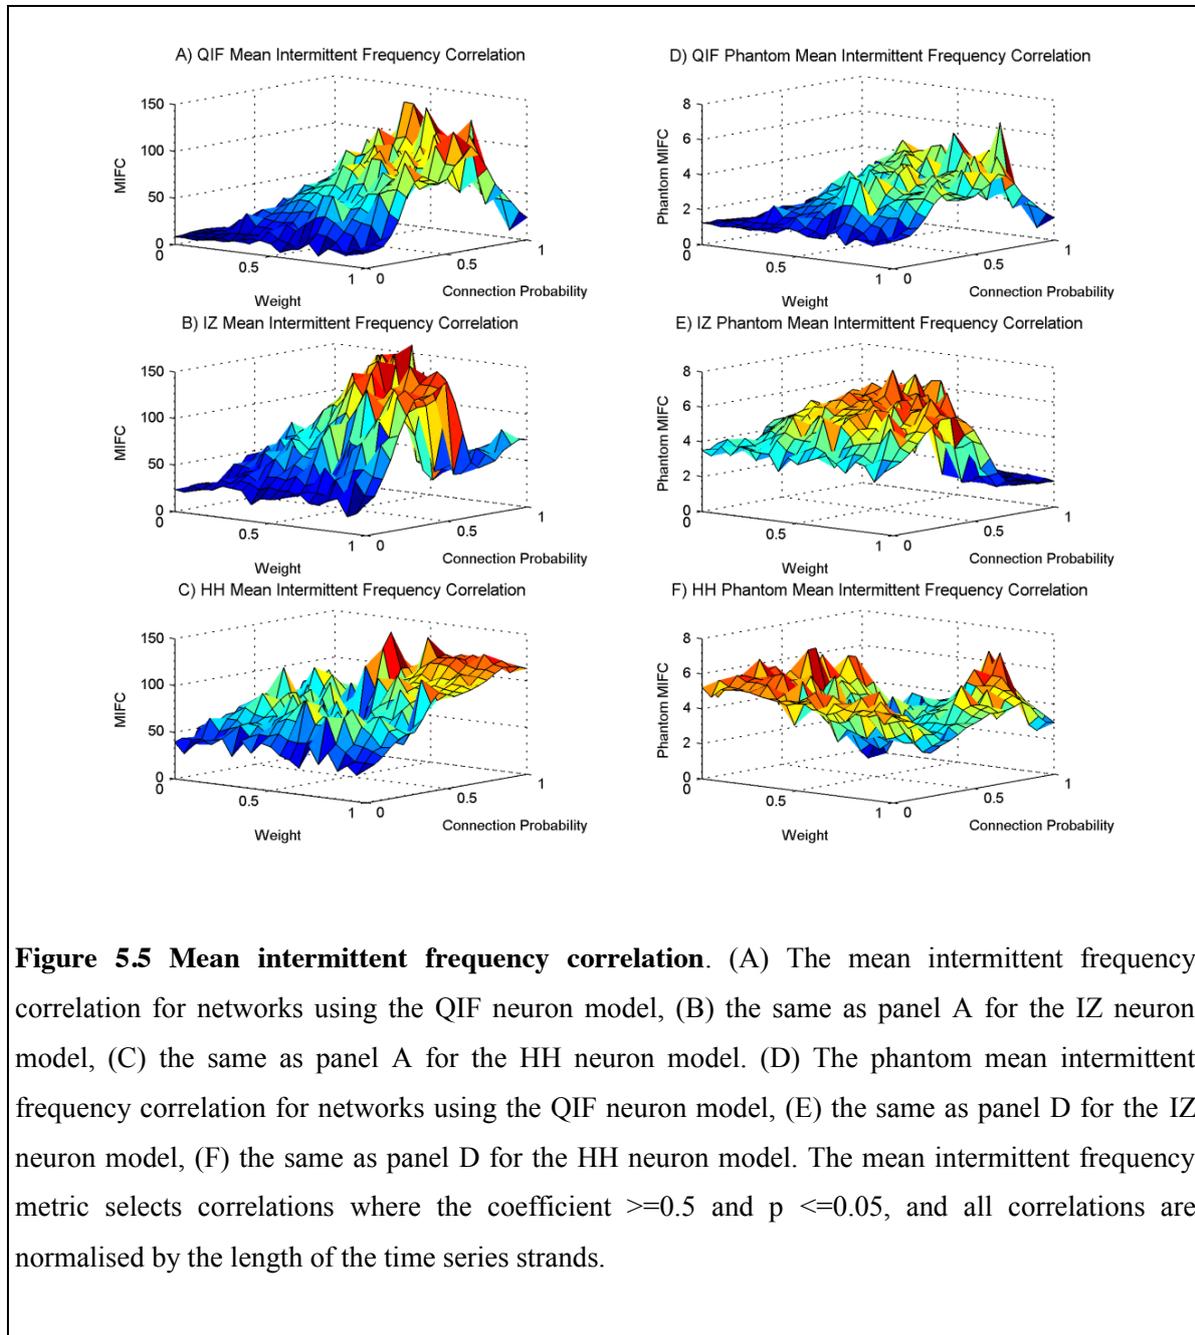
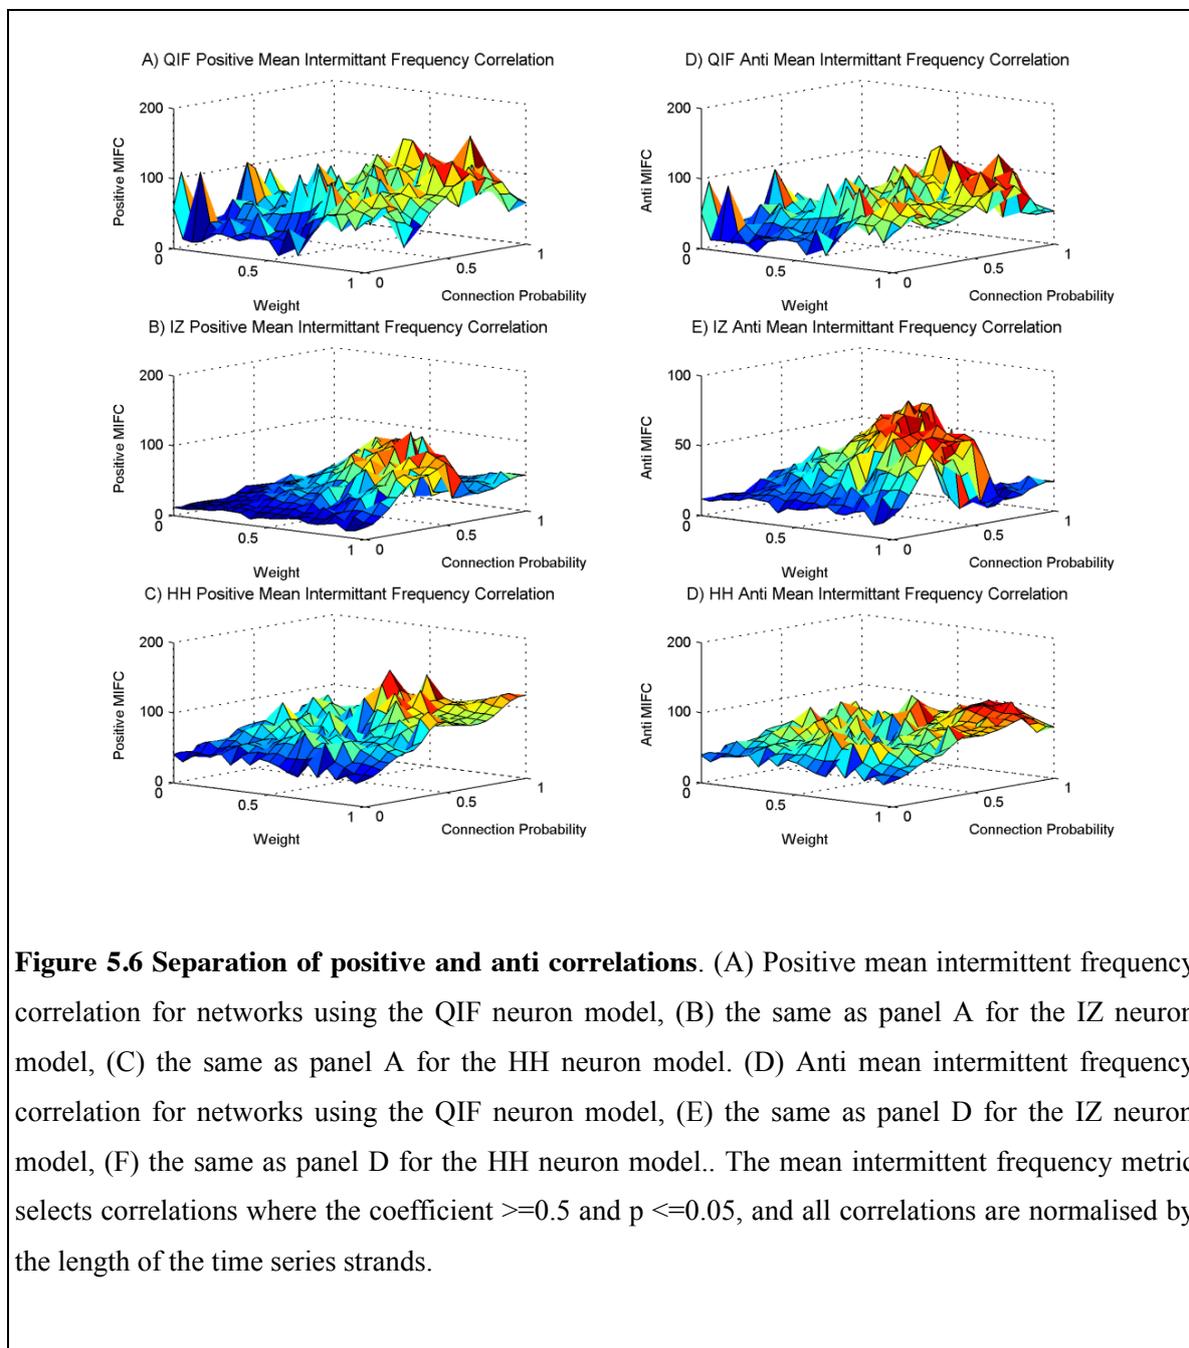


Figure 5.4 The number of correlations found. (A) The average number of mean intermittent frequency correlations found for networks using the QIF neuron model, (B) same as panel A for the IZ neuron model. (C) same as panel A for the HH neuron model. (D) The average number of phantom mean intermittent frequency correlations found for networks using the QIF neuron model, (E) the same as panel D for the IZ neuron model, (F) the same as panel D for the HH neuron model.

fluctuations in frequency imply more episodes of synchrony, suggesting that the fluctuating influences between oscillators are moving each other towards synchronous behaviour. As the number of significant correlations is so high, we can conclude that this influence works across frequency bands. To control for coincidental (“phantom”) correlations, the number of significant correlations found compared to the number found when randomising one of the time series before correlating is presented in figures 5.4D, 5.4E and 5.4F. We can see that for the QIF and IZ models, although many significant phantom correlations are found and they follow the same trend, the number found is an order of magnitude less than the number of correlations found in the original data. The HH model displays a similar ratio between real and phantom correlations. However, the region of low weight and connection probability shows large numbers of correlations in line with the greater number of coexisting frequencies found in that area in figure 5.3C.

