

UNIVERSITY OF CALIFORNIA
SANTA BARBARA

**Dynamical Characterization and Feedback Control of Oscillatory
Neural Systems**

by

Per Danzl

A Dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mechanical Engineering

Committee in charge:
Professor Jeffrey Moehlis, Chair
Professor Paul Atzberger
Professor Francesco Bullo
Professor João Hespanha
Professor Bradley Paden

June 2010

The dissertation of Per Danzl is approved:

Professor Paul Atzberger

Professor Francesco Bullo

Professor João Hespanha

Professor Bradley Paden

Professor Jeffrey Moehlis, Chair

June 2010

Dynamical Characterization and Feedback Control of Oscillatory Neural Systems

Copyright © 2010
by
Per Danzl

Dedicated to my wife, Susan

Acknowledgments

My sincere gratitude goes to my advisor, Jeff Moehlis. One could not ask for a more helpful and inspiring mentor. I would also like to thank the esteemed members of my committee: Paul Atzberger, Francesco Bullo, João Hespanha, and Brad Paden.

Through my graduate studies I had the pleasure of working with some truly exceptional collaborators. From the UCSB campus, Guillaume Bonnet, João Hespanha, Ali Nabi, and Robert Hansen contributed to the studies presented in this dissertation. Our collaboration with Tay Netoff and Tyler Stigen of the University of Minnesota has broadened our understanding of the experimental neuroscience research.

I am grateful for two years of generous funding from the National Science Foundation via the IGERT program in Computational Science and Engineering, led by Linda Petzold.

It has been a pleasure to be a part of the Moehlis research group. I would like to acknowledge the other researchers group: Lina Kim, Margot Kimura, Allison Kolpas, Barry DeMartini, Ali Nabi, Michael Busch, and Gabor Orosz.

My path to graduate school began when I was working as an engineer at Honeywell. There, my mentors Matt Arend and Corey Bourassa encouraged me to pursue a graduate degree, and have since been a continuing source of support.

Finally, I would like to thank my parents, John and Joni, and my wife, Susan. This endeavour would not have been possible without their encouragement, support and love.

Vita of Per Danzl

June 2010

EDUCATION

- **The University of California, Santa Barbara**, Santa Barbara, CA
Department of Mechanical Engineering
Center for Control Dynamical Systems and Computation
Ph.D. in Mechanical Engineering; Adviser: Jeffrey Moehlis
September 2005–June 2010; GPA: 3.93/4.00
Dissertation topic: *Dynamical Characterization and Feedback Control of Oscillatory Neural Systems*
- **The University of Minnesota**, Minneapolis, MN
Bachelor of Science in Mechanical Engineering
September 1998–December 2002; GPA: 3.55/4.00

HONORS AND AWARDS

- National Science Foundation *Integrative Graduate Education Research Traineeship* Fellowship, 2007–2008.
- Walter and Margaret Pierce Scholarship, 2000.
- Institute of Technology Deans List, 1999–2002.

Abstract

Dynamical Characterization and Feedback Control of Oscillatory Neural Systems

by

Per Danzl

Oscillatory spiking neurons have been shown to play an important role in dynamical diseases of the nervous system, particularly Parkinson's disease. This dissertation seeks to understand the population-level response dynamics of such neurons, and proposes several event-based feedback control strategies to control their spike timing and degree of synchronization. These control schemes are a significant departure from the existing open-loop electrical deep brain stimulation protocols.

In the first part of this dissertation, we consider population-level dynamics of oscillatory neurons. We study the ensemble response to independent spike trains with interspike intervals drawn from a Poisson distribution, which has been shown to be a reasonable representation of background spiking activity in the brain. The concept of partial phase synchronization, a quantitative tool with which to characterize the state of a population, is presented and shown to provide useful information that is distinct from measurements of spiking synchrony. Then, we investigate parametric resonance to sinusoidal stimulus, and how nonlinearity and coupling lead to a wide variety of stable and unstable solutions, which are categorized based on symmetry considerations and solution types.

In the second part of the dissertation, we consider several event-based feedback control schemes for spike timing control. The first approach uses biologically-inspired impulsive and quasi-impulsive stimulus protocols to drive a neural oscillator to spike in phase with a reference oscillator. We show how these control schemes can be used to desynchronize populations of neurons that, in absence of control, are driven to spike in synchrony by a pacemaker. The second approach uses a minimum-time-to-reach optimal control scheme, based on the Hamilton-Jacobi-Bellman framework, to drive a neural oscillator's state to a point at which its asymptotic phase is extremely sensitive to background noise, which in practice provides effective phase randomization. An extension of this control scheme is presented which prevents pathological synchronous spiking in a network of all-to-all coupled neurons. Finally, we propose a set of algorithms for stabilizing the interspike time interval between a pair of oscillatory neurons that is tailored to experimental implementation.

Contents

List of Figures	xii
List of Tables	xvii
1 Introduction	1
2 Neuron models	5
2.1 Introduction and background	5
2.2 Conductance-based models	6
2.2.1 Hodgkin-Huxley	7
2.2.2 2D Hodgkin-Huxley reduction	8
2.2.3 Fast-slow reduction: FitzHugh-Nagumo	10
2.2.4 Hindmarsh-Rose	10
2.3 Phase, isochrons, and phaseless sets	11
2.4 Phase-reduced models	12
2.4.1 Type II phase response curves	13
2.5 Parametric resonance models	14
2.6 Model identification	14
2.6.1 Example: FitzHugh-Nagumo	16
3 Population Dynamics: Response to Poisson Noise	18
3.1 Introduction and background	19
3.2 Stimulus model	20
3.3 Mechanism for partial phase synchronization	20
3.3.1 Intuitive description	20
3.3.2 Theoretical development	22
3.4 Simulation methods: phase sampling	26
3.4.1 Conductance-based models	26
3.4.2 Phase-reduced models	27
3.5 Results	28
3.5.1 Phase distribution dynamics	28
3.5.2 Parametric study	30
3.6 Discussion	31

3.6.1	Relationship between partial phase synchrony and spike synchrony	31
3.6.2	Effects of partial phase synchronization	36
4	Population Dynamics: Parametric Resonance	39
4.1	Introduction and background	40
4.2	A parametrically forced oscillator	41
4.3	Uncoupled parametrically forced oscillators	42
4.3.1	Combinatorial considerations	42
4.3.2	Symmetry considerations	44
4.4	Coupled parametrically forced oscillators	46
4.4.1	Periodic orbits for weakly coupled oscillators	46
4.4.2	$N = 2$ coupled oscillators	48
4.4.3	$N = 3$ coupled oscillators	50
4.4.4	N coupled oscillators	61
4.5	Discussion	62
5	Reference-tracking Spike Timing Control	65
5.1	Introduction and background	65
5.2	Control of individual neurons	67
5.2.1	Event-based control	67
5.2.2	Reference-phase tracking	67
5.2.3	Impulsive control	69
5.2.4	Quasi-impulsive control	71
5.3	Control of an Ensemble of Neurons	76
5.3.1	Desynchronizing an ensemble	76
5.3.2	Desynchronizing a pacemaker-driven ensemble	77
5.4	Discussion	82
6	Phase-randomizing Control	83
6.1	Introduction and background	84
6.2	Normalized model	85
6.3	Control objectives	86
6.4	Optimal stimulus design	86
6.5	Numerical methods	89
6.6	Results	90
6.6.1	Phase randomization: single neuron results	90
6.6.2	Phase randomization: ensemble results	94
6.6.3	Spike rate control	94
6.6.4	Connection to isochrons	95
6.6.5	FitzHugh-Nagumo model	97
6.7	Extension: globally coupled networks	98
6.8	Discussion	101

7	Master-slave Control: Pairwise Spike Timing	103
7.1	Introduction and background	103
7.2	Master-slave neuron models	104
7.3	Control objective	106
7.4	Algorithms	106
7.4.1	Pre-processing	106
7.4.2	Event-based algorithm	107
7.4.3	Numerical example	110
7.5	Discussion	113
8	Conclusion	114
	Bibliography	118

List of Figures

2.1	Hodgkin-Huxley I_b bifurcation diagram. The vertical axis displays voltage values for the fixed points and the maximum and minimum voltage values for the periodic orbits. Stable fixed points along the thick black line, unstable fixed points along dashed line. Stable periodic orbits indicated by filled circles, unstable periodic orbits as open circles. Bifurcation types as noted.	8
2.2	Stable periodic orbit for the reduced Hodgkin-Huxley model with $I_b = 10$ mA (red curve). Isochrons, spaced uniformly in phase, are shown as dashed lines. The location of the unstable equilibrium point, where the isochrons converge, is shown by a black X. The spike point (maximum membrane voltage) is shown by the blue square. . .	10
2.3	Phase response curve approximations for the Hindmarsh-Rose model with $I_b = 5$ mA (solid) and the Hodgkin-Huxley model with $I_b = 10$ mA (dashed).	13
2.4	Phase response curve calculated for the Hodgkin-Huxley neuron model for $I_b = 10$ mA [86]. Important points are labeled as follows: α is the phase at which the phase response curve takes on its minimal value, Z_{min} , β is where it takes on its maximal value Z_{max} , and γ is the interior zero-crossing.	14
2.5	(a) Simulated experimental voltage data (solid blue line) and voltage trace from estimated model (dashed green line). (b) Convergence of the cost function (2.15) for the parameter estimation procedure over 18 internal fmincon iterations.	17
3.1	Initial phases as black markers labeled i , phase after drift as open markers labeled i' , and phase after drift and stimulus as striped markers labeled i'' (a). Distribution of many neurons after exposure to stimulus (b).	21
3.2	Simulation results, $\rho(\theta, t)$, for Hindmarsh-Rose (a), (c) and Hodgkin-Huxley (b),(d) (resp., conductance-based model, phase-reduced model).	

3.3	Comparison of theoretical distribution estimate with averaged simulation data for both conductance-based and phase-reduced models for Hindmarsh-Rose (a) and Hodgkin-Huxley (b). Solid lines represent full conductance-based model simulation data, dashed lines represent phase-reduced model simulation data, and dotted lines represent the theoretical estimates.	30
3.4	Numerical (a) and theoretical (b) ρ_{max} over (α, \bar{I}) parameter space for Hindmarsh-Rose model.	31
3.5	Hodgkin-Huxley population-level response characteristics. Numerical ρ_{max} over (α, \bar{I}) parameter space (a). Numerical θ_{max} over (α, \bar{I}) parameter space (b). Comparison of numerical results (c) to theoretical predictions (d) over the range of (α, \bar{I}) for which the assumption $\omega/\alpha = \mathcal{O}(1)$ is reasonable.	32
3.6	Illustrations showing that partial phase synchrony and spike synchrony need not be strongly related. Plot (a) shows the scenario from Example 1 where there is a high degree of phase synchronization but no spike synchronization. Plot (b) shows the scenario from Example 2 where there is complete spike synchronization but only small time intervals of phase synchronization. The bottom figures show a quantification of phase synchrony, as represented by the magnitude of the Kuramoto order parameter. The results can be sharpened by increasing the number of neurons, and reducing the parameters Δ and ε	35
3.7	First Spike Time Histograms for Hindmarsh-Rose (a), (c) and Hodgkin-Huxley (b),(d). Plots (a) and (b) begin from a uniform distribution. Plots (c) and (d) begin from the partially synchronized distributions shown in Figure 3.3. This analysis tracks the time of the first spike after the unit step-function stimulus is turned on. 1000 realizations are used.	37
4.1	Bifurcation diagram for fixed $b = 0.2$ and $F = 0.5$. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.	43
4.2	Existence and stability of solutions for different regions of parameter space for $b = 0.2$, with bifurcations sets shown as lines.	43
4.3	Stable periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2$. The response is at half the frequency of the forcing. We will refer to this as the A solution.	44
4.4	Stable periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2$. This is shifted by one period of the forcing from the solution shown in Fig. 4.3. We will refer to this as the B solution.	45

4.5	Bifurcation diagram for fixed $b = 0.2, F = 0.5, \omega_f = 2$ with the coupling strength c treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. For small (positive) c , from top to bottom, the solutions of the branches approach the $AA, A0, AB, 0A,$ and 00 as $c \rightarrow 0$. Note, for example, that the AA and BB solutions have the same $\max(x_1)$ value, and thus appear to be on the same branch in this projection. A similar coincidence between symmetry-related solutions occurs for all other branches.	50
4.6	Bifurcation diagram for fixed $b = 0.2, F = 0.5, \omega_f = 2.25$ with the coupling strength c treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.	51
4.7	Partial bifurcation diagram for Eqs. (4.18)-(4.19) for fixed $b = 0.2, F = 0.5, c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.	52
4.8	Full bifurcation diagram for Eqs. (4.18)-(4.19) for fixed $b = 0.2, F = 0.5, c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.	53
4.9	Example large-small periodic orbit for $N = 2, b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$. In the lower panel, the dark line and light line correspond to the two different oscillators.	54
4.10	Time series for anti-synchronized chaos for $N = 2, b = 0.2, F = 0.5, \omega_f = 2,$ and $c = -0.9$, where the dark line and light line correspond to the two different oscillators.	54
4.11	Phase space representation of anti-synchronized chaos for $N = 2, b = 0.2, F = 0.5, \omega_f = 2,$ and $c = -0.9$, where the dark line and light line correspond to the two different oscillators.	55
4.12	$N = 3$ oscillator networks. Ring $[\tilde{c} = c]$ (a), morph $[0 < \tilde{c} < c](b),$ and line $[\tilde{c} = 0]$ (c).	55
4.13	Example of R_3 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with ring topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators, the latter two of which have identical solutions.	56
4.14	Example of R_4 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with ring topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators, the latter two of which have identical solutions.	56
4.15	$N = 3$ ring bifurcation diagram for fixed $b = 0.2, F = 0.5, c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.3, and have symmetry-related solutions tabulated in Table 4.2.	57

4.16	Time series for unsynchronized chaos for $N = 3$ ring $b = 0.2$, $F = 0.5$, $\omega_f = 2$, and $c = -0.9$, where the dark line, the dashed line, and light line correspond to the three different oscillators.	58
4.17	Phase space representation of unsynchronized chaos for $N = 3$ ring $b = 0.2$, $F = 0.5$, $\omega_f = 2$, and $c = -0.9$, where the dark line, the dashed line, and light line correspond to the three different oscillators.	59
4.18	Example of L_5 periodic orbit for $b = 0.2$, $F = 0.5$, $\omega_f = 2.25$, $c = 0.03$ with line topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators. . .	62
4.19	Example of L_9 periodic orbit for $b = 0.2$, $F = 0.5$, $\omega_f = 2.25$, $c = 0.03$ with line topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators. . .	63
4.20	$N = 3$ line bifurcation diagram for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter ($\max(x_1)$ projection). Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.5, and have symmetry-related solutions tabulated in Table 4.4.	64
4.21	$N = 3$ line bifurcation diagram for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter ($\max(x_2)$ projection). Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.5, and have symmetry-related solutions tabulated in Table 4.4.	64
5.1	Phase error sign convention examples. Reference trajectories are shown as dashed lines. The controlled neuron's phase is shown as a solid line. When the controlled neuron's phase reaches $2\pi = 0 \pmod{2\pi}$, it spikes and triggers assessment of phase error relative to the reference oscillator.	69
5.2	Quasi-impulsive control algorithm performance. The left plot shows the phase error gain. The right plot shows the $\Delta\theta \mapsto \Delta\theta^+$ map. Solid lines are results from the phase-reduced model to be compared with the white circle markers, which are results from the full-dimensional Hodgkin-Huxley system.	76
5.3	Pacemaker network. Each of the N neurons is driven by control signals u_i and the pacemaker voltage V_P . The controller observes spiking behavior of the individual neuron voltages V_i and the pacemaker voltage V_P	77
5.4	The map M_{HH} for $K_P = 2$	79

5.5	(a) Quasi-impulsive anti-pacemaker control scheme applied to a pacemaker-driven ensemble of $N = 10$ phase neurons. The spike timing of the ensemble, as well as the phase distribution, are desynchronized by the action of the controller. (b) When the controller is turned off, the pacemaker drives the ensemble back into spiking synchrony. The top panels show the time evolution of each neuron's phase. The middle panel tracks the most recent value of the ensemble's interspike interval (ISI). The bottom panels show the magnitude of the order parameter $ r_1 $	81
6.1	Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple circle). The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, n) using a control signal bounded by $\mu = 10$	91
6.2	(a) State-feedback optimal control law (grayscale background) with the periodic orbit (red line), optimal state trajectory (blue dashed line) and the target set (purple circle). Gray regions signify maximum control in the negative direction ($I = -10$ mA), and white regions signify maximum control in the positive direction ($I = +10$ mA). (b) Optimal voltage (black line) and gating variable (green line) trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, starting from the spike point, in minimal time ($t_{\text{end}} \approx 6.88\text{msec}$).	92
6.3	Ensemble asymptotic phase results for 10^5 uncoupled neurons, each starting from the spiking point. In (a), neurons are exposed only to unit-magnitude noise. In (b), the neurons are also exposed to the phase randomizing control signal.	93
6.4	Level sets of the numerical approximation of the value function \mathcal{V} for optimal spike-rate-increasing control. Superimposed are the periodic orbit (red line) and the target set (purple ellipse). Again, we consider $\mu = 10$	95
6.5	(a) State-feedback optimal control law for optimal spike rate increase. Notation in this figure is the same as Figure 6.2. (b) Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom). Here, the objective is to reach an ε -ball centered at the spike point in minimal time starting from the spike point. This control protocol decreases the interspike interval neuron from 11.84 msec to 7.74 msec, resulting the fastest spiking rate possible with a stimulus constrained by 10 mA.	96
6.6	Level sets approximating isochrons for the reduced Hodgkin-Huxley model generated using parameters $\varepsilon = 0.01, \mu = 1$	97

6.7	Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple line) for the FitzHugh-Nagumo model. The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, w) using a control signal bounded by $\mu = 0.2$	98
6.8	(a) State-feedback optimal control law, optimal state trajectory, and the target set for the FitzHugh-Nagumo model with $\mu = 0.2$. Notation in this figure is the same as in Figure 6.2. (b) Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, in minimal time ($t_{\text{end}} \approx 0.5$).	99
6.9	(a) Network results for $N = 100$ globally coupled Hodgkin-Huxley neurons, $k = 0.01$, $D = 1.0$, $\mu = 10$. Top panel shows voltage traces for each neuron. Middle panel shows the control signal applied to the population. Controller is activated at $t = 20$ msec indicated by the vertical dashed line. Bottom panel is the population-level observable $\bar{V}(t)$ with the threshold $V_{\text{thresh}} = 0$ shown as a red dashed line. (b) Raster plot of spike times for the network showing desynchronization.	102
7.1	Typical spike advance curve $f(I)$ invertible over (I_{\min}, I_{\max}) . This function characterizes the slave neuron's response to a stimulus of magnitude I applied at the maximal phase θ_c	105
7.2	Example master-slave control scenario. Slave phase, θ_s , shown as solid blue line with spike events as blue 'x' markers. Master phase, θ_m , shown as red dashed line with red '+' markers. Target slave spike times (master spike times offset by Δt_D) shown by green 'o' markers. The thin black dotted line is the optimal control phase, $\theta_c = 3\pi/2$, which is where we apply the stimulus pulse. In this example, the slave spikes converge to the desired phase offset (relative to the master spikes) after two periods of control. Then the controller simply adjusts the period of the slave each time to match that of the master.	111

List of Tables

3.1	Simulation parameters for Figures 3.2 and 3.3.	29
4.1	Periodic solutions for uncoupled two oscillator system	49
4.2	$N = 3$ ring symmetry classes	54
4.3	Solution representatives for $ c \approx 0$ for each symmetry class for ring topology	55
4.4	$N = 3$ line symmetry classes	60
4.5	Solutions representatives for $ c \approx 0$ for each symmetry class for line topology	61
5.1	Constraints on the tunable desired contraction factor K . If K is chosen to be greater than the value of the right column, the corner condition in the left column is satisfied.	71
5.2	Constraints on the minimum admissible stimulus constraint $C = C_{min}$. If C is chosen to be greater than the value of the right column, the condition in the left column is satisfied.	73
7.1	Values used for the numerical example in Figure 7.2.	110

Chapter 1

Introduction

This dissertation investigates the dynamical characteristics of oscillatory spiking neurons and proposes novel feedback-based stimulation protocols for controlling their spiking behavior. We restrict our study to relatively idealistic mathematical models, but the results presented here are best interpreted in the larger context of physiology and human disease.