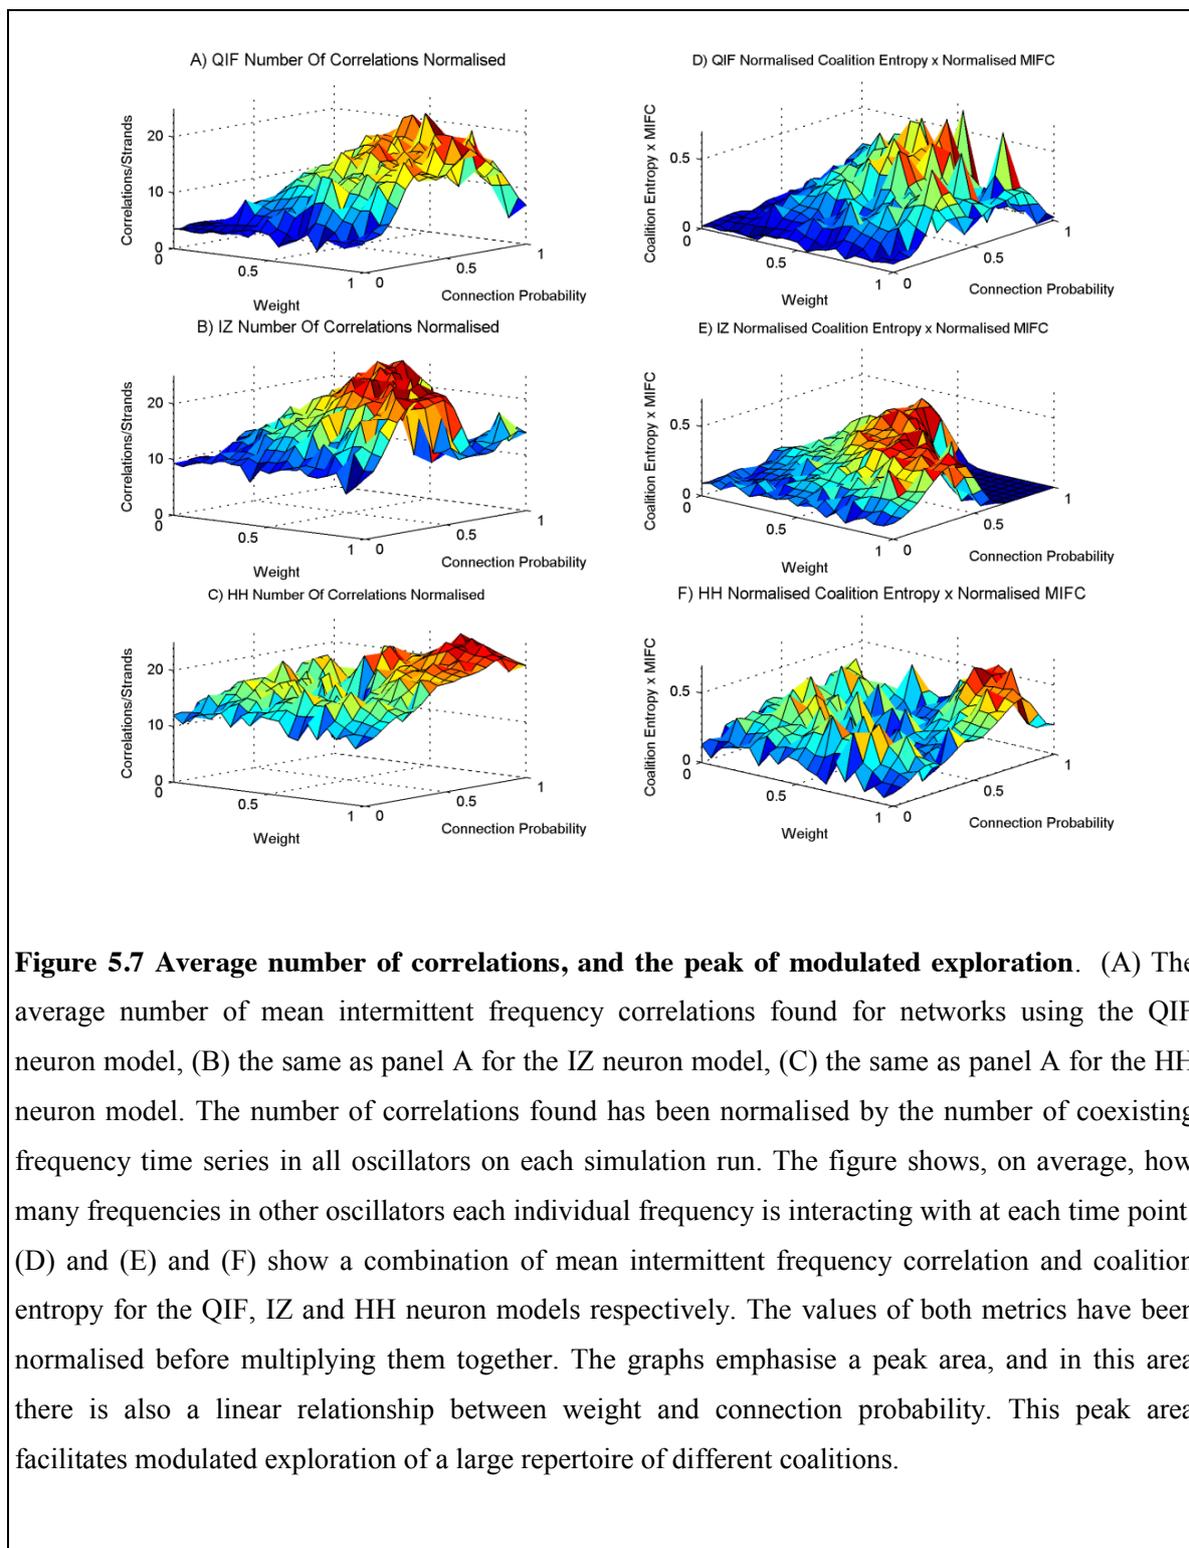
The mean intermittent frequency correlations are shown in figures 5.5A, 5.5B and 5.5C for QIF, IZ and HH models respectively. As causal influence between oscillators increases in the network this correlation measure increases, meaning that correlation directly reflects causation in this case. Figures 5.6A-5.6F show that when separating these data into positive correlations and anti-correlations both follow the same trend. The metric not only identifies the significant correlations but also normalises each of these correlations by the length of the intermittent fluctuating frequency time series. The resulting values are therefore always much less than the number of correlations. The mean intermittent frequency correlation shown for phantom correlations in figures 5.5D, 5.5E and 5.5F is very low, peaking at around 7 compared to real correlations, which peak at around 140. The ratio is double that found for the simple ‘number of significant correlations found’ of figure 5.4 and so is even stronger justification for the claim that the correlations found are significant. For the HH model, the area of low weight and connection probability shows small mean intermittent frequency correlation values, in contrast with the corresponding ‘number of correlations found’ shown in figure 5.4C due to greater number of coexisting frequencies seen in figure 5.3C. This is because, although significant correlations are found, they are being normalised to a lesser value by length, meaning that these are very short time series. This is further reason for discarding high coalition entropy values in this area, on the grounds that they are not due to any consequential interactions but are merely coincidental. The performance of HH networks around this low parameter region is therefore very erratic compared to the behaviour in the mid and high connectivity regions which exhibits stable and modulatory influences.





In figures 5.7A, 5.7B and 5.7C, the number of correlations found has been normalised by the number of coexisting frequency time series in all oscillators on each run. These figures show, on average, how many frequencies in other oscillators each individual frequency is interacting with at each time point. The mid area of figures 5.3A, 5.3B and 5.3C show ≈ 3 coexisting time series per oscillator at each time point. The mid and high parameter region in figures 5.7A-5.7C show many more than the 9 interactions we would expect if each frequency was only interacting with frequencies in other oscillators that are in the same frequency band. We can safely conclude from this that frequencies in different neural populations communicate across bands. This type of complexity is not manifest in simple oscillator models, a shortcoming that is most evident at high levels of synchrony when simple oscillator models, unlike systems of neurons, display only a single shared frequency.

Figures 5.7D-5.7F show a combination of mean intermittent frequency correlation with the coalition entropy. To obtain this combination the mean intermittent frequency correlation and the coalition entropy were normalised, and multiplied the results together. For QIF and IZ models, there is a range from weight value 0.35 and connection probability 1 to weight value 1 and connection probability 0.35 at which the amount of correlation between fluctuating frequencies across oscillators coincides with a measure of coalition entropy such that they are at a combined peak. This area, in which the two metrics are balanced, facilitates metastable dynamics in which there is a richness of influence and interaction between different oscillators and across frequency bands modulating each other's behaviour, enabling the exploration of a large repertoire of different coalitions. It is noteworthy that there is a linear relationship between weight and connection probability at which this is best facilitated. This indicates that both connection probability and synaptic weight have a complementary effect upon metastability. The medium-to-high level of synchrony in this area further suggests that the conditions for information transfer between populations are fulfilled (Buehlmann and Deco, 2010). These traits are desirable in order to facilitate a system versatile at exploration, integration and communication between functionally related areas during cognitive processing (Chialvo, 2010; Kelso, 2012; Shanahan, 2012).



5.5 Discussion

A general rule of thumb for oscillator systems is that greater connectivity produces more synchrony. Unlike simple oscillator models, systems of spiking neurons display great spectral complexity with many coexisting frequencies within a single oscillator at one time. The work presented in this chapter demonstrates that this complexity increases with connectivity, not only in the number of coexisting frequencies but also in the amount of interaction across frequency bands. As causal interactions increase, so does correlation between these fluctuating frequencies, as well as the tendency towards more and longer episodes of synchrony and information transfer, which implies that they are modulating each other towards communication.

The focus of this chapter is the theory of metastability in which neural behaviour produces episodes of synchronisation and desynchronisation between oscillating populations, for which the combined effect amongst a collection of oscillators is to explore many different coalitions over time. The results presented identify an area in the weight and connectivity space at which spiking neuron models are at a balance in which coalition entropy is exhibited due to the influential modulation between populations through their oscillatory behaviour. In this area of balance the neural systems influence each other across frequency bands in a way that promotes exploration of, and communication between, coalitions. This is because the fluctuating oscillatory frequencies in each neural population modulate each other so as to drive the system towards episodes of synchrony between different neural populations, enabling communication between them. Whilst doing this, the variation in synchronous coalitions of neural populations over time is kept very high, and hence this area of the connectivity space may be described as encouraging exploration.

A dynamical system whose component parts interact so as to direct the system through varieties of coalitions would form a good basis for contextual exploration as well as integration among, and communication between, functionally related areas during cognitive processing. Further to this, maintaining a large repertoire of synchronous coalitions promotes versatile exploration of novel functional combinations, a desirable trait when problem solving.

6 Metastability and Plasticity

The previous chapter demonstrated how metastable dynamics that support transient synchronous coalition formation naturally arise in networks of spiking neurons. Using computer simulations, this chapter extends this finding by establishing that networks of spiking neurons connected according to a modular small-world topology give rise to heightened levels of metastability. Further to this, the results presented show that such modular small-world topologies form as a natural consequence of synaptic plasticity being present in the neural network during the interaction between many different oscillating neural populations.

6.1 Transient Synchronisation and Brain Coordination

There is extensive empirical evidence that shows the tendency for several different brain areas to synchronise for short amounts of time, and that different synchronous groups appear at different times. For example, Betzel *et al* (2012) studied fluctuations in functional connectivity of human resting state networks from EEG recordings. They report that such fluctuations occurred on a time scale of tens to hundreds of milliseconds, and involved spatially local and remote sites, resulting in fast reconfigurations of network states. It has been suggested that such transient periods of synchronisation and desynchronisation provide a mechanism for the formation of coalitions of functionally related neural areas (Chialvo, 2010; Kelso, 2012; Shanahan, 2012). In addition, the ‘communication through coherence’ hypothesis proposes that synchronisation opens up communication channels between distant neuronal groups (Fries, 2005), providing optimal conditions for information transfer (Buehlmann and Deco, 2010). Computer models that reproduce the fluctuations in human resting state networks by using metastable systems of coupled oscillators lend further support to the hypothesis that transient periods of neural synchronisation and

desynchronisation are important (Cabral et al., 2011; Cabral et al., 2013; Hellyer et al., 2014). Whilst it remains an open question what the underlying mechanism is that gives rise to these observed phenomena, the previous chapter goes some way in demonstrating that intrinsic metastable dynamics that appear in simulation may be the cause of the observed phenomena *in vivo*. However, it is unclear what structural and functional characteristics are required to facilitate such apparently metastable dynamics. An investigation of the required functional and structural characteristics are the subject of this chapter.

6.2 Local and Global Processing

Kelso and Tognoli (2007) argue that the fundamental property of complex systems that operate in a ‘metastable’ dynamical regime is the duality of large-scale processing by sets of distributed, interconnected areas and local processing within those areas. The simulations presented in this chapter add weight to the intuitive idea that such local-global dynamics are promoted by modular small-world connectivity. Small-world networks are characterised as having a high clustering coefficient but with a short characteristic path length (Watts and Strogatz, 1998). A small-world network is also modular if it can be partitioned into sets of nodes (modules) that have dense within-module connections but are only sparsely connected to other modules. Clustering and modularity tend to favour localised processing, while a short characteristic path length enables easy connectivity between arbitrary sites. As the results presented confirm, a balance of these two properties promotes the sort of local-global dynamics that is the hallmark of metastability.

6.3 Plasticity Experiment

6.3.1 Experimental Setup

This chapter demonstrates not only that modular small-world structure promotes the transient formation of synchronous coalitions through metastable dynamics, but also that the modular small-world topology required to support this dynamics will naturally arise through synaptic plasticity. To do this, networks in which populations of spiking neurons are nodes are simulated. The networks are divided into populations, and each population is configured to collectively oscillate at a different intrinsic frequency. The interaction of these oscillating neural populations leads to complex dynamics. Acting in concert with synaptic plasticity, these dynamics naturally result in the restructuring of the

network topology to have modular small-world connectivity, which in turn heightens levels of metastability in the system.

6.3.2 Small-world index

The small-world index (Humphries, Gurney and Prescott, 2006) is used to quantify the degree of small-world connectivity within networks of oscillators. Within small-world networks most nodes are not neighbours, but most nodes can be reached from every other node by a path which consists of small number of hops. Small-world networks have a small characteristic path length, and a clustering coefficient that is significantly higher than expected by random chance.

The path length between any pair of nodes in a graph G is defined as the number of nodes in the shortest path connecting the two. The average path length over all pairs of nodes in G is λG . The clustering coefficient of a node $i \in G$ is defined as the ratio of actual edges between neighbours of i to all possible edges between those neighbours. The clustering coefficient γG of the graph G is the mean of the clustering coefficient over all $i \in G$. The small-world index of G (σG) is given by:

$$\sigma G = \frac{\gamma G / \gamma_{rand}}{\lambda G / \lambda_{rand}}$$

Where γ_{rand} and λ_{rand} are the clustering coefficient and average path length of a random graph with the same connectivity. The index is high if G exhibits comparable mean path length with a higher clustering coefficient to a randomly connected graph with the same number of nodes and edges.

6.3.3 Modularity

Many networks divide naturally into communities or modules with dense connections within communities but sparser connections between them. Modularity (Leicht and Newman, 2008) is used to quantify the degree of modular connectivity within networks of oscillators. Large positive values of modularity indicate when there are more edges within communities than would be expected by chance. For directed networks, the crucial point is that the expected positions of edges in the network depend on their direction. Consider two nodes, i and j , and suppose node i has high out-degree but low in-degree while node j has the reverse situation. In this example an edge is more likely to run from i to j than vice versa. As it is a bigger surprise to find an edge from j to i than from i to j , it

should make a bigger contribution to the modularity, since modularity should be high for statistically surprising configurations.

The probability of an edge from node j to node i is $(k_i^{\text{in}}k_j^{\text{out}})/m$, where k_i^{in} is the in degree of node i , k_j^{out} is the out degree of node j , and m is the total number of edges in the network. For a given partitioning of the nodes into communities, the modularity Q is defined as:

$$Q = \frac{1}{m} \sum_{ij} \left(A_{ij} - \frac{k_i^{\text{in}}k_j^{\text{out}}}{m} \right) \delta_{c_i, c_j}$$

where A_{ij} is an element of the adjacency matrix, δ_{ij} is the Kronecker delta (equal to 1 if i and j are the same and 0 otherwise), and c_i is the label of the community to which node i is assigned. The task of finding a partitioning of the nodes that maximises Q is known to be NP-complete, so practical methods based on modularity optimisation make use of approximate schemes and heuristics. The work presented uses the stochastic algorithm in the brain connectivity toolbox of (Rubinov and Sporns, 2010), where slightly differing results may be produced each time modularity is calculated. In the present work, the maximum over 100 samples of calculated modularity is used to obtain the final modularity measure.

6.3.4 Knotty centrality

The network measure called ‘knotty centrality’ quantifies the extent to which a given subset of a graph’s nodes constitutes a densely intra-connected topologically central connective core (Shanahan and Wildie, 2012). For a directed graph G with N nodes, the knotty centrality (KC) of a (non-empty, non-singleton) subset S of the nodes in G is given by:

$$KC(S) = \frac{E_s}{N_s(N_s - 1)} \sum_{i \in S} bc(i)$$

where E_s is the number of edges between nodes in S , and N_s is the number of nodes in S , and $bc(i)$ is the betweenness centrality of node i normalised with respect to the whole graph. The measure can be applied to either weighted or unweighted graphs by substituting weighted or unweighted variants of betweenness centrality (Brandes, 2001). Knotty centrality ranges from 0 to 1. It is 0 if none of the nodes in S is adjacent, and 1 if S is a clique.

For non-trivial graphs it is infeasible to calculate the knotty-centrality of all 2^N subsets of G and pick the one with the maximum value. Instead an exhaustive search of all subsets of G whose members fall in the top M nodes for betweenness centrality is performed. Given the subset with the highest knotty centrality, gradient ascent is then used to add nodes to it so as to maximise the knotty-centality.

6.3.5 Metastability

Shanahan (2010) introduced a measure of metastability (σ) based upon an estimate of the variance in synchrony amongst oscillators over time.

$$\sigma = \frac{1}{t_{\max} - 1} \sum_{t < T} (\varphi(t) - \varphi)^2$$

where $\varphi(t)$ is the synchrony of the system at time t , and φ is the mean synchrony of the system over the entire time of the simulation. The measures $\varphi(t)$ and φ are detailed in section 4.6. t_{\max} is the length of time of the simulation. This metastability metric has a shortcoming in that the distribution of synchrony may have a fairly high variance, resulting in high metastability, but be skewed towards desynchrony and therefore not have any stable episodes. To correct for this, the measure used in this chapter (σ_{fix}) scales the measure of metastability defined above according to the skewness of the distribution if and only if the distribution is skewed towards desynchrony.

$$\sigma_{fix} = \sigma(1 - s)$$

where s is 0 if the distribution of synchrony is not skewed towards dysynchrony, and is a value between 0 and 1 with 1 being maximum skewness if the distribution of synchrony is skewed towards desynchrony.

6.3.6 Results

The approach taken in this chapter is to build simulations of interacting neural oscillator populations and study the network structure and dynamics before and after plasticity is applied. As with the previous chapter, a number of different measures are then used to assess the network's overall dynamics. The frequency modulation between neural oscillators, averaged in the network as a whole, is measured using the mean intermittent frequency correlation, as defined in the previous chapter. Coalition entropy (Section 5.4.4) measures the variety of metastable states entered by a system of oscillators and is calculated from the number of distinct states the system can generate and the

probability of each state occurring. By identifying high levels of mean intermittent frequency correlation and coalition entropy it is possible to discern whether the neural populations in a network are modulating each others' oscillatory behaviour so as to explore many different coalitions. If the network also displays levels of synchrony that promote information transfer during coalition formation, then the network exhibits metastability. This is characterised by the tendency of a system to continuously migrate between a variety of synchronous states. This measure of metastability, which is called 'composite metastability', is contrasted with the measure of metastability defined in section 6.3.5 that is based on that introduced by Shanahan (2010).

A series of experimental simulations were performed in each of which 24 neural QIF PING oscillators were chosen, with frequencies drawn from a uniform distribution, from the set that had been evolved with intrinsic frequencies ranging from 30Hz to 50Hz. The probability of one oscillator providing neural input to another was determined with a given probability C . The probability C was the same for all oscillator to oscillator connections in the same experimental simulation. Given that a connection was established from oscillator n to oscillator m , the excitatory neurons in oscillator n synaptically connect to the excitatory neurons in oscillator m . To maintain consistency with the work in the previous chapter, the number of synaptic connections formed was 20 percent of the 40000 possible synaptic connections from the 200 excitatory neurons in oscillator n to the 200 excitatory neurons in oscillator m . For all synaptic connections formed the weight of the synapse was set to W . The values for W and C were randomly chosen at the beginning of each experimental simulation from a uniform distribution between 0 and 1 for C , and a uniform distribution between 0 and 0.38 for W . The present work uses 24 neural oscillators rather than the 10 used in the previous chapter. As a result, the maximum synaptic weight value W is 0.38 rather than the higher 1.0 used in the previous chapters experiment. However, the shape of the data within these weight and connection probability ranges presented in the results that follow is the same for both pieces of work. In the present chapter 480 simulations were performed. As the weight and connection probability for each of the 480 simulations were chosen at random, these data points are scattered throughout this parameter space. Figures 6.1-6.6 show various measures taken from these 480 simulations. These are analysed and discussed in detail below. In figures 6.1-6.4 a surface has been fitted to the underlying trend of the 480 data points for each measure depicted.

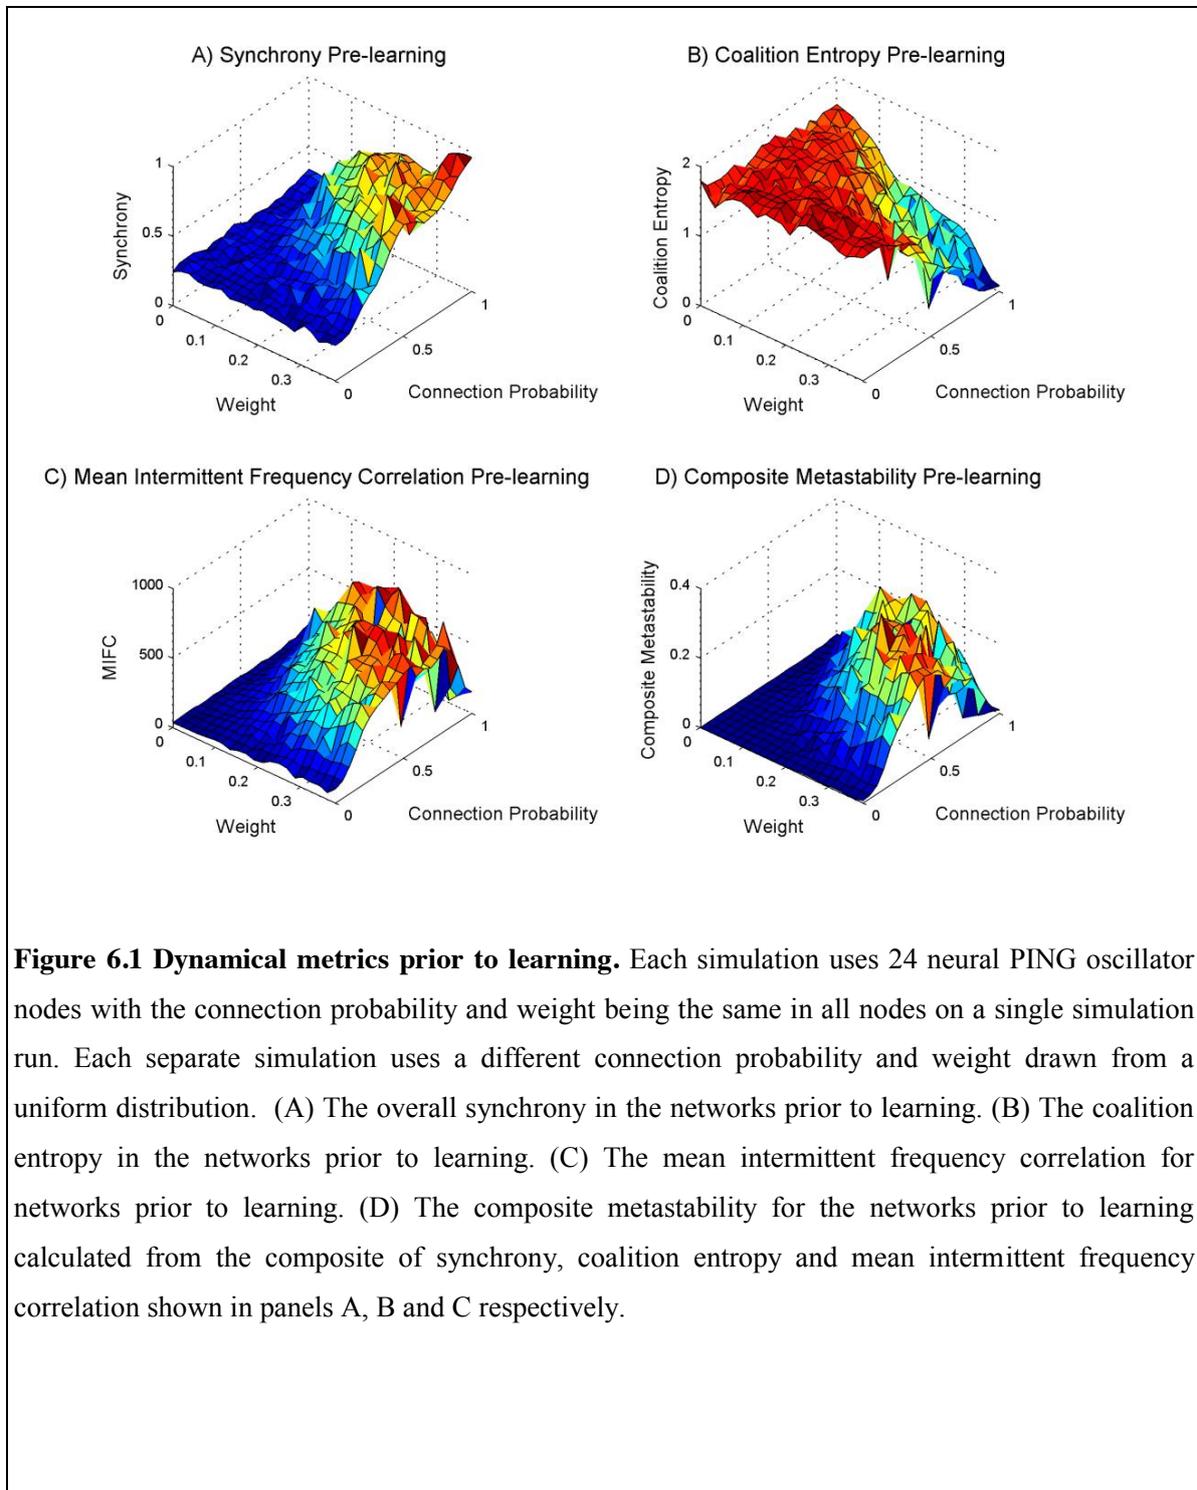
Throughout each simulation, external stimulus was provided to each neural oscillator by a Poisson process with parameter $\lambda = 4.375$. The inputs were scaled by 8 in order to provide sufficient stimulus

to induce firing. Each experiment was run for 2000 ms of simulated time with no STDP, after which the simulation continued running for 10000 ms using STDP, and finally for a further 2000 ms without STDP. After each experiment, the firing activity of the excitatory layers in each oscillator was used to calculate synchrony, coalition entropy, mean intermittent frequency correlation and composite metastability. This was performed on the simulation times before and after STDP was applied. The first 500 ms of these simulation times were discarded in the calculation of these metrics to eliminate initial transients.

Figure 6.1A shows the synchrony through the parameter sweep prior to STDP. Unsurprisingly, and in accord with previous findings, synchrony increases as connectivity increases. The neural systems never reach full synchrony, as on the simulations which attain the highest values of synchrony these systems regularly deviated from full synchrony over the course of the simulation. The value of 0.25 synchrony in the area of low weight and low connection probability represents only coincidental alignments in phase.

Figure 6.1B shows coalition entropy for the simulations prior to STDP. The trend for coalition entropy has the reverse form from synchrony, decreasing as synchrony increases, and the oscillators become more aligned in phase for more of the time. The measure of coalition entropy used cannot distinguish coincidentally synchronous coalitions from those that are genuinely coupled. However, when we contrast with the graphs of synchrony we can get an idea of what is happening. Regions of the parameter space with low weight and low connection probability exhibit high coalition entropy, but the same regions present low values for synchrony. This suggests that the many coalitions that appear are constituted by very short coincidental alignments in phase that are not capable of significant information transfer (Buehlmann and Deco, 2010). The mid parameter space area shows fairly high values for synchrony, indicating the capacity for substantial information transfer, as well as high coalition entropy indicating transfer between many different groups at different times. The region of the parameter space in which the synaptic weight and the inter-population connection ratio are high facilitates more information transfer but less variation in coalitions.

The mean intermittent frequency correlation prior to STDP is shown in figure 6.1C. As causal influence between oscillators increases in the network so does this correlation measure, suggesting that correlation directly reflects causation in this case. The previous chapter demonstrates the statistical significance of this metric, and shows comparable results for synchrony, coalition entropy and mean intermittent frequency correlation when using Izhikevich neurons and Hodgkin-Huxley



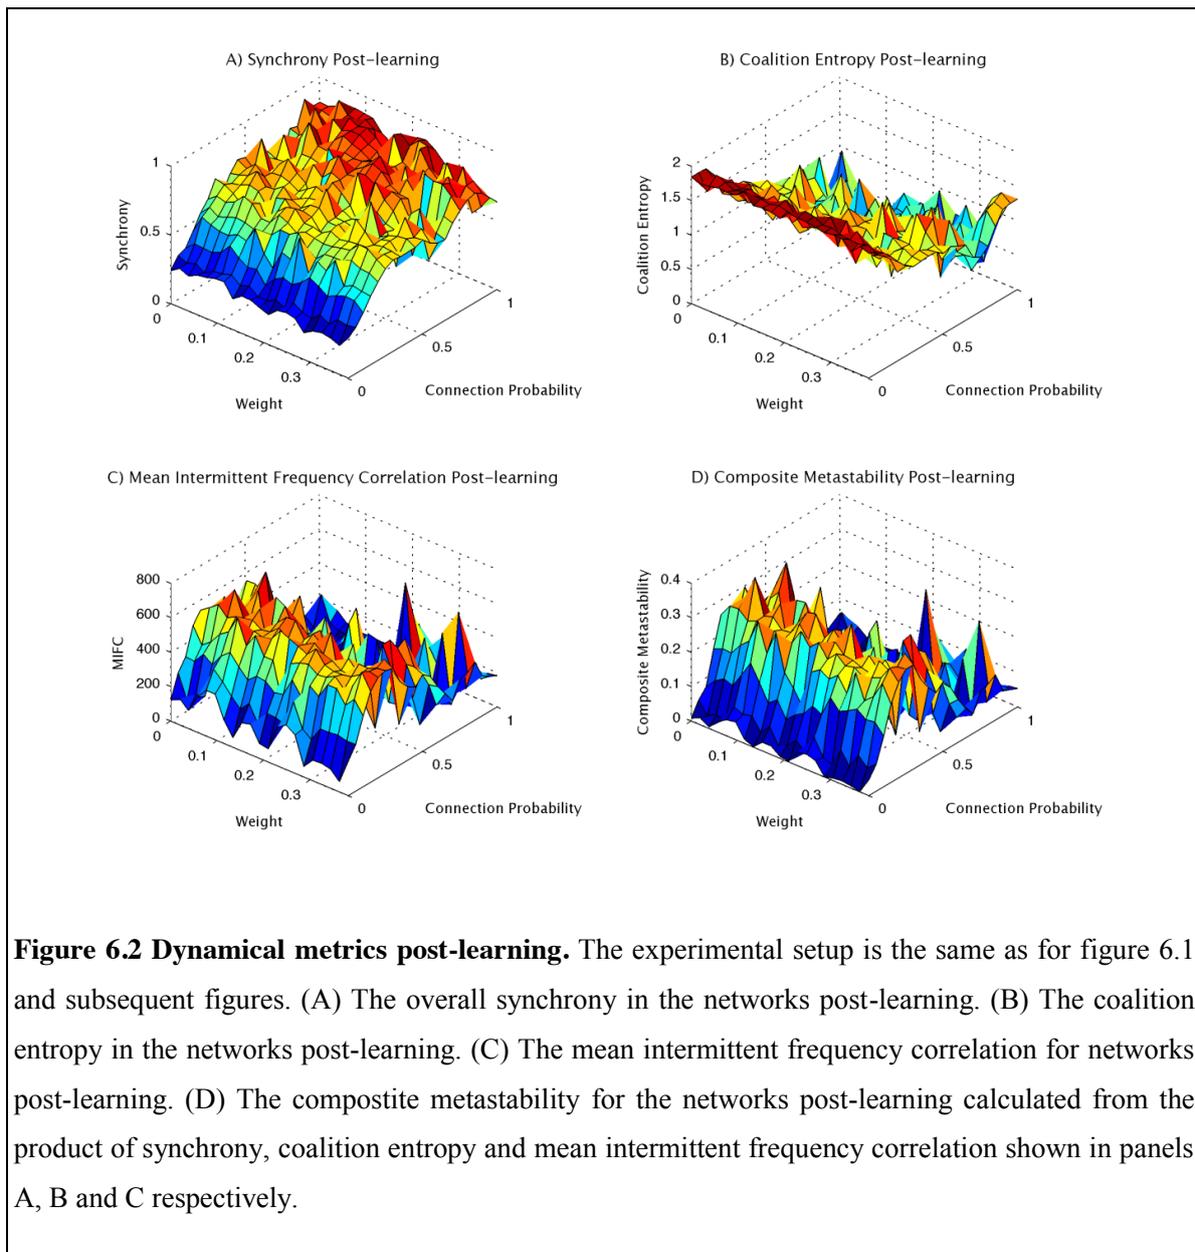
neurons. The previous chapter further demonstrates that fluctuating frequencies in different neural populations influence each other across frequency bands. The present chapter goes on to show the effect on these dynamics as well as on structure when STDP is applied to these networks.