Many human diseases have dynamical mechanisms involved in either their causes or symptoms, including those affecting the nervous system, endocrine system, cardiac and respiratory systems, circadian rhythms, hearing, vision, and hormone systems [54]. It has been noted that tools of dynamical systems, such as bifurcation and stability analysis, are naturally suited for understanding these dynamical diseases [7]. Synchronization is a particularly interesting dynamical behavior, especially in the context of neuronal systems where it has been observed to play a significant role in pathological tremor and epilepsy [98, 103, 15].

An important example of a dynamical disease is Parkinson's disease, a chronic progressive neurodegenerative disorder characterized by tremor, rigidity, and bradykinesia [12]. There is evidence to suggest that the pathological tremors associated with Parkinson's disease occur when groups of oscillatory neurons in the motor control region of the brain spike in synchrony [77, 34]. Such oscillatory spiking behavior has also been observed in extensive experimental investigations of the mammalian brain in the context of motor control, movement, and the sleep-wake cycle [24, 92].

Synchrony is an interesting phenomenon from a dynamical systems perspective. Most often in neuroscience, synchrony is discussed in terms of spike times [43, 72]. Phase synchronization is also an important characteristic that provides different information about the state of a population of oscillatory neurons. In this dissertation we will investigate the onset of spiking synchrony and phase synchrony as a function of background noise parameters with the goal of characterizing population level responses. By understanding the phase distribution, models of oscillatory neuronal populations involved in pathology, like Parkinsonian tremor, can be more accurately initialized and studied.

Parkinson's disease and its associated pathology, including neuronal synchrony, present interesting opportunities and challenges for controls engineering. Electrical

Deep Brain Stimulation (EDBS) is an FDA-approved treatment for Parkinson’s disease (which has also been shown to be effective for the treatment of depression and epilepsy). The current EDBS technology consists of a single electrode, implanted deep in the brain, that applies a high-frequency (on the order of 100 Hz) open-loop periodic stimulus waveform that can silence pathological tremors [10]. The efficacy of this treatment shows that it is possible to use implanted hardware to control pathological neuronal behavior.

Unfortunately, the current EDBS technology has several shortfalls. The efficacy of the EDBS system varies widely from patient to patient, and is governed by empirically derived stimulus waveforms with tunable parameters. Its open-loop nature means that the stimulus is always applied to the tissue, even if the patient is not currently experiencing symptoms. Over time, this can lead to excessive tissue damage which lessens efficacy and may ultimately lead to a situation where the EDBS system is no longer able to mitigate the symptoms.

There is much interest in designing EDBS control systems that use additional electrodes for feedback, thereby making the stimulus “demand-controlled” [98, 99]. A feedback-based approach is attractive from a clinical perspective in that the biological tissue is only stimulated when necessary, thereby reducing the overall accumulation of negative side effects of electrical stimulation, and also the amount of power required from the implanted battery.

The neuronal networks that constitute the motor control system in human brain are exceptionally complex, so we take a more modest approach by considering only single neurons or simple networks. From an experimental perspective, it is possible to isolate single neurons from a slice of excised rat brain in the laboratory, and equip them with sensory electrodes and stimulus electrodes. Using recently developed experimental techniques, one can gather data and implement novel closed-loop control algorithms on living cells in real time [94].

There are, however, challenges to implementing feedback control for neurons. First, the conductance-based formalism first proposed by Hodgkin and Huxley in [50], a general modeling methodology in neuroscience, yields systems of continuous-time ordinary differential equations that tend to be *highly* nonlinear. Many traditional feedback control system designs require the controlled system to be either linear or only *weakly* nonlinear. A second challenge is that a control scheme cannot stimulate biological tissue with arbitrarily large signals. The magnitude of the electrical stimulus must be constrained to acceptable ranges for hardware implementation and biological tissue tolerance. A third challenge is the fact that the only state that is directly observable is the neuron’s membrane voltage. A control system cannot measure the dynamic states of the many ion channels that play a critical role in the oscillatory behavior of neural spiking. This poses particular problems for implementing traditional nonlinear feedback control systems, which depend on continuous measurements of state [55]. The control strategies proposed in this dissertation are intended to at least partially overcome these challenges.

An outline of this dissertation is as follows. In Chapter 2, we review the modeling methodologies we will use to model individual neurons. This includes the

“full-dimensional” conductance-based models in the Hodgkin-Huxley formalism, as well as reduced dimensional models including fast-slow planar models, nonlinear phase oscillators, and parametric resonators. We will also review some central concepts of model identification, to enable the construction of appropriate models from experimental data. Chapter 3 presents results on the population-level dynamics. We investigate the response of ensembles of neurons to background noise in the form of spikes with interspike intervals drawn from a Poisson distribution. This chapter will introduce a phase sampling algorithm that allows conductance-based model data to be correctly interpreted in the context of equivalent phase models. Then, in Chapter 4, we consider small populations of coupled nonlinear resonators, addressing their bifurcation structure and symmetries, and their response to periodic input. In Chapters 5-7, we shift our focus to feedback control of neural oscillators. Spike timing control from a reference phase tracking perspective is discussed in Chapter 5. We develop event-based feedback schemes based on impulsive stimuli and extend our results to the quasi-impulsive case that is experimentally realizable. Results are presented that show how this scheme can desynchronize a population of spiking neurons, even when the ensemble is driven by an uncontrollable pacemaker neuron. Chapter 6 presents another event-based control scheme, based on the Hamilton-Jacobi-Bellman time-optimal framework, that effectively randomizes the asymptotic phase of the controlled neuron using constrained stimulus. We develop an extension that is shown to randomize phase in an all-to-all coupled network where every neuron receives the same stimulus. Some surprising connections are drawn between computational methods for optimal control and dynamical characteristics of the underlying neuron model, namely the isochrons. Chapter 7 details recent collaborations with experimentalists wherein we have developed a pair-wise spike timing algorithm that stabilizes a desired interspike interval even if the pair of neurons have different natural frequencies. In Chapter 8 we discuss the main conclusions drawn from this work, and also discuss future directions.

Publications

This work is supported by results originally documented in the following peer-reviewed publications:

- [28] P. Danzl, J. Hespanha, and J. Moehlis. Event-based minimum-time control of oscillatory neuron models. *Biological Cybernetics*, 101(5):387–399, 2009.
- [27] P. Danzl, R. Hansen, G. Bonnet, and J. Moehlis. Partial phase synchronization of neural populations due to random Poisson inputs. *Journal of Computational Neuroscience*, 25(1):141–157, 2008.
- [31] P. Danzl and J. Moehlis. Weakly coupled parametrically forced oscillator networks: existence, stability, and symmetry of solutions. *Nonlinear Dynamics*, 59(4): 661–680, 2010.

- P. Danzl, A. Nabi, and J. Moehlis. Charge-balanced spike timing control for phase models of spiking neurons. Submitted to *Discrete and Continuous Dynamical Systems*, 2010.
- [94] T. Stigen, P. Danzl, J. Moehlis, and T. Netoff. Linear control of neuronal spike timing using phase response curves. In Proceedings of the 31st Annual International Conference of the IEEE EMBS, pages 1541–1544, Minneapolis, MN, 2009.
- [30] P. Danzl and J. Moehlis. Spike timing control of oscillatory neuron models using impulsive and quasi-impulsive charge-balanced inputs. In *Proceedings of the 2008 American Control Conference*, pages 171–176, Seattle, WA, 2008.
- [29] P. Danzl and J. Moehlis. Event-based feedback control of nonlinear oscillators using phase response curves. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 5806–5811, New Orleans, LA, 2007.

Chapter 2

Neuron models

In this chapter, we introduce a set of dynamical models for the activity of oscillatory neurons and we provide some tools to fit such models to experimental observations. We focus on nonlinear ordinary differential equation (ODE) models based on the Hodgkin-Huxley formalism. Several techniques are used to generate reduced-order models that are useful from the perspective of dynamical characterization, numerical simulation, and control. Phase-reduction is introduced, which provides nonlinear phase oscillator models that will be of central importance to many of the results presented later in this dissertation.

Background information and context is presented in Section 2.1 along with a brief literature overview. In Section 2.2, we discuss an important approach to modeling neuronal dynamics, namely the conductance-based methodology first proposed by Nobel laureates Hodgkin and Huxley in their seminal work [50]. Also included in this section are two methods for reducing dimensionality of the resulting ODEs to planar systems. The concept of phase is introduced in Section 2.3, along with the related notions of isochrons and phaseless sets. Section 2.4 provides an overview of phase reduction methods, that allow us to model the dynamics of a higher dimensional nonlinear oscillatory ODEs as a one-dimensional phase oscillator. We also discuss the concepts of isochrons and phaseless sets. We introduce another type of oscillator model based on parametric resonance in Section 2.5, which extends the applicability of some of the results presented in this dissertation to microelectromechanical systems (MEMS). In Section 2.6, we briefly review optimization-based methods for fitting model parameters to experimental time-series data.

2.1 Introduction and background

Many methods of modeling single neurons have been developed to address the wide variety of perspectives and paradigms of interest in the neuroscience field [40]. Simple mathematical model types include integrate and fire, spike response, spike rate models, and mean field models like the Kuramoto model [59, 48, 96]. More detailed modeling paradigms include the conductance-based formalism and its stochastic extensions, as well as spatio-temporal and information theoretic

approaches. The voltage across the cell membrane is a fundamental measurable quantity that determines how neurons communicate with each other.

Nonlinear models of oscillatory neurons, such as the Hodgkin-Huxley model, display bifurcations in which oscillatory behavior can arise from a stable equilibrium due to the variation of a single parameter. Other oscillatory biological systems also display bifurcations [6] such as the development of abnormal tremor in tardive dyskinesia [74], respiratory arrest [78], and the focus of this study, Parkinsonian tremor [12, 13].

Networks of neurons present a wide variety of modeling paradigms as well [49] and pose additional challenges due to the fact that neural networks often have densely coupled architectures that are difficult, if not impossible, to experimentally map *in vivo*. The network topology also tends to change over time in living tissue.

We will focus on ODEs for modeling neural dynamics of individual neurons, and consider only the simplest of network topologies for the population-level results presented in this dissertation.

2.2 Conductance-based models

A neuron’s cellular membrane separates the ionic concentrations on the outside of the cell from those on the inside. By virtue of this separation of charge, the membrane may be viewed as a capacitor. Pores in the cell membrane allow the ions to migrate through, driven by the differences in concentrations on each side of the membrane. Collectively, these microscopic pores can be modeled as *gating variables* which determine how the set of ionic channels conduct ionic flow across the membrane. By viewing the cellular membrane as a capacitor in parallel with these ionic channels, one can construct a model of membrane voltage dynamics using electrical circuit analysis. The models arising from this methodology are known as “conductance-based” models.

A. L. Hodgkin and A. F. Huxley pioneered such a modeling framework in the early 1950’s. Using the squid *Loligo* giant axon, they were able to deduce the macroscopic behavior of the ionic channels and empirically fit curves describing the nonlinearities to such accuracy that simulations of the resulting models agreed well with experimental measurements. Since its inception, this modeling framework has been used by many researchers to characterize a wide variety of neurons from crustacean axons [25] to guinea pig pyramidal cells [102] to bursting pacemaker neurons responsible for respiratory rhythm generation in the pre-Botzinger complex in mammals [21]. In the case of long axons, Hodgkin and Huxley proposed a spatial extension of this framework [50], but for this work, we will consider the “space-clamped” ODE form; this formulation assumes spatially uniform membrane voltage, which is an appropriate simplification for modeling the general spiking behavior of oscillatory neurons. The general form for a space-clamped conductance-based ODE

system is

$$c\dot{V} = I_g(V, \mathbf{n}) + I_b + I(t), \quad (2.1)$$

$$\dot{\mathbf{n}} = \mathbf{G}(V, \mathbf{n}). \quad (2.2)$$

Here, $V \in \mathbb{R}$ is the voltage across the membrane, $\mathbf{n} \in \mathbb{R}_{[0,1]}^d$ is the vector of ionic channel gating variables, $c \in \mathbb{R}^+$ is the membrane capacitance, $I_g : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the sum of the membrane currents, $I_b \in \mathbb{R}$ is a constant baseline current, and $I_i : \mathbb{R} \rightarrow \mathbb{R}$ is the current stimulus.

We now describe several specific conductance-based models of this form for neurons which will be used throughout this dissertation.

2.2.1 Hodgkin-Huxley

A. L. Hodgkin and A. F. Huxley developed a four-dimensional model of the squid giant axon in 1952 [50]. Since its inception, this has served as a prototype for modeling membrane excitation of a class of neuron cells termed ‘‘Type II’’ [18] or ‘‘Class II’’ [51]. The system has also become important as a building block in systems of weakly connected neurons [51]. Following [18], the Hodgkin-Huxley equations are presented using several simplifications. The scaling of the model has been adjusted to give voltage in mV, current in mA, and time in msec. Although the following model was developed to characterize the membrane voltage dynamics of the squid giant axon, we will use it as a primary example of an oscillatory neuron for the simple reason that it is perhaps the most studied neuron model in mathematical neuroscience.

$$\begin{aligned} \dot{V} &= (I_b + I(t) - \bar{g}_{\text{Na}}h(V - V_{\text{Na}})m^3 - \bar{g}_{\text{K}}(V - V_{\text{K}})n^4 - \bar{g}_{\text{L}}(V - V_{\text{L}}))/c, \\ \dot{m} &= a_m(V)(1 - m) - b_m(V)m, \\ \dot{h} &= a_h(V)(1 - h) - b_h(V)h, \\ \dot{n} &= a_n(V)(1 - n) - b_n(V)n, \end{aligned}$$

$$\begin{aligned} \alpha_m(V) &= 0.1(V + 40)/(1 - \exp(-(V + 40)/10)), \\ \beta_m(V) &= 4 \exp(-(V + 65)/18), \\ \alpha_h(V) &= 0.07 \exp(-(V + 65)/20), \\ \beta_h(V) &= 1/(1 + \exp(-(V + 35)/10)), \\ \alpha_n(V) &= 0.01(V + 55)/(1 - \exp(-(V + 55)/10)), \\ \beta_n(V) &= 0.125 \exp(-(V + 65)/80). \end{aligned}$$

$$\begin{aligned} V_{\text{Na}} = 50 \text{ mV}, \quad V_{\text{K}} = -77 \text{ mV}, \quad V_{\text{L}} = -54.4 \text{ mV}, \quad \bar{g}_{\text{Na}} = 120 \text{ mS/cm}^2, \\ \bar{g}_{\text{K}} = 36 \text{ mS/cm}^2, \quad \bar{g}_{\text{L}} = 0.3 \text{ mS/cm}^2, \quad c = 1 \text{ }\mu\text{F/cm}^2. \end{aligned}$$

In this model, \bar{g}_{Na} , \bar{g}_{K} , and \bar{g}_{L} represent the conductances of the sodium, potassium, and leakage channels. The reversal potentials of these channels are V_{Na} ,

V_K , and V_L . The baseline current I_b will be treated as a bifurcation parameter. Figure 2.1 shows the bifurcations of the Hodgkin-Huxley system in the region of interest of I_b parameter space. For small values of I_b , the system has a single stable

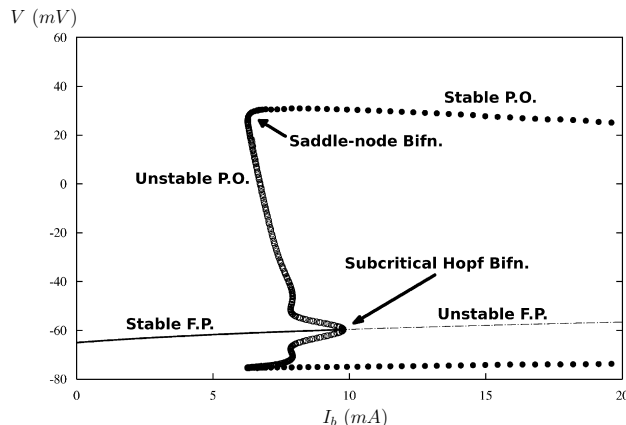


Figure 2.1: Hodgkin-Huxley I_b bifurcation diagram. The vertical axis displays voltage values for the fixed points and the maximum and minimum voltage values for the periodic orbits. Stable fixed points along the thick black line, unstable fixed points along dashed line. Stable periodic orbits indicated by filled circles, unstable periodic orbits as open circles. Bifurcation types as noted.

fixed point. This stable fixed point becomes unstable when the system undergoes a subcritical Hopf bifurcation that creates an unstable periodic orbit that then undergoes bifurcations and becomes stable through a saddle-node bifurcation. We note that at $I_b = 10$, the value used for much of the work presented in this dissertation, the fixed point is unstable and the large amplitude periodic orbit is stable.

2.2.2 2D Hodgkin-Huxley reduction

For the results presented later in Chapter 6, it will be useful to approximate the four-dimensional Hodgkin-Huxley system by a two-dimensional system. A planar system allows for complete visualization of the phase space on the plane, and allows the visualization of optimal cost functions as surfaces above this plane. We consider the two-dimensional reduction of the Hodgkin-Huxley model analyzed in, for example, [69, 54].

To motivate the dimension reduction, we note that the gating variables in m , h , and n in (2.3) can be represented by equations of the form:

$$\tau_j(V) \frac{dj}{dt} = j_\infty(V) - j, \quad (2.3)$$

where

$$j_\infty(V) \equiv \frac{\alpha_j(V)}{\alpha_j(V) + \beta_j(V)} \quad (2.4)$$

and

$$\tau_j(V) \equiv \frac{1}{\alpha_j(V) + \beta_j(V)}. \quad (2.5)$$

The n and h variables evolve much more slowly than the V and m states over the domain of interest. This is because the Na^+ channels activate quickly and the membrane potential changes quickly compared to the slow inactivation of the Na^+ channels and the slow activation of the K^+ ion channels. For this reason, we replace m by $m_\infty(V)$, which is the value to which m would relax if V were constant [54]. Furthermore, it has been noted that, for the full four-dimensional Hodgkin-Huxley model in the oscillatory regime, the following relationship holds between n and h :

$$n + h \approx 0.8. \quad (2.6)$$

We therefore make the substitution $h = 0.8 - n$.

The dynamics for this reduced model are represented by the following ODE system:

$$\begin{aligned} \dot{V} &= (I(t) + D\eta(t) + I_b - \bar{g}_{\text{Na}}[m_\infty(V)]^3(0.8 - n)(V - V_{\text{Na}}) \\ &\quad - \bar{g}_{\text{K}}n^4(V - V_{\text{K}}) - \bar{g}_{\text{L}}(V - V_{\text{L}})) / c \\ \dot{n} &= \alpha_n(V)(1 - n) - \beta_n(V)n, \end{aligned} \quad (2.7)$$

where

$$\begin{aligned} \alpha_n(V) &= \frac{0.01(V + 50)}{1 - \exp[-(V + 55)/10]} \\ \beta_n(V) &= 0.125 \exp[-(V + 65)/80] \\ m_\infty(V) &= \frac{\alpha_n(V)}{\alpha_n(V) + \beta_n(V)} \\ V_{\text{Na}} &= 50\text{mV}, \quad V_{\text{K}} = 77\text{mV}, \quad V_{\text{L}} = 54.4\text{mV}, \\ C &= 1\text{F/cm}^2, \quad \bar{g}_{\text{Na}} = 120\text{mS/cm}^2, \\ \bar{g}_{\text{K}} &= 36\text{mS/cm}^2, \quad \bar{g}_{\text{L}} = 0.3\text{mS/cm}^2. \end{aligned}$$

As stated earlier, we will use a baseline current $I_b = 10$ mA, for which the system has a stable periodic orbit shown in Figure 2.2. The term $\eta(t)$ is a zero-mean white noise signal, with strength D intended to represent background noise in the neurological context. Notice, here we have explicitly separated the noise process $D\eta(t)$ from the injected stimulus $I(t)$. Though it is a slightly different notation convention from the other ODE models presented here, it will greatly clarify the presentation of the results in Chapter 6.