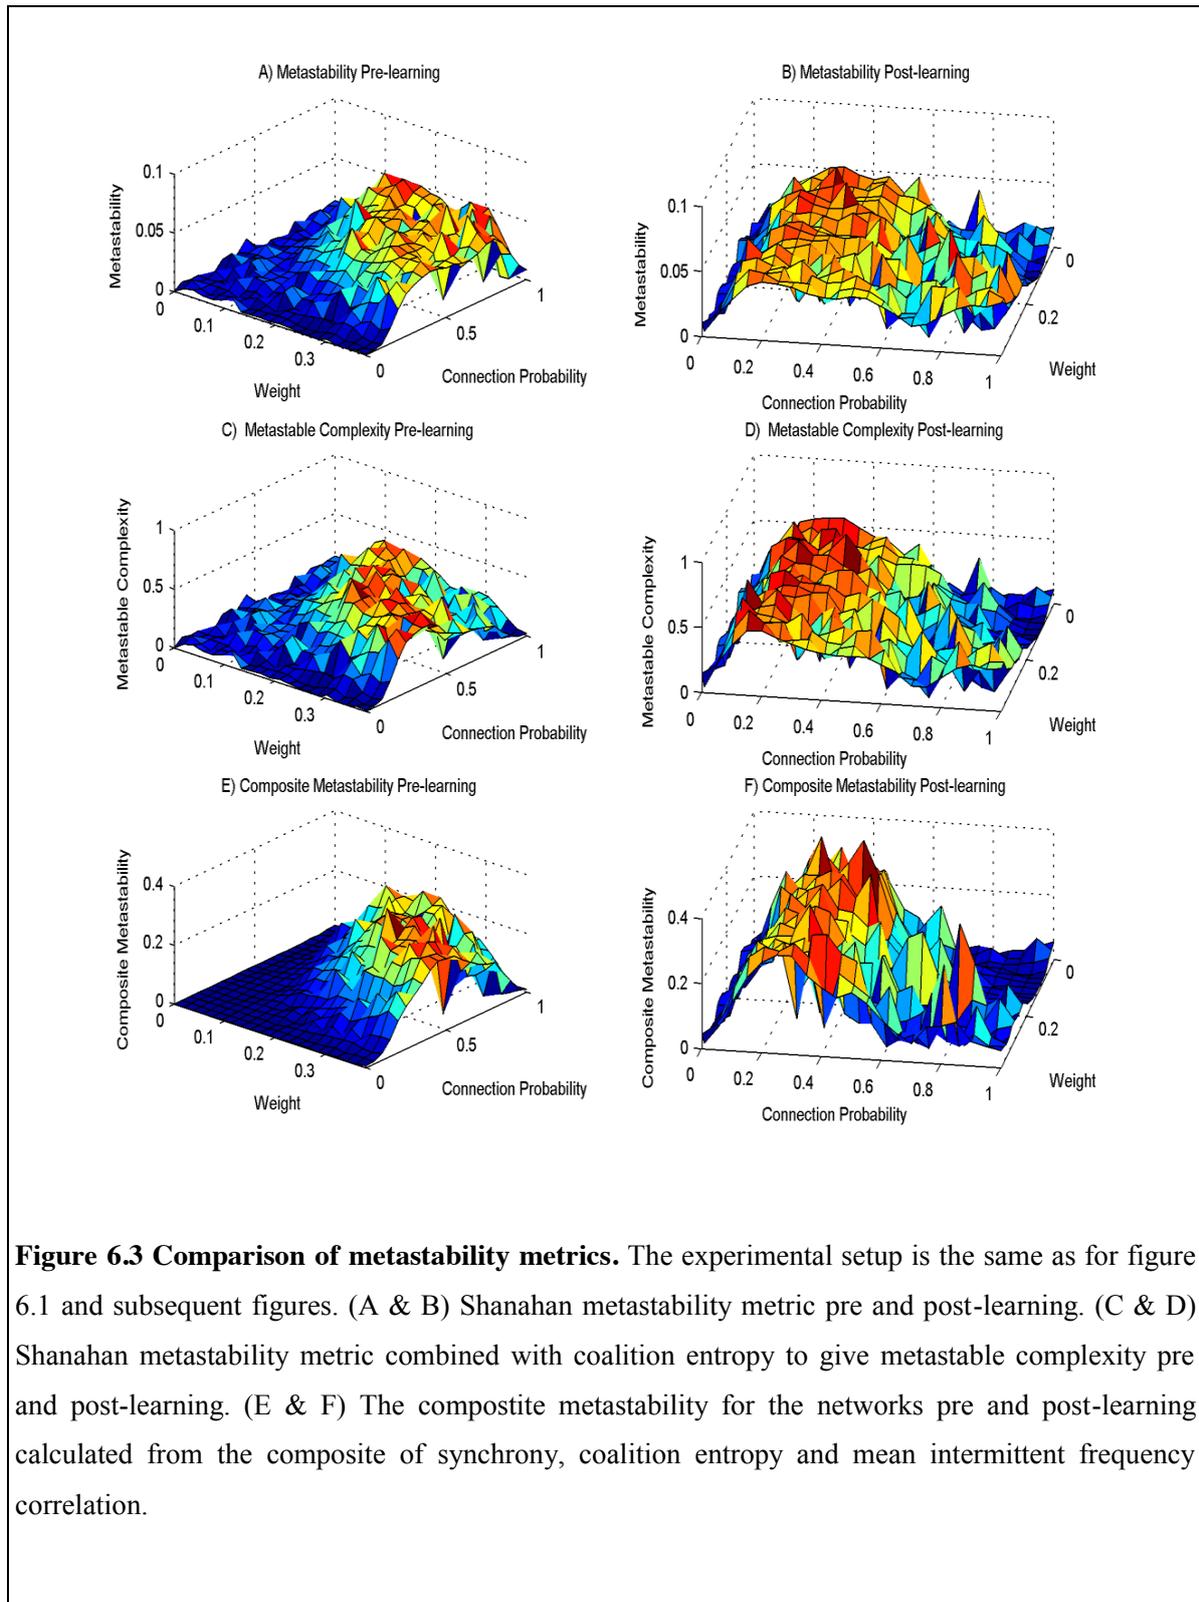
By normalising and combining intermittent frequency correlation, coalition entropy and synchrony we obtain a composite measure of metastability (figure 6.1D). There is a diagonal vector from weight value 0.2 and connection probability 1 to weight value 0.38 and connection probability 0.5 at which the composite metastability is at a peak. This area, in which the three metrics are balanced, facilitates metastable dynamics in which there is a richness of influence and interaction between different oscillators and across frequency bands modulating each other's behaviour, enabling the exploration of a large repertoire of different coalitions. The medium-to-high level of synchrony in this area further suggests that the conditions for information transfer between populations are fulfilled (Buehlmann and Deco, 2010). These traits are desirable in order to facilitate exploration, integration, and communication among functionally related areas during cognitive processing (Chialvo, 2010; Kelso, 2012; Shanahan, 2012).

Figure 6.2 show the same measures as figure 6.1 but for the period after STDP was applied. It is interesting to note that there is some slight variation given the inter-node connection weight prior to learning. However, the inter-node connection weight prior to learning has no major affect on the measures post-learning that have resulted from the modification of the network structure. It is also noteworthy that the connection probability prior to learning is the key variable affecting the changes that take place when plasticity is applied.

Figure 6.3 shows a comparison of metastability metrics taken from the simulations pre and post learning. Figures 6.3A and 6.3B show the metastability metric detailed in section 6.3.5 based on that introduced by Shanahan (2010). Comparing the pre-learning metric of Shanahan in figure 6.3A to that of composite metastability in figure 6.3E, the shape of the surfaces look similar. However, it can be seen that the area of high connection probability and high weight produces high metastability in figure 6.3A but low metastability in figure 6.3E. By taking the metastable metrics of figures 6.3A and 6.3B and normalising them, and then combining them with coalition entropy which is also normalised, the plots of figures 6.3C and 6.3D are obtained. These plots are termed 'metastable complexity' as they combine the variance in synchrony with a measure of the variation in different coalitions that appear. As can be seen in figure 6.3C the area of high connection probability and high weight produces low metastability much like that for composite metastability in figure 6.3E. It can be

concluded from this that the high metastability in the area of high connection probability and high weight in figure 6.3A is obtained because there is variation in synchrony but that there is less variation in coalitions and the later is not accounted for. Hence, the same coalitions are reappearing at different times in the simulations in this area of the plot. In accounting for coalition entropy, figures 6.3C and 6.3D appear very similar to figures 6.3E and 6.3F respectively.





It is interesting to note from figures 6.3B, 6.3D and 6.3F that when connection probability is high prior to learning the metastability post-learning is low. This behaviour can be elucidated further by examining what is happening with the coalitions that form throughout the post-learning time. Figure 6.4A shows the number of unique coalitions that appear post-learning regardless of the number of times they reappear. As can be seen in the area for high connection probability between 50 and 500 unique coalitions appear in this area. However, if we ignore coalitions that are around for 10 ms or less, so as to ignore coincidental alignments in phase and only consider those that are around for a longer time so that they promote information transfer, then the story is a little different. Figure 6.4B shows the number of unique coalitions that appear for more than 10 ms post-learning regardless of the number of times they reappear. The same high connection probability area exhibits between two and eight long lasting coalitions. Comparing the mean size of these coalitions in figure 6.4C we can see that the same area shows coalition sizes of between three and twelve oscillator nodes, as would be expected. The area of low weight and high connection probability shows coalition sizes of ≈ 12 . Figure 6.4D shows the mean duration that the coalitions are present for. The area of low weight and high connection probability show durations of ≈ 1450 ms indicating that the (typically) two coalitions of approximately twelve nodes are present for virtually all of the 1500 ms of post-learning simulation time. As would be expected, the same area in figure 6.3F exhibits very low metastability.

The behaviour of the coalitions in this area is indicative of multistability. Maistrenko *et al* (2007) explored analytically the behaviour of systems of Kuramoto oscillators that were uniformly connected and had an STDP-type rule applied to them. They found that the systems would tend towards multistability. The area of high connection probability in the simulations presented in this chapter is uniformly connected. It seems that when there is also a low weight then the systems tend towards multistability. However, it is unclear whether the areas of high connection probability and high weight would also tend towards multistability if given enough learning time. This is certainly an area for future exploration. Intuitively, it seems possible that a pre-existing high weight configuration would take longer to settle into a multistable regime than a low weight configuration. If this area of connectivity were proven also to be multistable then one would be inclined to conclude that uniformly connected neural oscillators display the same behaviour as the multistable Kuramoto oscillators described by Maistrenko *et al*. The areas of lower connection probability are more sparsely connected and are metastable rather than multistable. It may further be concluded that the inability to balance all possible connection pathways between nodes due to the sparse and non-

uniform connectivity keeps a tension in the system that prevents multistability and instead promotes metastability.

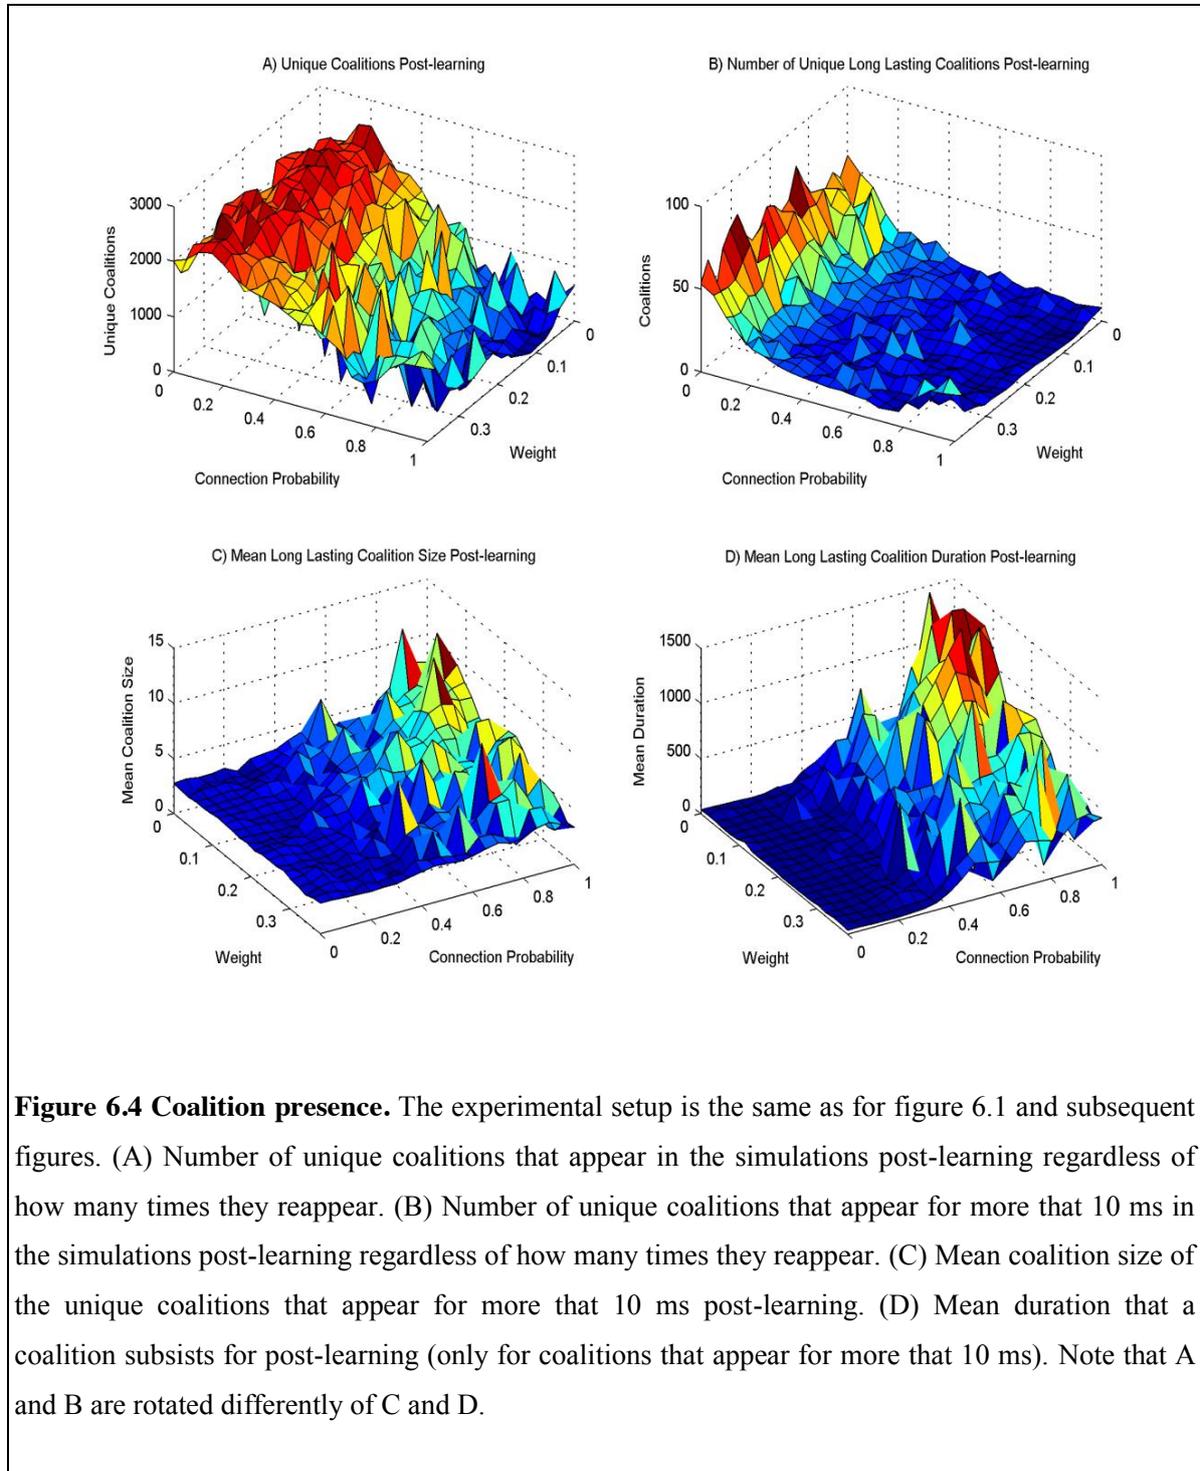


Figure 6.5A shows the connection probability prior to learning and the metastability after plasticity is applied and the network structure has changed. There is a rise in metastability as pre-learning connection probability rises, it peaks at around pre-learning connection probability 0.3, and then falls off. The metastable curve of interest lies between pre-learning connection probability 0.1 and 0.5. The elevated metastable area is not only in a different place to that of figure 6.1D, but also peaks at a value around 30% higher than its maximum in the pre-learning part of the simulation. This begs the question, what has changed in the structure of the networks that started in this connection probability area so as to alter their metastability? Figure 6.5B show the distribution of weights post-learning. The weights form a bimodal distribution in which they tend to polarise towards 0.1 and 1. This polarising behaviour is expected with an additive nearest neighbour STDP rule. The area below connection probability ≈ 0.1 does not adapt so much due to over sparseness in the network structure. Given that the area greater than connection probability ≈ 0.1 shows similar weight distribution, the question as to what alters metastability is naturally drawn away from the weights and to topological structure.

Figure 6.5C shows the small-world index of the networks post-learning plotted against pre-learning connection probability. To understand the significance of the plots requires a short discussion of the the definition of small-world index. The small-world index is calculated relative to random networks with the same degree distribution. Hence small-worldness is evident at index values greater than 1. A value of 1 would represent a random network. The values shown in figure 6.5C are above 1 but generally below 1.5. This may appear low for a small-world index. However, the value of the small-world index is not normalised for different network sizes, and as a result the small-world index can be larger for networks with more nodes compared to those with fewer nodes. This is due to the metric balancing particular ratios that are ultimately dependent on mean degree and the number of nodes in the network. (The definition of the small-world index is given in section 6.3.2.) To achieve small-worldness the characteristic path length of the network should be approximately the same as the characteristic path length of a random network with the same degree distribution. As a result the denominator of the small-world equation should be approximately 1. The numerator divides the clustering coefficient of the network by that of a random network with the same connectivity distribution. As the mean degree of the network increases towards the number of nodes in the network, the clustering coefficient of the network and its random analogue tend towards each other, hence reducing small-worldness. This implies that mean degree should be kept fairly low in order to achieve small-worldness. With a lower mean degree, the clustering coefficient of a small-world network will be higher than that of a random network with the same connectivity distribution,

resulting in a small-world index greater than 1. However, clustering coefficient is calculated from the number of triangles and possible triangles around a node given the degree of the node. The smaller the mean degree is relative to the number of nodes in the network then the fewer possible triangles there will be, particularly in the case of a random network. Such a random network will have a smaller the clustering coefficient and subsequently a greater small-world index can be achieved. It is important to bear in mind that the number of possible triangular pathways grows exponentially given the number of nodes in the network. Therefore the likelihood of a limited number of connections forming a triangle is equivalently reduced. As such the number of actual triangles around a node will be reduced as the number of nodes is increased unless the mean degree is increased exponentially. Increasing the mean degree to compensate, as mentioned above, counteracts small-worldness because the clustering coefficient of the network and its random version tend towards each other.

What results from this balancing of variables in a system in which there is exponential growth in possible triangles relative to the number of nodes, is a metric for which networks with more nodes can generate higher values of small-world index. Using the Watts-Strogatz (1998) technique for creating directed small-world networks and varying the connection and rewiring probabilities for a network with 24 nodes produces small-world indices up to ≈ 1.5 . For 48 nodes it produces small-world indices up to ≈ 2.5 , and for 98 nodes it produces indices up to ≈ 4.5 . Given that the networks in this chapter have 24 nodes and the idealistic Watts-Strogatz technique for generating small-world networks of the same size does not produce a small-world index greater than ≈ 1.5 , the values shown in figure 6.5C which range from 1 to 1.5 are significantly small-world, particularly considering that STDP can only alter the restricted number of pathways between oscillator nodes that were selected using a random connection probability at the beginning of each simulation.

Whilst small-worldness is generally high for connection probabilities greater than ≈ 0.2 , figure 6.5D shows that high connection probability displays lower modularity. The area of interest in which the pre-learning connection probability is greater than 0.1 and less than 0.5, and where metastability is elevated post-learning, falls in an area of high modularity. Hence, these networks are both modular and small-world.

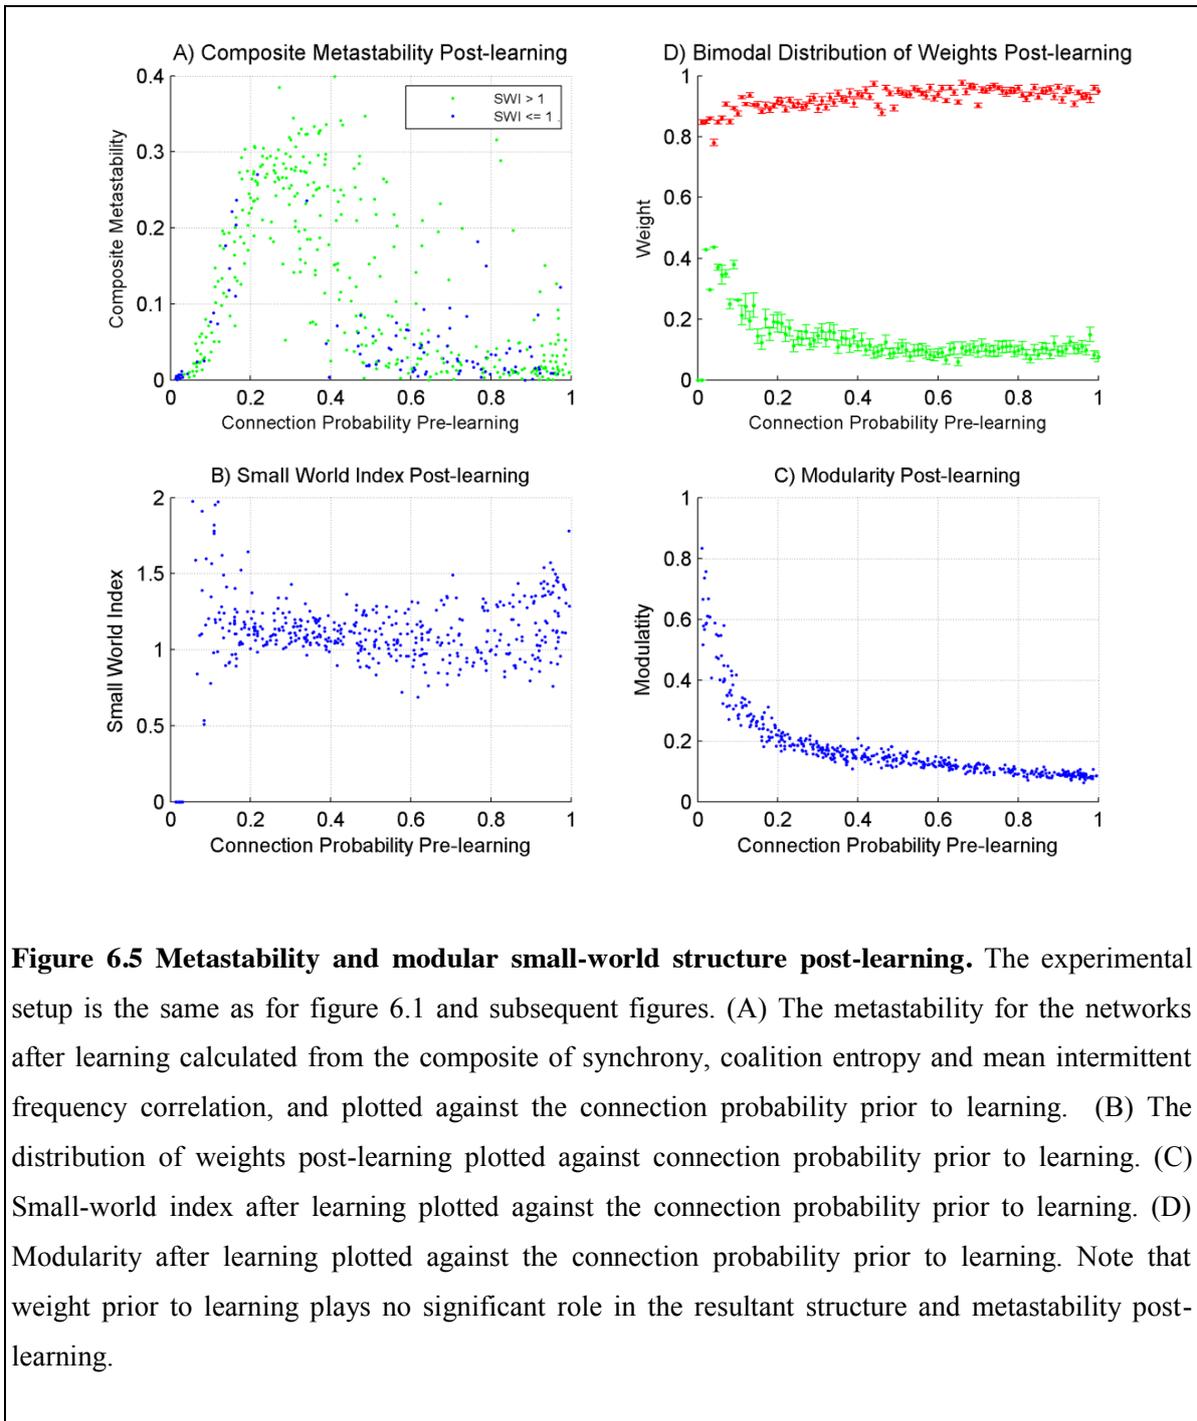
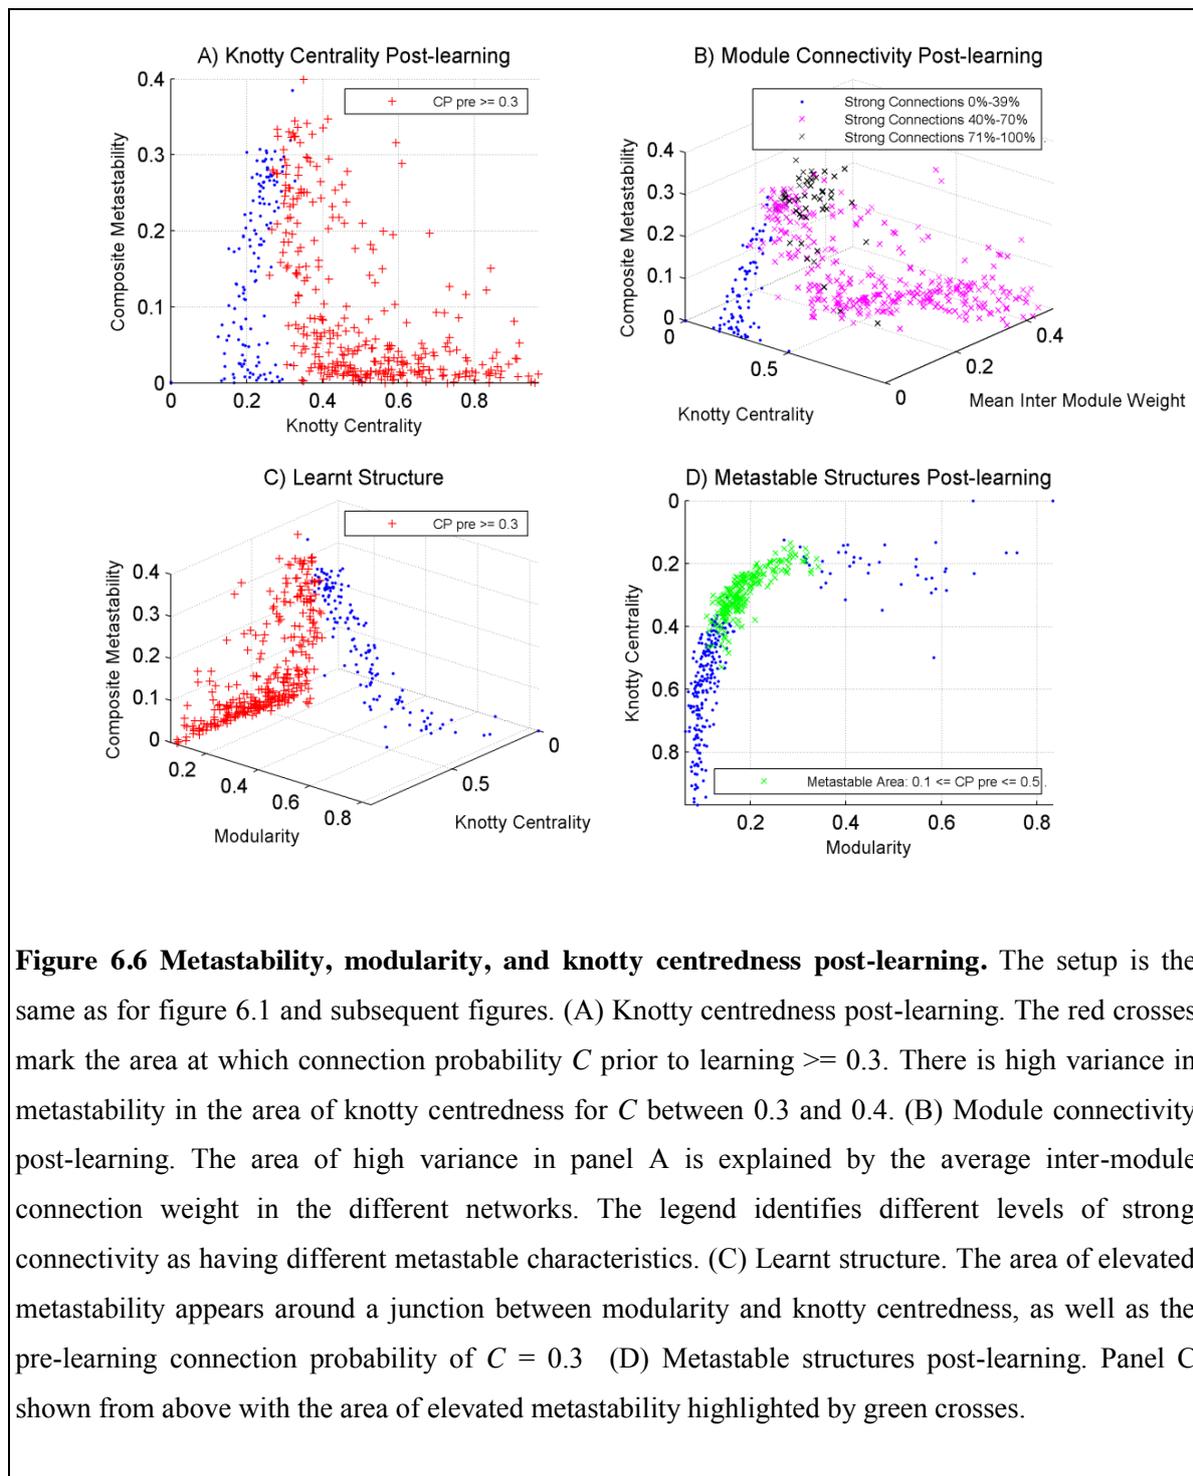