In the forthcoming theoretical development, we will consider the noiseless deterministic case when $D = 0$. We will return to the case of non-zero noise after we have computed the optimal control stimulus. The membrane capacitance is $c = 1.0$ in the units of (2.7), so for notational convenience in the following section, we will incorporate it into our input variable $u(t) = I(t)/c$.

Also, in Chapter 6 we will use the following shorthand notation for the right hand side of (2.7):

$$\begin{aligned}\dot{V} &= f_V(V, n) + u, \\ \dot{n} &= f_n(V, n).\end{aligned}\tag{2.8}$$




Figure 2.2: Stable periodic orbit for the reduced Hodgkin-Huxley model with $I_b = 10$ mA (red curve). Isochrons, spaced uniformly in phase, are shown as dashed lines. The location of the unstable equilibrium point, where the isochrons converge, is shown by a black X. The spike point (maximum membrane voltage) is shown by the blue square.

2.2.3 Fast-slow reduction: FitzHugh-Nagumo

Next we consider another classical neural model, the FitzHugh-Nagumo model [37, 71] in a dimensionless form taken from [54]:

$$\begin{aligned}\delta \dot{V} &= V(V + a)(1 - V) - w + u \\ \dot{w} &= V - 0.5w.\end{aligned}\tag{2.9}$$

Here, u represents the input to the neuron. The parameter a determines the location of the unstable equilibrium point inside the periodic orbit. In the results presented in this dissertation we will take $a = 0.6$. The fast-slow time scale separation is readily apparent, and scales with the δ parameter, which we will take to be 0.01. An advantage of this model is that it displays the fundamental dynamical characteristics of the Hodgkin-Huxley system, but with much simpler equations. The nullclines, surfaces for which a dimension of the right hand side of the ODE is equal to zero, are simply represented by polynomial functions. We view this model as the distillation of the most important parts of an oscillatory neuron ODE.

2.2.4 Hindmarsh-Rose

As a prototypical Type I (meaning that positive input always advances spike timing) neuron model, we consider the Hindmarsh-Rose equations [85], which represent a reduction of the Connor model for crustacean axons [25]:

$$\begin{aligned}\dot{V} &= [I_b + I_i(t) - \bar{g}_{\text{Na}}m_\infty(V)^3(-3(q - Bb_\infty(V)) + 0.85)(V - V_{\text{Na}}) \\ &\quad - \bar{g}_{\text{K}}q(V - V_{\text{K}}) - \bar{g}_{\text{L}}(V - V_{\text{L}})]/c, \\ \dot{q} &= (q_\infty(V) - q)/\tau_q(V),\end{aligned}$$

$$\begin{aligned}q_\infty(V) &= n_\infty(V)^4 + Bb_\infty(V), \quad b_\infty(V) = (1/(1 + \exp(\gamma_b(V + 53.3))))^4, \\ m_\infty(V) &= \alpha_m(V)/(\alpha_m(V) + \beta_m(V)), \quad n_\infty(V) = \alpha_n(V)/(\alpha_n(V) + \beta_n(V)), \\ \tau_q(V) &= (\tau_b(V) + \tau_n(V))/2, \quad \tau_n(V) = T_n/(\alpha_n(V) + \beta_n(V)), \\ \tau_b(V) &= T_b(1.24 + 2.678/(1 + \exp((V + 50)/16.027))), \\ \alpha_n(V) &= 0.01(V + 45.7)/(1 - \exp(-(V + 45.7)/10)), \\ \alpha_m(V) &= 0.1(V + 29.7)/(1 - \exp(-(V + 29.7)/10)), \\ \beta_n(V) &= 0.125 \exp(-(V + 55.7)/80), \quad \beta_m(V) = 4 \exp(-(V + 54.7)/18).\end{aligned}$$

$$\begin{aligned}V_{\text{Na}} &= 55 \text{ mV}, \quad V_{\text{K}} = -72 \text{ mV}, \quad V_{\text{L}} = -17 \text{ mV}, \quad \bar{g}_{\text{Na}} = 120 \text{ mS/cm}^2, \\ \bar{g}_{\text{K}} &= 20 \text{ mS/cm}^2, \quad \bar{g}_{\text{L}} = 0.3 \text{ mS/cm}^2, \quad g_A = 47.7 \text{ mS/cm}^2, \\ c &= 1 \text{ } \mu\text{F/cm}^2, \quad \gamma_b = 0.069 \text{ mV}^{-1}, \\ T_b &= 1 \text{ msec}, \quad T_n = 0.52 \text{ msec}, \quad B = 1.26.\end{aligned}$$

Type I models yield non-negative phase response curves, which translate into functionally different response dynamics.

2.3 Phase, isochrons, and phaseless sets

The *phase* of an oscillator is a measure of the time it has spent evolving since it last passed through the marker event which defines zero phase. For neuron models, the spiking point, which for the 2D Hodgkin-Huxley model is the blue square shown in Figure 2.2, is conventionally taken as the marker event defining zero phase. Phase is an important concept for studying periodically spiking neuron models because it also indicates when the neuron will spike next in the absence of stimulus. Following [58], this extended concept of phase is called *asymptotic phase* and can be defined as follows:

Consider an arbitrary autonomous ODE

$$\dot{x} = f(x), \quad x \in \mathbb{R}^p, \quad p \geq 2 \tag{2.10}$$

with a stable hyperbolic limit cycle $\gamma(t)$ with natural period T , i.e. $\gamma(t + T) = \gamma(t)$ for all t . For each point x in the basin of attraction of the limit cycle γ , there exists

a unique asymptotic phase $\phi(x)$ such that

$$\lim_{t \rightarrow \infty} |x(t) - \gamma(t + \phi(x))| = 0. \quad (2.11)$$

State space is foliated with $(p-1)$ -dimensional hypersurfaces called *isochrons*, which are level sets of asymptotic phase [45]. For planar ($p = 2$) systems, the isochrons are simply lines of equal asymptotic phase. A set of isochrons for the reduced Hodgkin-Huxley model are shown as dashed lines in Figure 2.2. Points in state space where phase is not defined are called *phaseless sets*.

In Figure 2.2, one can see that the isochrons are most dense near the unstable equilibrium point at $(V_s, n_s) = (-59.6, 0.403)$. This unstable equilibrium point does not have a well-defined phase because it is not technically in the basin of attraction of the limit cycle. It is the only phaseless set “inside” the periodic orbit and is an attractive target for phase randomization. Later, in Chapter 6, we will show how to construct a stimulus to reach an arbitrarily small neighborhood of this point, then the noise process will randomly perturb the state onto a nearby isochron, effectively randomizing the asymptotic phase of the neuron.

2.4 Phase-reduced models

We choose baseline current, I_b , values such that each neuron’s only attractor is a stable periodic orbit: for example, $I_b = 5$ mA for the Hindmarsh-Rose model, and $I_b = 10$ mA for the Hodgkin-Huxley model. Since each system has a stable periodic orbit at the prescribed I_b value, it is useful to map the system to phase coordinates [108, 45, 59, 109, 18]. Following [45], there exists a homeomorphism $h: \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{S}^1$ mapping any point on the periodic orbit γ to a unique point the unit circle $\theta \in [0, 2\pi)$. Furthermore, in a neighborhood of the periodic orbit, the phases are determined by isochrons, whose level-sets are the sets of all initial conditions for which the distance between trajectories starting on the same isochron goes to zero as $t \rightarrow \infty$ [58]. By convention, when $\theta = 0$ the neuron fires. In the limit of small perturbations, the stimulus spikes serve to nudge the state slightly off the periodic orbit. In this way, the state may be moved onto a different isochron resulting in a difference in phase, with a magnitude and direction determined by the phase response curve (PRC). This reduction enables the dynamics of an $(d + 1)$ -ODE neuron model to be represented by the evolution of a scalar phase variable. A recent discussion of this framework was communicated in [47].

We can therefore model a neural population by a set of one-dimensional phase models - one for each neuron. In the presence of individual stimuli $I_i(t)$, this phase reduction yields the following uncoupled N -dimensional system of equations for the phases θ_i of an uncoupled population of N identical neurons [18, 47]:

$$\frac{d\theta_i}{dt} = \omega + \frac{Z_V(\theta_i)}{c} I_i(t), \quad i = 1, \dots, N. \quad (2.12)$$

Here $\theta_i \in [0, 2\pi)$, $\omega = 2\pi/T$ where T is the period of the periodic orbit, and $\mathbf{Z}(\theta)$ is the PRC. Of particular interest are perturbations in the voltage direction, i.e. $Z_V(\theta) = \frac{\partial \theta}{\partial V}$. The software package XPPAUT [36] was used to numerically calculate PRCs for the Hindmarsh-Rose and Hodgkin-Huxley models. For computational convenience, we use approximations to these PRCs shown in Figure 2.3: the Hindmarsh-Rose PRC is approximated by a curve-fit of the form $Z_V(\theta) \approx \frac{K}{\omega}(1 - \cos(\theta))$ where $K \approx 0.0036 \text{ (mV-msec)}^{-1}$ for $I_b = 5 \text{ mA}$ [18, 19], cf. [35]; the Hodgkin-Huxley PRC is approximated as a Fourier series with 21 terms.




Figure 2.3: Phase response curve approximations for the Hindmarsh-Rose model with $I_b = 5 \text{ mA}$ (solid) and the Hodgkin-Huxley model with $I_b = 10 \text{ mA}$ (dashed).

2.4.1 Type II phase response curves

We will consider neurons which possess a so-called “Type II” phase response curve [51], like the Hodgkin-Huxley model, seen in Figure 2.3. These phase response curves arise for systems exhibiting a Hopf (or Bautin) bifurcation [18]. In Figure 2.4, we show the Hodgkin-Huxley phase response curve labeled with the following important points:

$$\begin{aligned} \alpha &= \operatorname{argmin}(Z(\theta)) \quad , \quad Z_{min} = Z(\alpha) \\ \beta &= \operatorname{argmax}(Z(\theta)) \quad , \quad Z_{max} = Z(\beta). \end{aligned}$$

In general, phase response curves for Type II neurons are characterized by the following properties:

$$\begin{aligned} Z(0) &= 0 \quad , \quad Z'(0) < 0 \\ Z(\gamma) &= 0 \quad , \quad Z'(\gamma) > 0 \\ Z_{max} &> 0 \quad , \quad Z_{min} < 0 \\ 0 &< \alpha < \gamma < \beta < 2\pi \quad . \end{aligned} \tag{2.13}$$

Type II phase response curves present special challenges to control, notably the fact that they are not sector bounded. A positive stimulus can either advance or retard the phase depending on the phase at which it is administered.




Figure 2.4: Phase response curve calculated for the Hodgkin-Huxley neuron model for $I_b = 10$ mA [86]. Important points are labeled as follows: α is the phase at which the phase response curve takes on its minimal value, Z_{min} , β is where it takes on its maximal value Z_{max} , and γ is the interior zero-crossing.

2.5 Parametric resonance models

We remark that there are many other types of oscillators besides neurons. Furthermore, their oscillation can be due to external periodic forcing. In Chapter 4, we will consider another class of ODE models that arise from a form of the nonlinear Mathieu equation for a damped, parametrically forced oscillator, shown below:

$$\ddot{x} + b\dot{x} + x + x^3 = xF \cos(\omega_f t). \quad (2.14)$$

Here the term $b\dot{x}$ represents damping (we assume $b > 0$), the term $x + x^3$ represents a nonlinear restoring force, and the term $xF \cos(\omega_f t)$ represents parametric excitation which can be viewed as a time-periodic modulation of the linear part of the restoring force.

This model presents a particularly interesting phenomenon of resonance due to a driving force with frequency ω_f . By studying models of this type, we can extend the applicability of some of our results to microelectromechanical oscillators that could have utility as motion sensors and energy harvesters.

2.6 Model identification

Given a parametrized ODE neuron model and a set of experimental input-output data, we will show how one can obtain reasonable estimates of the parameters using an optimization-based parameter estimation approach. The treatment presented here is simplified; we refer to [32] for a more complete overview of optimization-based approaches for parameter estimation.

We note that filter-based methods, an alternative to the methods presented here, have also been used to estimate parameters for neuroscience models. A particular application is reported in [87], and, for a general review, we refer to [106]. These types of methods are particularly well suited when noise is present in the experimental data – a significant factor that we do not consider here.

Most often in experimental neuroscience, the only available experimental data is a length- N time-series voltage trace, $V_{\text{exp}}(t_k)$, sampled uniformly in time ($t_1 = \Delta t$, $t_2 = 2\Delta t$, \dots , $N\Delta t$), where Δt is the sampling period. The input to the neuron, $I(t_k)$, is experimentally generated and therefore known (or at least measurable).

The first step in model identification is to choose a parametrized model that will be fit to the experimental data. The conductance-based models presented in Section 2.2 are admissible types, as well as many other ODE models in the neuroscience literature. One groups all the unknown parameters, including the initial value of all unobservable states into a vector $p \in \mathbb{R}^{N_p}$, where N_p is the cardinality of this vector. In most cases, the unknown parameters will be known to lie in some subset $\mathcal{P} \subset \mathbb{R}^{N_p}$. Given the nonlinearity of the underlying dynamical models, this parameter estimation approach works best when \mathcal{P} can be bounded reasonably close to the actual value of the unknowns.

The objective is to find the “true” value of the unknown parameter vector, p , which for clarity will be denoted by p_{true} . We approach this as a constrained optimization problem, and will use numerical simulation to generate voltage trajectories for models with different guesses for p , denoted by the variable p_{guess} . We define the cost function for a generic choice of p as:

$$\mathcal{C}(p) = \|V_{\text{sim}}(p) - V_{\text{exp}}\|^2 \quad (2.15)$$

where $V_{\text{sim}}(p)$ is a length- N time-series voltage trace computed from numerically simulating the underlying dynamical system using parameter value p . We choose the Euclidean 2-norm for the function $\|\cdot\|$. We note that the initial value of the voltage is simply the first value of the experimental data V_{exp} , and is therefore known. Recall from above, the initial values of the unobservable states are not known and are treated as additional parameters in p . So, the cost function $\mathcal{C}(p_{\text{guess}})$ is simply a measure of the distance between the simulated data and the experimental data for a given guess of the parameter value p_{guess} .

We employ the MATLAB (R) Optimization Toolbox constrained minimization routine **fmincon** to take the initial parameter guess $p_{\text{guess}} \in \mathcal{P}$, assess the cost $\mathcal{C}(p_{\text{guess}})$, and develop a refined guess p_{guess}^+ . The optimizer iterates on this procedure using the active set line search method until either a local minimum is found, or the problem is declared unsolvable.

Convergence of this optimization procedure is not guaranteed, but in practice the optimizer tends to find local minima as long as the subset \mathcal{P} is not too large. The local optimal estimated parameter vector p_{opt} is given by the following:

$$p_{\text{opt}} = \arg \min_{p \in \mathcal{P}} \{\mathcal{C}(p)\}.$$

We reiterate here the importance of the fact that this p_{opt} is a local optimum.

Since, for the models we consider, the underlying dynamics are highly nonlinear, it is advisable to iterate on this constrained optimization process, each time beginning from a different initial guess p_{guess}^j drawn from a uniform distribution over \mathcal{P} . For each of the N_{it} iterations, successful convergence produces a local minimum p_{opt}^j . From this set $\{p_{\text{opt}}^j\}$, we choose the parameter vector that yields the lowest cost by Equation (2.15) as our final estimate for the parameter estimation algorithm.

Though there is no mathematically formal guarantee of success, this method provides a practical means to develop ODE-based models for real neurons in the laboratory, which is an important experimental connection to the theoretical results presented in this dissertation.

2.6.1 Example: FitzHugh-Nagumo

We now present an example of how to fit a parametrized model to experimental data using the FitzHugh-Nagumo model (2.9). First, we generate a set of time series voltage data, shown as the solid blue line in Figure 2.5(a). This data has been generated by simulation, but in practice it would be measured experimentally. The goal is to identify the parameters a and δ , as well as the initial condition, w_0 , of the unobservable state w . So in this case $p = [a \ \delta \ w_0]^T$. The simulated experimental data was generated using

$$p_{\text{true}} = [0.6 \ 0.01 \ 0.1]^T. \quad (2.16)$$

We now seek to develop some reasonable bounds on the domain \mathcal{P} of the parameters. For the FitzHugh-Nagumo system, it is known that a will be positive and of order one, and that the time scale separation variable δ will likely be between 0.001 and 0.1. The domain of the unobservable variable w is roughly bounded in magnitude by 0.5. Using this information, we propose the following bounds:


$$\mathcal{P} = (0.1, 1.0) \times (0.001, 0.1) \times (-0.5, 0.5). \quad (2.17)$$

Twenty runs of the optimization procedure were conducted with different p_0 values chosen at random from the set \mathcal{P} . One of the convergent runs began at $p_0 = [0.4235 \ 0.0908 \ 0.2265]^T$ and convergence was achieved after 18 internal **fmincon** iterations. Figure 2.5(b) shows the convergence in cost function (2.15) of the optimizer. The cost of the initial guess is 4.23×10^3 . After 18 internal iterations, the cost is 3.36×10^{-6} . The final estimate of the parameters is:


$$p_{\text{opt}} = [0.60001 \ 0.09999 \ 0.10000],$$

which is extremely close to the parameter value used to generate the observed data.

This example is highly idealized in that the bounds on the parameter values were known, the experimental data was noiseless, and the model being fit to the data had the identical structure of the model that generated the data. It shows that



(a)



(b)

Figure 2.5: (a) Simulated experimental voltage data (solid blue line) and voltage trace from estimated model (dashed green line). (b) Convergence of the cost function (2.15) for the parameter estimation procedure over 18 internal **fmincon** iterations.

the proposed procedure can successfully identify parameters, but more advanced methods may be required if noise and model selection are significant issues in the experiment.

Chapter 3

Population Dynamics: Response to Poisson Noise

The study presented in this chapter shows that populations of identical uncoupled neurons exhibit partial phase synchronization when stimulated with independent, random unidirectional current spikes with interspike time intervals drawn from a Poisson distribution. We characterize this partial synchronization using the phase distribution of the population, and consider analytical approximations and numerical simulations of phase-reduced models and the corresponding conductance-based models of typical Type I (Hindmarsh-Rose) and Type II (Hodgkin-Huxley) neurons, showing quantitatively how the extent of the partial phase synchronization depends on the magnitude and mean interspike frequency of the stimulus. Furthermore, we present several simple examples that disprove the notion that phase synchrony must be strongly related to spike synchrony. Instead, the importance of partial phase synchrony is shown to lie in its influence on the response of the population to stimulation, which we illustrate using first spike time histograms.

Section 3.1 provides some introductory background and a brief overview of the existing literature. We precisely describe the set of random stimuli in Section 3.2. Section 3.3 gives an intuitive description of the mechanism by which partial phase synchronization occurs, and presents a detailed theory for the long-time probability distribution function for the phase. In Section 3.4, we describe our simulation methods for both conductance-based models and phase models. This section provides the numerical tools necessary to interpret conductance-based model observations from the phase-based perspective. Section 3.5 gives quantitative details of how the extent of the phase synchronization depends on the magnitude of the stimulus and the mean interspike frequency of the stimulus for Hodgkin-Huxley and Hindmarsh-Rose neurons. We show that the simulation results for a population of phase oscillators very closely match those for a population of conductance-based neuron models exposed to the same type of stimulus, and both sets of simulation data yield long-time phase distributions that are very similar to the theoretical predictions. This verifies that the concept of phase synchronization, although developed based on phase oscillators, is an effective and accurate tool for modeling

populations of conductance-based neuron models. We discuss the practical implications of partial phase synchrony to the response of the population in Section 3.6.