Figure 6.5A shows a lot of variation in metastability between pre-learning connection probability 0.3 and 0.4. This variability is not explained simply by the modular small-worldness. Figure 6.6A shows

the knotty centredness of the networks post-learning and their respective metastability. The red crosses indicate where the connection probability prior to learning is greater than or equal to 0.3 (i.e. where metastability peaks and falls post-learning). Although there is a slight separation of the data over the area of decline in metastability, in general the same variability can be seen in metastability around the peak and fall area as in figure 6.5A. Using the modules of oscillator nodes calculated for the modularity metric, figure 6.6B shows the mean weight of oscillator-to-oscillator connections in different modules and plots these against each network's knotty centrality. The figure shows the metastability for knotty centrality values around 0.25 increasing as the mean inter-module weight increases (blue dots). From the point where mean inter-module weight reaches around 0.2, as the knotty centredness and mean inter-module weight both increase together, the metastability drops rapidly (magenta crosses). This means that from this critical point, the stronger the connections are both within the knotty centre and between modules then the less metastable the network will be. However, if the knotty centredness remains around 0.25 and only the mean inter-module weight increases then higher levels of metastability can be reached. By thresholding inter-module connection weights above 0.1 we can identify strong inter-module connections. The blue dots show the percentage of strong connections between 0-39%, the red crosses between 40-70%, and the black crosses between 71-100%. It is clear that there is an increase in metastability as the percentage of strong connections increases, but in the mid area (40-70%) metastability can fall. However, it is interesting to note that the highest percentage of strong inter-module connections exhibits greater metastability albeit that the knotty centrality for these remains around 0.25. The conclusion is that this area of knotty centredness is important in achieving metastability. A knotty centrality value between 0.2 and 0.4 produces elevated metastability, and the rise and fall of metastability within this area can be explained by the strength of the mean inter-module connection weights.

Figure 6.6C shows metastability in relation to modularity and knotty centredness. The red crosses in this figure identify the pre-learning connection probabilities greater than or equal to 0.3, thus separating the data into the metastable climb and fall. Figure 6.6D shows the same data from above with the area of metastable elevation highlighted with green crosses. This area of interest appears at a juncture between modularity and knotty centrality. It was formed through plasticity acting on networks that had a pre-learning connection probability between 0.1 and 0.5. The resultant networks are modular small-world networks that have a particular level of knotty centrality and whose level of metastability is also reliant on the mean inter-module connection weight.



6.4 Discussion

The ability to orchestrate both local and global activity is of crucial importance for the brain, since it consists of numerous specialised regions interacting on multiple levels. The work presented in this chapter supports the intuitive notion that network topology needs to support such complex interaction, and that the required dynamics are promoted by modular small-world connectivity. Modular small-world structures are prevalent in the brain. Sporns *et al* (2000) investigated the structure of large-scale cortical systems and found evidence of localised clustering as well as a low characteristic path length between sites, the attributes of modular small-world connectivity, and these findings have since been substantiated and extended (Bullmore and Sporns, 2009). In addition, Sporns (2013) reports on a growing body of work that draws attention to how the balance between structural segregation and integration is essential for the operation of the brain networks underlying cognitive function.

Computer simulations have also contributed weight to the argument that modular small-world connectivity promotes functional complexity. Sporns *et al* (2000) compared networks optimised in simulation for functional complexity and cortical connection matrices derived from real datasets, and found a significant overlap in their structural and functional characteristics. Functional complexity captures the extent to which a system is both functionally segregated such that small subsets of the system behave independently, and functionally integrated such that large subsets tend to behave coherently. The heightened metastability in modular small-world networks demonstrated in this chapter supports the view that this type of connectivity promotes functional complexity, as well as placing emphasis on the necessity for high levels of modularity.

The question naturally arises as to how the right topology of neural connections is formed. Brain networks are shaped by evolution, ontogenetic development, and experience (Sporns *et al.*, 2004). Levy *et al* (2001) report that synaptic plasticity facilitates the formation of sub-assemblies within a network, each of which exhibits its own oscillatory dynamics, a phenomenon they refer to as *distributed synchrony*. Using computer simulations, Kwok *et al* (2007) explored the development of small-world network structures on a micro scale. They used a Hebbian rewiring rule similar to STDP that connects pairs of neurons that fire together and disconnects those that do not. Given spontaneous neural firing activity characteristic of early development, the resulting network exhibited modular small-world connectivity. All of these findings demonstrate the natural tendency of neural networks

to self-organise into modular and small-world structures at the micro scale, and are consistent with the present work. However, the concern of this chapter is the macro scale.

Modular brain structures are hierarchical but not fractal between scales. To the extent that the brain exhibits any fractal self similarity, this is statistical rather than exact (Meunier, Lambiotte and Bullmore, 2010). Therefore the causes of modular small-world networks at the micro scale do not necessarily apply to the macro scale. An explanation of the emergence of larger scale modular small-world topology is required. The present work demonstrates how small-world topology at a macro scale emerges due to the interaction between adaptive processes and naturally occurring dynamics. It is interesting to note that the network dynamics and adaptive processes drive the network towards forming modular small-world architectures that support *heightened* levels of metastability. Networks whose starting configuration did not result in metastability exhibited very low levels of mean intermittent frequency correlation, implying little or no influential interaction and cooperation between oscillator nodes. It may be suggested that as a consequence they are not capable of functionally complex behaviour.

A key finding of the results is that a particular type of connective core within the network structure is required in order for metastability to be manifest. This implies that the connective core plays an important part in balancing metastability. It has been argued that the spatial and topological centrality of brain regions that form the core of the human connectome plays a role in maintaining efficient global communication. (van den Heuvel and Sporns, 2011; Shanahan, 2012). When calculated, the knotty centrality for an 82 region human brain connectome (van den Heuvel and Sporns, 2011) and a 52 region pigeon brain connectome (Shanahan et al., 2013) both had a knotty centredness that falls inside the range, detailed in the results, within which metastability is elevated. Interestingly, they fall to the right (≈ 0.37) of the metastable peak, in the area of high metastable variance. Intuitively this seems appropriate. Higher metastability ought to correlate with more lateral, novel and creative thought processes. As far as everyday cognition is concerned, not being overly metastable would result in more balanced reasoning whilst still allowing originality of thought. However, the ability to vary one's creativity is also clearly advantageous.

To summarise, the results presented imply that the phenomenon of transient episodes of synchrony in neural systems is inherently related to the topological structure of the network. This is because the natural formation of a modular small-world topology gives rise to metastable dynamics, which in turn facilitate the empirically observed coordination phenomenon.

7 Conclusion

This thesis presents a model-based exploration of metastability as the cause of the phenomenon of transient coalition formation between neural areas empirically observed *in vivo*. The main focus has been on understanding the interaction between oscillations in different neural populations and the resulting global system dynamics, as well as the relationship between network structure and dynamics within the brain. Gaining an understanding of the way in which brain-scale dynamics arise has required several studies. Firstly, assessing the validity of existing neural models of metastability that use abstract oscillator models by comparing the collective behaviour of oscillating populations of neurons to the collective behaviour of Kuramoto oscillators. Next, quantifying the modulatory influences between oscillatory neural populations and relating them to other measures that quantify metastability. Finally, applying plasticity to a network of neural oscillators and evaluating the resulting network structures and the effect it has upon metastability.

Metastability is characterised as intrinsically driven movements between transient, attractor-like states. Although transient episodes of synchrony between different neural areas are observed in the brain, it has been a point of contention as to whether an underlying metastable mechanism gives rise to the observed phenomena or whether it is just an epiphenomenon. The results presented here demonstrate that metastability naturally emerges in networks of oscillating spiking neuron populations in simulation.

Other work has aimed bridge the gap between simulation and empirical data. For example, Cabral *et al* (2011) performed simulations in which simple Kuramoto oscillators modelled the intrinsic

oscillatory dynamics of functional neural areas set within a human connectome. They reported the presence of metastability, and demonstrated that the resulting phenomena of transient synchronous coalitions reproduced the fluctuations in human resting state networks obtained *in vivo* (Hellyer et al., 2014). Chapter 4 explored the relationship between this type of simple oscillator model and their neural population cousins by emulating neurally the Kuramoto critical coupling experiment that showed an increase in synchrony as connection strength is increased in a uniformly connected oscillator network. It was shown that the neural systems broadly conform the behaviour of Kuramoto oscillators, and so substantiates the above metastability work. However, it was also demonstrated that at the point of maximum synchrony the neural systems not only displayed several coexisting frequencies within an individual oscillator population but also showed deviations from a measure of full synchrony likely caused by these additional fluctuating influences.

Having shown that neural models display greater spectral complexity during synchronisation than more abstract oscillator models, with several oscillatory frequencies coexisting within an individual neural oscillator population, chapter 5 explored the interaction of these frequencies between neural oscillator populations and across frequency bands. It was shown that the fluctuation of frequencies in each oscillator modulated the frequencies of the other oscillators, and in doing so was able to drive the system as a whole towards transient episodes of synchrony. An area in the connectivity space was located in which metastability was at a peak. Interestingly, there was an almost equal influence upon metastability from synaptic weight and macro scale connectivity. This suggests that it is simply a particular level of mutual influence between oscillators as a whole that is required in order to generate metastable dynamics. Given that additive STDP polarises the synaptic weights, the implication is that future work on metastability can focus on topology rather than weights in order to investigate the phenomenon.

Chapter 6 explored the effect of plasticity in a network of neural oscillators. It was demonstrated that, given an appropriate initial connection probability between neural oscillators, the network would restructure into a modular small-world network that exhibits heightened levels of metastability. Plasticity, responding to the interactions between different neural areas, naturally forms modular small-world networks, which in turn promote metastability, and metastability further enhances these structural features. Another key finding was that a particular type of connective core within the network structure is required in order for metastability to be manifest. This implies that the connective core plays an important part in balancing metastability. This fits well with the argument

that the spatial and topological centrality of brain regions that form the core of the human connectome plays a role in maintaining efficient global communication (van den Heuvel and Sporns, 2011; Shanahan, 2012).

The present work motivates more research into metastable phenomena. An area for further research would be to explore the space of small-world, modular, and connective core configurations and identify the metastable characteristics within it. Considering where the human connectome falls within this space would provide further insights into the constraints on brain dynamics. In addition, the present study has presumed that each neural population is active, and has therefore not included the amplitude of the oscillation as a variable in the study. Whilst this fits well with the studies that use Kuramoto oscillators, which do not model amplitude either, a more detailed study that uses amplitude may be called for. Further to this, it would be invaluable to explore whether and how metastable dynamics can be used to promote cognitive functions, such as behavioural responses if embodied in a robot, or even the potential of using these exploratory dynamics for problem solving.