3.1 Introduction and background

As discussed in Chapter 1, synchronized neural activity is believed to be important for various brain functions, including visual processing [44, 105], odor identification [38], signal encoding [93], cortical processing [88], learning [90], and memory [56]. It can also be detrimental. For example, resting tremor in patients with Parkinson’s disease has been linked to synchronization of a cluster of neurons in the thalamus and basal ganglia [77]. Similarly, essential tremor and epileptic seizures are commonly associated with synchronously firing neurons [34, 103], which can become (partially or fully) synchronized due to coupling and/or stimulation by common inputs.

Even common inputs that are random or noisy can lead to synchronization [42, 43, 72]. This applies to a broad range neuron models [84, 100], with little constraint on intrinsic properties. It has also been shown experimentally that some neurons, in particular olfactory bulb mitral cells, can synchronize in this manner *in vitro* [39]. This is relevant to spike timing reliability experiments, which found that repeated injection of the same fluctuating current into a single cortical neuron leads to a more reproducible spiking pattern than injection of a constant current [63]. Indeed, an experiment in which multiple uncoupled neurons are subjected to a common input is equivalent to an experiment in which a single neuron is subjected to the same input over multiple trials, as in [63], so that spike timing reliability can be viewed as synchronization across trials.

Typically, as in the references cited above, synchrony in the context of neuroscience is discussed in terms of synchronization of action potentials (spikes). Synchronization of spike times is a natural way to quantify the dynamic behavior of a population of neurons, since typically the only observable quantities are voltages. There is, however, another form of synchrony that can play an important role in the dynamic response of a population of oscillatory neurons to stimulus.

In this chapter, we consider *partial phase synchronization*, a characteristic that provides information about the dynamical state of a population of oscillatory neurons not easily obtained by studying spike synchrony alone. The concept of partial phase synchronization applies to populations of oscillatory neurons, each of which evolves in time according to dynamics that can be represented by a one-dimensional phase oscillator (a “simple clock”, in the terminology of [109]). As we discussed earlier in Section 2.3, the phase of a neuron relates its state, in time, to the firing of an action potential (or other marker event on its periodic orbit), as described for example in [45] and [18]. The dynamical state of a population of phase oscillators can be characterized by the distribution of their phases (over the unit circle). Partial phase synchronization refers to the degree to which this phase distribution possesses a single dominant mode, meaning that there is a higher density of neu-

rons with phases near this region than anywhere else on the circle. Complete phase synchronization is the limiting case of partial phase synchronization where all the neurons have exactly the same phase, which yields a phase distribution in the form of a Dirac delta function.

We will now proceed to examine how an uncoupled ensemble of neurons can develop phase synchronization due to independent random stimuli. Such stimuli represent background activity of other neurons, based on the experimental evidence presented in [89].

3.2 Stimulus model

To model the independent random stimuli, we suppose that each neuron receives δ -function current inputs of strength \bar{I} at times determined by a Poisson process with mean frequency α :

$$I_i(t) = \bar{I} \sum_k \delta(t - t_i^k), \quad (3.1)$$

where t_i^k is the time of the k^{th} input to the i^{th} neuron. The times of these inputs are determined by drawing the interspike intervals from the distribution

$$p(\tau; \alpha) = \alpha e^{-\alpha\tau}. \quad (3.2)$$

We emphasize that the neurons in our models are not coupled to each other, and that each one receives the same baseline current I_b but a *different* current stimulus $I_i(t)$.

In our analysis, we are interested in determining how the population-level behavior depends on α and \bar{I} . Experimentalists have found that the mean excitatory spike frequency in certain brain neurons is about 83 Hz or $\alpha = .083 \approx 0.1$ spikes/msec [89]. This falls at the high end of the *gamma-rhythm* range [57], and sets the scale for biologically relevant α values. We base our \bar{I} scaling on a unit 1 mA, so that a spike gives an instantaneous voltage change of order $\bar{I}/c = 1$ mV. Since the dynamic range of a periodically firing neuron is approximately 100 mV, this corresponds to a *small* perturbation in the *depolarizing* direction [99].

3.3 Mechanism for partial phase synchronization

3.3.1 Intuitive description

We take partial phase synchronization to mean that there is a higher probability of a neuron having a certain range of phases than another range of equal size. The following illustrates qualitatively how partial phase synchronization can occur for a population of neurons subjected to the previously described stimuli. We consider neurons that are modeled as phase oscillators, per Section 2.4. Take, for example, the PRC $Z_V(\theta) = 1 - \cos(\theta)$. Recall the time evolution of a neuron is governed

by (2.12). Now imagine a set of four such neurons on the unit circle with this PRC starting at $\{\theta_1, \dots, \theta_4\} = \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$. These are shown as black markers on Figure 3.1(a) and are labeled according to their i^{th} indices. Figure 3.1(a) also shows the phase of the four neurons in the absence of stimuli after some time interval Δt has elapsed: each neuron has advanced in phase by $\Delta\theta = \omega\Delta t$ to positions indicated with open markers, labeled i' . This behavior is termed *drift* and is simply due to the natural frequency of the neurons. If each of the neurons was exposed to a unit stimulus during this same time interval Δt , their phases advance additionally, as determined by the PRC, by $\Delta\theta = Z_V(\theta)\bar{I}/c$ (if the PRC contains negative values, the phase can also be retarded). The position of each of the neurons after drift and stimulus is indicated by striped markers labeled by i'' . Notice that the position of neuron 1 (starting at $\theta = 0$) with stimulus ($1''$) is the same as without stimulus ($1'$). This is because $Z_V(0) = 0$ for this example.




Figure 3.1: Initial phases as black markers labeled i , phase after drift as open markers labeled i' , and phase after drift and stimulus as striped markers labeled i'' (a). Distribution of many neurons after exposure to stimulus (b).

If all four neurons are subjected to stimulus, it is apparent from the i'' locations in Figure 3.1(a) that the distribution of neurons has changed. What began as uniform has become asymmetric. The situation becomes more apparent if many more neurons are displayed, as in Figure 3.1(b). Here we see a definite “bunching” around $\theta = 0$. Many neurons have similar phases, hence phase synchronization has occurred.

The above description assumed each neuron receives the same input. Conceptually, one can envision a similar argument if all neurons receive unidirectional independent Poisson-type stimuli with the same statistics, i.e. mean interspike frequency α and magnitude \bar{I} .

3.3.2 Theoretical development

Suppose that each neuron in a population received the same *non-random* stimulus $I(t)$. For example, $I(t)$ might be a step function stimulus [18] or a sinusoidal stimulus. We are interested in the probability that a neuron will have a phase between θ and $\theta + d\theta$ at time t , given by $\rho(\theta, t)d\theta$, where $\rho(\theta, t)$ is the probability distribution function for the population of neurons. The phase synchronization can be represented by the shape of $\rho(\theta, t)$ over $\theta \in [0, 2\pi)$ at time t . The following partial differential equation can be derived for $N \rightarrow \infty$ for such a population of neurons (e.g., [18]):

$$\frac{\partial \rho}{\partial t} = -\frac{\partial}{\partial \theta} [(\omega + Z(\theta)I(t)/c)\rho(\theta, t)]. \quad (3.3)$$

For the present problem with independent *random* inputs $I_i(t)$, the situation is more complicated. We proceed by deriving an expansion of the density evolution that is conducive to perturbation methods consistent with the neuron models and random inputs under consideration. For the values of mean spike frequency, α , and neuron natural firing frequency, ω , used in this chapter, it is reasonable to consider the ratio ω/α as $\mathcal{O}(1) = \mathcal{O}(\epsilon^0)$ for the Hodgkin-Huxley model and $\mathcal{O}(\epsilon)$ for the Hindmarsh-Rose model, where ϵ is a small parameter.

Kramers-Moyal Expansion

The probability distribution $\rho(\theta, t)$ obeys the Kramers-Moyal expansion [23]

$$\frac{\partial \rho(\theta, t)}{\partial t} = \sum_{n=1}^{\infty} \left(-\frac{\partial}{\partial \theta}\right)^n [D^{(n)}(\theta)\rho(\theta, t)], \quad (3.4)$$

where

$$D^{(n)}(\theta_0) = \frac{1}{n!} \lim_{\Delta t \rightarrow 0} \frac{\mathbb{E}[(\theta(\Delta t) - \theta_0)^n]}{\Delta t}. \quad (3.5)$$

Here \mathbb{E} denotes the expected value, and $\theta(0) = \theta_0$ for the realizations used to calculate $D^{(n)}(\theta_0)$. Now

$$\theta(\Delta t) = \theta_0 + \omega\Delta t + \Delta\theta,$$

where $\Delta\theta$ is the total jump in θ due to the random Poisson inputs in $[0, \Delta t)$. For sufficiently small Δt , we expect that either no inputs occur, in which case $\Delta\theta = 0$, or one input occurs, giving

$$\Delta\theta \approx \bar{I}Z(\theta_0)/c \equiv (\Delta\theta)_{jump}. \quad (3.6)$$

For $n = 1$ and sufficiently small Δt , we have

$$\mathbb{E}[\theta(\Delta t) - \theta_0] = \omega\Delta t + \mathbb{E}[\Delta\theta|\text{input}] \times p(\text{input}) + \mathbb{E}[\Delta\theta|\text{no input}] \times p(\text{no input}). \quad (3.7)$$

Using

$$\mathbb{E}[\Delta\theta|\text{input}] = (\Delta\theta)_{jump}, \quad \mathbb{E}[\Delta\theta|\text{no input}] = 0,$$

$$p(\text{input}) = \alpha\Delta t, \quad p(\text{no input}) = 1 - \alpha\Delta t$$

gives

$$\mathbb{E}[\theta(\Delta t) - \theta_0] = \omega\Delta t + \alpha\Delta t \frac{Z(\theta_0)\bar{I}}{c}, \quad (3.8)$$

so that

$$D^{(1)}(\theta_0) = \omega + \alpha \frac{Z(\theta_0)\bar{I}}{c}. \quad (3.9)$$

Similarly for $n > 1$,

$$\mathbb{E}[(\theta(\Delta t) - \theta_0)^n] = \mathbb{E}[(\omega\Delta t + \Delta\theta)^n] = \alpha\Delta t \left(\frac{\bar{I}Z(\theta_0)}{c} \right)^n + \mathcal{O}((\Delta t)^2), \quad (3.10)$$

so that

$$D^{(n)}(\theta_0) = \frac{\alpha}{n!} \left(\frac{\bar{I}Z(\theta_0)}{c} \right)^n \quad n > 1. \quad (3.11)$$

Finally, letting

$$Z(\theta) = Z_d\psi(\theta), \quad \epsilon = \frac{Z_d\bar{I}}{c}, \quad (3.12)$$

where Z_d and ϵ are nondimensional, we obtain the following equation for the steady ($\partial/\partial t = 0$) density $\rho_s(\theta)$:

$$0 = -\frac{\omega}{\alpha} \frac{d\rho_s}{d\theta} + \sum_{n=1}^{\infty} \epsilon^n \frac{(-1)^n}{n!} \left(\frac{d}{d\theta} \right)^n [(\psi(\theta))^n \rho_s(\theta)]. \quad (3.13)$$

We now consider the small ϵ limit, corresponding to random inputs with small amplitude. We substitute

$$\rho_s(\theta) = \rho_0(\theta) + \epsilon\rho_1(\theta) + \epsilon^2\rho_2(\theta) + \dots \quad (3.14)$$

into (3.13), and to solve at successive orders of ϵ . We divide this into two cases based on the relative size of ω/α .

Hodgkin-Huxley: $\frac{\omega}{\alpha} = \mathcal{O}(1)$

This case applies when the ratio of natural frequency to mean spike frequency is of order 1, which corresponds to the Hodgkin-Huxley model in our analysis.

At $\mathcal{O}(\epsilon^0)$, we find that

$$\frac{d\rho_0}{d\theta} = 0 \quad \Rightarrow \quad \rho_0 = \text{constant} = \frac{1}{2\pi}, \quad (3.15)$$

where the value of the constant follows from the normalization

$$1 = \int_0^{2\pi} \rho_0 d\theta. \quad (3.16)$$

At $\mathcal{O}(\epsilon)$, we obtain

$$-\frac{\omega}{\alpha} \frac{d\rho_1}{d\theta} - \frac{d\psi}{d\theta} \rho_0 = 0, \quad (3.17)$$

which has solution

$$\rho_1(\theta) = -\frac{\alpha}{2\pi\omega} \psi(\theta) + k_1. \quad (3.18)$$

The value of the constant k_1 is determined by the normalization condition

$$\int_0^{2\pi} \rho_1(\theta) d\theta = 0 \quad \Rightarrow \quad k_1 = \frac{\alpha}{(2\pi)^2 \omega} \int_0^{2\pi} \psi(\theta) d\theta. \quad (3.19)$$

Thus, to $\mathcal{O}(\epsilon)$,

$$\rho_s(\theta) \approx \frac{1}{2\pi} + \frac{\alpha \bar{I}}{2\pi\omega c} \left(-Z(\theta) + \frac{1}{2\pi} \int_0^{2\pi} Z(\theta) d\theta \right). \quad (3.20)$$

At $\mathcal{O}(\epsilon^2)$, we obtain

$$-\frac{\omega}{\alpha} \frac{d\rho_2}{d\theta} - \frac{d}{d\theta} [\psi(\theta) \rho_1(\theta)] + \frac{\rho_0}{2} \frac{d^2}{d\theta^2} [(\psi(\theta))^2] = 0. \quad (3.21)$$

This has solution

$$\rho_2(\theta) = -\frac{\alpha}{\omega} \psi(\theta) \rho_1(\theta) + \frac{\alpha}{4\pi\omega} \frac{d}{d\theta} [(\psi(\theta))^2] + k_2. \quad (3.22)$$

Using the normalization condition

$$\int_0^{2\pi} \rho_2(\theta) d\theta = 0, \quad (3.23)$$

we obtain the following approximation for $\rho_s(\theta)$ accurate to $\mathcal{O}(\epsilon^2)$:

$$\begin{aligned} \rho_s(\theta) \approx & \frac{1}{2\pi} + \frac{\alpha \bar{I}}{2\pi\omega c} \left(-Z(\theta) + \frac{1}{2\pi} \int_0^{2\pi} Z(\theta) d\theta \right) \\ & + \frac{\alpha^2 \bar{I}^2}{2\pi\omega^2 c^2} \left\{ [Z(\theta)]^2 - \frac{1}{2\pi} Z(\theta) \int_0^{2\pi} Z(\theta) d\theta \right\} \\ & + \frac{\alpha \bar{I}^2}{4\pi\omega c^2} \frac{d}{d\theta} [(Z(\theta))^2] + k'_2, \end{aligned} \quad (3.24)$$

where

$$k'_2 = \frac{1}{2\pi} \int_0^{2\pi} \left\{ \frac{\alpha \bar{I}}{\omega c} Z(\theta) \left[\frac{\alpha \bar{I}}{2\pi\omega c} \left(-Z(\theta) + \frac{1}{2\pi} \int_0^{2\pi} Z(\theta) d\theta \right) \right] - \frac{\alpha}{4\pi\omega} \left(\frac{\bar{I}}{c} \right)^2 \frac{d}{d\theta} [(Z(\theta))^2] \right\} d\theta. \quad (3.25)$$

The truncation at $\mathcal{O}(\epsilon^2)$ is equivalent to formulating the problem in terms of a Fokker-Plank equation, the same form of which is known as the diffusion equation.

Hindmarsh-Rose: $\frac{\omega}{\alpha} = \mathcal{O}(\epsilon)$

This case applies when the natural frequency is much smaller than the mean spike frequency, which corresponds to the Hindmarsh-Rose model in the regime studied in this chapter.

We let

$$\frac{\omega}{\alpha} = \epsilon \tilde{\omega}, \quad (3.26)$$

where $\tilde{\omega} = \mathcal{O}(1)$. At $\mathcal{O}(\epsilon)$ we obtain

$$\tilde{\omega} \frac{d\rho_0}{d\theta} = -\frac{d}{d\theta} [\psi(\theta) \rho_0(\theta)]. \quad (3.27)$$

This has solution

$$\rho_0(\theta) = \frac{c_0}{\omega + \frac{\alpha \bar{I}}{c} Z(\theta)}, \quad (3.28)$$

where

$$1 = \int_0^{2\pi} \rho_0(\theta) d\theta \quad \Rightarrow \quad c_0 = \left(\int_0^{2\pi} \frac{d\theta}{\omega + \frac{\alpha \bar{I}}{c} Z(\theta)} \right)^{-1}. \quad (3.29)$$

A useful interpretation for this result is that, on average, the current $\alpha \bar{I}$ enters each neuron during every time unit. If such current came in uniformly, so that $I(t) = \alpha \bar{I}$, equation (3.3) would have a steady distribution given by (3.28).

At $\mathcal{O}(\epsilon^2)$ we obtain

$$\tilde{\omega} \frac{d\rho_1}{d\theta} = -\frac{d}{d\theta} [\psi(\theta) \rho_1(\theta)] + \frac{1}{2} \frac{d^2}{d\theta^2} [(\psi(\theta))^2 \rho_0(\theta)]. \quad (3.30)$$

This has solution

$$\rho_1(\theta) = \frac{\alpha \bar{I} Z_d \frac{d}{d\theta} [(\psi(\theta))^2 \rho_0(\theta)] + c_1}{2c \left(\omega + \frac{\alpha \bar{I}}{c} Z(\theta) \right)}. \quad (3.31)$$

Using the normalization condition

$$0 = \int_0^{2\pi} \rho_1 d\theta, \quad (3.32)$$

we obtain the following approximation for $\rho_s(\theta)$ accurate to $\mathcal{O}(\epsilon)$:

$$\rho_s(\theta) \approx \frac{c_0}{\omega + \frac{\alpha \bar{I}}{c} Z(\theta)} + \frac{\alpha \bar{I}^2 \frac{d}{d\theta} [(Z(\theta))^2 \rho_0(\theta)] + c'_1}{2c^2 \left(\omega + \frac{\alpha \bar{I}}{c} Z(\theta) \right)}, \quad (3.33)$$

where

$$c'_1 = -c_0 \int_0^{2\pi} \frac{d}{d\theta} [(Z(\theta))^2 \rho_0(\theta)] \frac{d\theta}{\omega + \frac{\alpha \bar{I}}{c} Z(\theta)}. \quad (3.34)$$

A different approach for the derivation of the steady probability distribution for a related problem is given in [72], which is primarily concerned with showing

phase synchronization in the case in which all $I_i(t)$ in (2.12) are identical. Here the dynamics are reduced to a random phase map, and the evolution of the density associated with this map is described by a generalized Frobenius-Perron equation. The steady distribution can be found numerically, or analytically in certain limits. We find that our theory described above gives good agreement with our numerical results, and provides a straightforward alternative to the approach of [72].

3.4 Simulation methods: phase sampling

3.4.1 Conductance-based models

A numerical routine was constructed to simulate large populations of uncoupled neurons described by $(m + 1)$ -dimensional conductance-based models. Since the neurons are not coupled, the simulations can be conducted on a neuron-by-neuron basis. A second-order Runge-Kutta method with small fixed time step Δt was used for $\mathcal{O}((\Delta t)^2)$ accuracy and compatibility with the Poisson-distributed stimuli. In our numerical models, we approximate the δ -function by a rectangular spike of duration Δt and magnitude $1/\Delta t$.

To obtain a uniform initial phase distribution across the population, the $(m+1)$ -dimensional periodic orbit is calculated at run-time as a set of $(m+2)$ -tuples of the states and time. The time is then scaled to $[0, 2\pi)$, and initial states are assigned by drawing a scalar θ_0 from a uniform random number generator and using the periodic orbit to interpolate the initial data. Once the initial condition has been set, the system is integrated. Periodically throughout the time evolution of each neuron, the program computes its phase using a routine which will presently be described, yielding a sampled time-series of the neuron's phase even though the equations are not in phase-reduced form.