Before this can be done a greater understanding of how and why coalitions are formed relative to network structure and state is required. One possible way of exploring this would be to locate different coalitions that form in a network and use Granger causality to identify the network characteristics that give rise to them. If this is performed over many different network configurations then it may be possible to use data mining on the results from the Granger causality analysis to identify key conditions and principles that give rise to specific types of coalition formation. Such a study may lead to the ability to design network topologies that facilitate the formation of coalitions between functional groups of neurons that are advantageous for cognitive processing.

A further issue raised in the study was how the small-world index is not normalised over different network sizes. This means that a direct comparison between networks is presently not possible. It would be a very worthy task to adjust the small-world index so as to normalise over network sizes as this would assist researchers a great deal.

Given that metastability emerges naturally in neural simulations and that such dynamics map well to empirical data, the suggestion made here is that metastability is also present in biological brains. If this is not the case then researchers are faced with an important question as to why this complex interaction between oscillators, which is ubiquitous in dynamical systems in nature and seemingly also present in simulations of neural oscillators, is not present in the brain. The simulations presented

in this work suggest that modular small-world network characteristics akin to those found in the brain, and the metastable dynamics they promote, facilitate versatile exploration, integration, and communication between functionally related neural areas, and thereby support sophisticated cognitive processing in the brain. If metastability is present in the brain then, at the very least, further investigation will be required to better understand how these metastable dynamics operate in such a way that human cognition is not disrupted, but instead is stable and effective. However, I hypothesise that such dynamics form a good basis for contextual exploration, integration and communication between functionally related areas during cognitive processing. In an embodied system in which different brain areas are processing information relevant to the ongoing environmental situation this potentially provides a mechanism for problem solving. Presently we do not understand enough about the phenomena in order to make predictions that can be tested *in vivo* so as to confirm the hypothesis of metastability in the brain, let alone understand its cognitive purpose. Nevertheless, the present work demonstrates the presence of metastability in neural simulations, providing a strong indication that such phenomena should be present in the brain, and in doing so warrants much further investigation and research.

8 Bibliography

Abbott, L.F. and Nelson, S.B. (2000) 'Synaptic plasticity: taming the beast', *Nature neuroscience*, vol. 3, november, pp. 1178–83.

Abrams, D. and Strogatz, S. (2004) 'Chimera states for coupled oscillators', *Physical Review Letters*, vol. 93, no. 17, p. 174102.

Acebrón, J., Bonilla, L., Pérez-Vicente, C., Ritort, F. and Spigler, R. (2005) 'The Kuramoto model: A simple paradigm for synchronization phenomena', *Reviews of Modern Physics*, vol. 77, no. 1, pp. 137-185.

Alger, B.R.N.R.A. (1979) 'GABA-mediated biphasic inhibitory responses in hippocampus', *Nature*, vol. 281, pp. 315 - 317.

Andersen, P., Dingledine, R., Gjerstad, L., Langmoen, I.A. and Laursen, A.M. (1980) 'Two different responses of hippocampal pyramidal cells to application of gamma-amino butyric acid', *Journal of Physiology*, vol. 305, August, pp. 279-96.

Axmacher, N., Mormann, F., Fernández, G., Elger, C.E. and Fell, J. (2006) 'Memory formation by neuronal synchronization', *Brain research reviews*, vol. 52, no. 1, pp. 170-82.

Barlow, J.S. (1993) *The Electroencephalogram: Its Patterns and Origins*, MIT Press.

Bartos, M., Vida, I., Frotscher, M., Meyer, A., Monyer, H., Geiger, J.R.P. and Jonas, P. (2002) 'Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks', *Proceeding National Academy of Science USA*, vol. 99, pp. 13222–13227.

Basar, E., Basar-Eroglu, C., Karakas, S. and Schurmann, M. (2000) 'Brain oscillations in perception and memory', *International Journal of Psychophysiology*, vol. 35, no. 23, pp. 95–124.

Bear, F., Connors, W. and Paradiso, M. (1996) *Neuroscience: Exploring the Brain*, Williams and Wilkins.

- Belluscio, M.A., Mizuseki, K., Schmidt, R., Kempster, R. and Buzsáki, G. (2012) 'Cross-Frequency Phase-Phase coupling between theta and gamma oscillations in the hippocampus', *The Journal of Neuroscience*, 32(2):423–435., vol. 32, no. 2, pp. 423–435.
- Betz, R.F., Erickson, M., Abell, M., O'Donnell, B.F., Hetrick, W.P. and Sporns, O. (2012) 'Synchronization dynamics and evidence for a repertoire of network states in resting EEG', *Frontiers in Computational Neuroscience*, vol. 6, September, pp. 1-13.
- Börgers, C. and Kopell, N. (2005) 'Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons', *Neural Computation*, vol. 17, no. 3, pp. 557–608.
- Brandes, U. (2001) 'A faster algorithm for betweenness centrality', *Journal of Mathematical Sociology*, vol. 25, no. 2, pp. 163–177.
- Breakspear, M., Heitmann, S. and Daffertshofer, A. (2010) 'Generative models of cortical oscillations: Neurobiological implications of the kuramoto model', *Frontiers in Human Neuroscience*, vol. 4, p. 190.
- Bressler, S.L. (1995) 'Large-scale cortical networks and cognition', *Brain Research Reviews*, vol. 20, no. 3, pp. 288–304.
- Brunel, N. (2000) 'Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons', *Journal of Computational Neuroscience*, vol. 8, no. 3, pp. 183–208.
- Brunel, N. and Wang, X. (2003) 'What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance', *Journal of Neurophysiology*, vol. 90, no. 1, pp. 415–430.
- Buckner, R.L., Andrews-Hanna, J.R. and Schacter, D.L. (2008) 'The brain's default network: anatomy, function, and relevance to disease', *Annals of the New York Academy of Sciences*, vol. 1124, pp. 1–38.
- Buehlmann, A. and Deco, G. (2010) 'Optimal Information Transfer in the Cortex through Synchronization', *PLoS Computational Biology*, vol. 6, no. 9.
- Bullmore, E. and Sporns, O. (2009) 'Complex brain networks: graph theoretical analysis of structural and functional systems', *Nature reviews. Neuroscience*, vol. 10, no. 3, pp. 186-98.

Buzsáki, G. and Draguhn, A. (2004) 'Neuronal oscillations in cortical networks', *Science*, vol. 304, no. 5679, pp. 1926–1929.

Buzsáki, G., Buhl, D., Harris, K., Csicsvari, J., Czeh, B. and Morozov, A. (2003) 'Hippocampal network patterns of activity in the mouse', *Neuroscience*, vol. 116, no. 1, pp. 201–211.

Cabral, J., Hugues, E., Sporns, O. and Deco, G. (2011) 'Role of local network oscillations in resting-state functional connectivity', *NeuroImage*, vol. 57, pp. 130–139.

Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., Baker, A. and ... Deco, G. (2013) 'Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations', *NeuroImage*, vol. 90C, pp. 423–435.

Canolty, R.T., Edwards, E., Dalal, S.S., Soltani, M., Nagarajan, S.S., Kirsch, H.E., Berger, M.S., Barbaro, N.M. and Knight, R.T. (2006) 'High gamma power is phase-locked to theta oscillations in human neocortex', *Science*, vol. 313, no. 5793, pp. 1626–1628.

Cheng, H., Wang, Y., Sheng, J., Kronenberger, W.G., Mathews, V.P., Hummer, T.A. and Saykin, A.J. (2012) 'Characteristics and variability of structural networks derived from diffusion tensor imaging', *NeuroImage*, vol. 61, no. 4, pp. 1153–1164.

Chialvo, D.R. (2010) 'Emergent complex neural dynamics', *Nature Physics*, vol. 6, no. 10, pp. 744–750.

Colgin, L.L..D.T..F.M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M. and Moser, E.I. (2009) 'Frequency of gamma oscillations routes flow of information in the hippocampus', *Nature*, vol. 462, no. 7271, pp. 353–357.

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A. and Buzsáki, G. (1999) 'Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat', *The Journal of Neuroscience*, vol. 19, no. 1, pp. 274–287.

Daubechies, I. and Maes, S. (1996) 'A nonlinear squeezing of the continuous wavelet transform based on auditory nerve models', *Wavelets in Medicine and Biology*, CRC Press, pp. 527-544.

-
- Doesburg, S.M., Roggeveen, A.B., Kitajo, K. and Ward, L.M. (2008) 'Large-Scale Gamma- Band phase synchronization and selective attention', *Cerebral Cortex*, vol. 18, no. 2, pp. 386–396.
- Engel, A.K. and Fries, P. (2010) 'Beta-band oscillations signalling the status quo?', *Current Opinion in Neurobiology*, vol. 20, no. 2, pp. 156–165.
- Engel, A.K. and Singer, W. (2001) 'Temporal binding and the neural correlates of sensory awareness', *Trends in Cognitive. Science*, vol. 5, pp. 16–25.
- Ermentrout, B. (1996) 'Type I membranes, phase resetting curves, and synchrony', *Neural Computation*, vol. 8, no. 5, pp. 979–1001.
- Ermentrout, G.B. and Kopell, M. (1986) 'Parabolic bursting in an excitable system coupled with a slow oscillation', *SIAM Journal on Applied Mathematics*, vol. 46, no. 2, p. 233.
- Fell, J., Fernández, G., Klaver, P., Elger, C.E. and Fries, P. (2003) 'Is synchronized neuronal gamma activity relevant for selective attention', *Brain Research Reviews* , vol. 42, pp. 265–272.
- Fidjeland, A.K. and Shanahan, M. (2010) 'Accelerated Simulation of Spiking Neural Networks Using GPUs', *Proceedings IJCNN 2010*, pp. 1-8.
- Fox, P.T. and Friston, K.J. (2012) 'Distributed processing; distributed functions?', *Neuroimage*, vol. 61, pp. 407-426.
- Freeman, L. (1977) 'A set of measures of centrality based on betweenness', *Sociometry*, vol. 40, no. 1, pp. 35– 41.
- Freeman, L. (1979) 'Centrality in social networks conceptual clarification', *Social Networks*, vol. 1, no. 3, pp. 215–239.
- Fries, P. (2005) 'A mechanism for cognitive dynamics: neuronal communication through neuronal coherence', *Trends in cognitive sciences*, vol. 9, no. 10, pp. 474-80.
- Fries, P. (2009) 'Neuronal gamma-band synchronization as a fundamental process in cortical computation', *Annual review of neuroscience*, vol. 32, pp. 209-24.
- Fries, P., Reynolds, J.H., Rorie, A.E. and Desimone, R. (2001) 'Modulation of oscillatory neuronal synchronization by selective visual attention', *Science* , vol. 291, no. 5508, pp. 1560–1563.

Fries, P., Schroder, J., Roelfsema, P.R., Singer, W. and Engel, A.K. (2002) 'Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection', *Journal of Neuroscience*, vol. 22, no. 9, pp. 3739–3754.

Geisler, C. (2005) 'Contributions of intrinsic membrane dynamics to fast network oscillations with irregular neuronal discharges', *Journal of Neurophysiology*, vol. 94, no. 6, pp. 4344–4361.

Gerstner, W. and Kistler, W. (2002) *Spiking neuron models*, Cambridge University Press.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C.J., Wedeen, V.J. and Sporns, O. (2008) 'Mapping the structural core of human cerebral cortex', *PLoS Biology*, vol. 6, no. 7, p. e159.

Hansen, P., Kringelbach, M. and Salmelin, R. (2010) *MEG: An Introduction to Methods*, Oxford University Press.

Hasenstaub, A., Shu, Y. and Haider, B. (2005) 'Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks', *Neuron*, vol. 47, no. 3, pp. 423-35.

Hellyer, P.J., Shanahan, M., Scott, G., Wise, R.J.S., Sharp, D.J. and Leech, R. (2014) 'The control of global brain dynamics: opposing actions of frontoparietal control and default mode networks on attention', *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 34, no. 2, pp. 451–61.

He, Y., Wang, J., Wang, L., Chen, Z.J., Yan, C., Yang, H., Tang, H., Zhu, C., Gong, Q., Zang, Y. and Evans, A.C. (2009) 'Uncovering intrinsic modular organization of spontaneous brain activity in humans', *PLoS ONE*, vol. 4, no. 4, p. e5226.

Hodgkin, A.L. (1948) 'The local electric changes associated with receptive action in a non-mediated axon', *Journal of Physiology*, vol. 107, pp. 165-181.

Hodgkin, A.L. and Huxley, A.F. (1952) 'A quantitative description of ion currents and its applications to conduction and excitation in nerve membranes', *Journal of Physiology*, vol. 117, pp. 500-544.

Holland, J. (1975) *Adaptation in Natural and Artificial Systems*, The University of Michigan Press.

-
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R. and Hagmann, P. (2009) 'Predicting human resting-state functional connectivity from structural connectivity', *Proceedings of the National Academy of Sciences*, vol. 106, no. 6, pp. 2035–2040.
- Honey, C.J., Thivierge, J. and Sporns, O. (2010) 'Can structure predict function in the human brain?', *NeuroImage*, vol. 52, no. 3, pp. 766–776.
- Hosaka, R., Ikegnuci, T., Nakamura, H. and Akaki, O. (2004) 'Information Transformation from a Spatiotemporal Pattern to Synchrony through STDP Network', *Proceedings IJCNN 2004*, pp. 1475-1480.
- Huettel, S.A., Song, A.W. and McCarthy, G. (2009) *Functional Magnetic Resonance Imaging*, Massachusetts: Sinauer.
- Humphries, M., Gurney, K. and Prescott, T. (2006) 'The brainstem reticular formation is a small-world, not scale-free, network', *Proceedings of the Royal Society B*, vol. 273, no. 1585, pp. 503–511.
- Izhikevich, E. (2003) 'Simple model of spiking neurons', *IEEE transactions on neural networks*, vol. 14, no. 6, pp. 1569-72.
- Izhikevich, E. (2007) *Dynamical Systems In Neuroscience*, Cambridge University Press.
- Jadbabaie, A., Motee, N. and Barahona, M. (2004) 'On the Stability of the Kuramoto Model of Coupled Nonlinear Oscillators', *Systems Engineering*, pp. 1-8.
- Jensen, O., Kaiser, J. and Lachaux, J. (2007) 'Human gamma-frequency oscillations associated with attention and memory', *Trends in Neurosciences*, vol. 30, no. 7, pp. 317–324.
- Jirsa, V.K., Jantzen, K.J., Fuchs, A. and Kelso, J.A.S. (2002) 'Spatiotemporal forward solution of the EEG and MEG using network modelling', *IEEE Trans Med Imaging*, vol. 21, no. 5, pp. 493-504.
- Joyce, K.E., Laurienti, P.J., Burdette, J.H. and Hayasaka, S. (2010) 'A new measure of centrality for brain networks', *PLoS ONE*, vol. 5, no. 8, p. e12200.
- Kelso, J.A.S. (2012) 'Multistability and metastability: understanding dynamic coordination in the brain', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, vol. 367, no. 1591, pp. 906-18.