Intuitively, the phase should be related to the time it would take the neuron to fire, beginning at its present state, in the absence of stimuli. Computationally, some additional care is required. We use the Hodgkin-Huxley equations as an example to illustrate some of the computational complexities. For the I_b values considered, the vector field of this system contains a spiral source near the periodic orbit [83], an example of a phase singularity [109]. Therefore, it is possible for a trajectory on the periodic orbit passing near this source to be pushed arbitrarily close to it by stimulus. A neuron with a state very near the spiral source may take a significant amount of time to evolve back to a neighborhood close to the periodic orbit, especially if the (positive) real parts of the unstable eigenvalues are small. This situation is extremely rare for the models and parameter ranges we consider. In fact, it did not occur for any of the simulations which generated the data presented in this chapter. However, to be useful in a general setting, the phase calculation scheme must be able to identify this case and alert the program that a phase singularity has been reached. In most settings dealing with behavior of large populations, these cases can be dismissed without loss of statistical significance,

since they appear to occur on a set of extremely small measure.

To compute phase, we use the current state vector as the initial ($t = 0$) data for a new simulation which is run in absence of stimulus. This simulation is integrated until two spikes are observed. Let the time it takes for a spike (above a set voltage threshold) to occur be T_1 (a large upper bound on the allowable T_1 is used to identify possible phase singularities). Let the time at which the next spike occurs be T_2 . The interspike interval between the two is then $T_{int} = T_2 - T_1$. Then phase can be computed as

$$\theta = 2\pi \left(1 - \frac{T_1 \pmod{T_{int}}}{T_{int}} \right). \quad (3.35)$$

An important component of this algorithm is its ability to correctly identify a spike. We use a selection of logical checks which test the voltage time series for maxima above a set threshold, and have a prescribed minimum interspike time interval. Obviously, some knowledge of the spike magnitude and period is essential, but this is readily available since the computation of the periodic orbit at run-time provides all the necessary information. Using this method, one can simulate N neurons for P sample steps and receive, as output, an $N \times P$ matrix containing the discrete time evolution of θ_i as the values in row i .

The above method provides an accurate and robust way of simulating and determining phase information for large populations of uncoupled neurons of arbitrary dimension with a wide range of forcing functions and environmental parameters. Not limited to *small* perturbations, the utility of this method is constrained only by the relative stiffness of the neuron model and the geometry of the basin of attraction of the periodic orbit. However, such flexibility comes at the price of computational speed, which is strongly dependent on the dimension of the neuron model and the type of solver used.

3.4.2 Phase-reduced models

We are interested in the population dynamics over as wide a (α, \bar{I}) range as possible. In order to more efficiently map the parameter space, we use the phase-reduced form of the neuron equations. We have implemented an algorithm which very efficiently simulates (2.12) for large N and for long times. It exploits the fact that equation (2.12) can be solved exactly for the time intervals for which no inputs are present. One gets times at which inputs occur by recognizing that the time interval τ between subsequent inputs for an individual neuron can be obtained by sampling the distribution (3.2). When an input comes for the i^{th} neuron, we instantaneously let

$$\theta_i \rightarrow \theta_i + Z_V(\theta_i)\bar{I}/c.$$

Again, the uncoupled assumption of the population allows each θ_i to be simulated independently, and the program creates an $N \times P$ output matrix of the discrete time evolution of θ_i .

3.5 Results

We simulate populations of uncoupled neurons each subjected to independent spike trains of set magnitude, \bar{I} , with interspike intervals drawn from a Poisson distribution parametrized by mean spike frequency α . Our results include grayscale plots of $\rho(\theta, t)$ showing the time evolution of the phase distribution of the population, time-averaged distribution curves, and mappings of the distribution peak value and location as a function of \bar{I} and α . Comparisons are made between simulations of the full conductance-based models, simulations of the phase models, and theoretical estimates.

For notational convenience, we will refer to the $N \times P$ output matrix as Θ . We remind the reader that the i^{th} row of Θ represents the time-series values of θ_i , the phase of neuron i , as it evolves from initial time t_0 to final time t_f . While the differential equations are solved using a small time step Δt , we report phase using a larger time step Δp . This ensures the displayed data will not have spurious characteristics due to limitations of graphics resolution, and results in a dramatic performance increase in the simulations of the full conductance-based models. Therefore, the P columns of Θ represent the $(t_f - t_0)/\Delta p$ sample points, plus a left-concatenated column of the initial phases, i.e. $P = 1 + (t_f - t_0)/\Delta p$.

3.5.1 Phase distribution dynamics

The phase density of the population at each sample point, $\rho(\theta, t)$ is computed from Θ by taking an appropriately scaled histogram of the corresponding column of Θ . For example, $\rho(\theta)$ at the k^{th} sample point is calculated by

$$\rho(\theta, t = k\Delta p) = \text{hist}(\Theta \mathbf{e}_k, n_{bins}) \frac{n_{bins}}{2\pi N}, \quad (3.36)$$

where $\text{hist}(\cdot, \cdot)$ is the standard histogram binning function, n_{bins} is the number of bins dividing the interval $[0, 2\pi)$, and \mathbf{e}_k is an $P \times 1$ vector of zeros with a 1 in the k^{th} entry. The argument $\Theta \mathbf{e}_k$ is simply the k^{th} column of Θ , representing the phase of all the neurons at time $t = (k + 1)\Delta p$, where $k = 1$ implies $t = t_0$. The scaling used in (3.36) gives the normalizing condition

$$\sum_{j=1}^{n_{bins}} \rho\left(\theta = \frac{j}{n_{bins}2\pi}, t\right) \Delta\theta = 1, \quad \forall t, \quad (3.37)$$

where $\Delta\theta$ is the bin size $2\pi/n_{bins}$. Equation (3.36) returns a $n_{bins} \times 1$ vector discretization of $\rho(\theta, t)$.

Populations of 1000 neurons were simulated, beginning with a uniform phase distribution at $t_0 = 0$, and subjected to random sequences of spikes with magnitude $\bar{I} = 1$ mA. An integration time step of $\Delta t = 0.01$ msec was used. Since the time scales of the Hindmarsh-Rose and Hodgkin-Huxley equations are significantly different, separate simulation parameters are necessary, as given in Table 3.1.

Parameter	Hindmarsh-Rose	Hodgkin-Huxley
I_b (mA)	5	10
α (spikes/msec)	0.1	1
t_f (msec)	1000	100
Δp (msec)	1	0.1

Table 3.1: Simulation parameters for Figures 3.2 and 3.3.

Figures 3.2(a) and 3.2(b) show our results for the simulations of the full conductance-based models, to be compared with Figures 3.2(c) and 3.2(d), which are the results from the phase-reduced model simulations. It is apparent that the phase reduction yields both qualitatively and quantitatively similar results. Since the phase response curve for the Hodgkin-Huxley system is of relatively small magnitude, we have used an α value of 1, rather than 0.1, to more clearly illustrate the dynamic behavior of the phase distribution.




Figure 3.2: Simulation results, $\rho(\theta, t)$, for Hindmarsh-Rose (a), (c) and Hodgkin-Huxley (b),(d) (resp., conductance-based model, phase-reduced model).

One notices in these figures a behavior that can be described as “breathing”, i.e. there is an oscillation (with period of approximately 200 ms for the Hindmarsh-

Rose model and approximately 15 ms for the Hodgkin-Huxley model) about a mean distribution. We found numerically that such oscillations persist for long times of the order of 10^6 msec. It is worth noting that this oscillatory behavior is exhibited in both the phase-reduced and conductance-based model simulations, suggesting that it may not be an artifact of the numerics, but rather something inherent in the dynamics which is not predicted by the theoretical framework. For the purposes of our subsequent analysis, we view these oscillations as a secondary effect, to be pursued in future work. To illustrate the average shape of the $\rho(\theta, t)$ distribution curves, we average each simulation over the last 25% of the integration time interval, which filters the oscillatory effects of the “breathing” behavior.




Figure 3.3: Comparison of theoretical distribution estimate with averaged simulation data for both conductance-based and phase-reduced models for Hindmarsh-Rose (a) and Hodgkin-Huxley (b). Solid lines represent full conductance-based model simulation data, dashed lines represent phase-reduced model simulation data, and dotted lines represent the theoretical estimates.

Figure 3.3(a) shows that theoretical predictions match the numerical data very closely for the Hindmarsh-Rose system. The results for the Hodgkin-Huxley system are qualitatively similar, as illustrated in Figure 3.3(b), but display a slight mismatch of the θ location of the distribution peak.

3.5.2 Parametric study

The computationally efficient phase-reduced model allows for a more complete mapping of the population response in parameter space. As a characterization of the population response, we consider the magnitude, ρ_{max} , and location, θ_{max} , of the peak of the $t \rightarrow \infty$ averaged probability distribution. As in Section 3.5.1, we consider populations of 1000 neurons. The Hindmarsh-Rose neurons are simulated for 1000 msec (natural period ≈ 312 msec). The Hodgkin-Huxley neurons are simulated for 100 msec (natural period ≈ 14.6 msec).

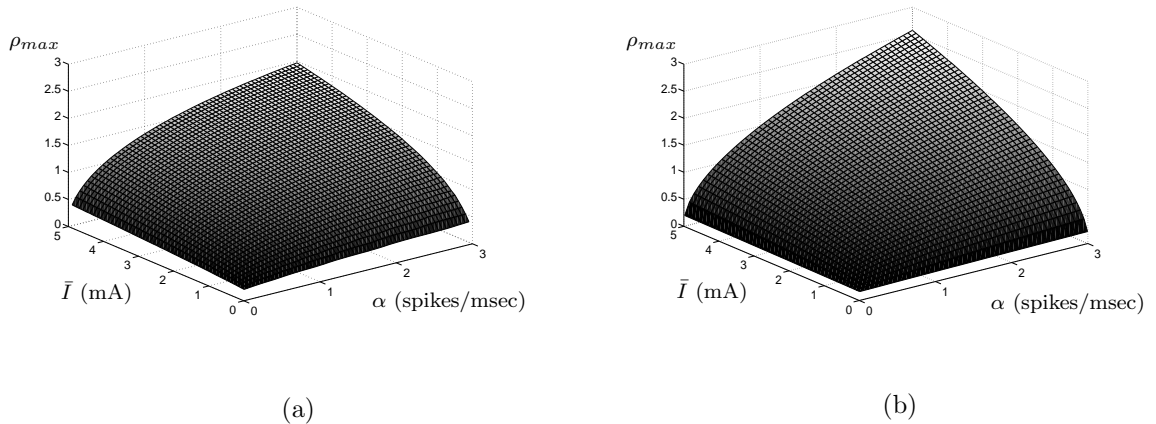


Figure 3.4: Numerical (a) and theoretical (b) ρ_{max} over (α, \bar{I}) parameter space for Hindmarsh-Rose model.

Figure 3.4 compares the phase model results to the theoretical prediction using (3.33) for the Hindmarsh-Rose system over $0 < \alpha \leq 3$ spikes/msec and $0 < \bar{I} \leq 5$ mA, for the ρ_{max} characteristic. We find qualitative and reasonable quantitative agreement. For simulations of the full conductance-based model, simulations of the phase model, and theoretical predictions, the location θ_{max} of the probability distribution function peak ρ_{max} stays at $\theta_{max} = 0$. The plots in Figure 3.5 show our results over the same ranges of α and \bar{I} for the Hodgkin-Huxley system. Figure 3.5(a) illustrates the more complex parameter dependence of ρ_{max} . While the theory predicts that θ_{max} will be independent of α and \bar{I} , Figure 3.5(b) shows that it is weakly dependent on them. Since our theory is based on the ratio ω/α being order ϵ , we have plotted the theoretical results over a reduced region of α and \bar{I} parameter space, as shown in Figure 3.5(d) to be compared with the simulation results reproduced over this reduced range in Figure 3.5(c). We see that the theoretical results qualitatively capture the trend of the data, but start to differ quantitatively as α and \bar{I} increase. Importantly, for both the Hindmarsh-Rose and Hodgkin-Huxley models, we have shown that the degree of synchronization, as measured by the peak of the probability distribution function of the phase, increases with both the size and frequency of the stimuli, and can become quite substantial.

3.6 Discussion

3.6.1 Relationship between partial phase synchrony and spike synchrony

The importance of studying partial phase synchronization lies not in its relationship to spike synchrony, but in its characterization of the dynamical state of the population, which governs the response of the population to future stimulus.

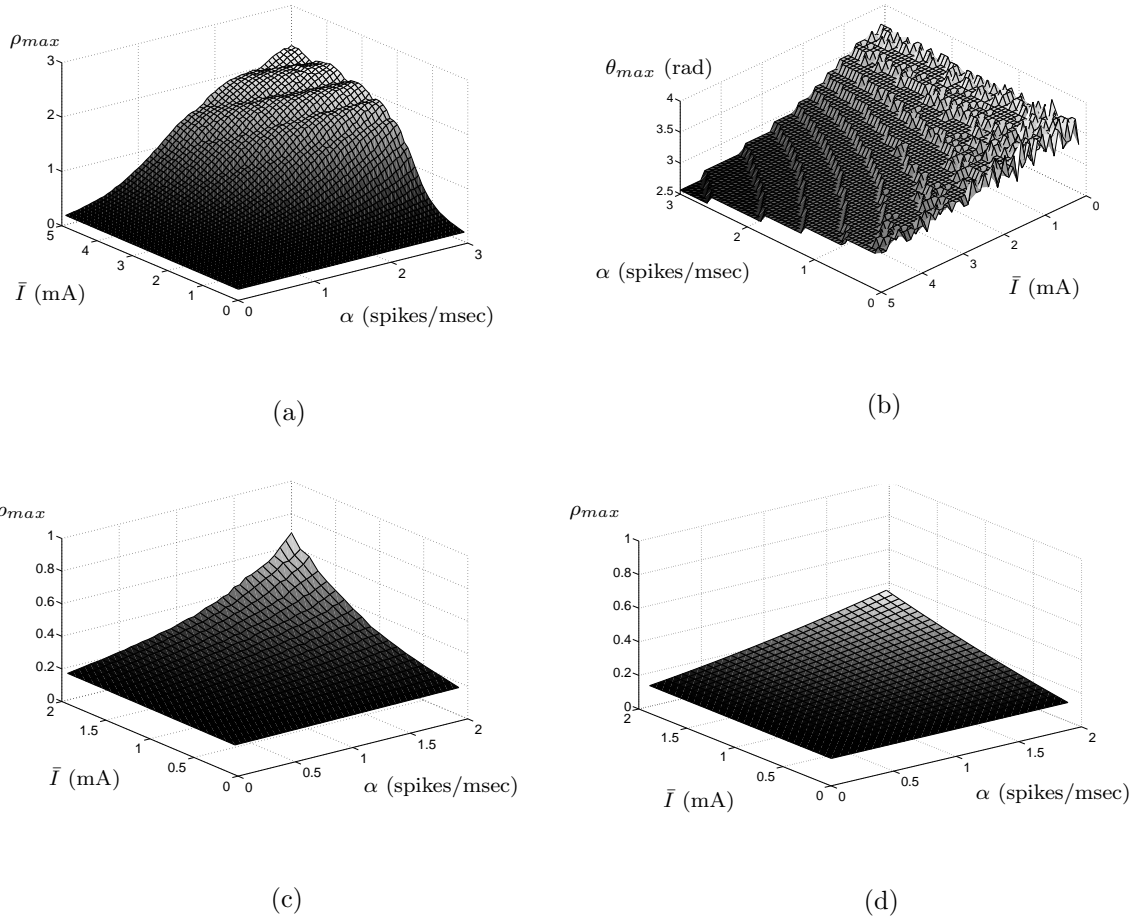


Figure 3.5: Hodgkin-Huxley population-level response characteristics. Numerical ρ_{max} over (α, \bar{I}) parameter space (a). Numerical θ_{max} over (α, \bar{I}) parameter space (b). Comparison of numerical results (c) to theoretical predictions (d) over the range of (α, \bar{I}) for which the assumption $\omega/\alpha = \mathcal{O}(1)$ is reasonable.

Before addressing the practical implications of partial phase synchrony, we will illustrate some important points regarding the relationship of phase synchrony and spike synchrony by discussing several simple examples. We will show that, in the general case of a population of stimulated neurons, the degree of partial phase synchronization is only weakly related to spike synchrony.

In the following scenarios, we consider a population of N identical uncoupled phase oscillators each obeying the following ODE:

$$\dot{\theta}_i = \omega + Z(\theta)I_i(t)/c.$$

In the *absence* of input, i.e. $I_i(t) \equiv 0$, there is a *strong* relationship between phase synchrony and spike synchrony for populations of uncoupled identical oscillators. If all oscillators have identical phase, it is obvious that they will cross the $0/2\pi$ spiking threshold at identical time. At the opposite extreme, if the distribution of

the population is uniform, it follows that the spike times will be uniform in time (and thus desynchronized).

In the presence of input, coupling, or large distributions of natural frequency, this strong relationship no longer holds. Since we have restricted our attention to populations of identical uncoupled neurons in this chapter, we will develop several simple examples that illustrate how this relationship breaks down in the presence of input stimulus.

Our primary concern is showing the mathematical relationship between phase synchrony and spike synchrony, so we will assume that there are no constraints on the inputs we can use, and the oscillators themselves are governed by phase response curves that allow them to be controlled to achieve any prescribed dynamics:

$$\dot{\theta}_i = f(\theta_i).$$

This assumption is not true in general for real neurons (since the phase response curves are often nonlinear and non-invertible). The question of whether population of neurons with a given phase response curve can be driven to achieve specific dynamics when stimulated by inputs from a restricted set is a control theoretic problem and an aim of our future research. For the time being, we will follow our controllability assumption, which will allow us to develop examples that are simple and easy to verify.

In order to quantify partial phase synchronization, we introduce the Kuramoto order parameter [59]:

$$r(t)e^{\sqrt{-1}\psi(t)} = \frac{1}{N} \sum_{i=0}^N e^{\sqrt{-1}\theta_i(t)}. \quad (3.38)$$

In particular, we consider the magnitude of the order parameter, $|r(t)|$, which ranges from zero for a uniform phase distribution to one for a completely synchronized phase distribution (a Dirac delta function).

Example 1: Phase synchrony without spike synchrony

Consider a population of N (uncoupled) phase oscillators, labeled $\theta_1, \dots, \theta_N$, evolving according to the following ODE:

$$\dot{\theta}_i = f(\theta_i) = \begin{cases} \frac{(\pi-\Delta)\omega N}{\frac{\Delta\omega N}{\pi(N-2)}} & \text{if } \theta_i \in \{[0, \pi - \Delta) \cup [\pi + \Delta, 2\pi)\} \\ \text{otherwise} & \end{cases} \quad (3.39)$$

with the following initial conditions:

$$\theta_i(0) = \begin{cases} 0 & \text{for } i = 1 \\ \pi - \Delta + \frac{2\Delta(i-2)}{N-2} & \text{for } i \neq 1. \end{cases} \quad (3.40)$$

This population will have uniform interspike intervals, i.e. no spike synchrony whatsoever, while possessing an order parameter $|r| \rightarrow 1$ as $\Delta \rightarrow 0$ and $N \rightarrow \infty$, which indicates an arbitrarily high degree of phase synchronization.

What has been done here is to define a neighborhood around $\theta = \pi$, with interval measure equal to 2Δ , where the neurons move very slowly. Outside this neighborhood the neurons advance very quickly. We have set up the initial conditions such that most of the oscillators (all but two) are in the small region where they move slowly. The one that starts at zero quickly moves to the starting edge of the π -neighborhood phase interval. During the same time interval, the last oscillator, θ_N , advances quickly from the end of the π -neighborhood phase interval around the circle to the $0/2\pi$ spiking threshold. Also during this time interval, all the other neurons slowly advance in order along the π -neighborhood phase interval. After the interspike time interval, the neurons have rotated positions, but the overall population distribution is identical to where it started from, so after relabeling the oscillators, we are back to the starting configuration. Now if we make Δ small, we increase the level of phase synchronization, but the firing continues to be completely uniform (desynchronized). In fact, as Δ approaches zero, all the neurons except one will be arbitrarily close to π , so $|r| \rightarrow \frac{N-1}{N} = 1 - \frac{1}{N}$. Furthermore, as we increase the number of neurons, $|r| \rightarrow 1$ asymptotically. The value of $|r|$ will, of course, never reach 1, but it can be made arbitrarily close.

We illustrate this with a numerical example consisting of a set of $N = 6$ phase oscillators with natural frequency $\omega = 1$. We set $\Delta = 0.5$ so that the trajectories will not be overly crowded in the π -neighborhood and the plotted results can be easily understood. Figure 3.6 (a) shows the trajectories generated by (3.39) along with the magnitude of the order parameter $|r(t)|$. The interspike time intervals here are uniformly equal to ratio of the natural period T to the number of neurons N , so there is no spike synchrony whatsoever. Yet the magnitude of the order parameter is at all times above 0.5. This value can be driven arbitrarily close to one by increasing the number of neurons and decreasing Δ . We note that all the phase trajectories are in solid black, because it is not important to be able to differentiate between them. Any set of trajectories following such paths, regardless of which neuron is on which trajectory at a given time, will generate the same population-level results. This convention will be used in the next example, as well.

Example 2: Spike synchrony with minimal phase synchrony

To further simplify the following example, we will assume that we can prescribe the phase trajectory for the i^{th} oscillator directly, rather than via an ODE, according to the following equation:

$$\theta_i(t) = \begin{cases} \frac{\pi+2\pi(i-1)}{n\epsilon} t \pmod{T} & \text{if } t \pmod{T} \in [0, \epsilon) \\ \frac{\pi+2\pi(i-1)}{n} & \text{if } t \pmod{T} \in [\epsilon, T - \epsilon) \\ \frac{\pi+2\pi(i-1)}{n} + \frac{\pi+2\pi(n-i)}{n\epsilon} (t \pmod{T} - (T - \epsilon)) & \text{if } t \pmod{T} \in [T - \epsilon, T) \end{cases} \quad (3.41)$$

The population of oscillators evolving according to (3.41) has complete spike synchrony, but the trajectories have been configured such that the distribution will

be nearly uniform except near the spike times, where the distribution will collapse toward a Dirac delta distribution, then expand again to be nearly uniform after the spike. Thus we have shown that, even in the presence of complete spike synchrony, the phase synchrony can be confined to intervals of arbitrarily small measure. To quantify this, we note that the integral of the order parameter magnitude

$$R(t) \equiv \int_0^t |r(\tilde{t})| d\tilde{t}$$

can be made arbitrarily small by letting $\epsilon \rightarrow 0$. This means that the phase synchrony of this population is zero except on a set of arbitrarily small (but still positive) measure. We conclude that, although spike synchrony must be accompanied by phase synchrony at the instants of the spikes, we cannot say more about their relationship without considering controllability of the oscillators with respect to possible trajectories.

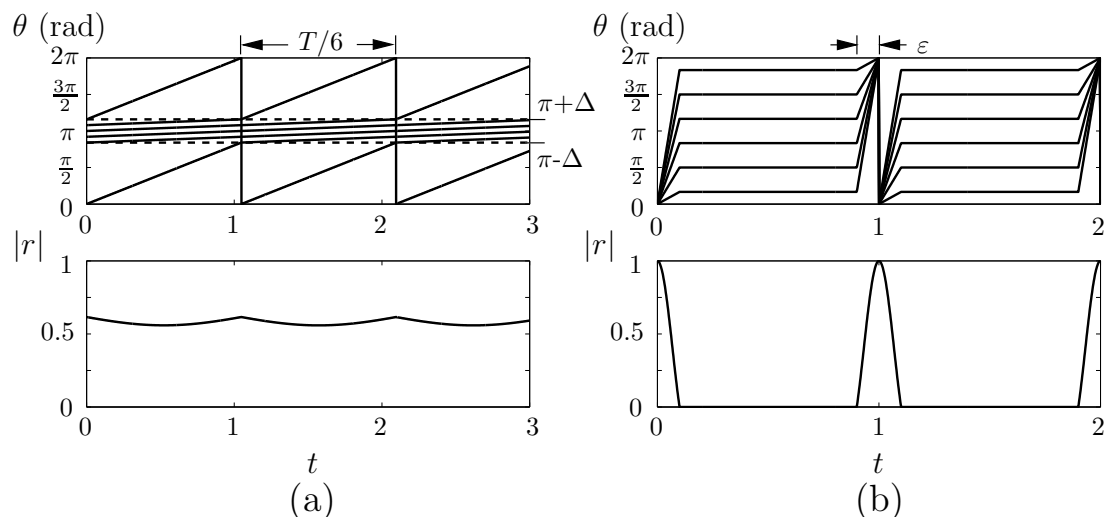


Figure 3.6: Illustrations showing that partial phase synchrony and spike synchrony need not be strongly related. Plot (a) shows the scenario from Example 1 where there is a high degree of phase synchronization but no spike synchronization. Plot (b) shows the scenario from Example 2 where there is complete spike synchronization but only small time intervals of phase synchronization. The bottom figures show a quantification of phase synchrony, as represented by the magnitude of the Kuramoto order parameter. The results can be sharpened by increasing the number of neurons, and reducing the parameters Δ and ϵ .

We illustrate this point with a small population of 6 oscillators with natural period $T = 1$. Figure 3.6(b) shows the trajectories generated by (3.41) as well as the associated $|r(t)|$. For clarity, we choose $\epsilon = 0.1$. The magnitude of the order parameter is zero everywhere except in an ϵ -neighborhood of the spike times. All of the oscillators spike together, which means that there is complete spiking synchrony. By reducing ϵ , we can confine the periodic phase synchrony to arbitrarily small time intervals centered at the spike time.

Poisson inputs

For a population of neurons subjected to independent Poisson inputs as described in Section 3.2, the independence of the inputs implies that the spike times will be uncorrelated. It is intuitively clear from equation (3.38) and Figure 3.3, the magnitude of the order parameter r can be non-zero when there are such inputs.

From a mathematical standpoint, we can illustrate this by considering the continuum limit as $N \rightarrow \infty$ and asymptotically as $t \rightarrow \infty$,

$$|r(t)| \rightarrow \left| \int_0^{2\pi} \rho_s(\theta) e^{\sqrt{-1}\theta} d\theta \right| \sim \frac{\pi}{2} (\rho_{max} - \rho_{min}). \quad (3.42)$$

The last relation follows from the fact that the absolute value of the integral is equal to the magnitude of the first Fourier component: when ρ_s is unimodal and not too sharply peaked, we expect that

$$\rho_s(\theta) \sim \frac{1}{2\pi} + \frac{1}{2} (\rho_{max} - \rho_{min}) \sin(\theta - \theta_0)$$

for some θ_0 , from which the relation follows. The values of ρ_{max} are precisely the results computed in Section 3.5.2 and displayed in Figures 3.4 and 3.5. We note in this case that $\rho_{min} \approx \frac{1}{\pi} - \rho_{max} \leq \frac{1}{2\pi}$. The neurons subjected to Poisson inputs thus have partial phase synchronization, with an order parameter increasing with both the magnitude of the stimulus and the mean spike frequency, although there is no spike synchronization. In the next section, we will discuss how a population that is partially phase-synchronized reacts differently to stimulus than if its phases were uniformly distributed. This will illustrate the importance of considering phase synchronization in addition to spike synchronization when studying the dynamical properties of populations of oscillatory neurons.

3.6.2 Effects of partial phase synchronization

The degree of partial phase synchronization affects the manner in which the population responds to stimulus, as can most clearly be seen by constructing a histogram of the first spike times of each neuron after receiving a step stimulus of 1 mA. For both the Hindmarsh-Rose and Hodgkin-Huxley systems, populations of 1000 phase models are initialized with uniform phase distribution. We expose the population to the step input and track when the next zero-crossing (firing) occurs. Figure 3.7(a) shows the first spike time histogram (FSTH) for Hindmarsh-Rose and Figure 3.7(b) shows the FSTH for Hodgkin-Huxley. To illustrate the effect of partial phase synchronization, we repeat the above simulations starting from the partially synchronized distribution shown in Figures 3.3(a) and (b). The results are shown below the FSTH plots for the uniform cases, in Figures 3.7(c) and (d). We see that for the Hindmarsh-Rose system, the partial phase synchronization tends to flatten the FSTH somewhat. For the Hodgkin-Huxley system, using the curve from the $\alpha = 1$ case for consistency, we see a very pronounced change in shape of

the FSTH. All four FSTH plots are tabulated over 1000 realizations, in order to show the character of the distributions independent of the randomness inherent in any individual realization. These results show that while partial phase synchrony does not imply firing synchrony, it can play an important role in determining how the population responds to a common stimulus.

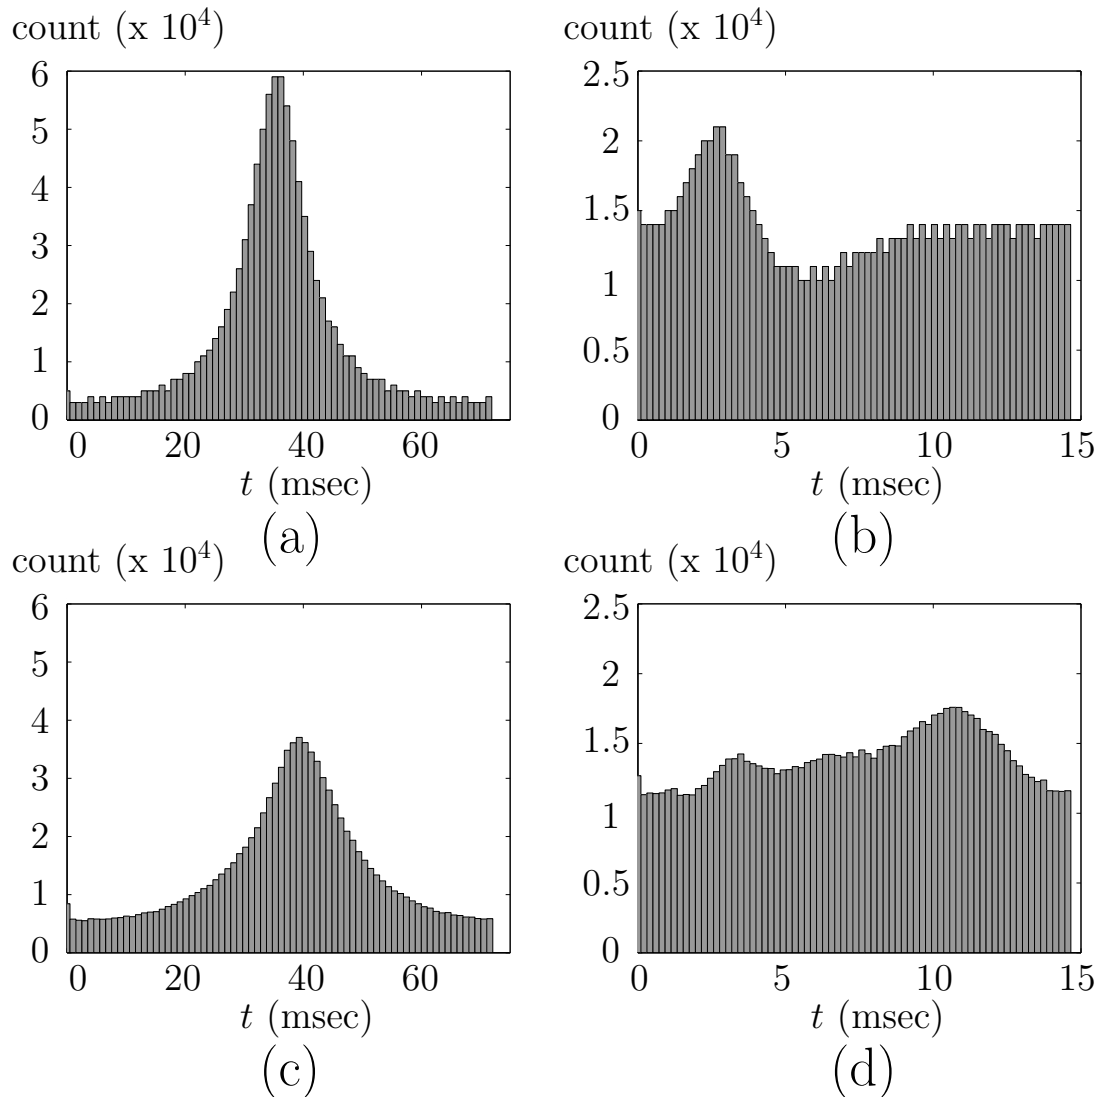


Figure 3.7: First Spike Time Histograms for Hindmarsh-Rose (a), (c) and Hodgkin-Huxley (b),(d). Plots (a) and (b) begin from a uniform distribution. Plots (c) and (d) begin from the partially synchronized distributions shown in Figure 3.3. This analysis tracks the time of the first spike after the unit step-function stimulus is turned on. 1000 realizations are used.

We have shown that populations of identical uncoupled neurons exhibit partial phase synchronization when stimulated with independent unidirectional current spikes with interspike time intervals drawn from a Poisson distribution. We characterized this partial synchronization by the phase distribution for the population,

using analytical approximations and numerical simulations of phase-reduced models and conductance-based models of typical Type I (Hindmarsh-Rose) and Type II (Hodgkin-Huxley) neurons. The results from the different approaches agree well with each other. We found that the degree of partial phase synchronization, as measured by the peak of the phase distribution, increases with both the size and frequency of the stimuli, and can become quite substantial.

We have shown that partial phase synchronization, a distinct phenomenon from spike synchronization, is an important consideration when using phase reduced models to infer dynamical characteristics of spiking neurons. Our results show that neural populations subjected to background activity from other neurons do not have a uniform distribution of phases, as is sometimes assumed in simulation studies. We show that such non-uniformity leads to different population-level response to other stimuli, suggesting that noisy inputs must be carefully incorporated into simulation studies in order to obtain biologically realistic results.

Chapter 4

Population Dynamics: Parametric Resonance

In this chapter, we discuss existence, stability, and symmetry of solutions for networks of parametrically-forced oscillators. In particular, we consider a nonlinear oscillator model with strong 2:1 resonance via parametric excitation. For uncoupled systems, the 2:1 resonance property results in sets of solutions that we classify using a combinatorial approach. The symmetry properties for solution sets are presented as are the group operators that generate the isotropy subgroups. We then impose weak coupling and prove that solutions from the uncoupled case persist for small coupling by using an appropriate Poincaré map and the Implicit Function Theorem. Solution bifurcations are investigated as a function of coupling strength and forcing frequency using numerical continuation techniques. We find that the characteristics of the single oscillator system are transferred to the network under weak coupling. We explore interesting dynamics that emerge with larger coupling strength, including anti-synchronized chaos and unsynchronized chaos. A classification for the symmetry-breaking that occurs due to weak coupling is presented for a simple example network.

The study presented here is an extension of results investigated in [70], and is organized as follows. We present introductory background material in Section 4.1. In Section 4.2 we consider the dynamics of a specific single parametrically forced oscillator and show that the origin is a stable equilibrium in the absence of forcing, and that for certain ranges of forcing frequency there exists a stable periodic orbit. Then, in Section 4.3.1, we consider a set of N uncoupled parametrically forced oscillators, identifying different periodic states for such systems. In Section 4.3.2, we study the symmetry properties of these models.

Section 4.4.1 shows that provided the periodic orbits for the uncoupled system are hyperbolic, there will be periodic orbits for the weakly coupled system close to the periodic states identified for the uncoupled system. In Section 4.4.2, we investigate the $N = 2$ case in the presence of bidirectional coupling. We explore the bifurcation structure of solutions with respect to both coupling strength and forcing frequency. Interesting dynamics, such as anti-synchronized chaos, are found

when the magnitude of the coupling strength is increased. Section 4.4.3 provides a detailed analysis of the $N = 3$ case in the presence of bidirectional coupling for both the ring topology and the line topology. We perform a set of numerical bifurcation analyses to show how the magnitude and stability of the periodic oscillations vary with forcing frequency. For the $N = 3$ case, we classify how the symmetry present in the uncoupled solution is either preserved or destroyed due to coupling for a simple example system for both the ring and the line topologies. We then relate the two topologies using an intermediate morphing topology (a ring with a single variable connection) and discuss how symmetry-related solutions organize differently depending on the network structure. We describe how these results can be generalized to N coupled parametrically forced oscillators in Section 4.4.4. Finally, we give conclusions in Section 4.5.

4.1 Introduction and background

The scientific study of coupled oscillators started with Christian Huygens' observations in the seventeenth century of mutual synchronization of pendulum clocks connected by a beam [52, 11]. More recently, it has been recognized that mutual synchronization of coupled oscillators - the adjustment of rhythms of oscillating objects due to their weak interactions - occurs in many biological systems, including neurons during epileptic seizures [103] and pacemaker cells in the human heart [65]. Coupled oscillators have also been studied in detail for technological systems, such as arrays of lasers and superconducting Josephson junctions: see [96], [79], and [95], a recent popular book on the topic, for many biological and technological examples of synchronization for coupled oscillators.

We classify as *autonomous oscillators* those for which the stable oscillations occur for an autonomous dynamical system, that is one for which there are no explicit time-dependent terms in the evolution equation. For example, the oscillations might arise through a Hopf bifurcation. In the limit of weak coupling, it is possible to reduce the dynamics of coupled autonomous oscillators to a phase model, with a single variable describing the phase of each oscillator with respect to some reference state (see, e.g., [45, 59, 109, 18]). This typically leads to models for which the dynamics depend only on the phase differences between different oscillators.

It is possible to show that several types of phase-locked solutions, for which the phase of all oscillators increases at the same constant rate, are guaranteed to exist in the weak coupling limit for *any* generic coupling function when the coupling topology has appropriate symmetry properties [1, 75, 26, 17]; for the case of identical all-to-all coupling for N oscillators, these are (i) in phase solution: all N oscillators have the same phase; (ii) two-block solutions: there are two blocks of oscillators, one in which p oscillators share the same phase, and one in which $N - p$ oscillators share the same phase; (iii) rotating block solutions: for $N = mk$, there are m blocks with k oscillators in each block sharing the same phase, with neighboring blocks differing in phase by $2\pi/m$; (iv) double rotating block solutions:

for $N = m(k_1 + k_2)$, there are two rotating block solutions, one with m blocks with k_1 oscillators in each block sharing the same phase and with neighboring blocks differing in phase by $2\pi/m$, another with m blocks with k_2 oscillators in each block sharing the same phase and with neighboring blocks differing in phase by $2\pi/m$, where there is a phase difference $0 < \phi < 2\pi/m$ between a block with k_1 oscillators and the closest phase-advanced block with k_2 oscillators.

On the other hand, we classify as *non-autonomous oscillators* those for which the stable oscillations only occur for a non-autonomous dynamical system, that is one for which there are explicit time-dependent terms such as time-periodic forcing. We will focus on parametrically forced oscillators, which are non-autonomous oscillators for which the forcing enters as a time-varying system parameter. Coupled parametrically forced oscillators arise in MEMS [20, 60, 16] and other application areas [101, 9, 8, 41], but have not received as much theoretical research attention as coupled autonomous oscillator systems. The work presented in this chapter represents the first steps in developing a comprehensive theory of the dynamics of general weakly coupled non-autonomous oscillators, in the spirit of the theory of general weakly coupled autonomous oscillators described in [1, 17]. We hope that such a theory will ultimately lead to novel sensing mechanisms using MEMS devices; for simplicity, here we will consider a model system which represents only a caricature of such devices.

The non-autonomous oscillator models presented in this chapter could also represent oscillatory neurons driven by periodic inputs. The current EDBS technology uses periodic waveforms that, if tuned correctly, could realize similar phenomena *in vivo*.