-
- Kelso, J.A.S. and Tognoli, E. (2007) 'Toward a Complementary Neuroscience: Metastable Coordination Dynamics of the Brain. Neurodynamics of Cognition and Consciousness', in Perlovsky, L.I. and Kozma, R. *Understanding Complex Systems*, Springer.
- Klimesch, W. (1999) 'EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis', *Brain research. Brain research reviews*, vol. 29, no. 2-3, pp. 169–95.
- Kondgen, H., Geisler, C., Fusi, S., Wang, X., Luscher, H. and Giugliano, M. (2007) 'The dynamical response properties of neocortical neurons to temporally modulated noisy inputs in vitro', *Cerebral Cortex*, vol. 18, no. 9, pp. 2086–2097.
- Kopell, N., Ermentrout, G.B., Whittington, M.A. and Traub, R.D. (2000) 'Gamma rhythms and beta rhythms have different synchronization properties', *Proceedings of the National Academy of Sciences*, vol. 97, no. 4, pp. 1867–1872.
- Kuhnert, M.T., Geier, C., Elger, C.E. and Lehnertz, K. (2012) 'Identifying important nodes in weighted functional brain networks: A comparison of different centrality approaches', *Chaos*, vol. 22, no. 2, p. 023142.
- Kuramoto, Y. (1984) *Chemical Oscillations, Waves and Turbulence*, Berlin: Springer-Verlag.
- Kwok, H.F., Jurica, P., Raffone, A. and van Leeuwen, C. (2007) 'Robust emergence of small-world structure in networks of spiking neurons', *Cognitive neurodynamics*, vol. 1, no. 1, pp. 39-51.
- Lachaux, J.P., Chavez, M. and Lutz, A. (2003) 'A simple measure of correlation across time, frequency and space between continuous brain signals', *Journal of neuroscience methods*, vol. 123, no. 2, pp. 175-88.
- Latham, P.E., Richmond, B.J., Nelson, P.G. and Nirenberg, S. (2000) 'Intrinsic dynamics in neuronal networks. i. theory', *Journal of Neurophysiology*, vol. 83, no. 2, pp. 808–827.
- Le Van Quyen, M. (2005) 'Anticipating epileptic seizures: from mathematics to clinical applications', *Comptes Rendus Biologies*, vol. 328, no. 2, pp. 187–198.
- Leicht, E.A. and Newman, M.E.J. (2008) 'Community structure in directed networks', *Physical Review Letters*, vol. 100, p. 118703.

-
- Levy, N., Horn, D., Meilijson, I. and Ruppin, E. (2001) 'Distributed synchrony in a cell assembly of spiking neurons', *Neural Networks*, vol. 14, no. 6-7, pp. 815-24.
- Lisman, J. (2005) 'The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme', *Hippocampus*, vol. 15, no. 7, pp. 913-22.
- Lisman, J.E. and Idiart, M.A. (1995) 'Storage of 7 +/- 2 short-term memories in oscillatory subcycles', *Science*, vol. 267, no. 5203, pp. 1512–1515.
- Lohmann, G., Margulies, D.S., Horstmann, A., Pleger, B., Lepsien, J., Goldhahn, D., Schloegl, H., Stumvoll, M., Villringer, A. and Turner, R. (2010) 'Eigenvector centrality mapping for analyzing connectivity patterns in fMRI data of the human brain', *PLoS ONE*, vol. 5, no. 4, p. e10232.
- Maistrenko, Y., Lysyansky, B., Hauptmann, C., Burylko, O. and Tass, P. (2007) 'Multistability in the Kuramoto model with synaptic plasticity', *Physical Review E*, vol. 75, no. 6, pp. 1–8.
- Markram, H., Lübke, J., Frotscher, M., Roth, A. and Sakmann, B. (1997) 'Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex', *Journal of physiology*, vol. 500, no. Pt 2, pp. 409-40.
- Masquelier, T., Guyonneau, R. and Thorpe, S.J. (2008) 'Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains', *PLoS ONE*, vol. 1, p. e1377.
- Masuda, N. and Aihara, K. (2004) 'Self-Organizing Dual Coding Based on input time population rate code mode', *Neural Computation*, vol. 16, pp. 627-663.
- McCormick, D.A., Sejnowski, T.J. and Steriade, M. (1993) 'Thalamocortical oscillations in the sleeping and aroused brain', *Science*, vol. 262, pp. 679-85.
- Meunier, D., Lambiotte, R. and Bullmore, E.T. (2010) 'Modular and hierarchically modular organization of brain networks', *Frontiers in neuroscience*, vol. 4, December, p. 200.
- Miltner, W.H.R., Braun, C., Matthias, A., Witte, H. and Taub, E. (1999) 'Coherence of gamma-band EEG activity as a basis for associative learning', *Nature*, vol. 397, no. 6718, pp. 434–436.
- Morrison, A., Diesmann, M. and Gerstner, W. (2008) 'Phenomenological models of synaptic plasticity based on spike timing', *Biological cybernetics*, vol. 98, no. 6, pp. 459-78.

Niedermeyer, E. and da Silva, F.L. (2004) *Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*, Lippincot Williams & Wilkins.

Nyhus, E. and Curran, T. (2010) 'Functional role of gamma and theta oscillations in episodic memory', *Neuroscience and biobehavioral reviews*, vol. 34, no. 7, pp. 1023-35.

Palva, J.M., Palva, S. and Kaila, K. (2005) 'Phase synchrony among neuronal oscillations in the human cortex', *The Journal of Neuroscience*, vol. 25, no. 15, pp. 3962–3972.

Papp, E., Leinekugel, X., Henze, D.A., Lee, J. and Buzsáki, G. (2001) 'The apical shaft of CA1 pyramidal cells is under GABAergic interneuronal control', *Neuroscience*, vol. 102, no. 4, pp. 715-721.

Park, H.J. and Friston, K. (2013) 'Structural and functional brain networks: from connections to cognition', *Science*, vol. 342, no. 6158, p. 1238411.

Parra, J., Kalitzin, S.N., Iriarte, J., Blanes, W., Velis, D.N. and Lopes da Silva, F.H. (2003) 'Gamma-band phase clustering and photosensitivity: is there an underlying mechanism common to photosensitive epilepsy and visual perception?', *Brain*, vol. 126, no. 5, pp. 1164–1172.

Penttonen, M. and Buzsáki, G. (2003) 'Natural logarithmic relationship between brain oscillators', *Thalamus & Related Systems*, vol. 2, no. 2, pp. 145–152.

Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P. and Andersen, R.A. (2002) 'Temporal structure in neuronal activity during working memory in macaque parietal cortex', *Nature Neuroscience*, vol. 5, no. 8, pp. 805–811.

Pfurtscheller, G., Stancak Jr, A. and Neuper, C. (1996) 'Post-movement beta synchronization. correlate of an idling motor area?', *Electroencephalography and Clinical Neurophysiology*, vol. 98, no. 4, pp. 281–293.

Roopun, A.K., Kramer, M., Carracedo, L.M., Kaiser, M., Davies, C.H., Traub, R.D. and Kopell, N. (2008) 'Period concatenation underlies interactions between gamma and beta rhythms in neocortex', *Frontiers in cellular neuroscience*, vol. 2, April, p. 1.

Rubinov, M. and Sporns, O. (2010) 'Complex network measures of brain connectivity: Uses and interpretations', *NeuroImage*, vol. 52, pp. 1059-69.

-
- Santos, B.A., Barandiaran, X.E. and Husbands, P. (2011) 'Metastable dynamical regimes in an oscillatory network modulated by an agents sensorimotor loop', *Proceedings IEEE Symposium on Artificial Life 2011*.
- Shanahan, M. (2010) 'Metastable chimera states in community structured oscillator networks', *Chaos*, vol. 20, no. 1, p. 013108.
- Shanahan, M. (2012) 'The brain's connective core and its role in animal cognition', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, vol. 367, no. 1603, pp. 2704-14.
- Shanahan, M., Bingman, V.P., Shimizu, T., Wild, M. and Güntürkün, O. (2013) 'Large-scale network organization in the avian forebrain: a connectivity matrix and theoretical analysis', *Frontiers in computational neuroscience*, vol. 7, July, p. 89.
- Shanahan, M. and Wildie, M. (2012) 'Knotty-centrality: finding the connective core of a complex network', *PLoS ONE*, vol. 7, no. 5, p. e36579.
- Shirvalkar, P.R., Rapp, P.R. and Shapiro, M.L. (2010) 'Bidirectional changes to hippocampal Theta-Gamma comodulation predict memory for recent spatial episodes', *Proceedings of the National Academy of Sciences*, vol. 107, no. 15, pp. 7054–7059.
- Siegel, M., Warden, M.R. and Miller, E.K. (2009) 'Phase-dependent neuronal coding of objects in short-term memory', *Proceedings of the National Academy of Sciences of the United States of America*, vol. 106, no. 50, pp. 21341-6.
- Song, S., Miller, K.D. and Abbott, L.F. (2000) 'Competitive Hebbian learning through spike-timing-dependent synaptic plasticity', *Nature neuroscience*, vol. 3, no. 9, pp. 919-26.
- Sporns, O. (2013) 'Network attributes for segregation and integration in the human brain', *Current opinion in neurobiology*, vol. 23, no. 2, pp. 162-71.
- Sporns, O., Chialvo, D.R., Kaiser, M. and Hilgetag, C.C. (2004) 'Organization, development and function of complex brain networks.', *Trends in cognitive sciences*, vol. 8, no. 9, pp. 418-25.
- Sporns, O., Honey, C.J. and Kötter, R. (2007) 'Identification and classification of hubs in brain networks', *PLoS ONE*, vol. 2, no. 10, p. e1049.

Sporns, O. and Kötter, R. (2004) 'Motifs in brain networks', *PLoS biology*, vol. 2, no. 11, p. e369.

Sporns, O. and Tononi, G. (2001) 'Classes of network connectivity and dynamics', *Complexity*, vol. 7, no. 1, pp. 28–38.

Sporns, O., Tononi, G. and Edelman, G.M. (2000) 'Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices', *Cerebral cortex*, vol. 10, no. 2, pp. 127-41.

Steriade, M. (2001) 'Impact of network activities on neuronal properties in corticothalamic systems', *Journal of Neurophysiology*, vol. 86, no. 1, pp. 1–39.

Steriade, M., Jones, E.G. and Llinás, R.R. (1990) *Thalamic Oscillations and Signalling*, John Wiley and Sons.

Strogatz, S. (2000) 'From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators', *Physica D: Nonlinear Phenomena*, vol. 143, no. 1-4, pp. 1-20.

Tognoli, E. and Kelso, J.A.S. (2011) 'Brain coordination dynamics: True and false faces of phase synchrony and metastability', *Progress in Neurobiology*, vol. 87, no. 1, pp. 31–40.

Tononi, G., Sporns, O. and Edelman, G.M. (1994) 'A measure for brain complexity: relating functional segregation and integration in the nervous system', *Proceeding National Academy of Science USA*, vol. 91, pp. 5033-5037.

Tort, A.B.L., Komorowski, R.W., Manns, J.R., Kopell, N.J. and Eichenbaum, H. (2009) 'Theta-Gamma coupling increases during the learning of Item-Context associations', *Proceedings of the National Academy of Sciences*, vol. 106, no. 49, pp. 20942–20947.

van den Heuvel, M.P., Mandl, R.C., Kahn, R.S. and Pol, H.E.H. (2009) 'Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain', *Human Brain Mapping*, vol. 30, no. 10, pp. 3127–3141.

van den Heuvel, M.P. and Sporns, O. (2011) 'Rich-club organization of the human connectome', *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 31, no. 44, pp. 15775-86.

-
- Varela, F., Lachaux, J., Rodriguez, E. and Martinerie, J. (2001) 'The Brainweb: Phase synchronization and large-scale integration', *Nature Reviews Neuroscience*, vol. 2, no. 4, pp. 229–239.
- Vida, I., Bartos, M. and Jonas, P. (2006) 'Shunting inhibition improves robustness of gamma oscillations in hippocampal interneuron networks by homogenizing firing rates', *Neuron*, vol. 49, no. 1, pp. 107–17.
- Wang, X. (2010) 'Neurophysiological and computational principles of cortical rhythms in cognition', *Physiological Reviews*, vol. 90, no. 3, pp. 1195–1268.
- Watts, D.J. (1999) 'Small worlds: the dynamics of networks between order and randomness' Princeton University Press.
- Watts, D.J. and Strogatz, S.H. (1998) 'Collective dynamics of 'small-world' networks', *Nature*, vol. 393, no. 6684, pp. 440–442.
- Whittington, M.A., Traub, R.D., Kopell, N., Ermentrout, B. and Buhl, E.H. (2000) 'Inhibition-based rhythms: experimental and mathematical observations on network dynamics', *International journal of psychophysiology: official journal of the International Organization of Psychophysiology*, vol. 38, no. 3, pp. 315–36.
- Wildie, M. and Shanahan, M. (2011) 'Metastability and chimera states in modular pulse-coupled oscillator networks.', *Chaos*, vol. 22, no. 4, p. 043131.
- Wildie, M. and Shanahan, M. (2012) 'Establishing Communication between Neuronal Populations through Competitive Entrainment', *Frontiers in computational neuroscience*, vol. 5, January, p. 62.
- Wilson, H.R. and Cowan, J.D. (1972) 'Excitatory and inhibitory interactions in localized populations of model neurons', *Biophysical Journal*, vol. 12, pp. 1–24.
- Womelsdorf, T., Schoffelen, J.M., Oostenveld, R., Singer, W., Desimone, R. and Engel, A.K.F.P. (2007) 'Modulation of neuronal interactions through neuronal synchronization', *Science*, vol. 316, no. 5831, pp. 1609-12.

Yu, W., Sommer, G., Daniilidis, K. and Duncan, J.S. (2005) 'Using skew Gabor filter in source signal separation and local spectral orientation analysis', *Image and Vision Computing*, vol. 23, pp. 377-392.

Zuo, X.N., Ehmke, R., Mennes, M., Imperati, D., Castellanos, F.X., Sporns, O. and Milham, M.P. (2012) 'Network centrality in the human functional connectome', *Cerebral Cortex*, vol. 22, no. 8, pp. 1862–1875.