Specifically, in this chapter we describe interesting synchronization phenomena that are possible for coupled parametrically forced oscillators. For example, consider two *uncoupled* oscillators whose response is at half the frequency of the driving voltage, as is common for MEMS devices [104]. Both oscillators could identically lock to the forcing, or they could lock one forcing period apart - both situations are allowable due to a discrete time-translation symmetry for the problem. We will show that different combinations of these states will persist if the oscillators are weakly *coupled*, with stability inherited from the stability properties of the periodic orbits which exist for the uncoupled system.

4.2 A parametrically forced oscillator

Consider the equation for a damped, parametrically forced oscillator

$$\ddot{x} + b\dot{x} + x + x^3 = xF \cos(\omega_f t). \quad (4.1)$$

which was first presented in Section 2.5. Here the term $b\dot{x}$ represents damping (we assume $b > 0$), the term $x + x^3$ represents a nonlinear restoring force, and the term $xF \cos(\omega_f t)$ represents parametric excitation which can be viewed as a time-periodic modulation of the linear part of the restoring force. For this system,

if $F = 0$ then $x \rightarrow 0$ as $t \rightarrow \infty$, as follows. Letting

$$V(x, \dot{x}) = \frac{1}{2}x^2 + \frac{1}{4}x^4 + \frac{1}{2}\dot{x}^2, \quad (4.2)$$

we find that $\frac{dV}{dt} = -b\dot{x}^2 \leq 0$, with equality only if $\dot{x} = 0$. The only point in phase space which starts in the set of points for which $V(x, \dot{x}) = 0$ and remains in this set for all time is $(x, \dot{x}) = (0, 0)$; by the LaSalle Invariance Principle [107], all trajectories thus approach this point as $t \rightarrow \infty$.

For appropriate F and ω_f , the system has a periodic response. Indeed, treating ω_f as a bifurcation parameter for fixed F and b , we obtain the bifurcation diagram shown in Fig. 4.1. (This and other numerical bifurcation analyses were done using AUTO [33] or XPP [36].) The 0 solution is characterized by $x = \dot{x} = 0$ for all time; it exists for all ω_f , being unstable near $\omega_f = 2$ and stable otherwise for the range shown. It loses stability in a bifurcation to a periodic orbit, with the periodic orbit branch turning around in a saddle-node bifurcation so that there is a region of bistability between the periodic orbit and the 0 solution. Such a bifurcation structure is common for MEMS devices, see e.g. [111, 82].

Figure 4.2 indicates the types of dynamics which occur in different parameter ranges, with the “parabola” corresponding to the loss of stability of the 0 state, and the “straight line” corresponding to the saddle-node bifurcation of the periodic orbit branch. As shown in Fig. 4.3, the response of this periodic orbit is at half the frequency of the forcing, as is common for parametrically forced oscillators [73]. We note that an equally valid periodic orbit for this forcing is shown in Fig. 4.4, which is shifted by one period of the forcing from the solution shown in Fig. 4.3; clearly these solutions are related by a time-translation symmetry. Indeed, our example oscillator governed by Equation (4.1) has a very convenient form of discrete time translation symmetry which is also related to phase space symmetry: translating by one forcing period is equivalent to multiplying both states at time t by -1 , so $(x(t+T), \dot{x}(t+T)) = (-x(t), -\dot{x}(t))$. These symmetry-related solutions will be crucial for understanding the different possible solutions when such oscillators are weakly coupled. We will find it convenient to distinguish the solutions shown in Figs. 4.3 and 4.4 by referring to one of them as the A solution and the other as the B solution. We introduce a new solution label $z \in \{A, B, 0\}$. The A and B solutions have the same $\max(x)$ value, and the periodic orbit branch in Fig. 4.1 corresponds to both (symmetry-related) solutions. The $z = 0$ label corresponds to the previously defined 0 solution, $x = \dot{x} = 0$. This notation will greatly simplify the following presentation.

4.3 Uncoupled parametrically forced oscillators

4.3.1 Combinatorial considerations

Suppose $b = 0.2, F = 0.5, \omega_f = 2$, so that when $c = 0$ each oscillator could be in a stable periodic state given by the A and B solutions, or it could be in the




Figure 4.1: Bifurcation diagram for fixed $b = 0.2$ and $F = 0.5$. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.




Figure 4.2: Existence and stability of solutions for different regions of parameter space for $b = 0.2$, with bifurcations sets shown as lines.

unstable 0 solution. There will be 3^N distinct periodic orbits for $c = 0$: oscillator 1 could be in A , B , or 0 , oscillator 2 could be in A , B , or 0 , etc. Of these solutions, a total of

$$\frac{N!}{p_A!p_B!p_0!} \tag{4.3}$$

solutions will have p_z oscillators in each state $z = A, B$, or 0 , where $p_A + p_B + p_0 = N$. This follows from the following combinatorial argument. Suppose we make a list of N symbols such that the i^{th} symbol is A, B , or 0 according to whether the i^{th} oscillator is in the A, B , or 0 state, respectively. In the N slots, there are

$$\binom{N}{p_A} = \frac{N!}{p_A!(N - p_A)!}$$




Figure 4.3: Stable periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2$. The response is at half the frequency of the forcing. We will refer to this as the A solution.

different ways to put the symbol A in p_A of the slots. Of the remaining $(N - p_A)$ slots, there are

$$\binom{N - p_A}{p_B} = \frac{(N - p_A)!}{p_B!(N - p_A - p_B)!}$$

different ways to put the symbol B in p_B of the slots. The remaining $(N - p_A - p_B)$ slots will have the symbol 0. The product of these is

$$\frac{N!}{p_A!(N - p_A)!} \times \frac{(N - p_A)!}{p_B!(N - p_A - p_B)!} = \frac{N!}{p_A!p_B!p_0!},$$

as in Eq. (4.3). As an illustration, suppose $N = 4, p_A = 2, p_B = 1$, and $p_0 = 1$. The different possible lists of symbols are

$$\begin{array}{cccccc} AAB0 & AA0B & ABA0 & A0AB & AB0A & A0BA \\ BAA0 & 0AAB & BA0A & 0ABA & B0AA & 0BAA, \end{array}$$

giving a total of

$$\frac{4!}{2!1!1!} = 12$$

possibilities. This combinatorial argument readily generalizes to situations in which more than three states are possible for each oscillator, for example when $b = 0.2, F = 0.5$, and $\omega_f = 2.25$.

4.3.2 Symmetry considerations

Periodic solution trajectories of an uncoupled N oscillator system are described by the population labeling variable \mathbf{z} , a length- N set of single oscillator solution labels z_i for $i = 1, \dots, N$. Since the order of the entries is irrelevant for an uncoupled




Figure 4.4: Stable periodic orbit for $b = 0.2$, $F = 0.5$, $\omega_f = 2$. This is shifted by one period of the forcing from the solution shown in Fig. 4.3. We will refer to this as the B solution.

system, we can group any state string into successive substrings of A 's, B 's, and 0 's. For example, a $00A0AB0AB$ state is equivalent to an $AAABB0000$ state. In any state string, there are N_A A 's, N_B B 's, and N_0 0 's, where $N_A + N_B + N_0 = N$.

It is possible to classify solutions of dynamical systems with symmetry according to their isotropy subgroup, which is the subgroup of all symmetry operations which leave the solution unchanged. It is convenient to order the solution into the substrings of A 's, B 's, and 0 's as described above. Then, we see that the isotropy subgroup of the system is

$$S_{(N_A+N_B)} \times S_{N_0}, \quad (4.4)$$

where

$$S_{(N_A+N_B)} = \langle \kappa_1, \kappa_2 \rangle, \quad S_{N_0} = \langle \kappa_3, \kappa_4 \rangle,$$

$$\begin{aligned} \kappa_1 & \cdot (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\ & \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\ & \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N) \\ \rightarrow & (z_2, z_1, \dots, z_{N_A-1}, z_{N_A}, \\ & \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\ & \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N), \end{aligned}$$

$$\begin{aligned} \kappa_2 & \cdot (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\ & \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\ & \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N) \\ \rightarrow & (-z_{N_A+N_B}, z_1, z_2, \dots, z_{N_A-1}, \\ & \quad -z_{N_A}, z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, \\ & \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N), \end{aligned}$$

$$\begin{aligned}
\kappa_3 & \cdot (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\
& \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\
& \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N) \\
& \rightarrow (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\
& \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\
& \quad z_{N_A+N_B+2}, z_{N_A+N_B+1} \dots, z_{N-1}, z_N), \\
\kappa_4 & \cdot (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\
& \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\
& \quad z_{N_A+N_B+1}, z_{N_A+N_B+2}, \dots, z_{N-1}, z_N) \\
& \rightarrow (z_1, z_2, \dots, z_{N_A-1}, z_{N_A}, \\
& \quad z_{N_A+1}, z_{N_A+2}, \dots, z_{N_A+N_B-1}, z_{N_A+N_B}, \\
& \quad z_N, z_{N_A+N_B+1}, z_{N_A+N_B+2} \dots, z_{N-1}).
\end{aligned}$$

In words, κ_1 swaps the first two A solutions, κ_2 is a cyclic permutation of the A and B solutions with minus signs appropriately included, κ_3 swaps the first two 0 solutions, and κ_4 is a cyclic permutation of the 0 solutions. The intuition behind this result is that the solution is unchanged by permuting the first $N_A + N_B$ states (the A and B states) however one desires, provided the A and B states are transformed into each other through appropriate multiplications by -1 . The N_0 remaining 0 solutions can be permuted however desired without changing the solution.

Let us again consider our example system $AAABB0000$. This solution has isotropy subgroup $S_5 \times S_4$, being unchanged by multiplication by any product of $\kappa_1, \kappa_2, \kappa_3, \kappa_4$. For example, it is unchanged under

$$\begin{aligned}
\kappa_3 \kappa_2^2 \kappa_1 \kappa_2 \kappa_1 \kappa_2^3 & \cdot (z_1, z_2, z_3, z_4, z_5, z_6, z_7, z_8, z_9) \\
& \rightarrow (-z_5, z_1, -z_4, -z_2, -z_3, z_7, z_6, z_8, z_9).
\end{aligned}$$

This discussion readily generalizes to conditions where the individual oscillators have more than three possible states by including additional κ operators to swap and permute the additional possible states.

4.4 Coupled parametrically forced oscillators

4.4.1 Periodic orbits for weakly coupled oscillators

When $c \neq 0$ but is small, we expect solutions analogous to those for $c = 0$ to exist, as follows. We write our system as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) + c\mathbf{g}(\mathbf{x}), \tag{4.5}$$

where

$$\mathbf{x} = (x_1, \dot{x}_1, \dots, x_N, \dot{x}_N)^T, \tag{4.6}$$

and $\mathbf{f}(\mathbf{x}, t)$ captures the terms which are independent of c , and $\mathbf{g}(\mathbf{x})$ captures the coupling terms. Let $T = 2\pi/\omega_f$ be the period of the forcing. We define P_c to be the time- $2T$ map, that is the map which takes an initial condition (the state at $t = 0$) to the state obtained by evolving for a time equal to twice the period of the forcing. Now, let

$$h(\mathbf{x}, c) = P_c(\mathbf{x}) - \mathbf{x}, \quad (4.7)$$

and let \mathbf{q}_0 be a point on one of the periodic solutions of the uncoupled problem, for example the AA solution in the $N = 2$ case. We see that

$$h(\mathbf{q}_0, 0) = 0. \quad (4.8)$$

The Implicit Function Theorem (see, e.g. the appendix of [31]) then implies that, provided the matrix $D_{\mathbf{x}}h(\mathbf{q}_0, 0)$ is invertible, there is a unique solution $\mathbf{q}(c)$ close to \mathbf{q}_0 , for any sufficiently small c , such that $h(\mathbf{q}(c), c) = 0$. This implies that

$$P_c(\mathbf{q}(c)) = \mathbf{q}(c), \quad (4.9)$$

that is, $\mathbf{q}(c) \approx \mathbf{q}_0$ is a fixed point of the time- $2T$ map, which means that it is a point on a periodic orbit with period $2T$ which is close to a periodic orbit of the uncoupled system. (A related argument is used to prove part (ii) of Theorem 4.1.1 in [46].)

It is instructive to consider an alternative, but equivalent argument. We know that \mathbf{q}_0 is a fixed point of P_0 , that is,

$$P_0(\mathbf{q}_0) = \mathbf{q}_0. \quad (4.10)$$

We will determine, to leading order in c , the condition which must be met for $\mathbf{q}(c)$ to be a fixed point of P_c . Consider the asymptotic expansions

$$\mathbf{q}(c) = \mathbf{q}_0 + c\mathbf{q}_1 + \dots \quad (4.11)$$

$$P_c(\mathbf{x}) = P_0(\mathbf{x}) + cp_1(\mathbf{x}) + \dots. \quad (4.12)$$

Setting $P_c(\mathbf{q}(c)) = \mathbf{q}(c)$, we obtain

$$\begin{aligned} \mathbf{q}_0 + c\mathbf{q}_1 + \dots &= P_c(\mathbf{q}_0 + c\mathbf{q}_1 + \dots) \\ &= P_0(\mathbf{q}_0 + c\mathbf{q}_1 + \dots) + p_1(\mathbf{q}_0 + c\mathbf{q}_1 + \dots) \\ &= P_0(\mathbf{q}_0) + cDP_0(\mathbf{q}_0)\mathbf{q}_1 + cp_1(\mathbf{q}_0) + \dots. \end{aligned}$$

This is valid at $\mathcal{O}(c^0)$ from Eq. (4.10). At $\mathcal{O}(c^1)$, we need

$$\mathbf{q}_1 = DP_0(\mathbf{q}_0)\mathbf{q}_1 + p_1(\mathbf{q}_0). \quad (4.13)$$

Solving for \mathbf{q}_1 ,

$$\mathbf{q}_1 = [Id - DP_0(\mathbf{q}_0)]^{-1}p_1(\mathbf{q}_0), \quad (4.14)$$

where Id is the identity matrix. In order to solve for \mathbf{q}_1 , it is necessary that $[Id - DP_0(\mathbf{q}_0)]$ be invertible. This is equivalent to the above condition for the Implicit Function Theorem to hold that $D_{\mathbf{x}}h(\mathbf{q}_0, 0)$ be invertible.

We now show that this matrix is invertible provided the periodic orbit for the uncoupled system is hyperbolic. Suppose that v is an eigenvector of $DP_0(\mathbf{q}_0)$ with eigenvalue λ , so that

$$[DP_0(\mathbf{q}_0)]v = \lambda v. \quad (4.15)$$

Then

$$(Id - [DP_0(\mathbf{q}_0)])v = (1 - \lambda)v. \quad (4.16)$$

Thus, the matrix $(Id - [DP_0(\mathbf{q}_0)])$ only has a zero eigenvalue if $\lambda = 1$. But the eigenvalues of $DP_0(\mathbf{q}_0)$ give the stability of the periodic orbit for the uncoupled problem; in particular, if it is a hyperbolic periodic orbit, none of the eigenvalues are on the unit circle. The hyperbolicity condition only needs to be checked for a single uncoupled oscillator, since we are assuming that the oscillators are identical.

Summarizing, provided the periodic orbit for the uncoupled system is hyperbolic, there will be a nearby periodic orbit for the system with sufficiently small coupling.

Furthermore, we expect that since the $c \rightarrow 0$ system limits to the $c = 0$ system, the periodic orbit for the weakly coupled system will “inherit” the stability properties from the periodic orbit for the uncoupled system. This follows from the continuity of the Poincaré map with respect to c , giving

$$\lim_{c \rightarrow 0} DP_c(\mathbf{q}(c)) = DP_0(\mathbf{q}_0). \quad (4.17)$$

This implies that the eigenvalues corresponding to the stability of the $\mathbf{q}(c)$ periodic orbit for the coupled system tend toward the eigenvalues corresponding to the stability of the \mathbf{q}_0 periodic orbit for the uncoupled system.

4.4.2 $N = 2$ coupled oscillators

Now consider $N = 2$ parametrically forced oscillators which are coupled linearly:

$$\ddot{x}_1 + b\dot{x}_1 + x_1 + x_1^3 = x_1 F \cos(\omega_f t) + c(x_2 - x_1), \quad (4.18)$$

$$\ddot{x}_2 + b\dot{x}_2 + x_2 + x_2^3 = x_2 F \cos(\omega_f t) + c(x_1 - x_2), \quad (4.19)$$

where x_i is the position of the i^{th} oscillator, $i = 1, 2$.

If $c = 0$, these are independent parametrically forced oscillators. Thus, for $b = 0.2$, $F = 0.5$, and $\omega_f = 2$, each oscillator could be in a stable periodic state given by the A or B solutions; each oscillator also has an unstable 0 solution. The periodic solutions for the uncoupled two oscillator system are given in Table 4.1.

Here the first symbol characterizes the state of the first oscillator, and the second symbol characterizes the state of the second oscillator. *In phase* solutions have both oscillators responding identically to the forcing, while *out of phase* solutions correspond to each oscillator undergoing an oscillation which is shifted by one period of the forcing relative to the other. *Large-small* solutions have one oscillator undergoing oscillations locked to the forcing while the other oscillator is stationary. The name comes from the result that when weak coupling is introduced, the former

Table 4.1: Periodic solutions for uncoupled two oscillator system

Type	Equivalent States
In-phase	AA, BB
Out-of-phase	AB, BA
Large-small	$A0, B0, 0A, 0B$
00 solution	00

oscillator will undergo relatively *large* oscillations, while the latter will undergo relatively *small* oscillations. For the 00 solution, both oscillators are stationary. The distinct solutions within a given class (in phase, out of phase, or large-small) are related by symmetry.

Results for $c \neq 0$ are illustrated for Eqs. (4.18,4.19) for $b = 0.2, F = 0.5, \omega_f = 2$ in Fig. 4.5, which show that for small $c > 0$, periodic orbits of the expected stability type exist and are close to the periodic orbits for the uncoupled system. Indeed, for the uncoupled system the A and B solutions are stable and the 0 solution is unstable (see Fig. 4.1); thus, the AA and AB solutions (and their symmetry-related counterparts BB and BA , respectively) are expected to be stable for small $|c|$, while the others are expected to be unstable. Figure 4.5 shows that as c increases, the out of phase AB solution loses stability. This illustrates that our arguments above are only valid for small $|c|$. We note that, in this figure, the fact that the branches come together at $c = 0$ is an artifact of the projection. The solutions are actually separated in phase space: even though they share the same value for $\max(|x_1|)$, the second oscillator has different behavior. Therefore, the uniqueness property from the Implicit Function Theorem argument is not violated.

For $b = 0.2, F = 0.5, \omega_f = 2.25$, Fig. 4.1 shows that the 0 solution is stable, and that there are stable periodic orbits (which are analogs of the A and B solutions discussed above) and unstable periodic orbits. Figure 4.6 shows that for small $c > 0$, periodic orbits of the expected stability type exist and are close to the periodic orbits for the uncoupled system. (To aid in interpreting this plot, we note that the maximum x values for the stable and unstable periodic orbits are approximately 0.72 and 0.45, respectively.) Indeed, for the uncoupled system the A , B , and 0 solutions are stable; thus, the in-phase AA solutions, out-of-phase AB solutions, large-small solutions (which in the limit $c \rightarrow 0$ approach the $A0$ solutions), and 00 solution are all stable for small $|c|$. All solutions which involve an unstable periodic orbit for the uncoupled system as $c \rightarrow 0$ are unstable.

We now fix the coupling strength as $c = 0.03$ and take $b = 0.2, F = 0.5$, and treat ω_f as a bifurcation parameter. Figure 4.7 shows the corresponding bifurcation diagram for the in-phase, out-of-phase, and 00 solutions. We see that the in-phase and out-of-phase solutions bifurcate from the 00 solution at different values of ω_f . As expected from the discussion above, for $\omega_f = 2$ both the in-phase and out-of-phase solutions are stable, while for $\omega_f = 2.25$ the in-phase, out-of-phase, and 00 solutions are all stable. We expect at $\omega_f = 2.25$ there will be large-small




Figure 4.5: Bifurcation diagram for fixed $b = 0.2, F = 0.5, \omega_f = 2$ with the coupling strength c treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. For small (positive) c , from top to bottom, the solutions of the branches approach the $AA, A0, AB, 0A,$ and 00 as $c \rightarrow 0$. Note, for example, that the AA and AB solutions have the same $\max(x_1)$ value, and thus appear to be on the same branch in this projection. A similar coincidence between symmetry-related solutions occurs for all other branches.

solutions; this is verified in Fig. 4.8 which identifies them as being on a branch which bifurcates from the out-of-phase solution branch. The large-small solution at $\omega_f = 2.25$ is shown in Fig. 4.9.

Although our analytical results only apply in the weak coupling limit, we note that interesting dynamics occur for larger $|c|$, such as anti-synchronized chaotic behavior for $b = 0.2, F = 0.5, \omega_f = 2$, and $c = -0.9$ shown in Figs. 4.10 and 4.11.

4.4.3 $N = 3$ coupled oscillators

For the identical bidirectional coupling considered in this chapter, the $N = 3$ case is the smallest network wherein network topology plays a significant role. We consider

$$\ddot{x}_1 + b\dot{x}_1 + x_1 + x_1^3 = x_1 F \cos(\omega_f t) + c(x_2 - x_1) + \tilde{c}(x_3 - x_1), \quad (4.20)$$

$$\ddot{x}_2 + b\dot{x}_2 + x_2 + x_2^3 = x_2 F \cos(\omega_f t) + c(x_1 - x_2) + c(x_3 - x_2), \quad (4.21)$$

$$\ddot{x}_3 + b\dot{x}_3 + x_3 + x_3^3 = x_3 F \cos(\omega_f t) + \tilde{c}(x_1 - x_3) + c(x_2 - x_3). \quad (4.22)$$

The cases $\tilde{c} = c$ and $\tilde{c} = 0$ are referred to *ring* and *line* coupling topologies, respectively, and are shown in Figure 4.12. We will consider the limit of weak coupling and $|\tilde{c}| \ll 1$, and we will show how the network topology will partition the set of all possible solutions into sets of symmetry-related solutions. We note that, due to the form of coupling, there will always exist a zero solution 000 with stability inherited from the uncoupled case, and a class of stable synchronous solutions $\{AAA, BBB\}$, regardless of network topology. In the following analysis, we will be




Figure 4.6: Bifurcation diagram for fixed $b = 0.2$, $F = 0.5$, $\omega_f = 2.25$ with the coupling strength c treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.

specifically considering the case where $b = 0.2$, $\omega_f = 2$, $F = 0.5$, which corresponds to the region in bifurcation space where, for a single oscillator, only the 0 solution and the large periodic orbit exist. The numerics were done using AUTO [33], which enables the analysis of stable *and* unstable orbits. We note that the representative solutions we consider are particular to the system described by Equations (4.20)-(4.22).

Notation

For convenience, we will introduce additional notation to describe the solution types and indicate solution symmetries. The $z \in \{A, B, 0\}$ solution label may now be indexed by subscripts, i.e. z_{XY} . The X subscript variable indicates two different qualifications of the solution, depending on the z solution type. If the solution is near or identically zero, i.e. $z = 0$, the subscript $X \in \{A, B, 0\}$ indicates if the solution is identically zero ($X = 0$) or if it is a small (near zero) periodic orbit with phase the same as either the A or B large periodic orbits ($X = A$ or $X = B$, respectively). For solutions near or identically zero, the Y superscript on the X subscript is not used.

For the large periodic orbit states $z \in \{A, B\}$, the subscript $X \in \{L, S, 0\}$ denotes the relative size of the solution and is intended to provide a qualification of the symmetry when coupling is included. If the solution is identical to the uncoupled stable periodic orbit, we will denote this by $X = 0$. If, on the other hand, the solution's magnitude (i.e. $\max_{t \in (0, T]}(x(t))$) is larger (resp., smaller) than the uncoupled stable periodic orbit, we denote this by $X = L$ (resp., $X = S$). When two or more oscillators both exhibit a large periodic orbit close in magnitude to the uncoupled solution, we will use our notation to indicate symmetry. For example, if there are two identical A solutions that are larger than the uncoupled orbit, we will indicate this by $A_L A_L$. If there are two A solutions that are larger




Figure 4.7: Partial bifurcation diagram for Eqs. (4.18)-(4.19) for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.

than the uncoupled solution, but not equal to each other, we will indicate this using the Y superscript as $A_{L^+}A_{L^-}$. The $Y \in \{-, +\}$ superscript will be used with the 0_X labels as well, also to indicate non-identical orbits. The ‘+’ solution is so marked because it has a larger magnitude than the ‘-’ solution (again, quantifying magnitude by $\max_{t \in (0, T]}(x(t))$). It is worthwhile to note here that these orbits and their magnitudes are all well-defined in the limit of small $|c|$.

For the rest of our discussion in this chapter, when we refer to oscillators being *synchronized*, we mean phase-synchrony, which does not imply any relationship between amplitudes. When amplitudes are equal, we will use the terms *identical* if they are indistinguishable, or *anti-synchronized* if they are in anti-phase.

Ring topology

For the $N = 3$ case, the ring topology is the same as the all-to-all topology. Each oscillator is coupled identically to the others. Even in such a simple network as shown below, the dynamics are rich, both for small coupling and as the coupling strength becomes largely negative.

The ring topology gives six symmetry classes, which are summarized in Table 4.2, including the zero (R_1) and synchronous (R_2) solutions.

The R_3 class is the set of all solutions that have two oscillators with synchronous large periodic orbits and one with a small near-zero magnitude orbit. For $c \lesssim 0$, the large periodic orbits are slightly larger in magnitude than the uncoupled orbit and the near-zero orbit is in-phase. For $0 \lesssim c$, the two synchronous large periodic orbits are slightly smaller than the uncoupled orbit, and the near-zero orbit is out-of-phase. This solution is depicted in Figure 4.13. All solutions in the R_3 class are unstable for small $|c|$.

The R_4 class consists of all solutions wherein two oscillators have synchronized large periodic orbits and the other oscillator has a similar large periodic orbit, but




Figure 4.8: Full bifurcation diagram for Eqs. (4.18)-(4.19) for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions.

shifted in time by one forcing frequency. With non-zero coupling, the S_3 symmetry of the AAB solution is broken. For $c \lesssim 0$, the two synchronous periodic orbits are smaller in magnitude than the anti-synchronous orbit, resulting in an S_2 symmetry. The relative magnitudes of the periodic orbits switch for the positive small coupling case, $0 \lesssim c$, shown in Figure 4.14. The R_4 solutions are all stable for small $|c|$.

The R_5 class is the set of all solutions that have a single oscillator with a large periodic orbit and two oscillators with near-zero orbits. For $c \lesssim 0$, the large periodic orbit is larger than the uncoupled orbit and the two near-zero solutions are identical (retaining an analog of the symmetry of the uncoupled solution) and in-phase with the other oscillator. As c is increased through zero, the magnitude of the A solution becomes smaller than the uncoupled orbit, and the two near-zero solutions switch to become out-of-phase with the A solution (but still synchronous and identical to each other). All solutions in the R_5 class are unstable for small $|c|$.

The R_6 class contains all solutions that have two oscillators with anti-synchronized large periodic orbits while the third oscillator has a near-zero orbit. For $c \lesssim 0$, the two large periodic orbits have identical magnitudes and are slightly larger than the uncoupled solution. For $0 \lesssim c$, the two large periodic orbits have identical magnitudes and are slightly smaller than the uncoupled solution. In all the cases, the near-zero solution is, in fact, identically zero. Thus, the symmetry present in the uncoupled solution is preserved for sets of solutions in the R_6 class. All solutions in the R_6 class are unstable for small $|c|$.

The properties of the ring topology solutions for $c = 0$ and $|c| \approx 0$ are summarized in Tables 4.2 and 4.3.




Figure 4.9: Example large-small periodic orbit for $N = 2$, $b = 0.2$, $F = 0.5$, $\omega_f = 2.25$, $c = 0.03$. In the lower panel, the dark line and light line correspond to the two different oscillators.

Table 4.2: $N = 3$ ring symmetry classes

Label	Symmetry-related States	Stability
R_1	000	Unstable
R_2	AAA, BBB	Stable
R_3	AA0, A0A, 0AA, BB0, B0B, 0BB	Unstable
R_4	AAB, ABA, BAA, BBA, BAB, ABB	Stable
R_5	A00, 0A0, 00A, B00, 0B0, 00B	Unstable
R_6	AB0, A0B, B0A, 0BA, 0AB, BA0	Unstable




Figure 4.10: Time series for anti-synchronized chaos for $N = 2$, $b = 0.2$, $F = 0.5$, $\omega_f = 2$, and $c = -0.9$, where the dark line and light line correspond to the two different oscillators.




Figure 4.11: Phase space representation of anti-synchronized chaos for $N = 2, b = 0.2, F = 0.5, \omega_f = 2,$ and $c = -0.9,$ where the dark line and light line correspond to the two different oscillators.




Figure 4.12: $N = 3$ oscillator networks. Ring [$\tilde{c} = c$] (a), morph [$0 < \tilde{c} < c$] (b), and line [$\tilde{c} = 0$] (c).

Table 4.3: Solution representatives for $|c| \approx 0$ for each symmetry class for ring topology

Symmetry Class	$c \lesssim 0$	$c = 0$	$0 \lesssim c$
R_1	$0_0 0_0 0_0$	$0_0 0_0 0_0$	$0_0 0_0 0_0$
R_2	$A_0 A_0 A_0$	$A_0 A_0 A_0$	$A_0 A_0 A_0$
R_3	$A_L A_L 0_A$	$A_0 A_0 0_0$	$A_S A_S 0_B$
R_4	$A_S A_S B_L$	$A_0 A_0 B_0$	$A_L A_L B_S$
R_5	$A_L 0_A 0_A$	$A_0 0_0 0_0$	$A_S 0_B 0_B$
R_6	$A_L B_L 0_0$	$A_0 B_0 0_0$	$A_S B_S 0_0$




Figure 4.13: Example of R_3 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with ring topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators, the latter two of which have identical solutions.




Figure 4.14: Example of R_4 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with ring topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators, the latter two of which have identical solutions.




Figure 4.15: $N = 3$ ring bifurcation diagram for fixed $b = 0.2, F = 0.5, c = 0.03$ with ω_f treated as a bifurcation parameter. Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.3, and have symmetry-related solutions tabulated in Table 4.2.

Forcing-frequency bifurcation diagram: ring topology

We again fix the coupling strength as $c = 0.03$ and take $b = 0.2, F = 0.5$, and treat ω_f as a bifurcation parameter. Figure 4.15 shows the $\max(|x_1|)$ -projection of the corresponding bifurcation diagram. This projection is identical in appearance to the $\max(|x_2|)$ and $\max(|x_3|)$ bifurcation diagrams due to the label permutation symmetry of the ring topology. The in phase AAA solution and the $A0B$ solution bifurcate from the 000 solution at different values of ω_f . This bifurcation diagram is similar to the $N = 2$ case shown in Figures 4.7 and 4.8. The main features are preserved, such as the large-small solution, which is now a large-small-small solution (or in our new notation 0_A0_AB).

Chaotic dynamics

Similar to the $N = 2$ case, interesting dynamics also emerge as $|c|$ is increased. Figure 4.16 shows unsynchronized chaotic behavior for $b = 0.2, F = 0.5, \omega_f = 2$, and $c = -0.9$. The phase portrait of the $N = 3$ chaotic system shown in Figure 4.17 displays a double-well potential structure similar to that of the $N = 2$ case shown previously in Figure 4.11, although the trajectories of the three oscillators are not synchronized.

Line topology

The line topology breaks the symmetry of the ring and introduces new symmetry classes, which are summarized in Table 4.4. There still exists an unstable zero solution (L_1) and a stable synchronous solution (L_2). However, the line topology means that the “center” oscillator gets input from both of the two “end” oscillators,




Figure 4.16: Time series for unsynchronized chaos for $N = 3$ ring $b = 0.2$, $F = 0.5$, $\omega_f = 2$, and $c = -0.9$, where the dark line, the dashed line, and light line correspond to the three different oscillators.

while the end oscillators only receive input from center oscillator. This has dramatic effect on the symmetry class partitions.

The L_3 class represents solutions where the center oscillator is synchronized in a large periodic orbit with one of the end oscillators. The other end oscillator is near-zero. Solutions in this class exhibit interesting bifurcation behavior. One of the end oscillators retains the uncoupled orbit amplitude for all sufficiently small c , regardless of sign. This can be interpreted as a region of the amplitude versus coupling strength curve where the slope is zero. For $c \lesssim 0$, the center oscillator has a periodic orbit that is slightly larger in amplitude than the uncoupled orbit and the other end oscillator has an in-phase near-zero orbit. For small positive coupling, the center oscillator has a slightly smaller magnitude orbit, and the near-zero orbit of the other end oscillator is out-of-phase. This is an example of broken symmetry due to coupling, since for small non-zero c , the permutation symmetry of the AA pair is destroyed - the changing amplitude of the center oscillator prevents the S_{N_A} permutation operators from mapping the solutions identically back to themselves. This class of solutions is unstable for small $|c|$.

The L_4 class is the set of solutions where the center oscillator is synchronized in a large periodic orbit with one of the end oscillators and the other oscillator has an anti-synchronized large periodic orbit. Similar to the L_3 solutions, this class also has a bifurcation pattern that leaves one of the end oscillators unchanged in a larger periodic orbit. The center oscillator and the other end oscillator are identical but out-of phase, and have larger magnitudes than the uncoupled case for negative coupling, and smaller magnitudes for positive coupling. Although the permutation symmetry between the A solutions is broken by non-zero c , the symmetry between the center A solution and the end B solution remains for small c . The L_4 solutions are all stable for small $|c|$.

The L_5 class is the set of solutions where the center and one of the end oscillators




Figure 4.17: Phase space representation of unsynchronized chaos for $N = 3$ ring $b = 0.2$, $F = 0.5$, $\omega_f = 2$, and $c = -0.9$, where the dark line, the dashed line, and light line correspond to the three different oscillators.

are near zero and the other end oscillator has a large periodic orbit. Symmetries in this class are completely broken for non-zero coupling. For $c \lesssim 0$, the large periodic orbit has larger magnitude than in the uncoupled system, and the two zero solutions become distinct near-zero periodic orbits in phase with the A solution. A surprising behavior found numerically is that for small positive coupling, the zero solution of the center oscillator switches to be out-of-phase with respect to the other two oscillators, even though this phase relationship is not a mirror image of the negative coupling case. For both coupling cases, the amplitude of the center oscillator's near-zero orbit is slightly larger than that of the end oscillator's near-zero orbit. We show this solution in Figure 4.18. These solutions are unstable for small $|c|$.

The L_6 class consists of the solutions where the center oscillator is the 0 solution and the two end oscillators have large periodic orbits that are in anti-synchrony. The symmetry of the uncoupled case is preserved for small coupling for solutions in this class. When the coupling is negative, the oscillatory solutions are larger than the uncoupled case, and for positive coupling the reverse is true, as shown in Table 4.5. Interestingly, the center oscillator's zero solution remains identically zero. The L_6 solutions are unstable for small $|c|$.

The L_7 class is the set of solutions where the center oscillator is near zero and the two end oscillators are synchronized with large periodic orbits. Solutions in this class also retain their symmetry for small non-zero coupling. For $c \lesssim 0$, the two end oscillators follow orbits that are slightly larger than the uncoupled A solution, and the zero solution becomes a near-zero solution in-phase with the other two oscillators. For $0 \lesssim c$, the magnitude of the end oscillators' larger periodic orbit becomes slightly smaller than the uncoupled case, and the near-zero solution is out-of-phase. These solutions are also unstable for small $|c|$.

The L_8 class is the set of solutions where the two end oscillators have large syn-

chronous periodic orbits in anti-synchrony with the center oscillator's large periodic orbit. Symmetries of solution trajectories in this class are modified by coupling in an interesting way. The two end oscillators retain their permutation symmetry, but the symmetries with the center oscillator are lost, because the amplitudes no longer match. When the small coupling is negative, the center oscillator has a larger orbit than the two end oscillators which themselves have a slightly larger orbit than the uncoupled case. The situation is reversed for the case of small positive coupling, as shown in Table 4.5. These two solutions are stable for small $|c|$.

The L_9 class is composed of the solutions for which the center oscillator has a large periodic orbit in anti-synchrony with one of the end oscillators, and the other end oscillator has a near zero orbit. The solution symmetries for this class are broken by coupling. For $c \lesssim 0$, the two non-zero solutions become larger than the uncoupled periodic orbit, but are not identical, and the zero solution becomes a near-zero solution in-phase with the center oscillator. Small positive coupling yields a similar situation, with the two periodic solutions becoming slightly smaller than the uncoupled orbit, and non-identical, while the near-zero solution is in-phase with the other end oscillator. This solution is shown in Figure 4.19. The L_9 class of solutions are all unstable for small $|c|$.

The L_{10} class is the set of solutions where the two end oscillators have near-zero orbits and the center oscillator has a large periodic orbit. This class of solutions retains a version of its uncoupled symmetry. For small negative coupling, the center oscillator has a slightly larger orbit than the uncoupled trajectory, and the two end oscillators follow identical synchronous near-zero orbits. The opposite is true for small positive coupling. This class of solutions is unstable for small $|c|$.

The properties for the solutions for $|c| \approx 0$ are summarized for reference in Table 4.5.

Table 4.4: $N = 3$ line symmetry classes

Label	Symmetry-related States	Stability
L_1	000	Unstable
L_2	AAA, BBB	Stable
L_3	$AA0, 0AA, BB0, 0BB$	Unstable
L_4	AAB, BAA, BBA, ABB	Stable
L_5	$A00, 00A, B00, 00B$	Unstable
L_6	$A0B, B0A$	Unstable
L_7	$A0A, B0B$	Unstable
L_8	ABA, BAB	Stable
L_9	$AB0, BA0, 0BA, 0AB$	Unstable
L_{10}	$0A0, 0B0$	Unstable

Table 4.5: Solutions representatives for $|c| \approx 0$ for each symmetry class for line topology

Symmetry Class	$c \lesssim 0$	$c = 0$	$0 \lesssim c$
L_1	$0_0 0_0 0_0$	$0_0 0_0 0_0$	$0_0 0_0 0_0$
L_2	$A_0 A_0 A_0$	$A_0 A_0 A_0$	$A_0 A_0 A_0$
L_3	$A_0 A_L 0_A$	$A_0 A_0 0_0$	$A_0 A_S 0_B$
L_4	$A_0 A_L B_L$	$A_0 A_0 B_0$	$A_0 A_S B_S$
L_5	$A_L 0_{A^+} 0_{A^-}$	$A_0 0_0 0_0$	$A_S 0_{B^+} 0_{A^-}$
L_6	$A_L 0_0 B_L$	$A_0 0_0 B_0$	$A_S 0_0 B_S$
L_7	$A_L 0_A A_L$	$A_0 0_0 A_0$	$A_S 0_B A_S$
L_8	$A_L^- B_L^+ A_L^-$	$A_0 B_0 A_0$	$A_S^+ B_S^- A_S^+$
L_9	$A_L^- B_L^+ 0_B$	$A_0 B_0 0_0$	$A_S^+ B_S^- 0_A$
L_{10}	$0_A A_L 0_A$	$0_0 A_0 0_0$	$0_B A_S 0_B$

Forcing-frequency bifurcation diagram: Line topology

We again fix the coupling strength as $c = 0.03$ and take $b = 0.2$, $F = 0.5$, and treat ω_f as a bifurcation parameter. Figure 4.20 shows the $\max(|x_1|)$ -projection of the corresponding bifurcation diagram and Figure 4.21 shows the $\max(|x_2|)$ -projection. We note that the $\max(|x_1|)$ -projection is visually identical to the $\max(|x_3|)$ -projection, as the two end oscillators (numbers 1 and 3) differ only by a label permutation. As in the previous bifurcation diagrams, the in-phase AAA solution and the $A0B$ solution bifurcate from the 000 solution at different values of ω_f . The structure of these bifurcation diagrams is similar to the $N = 2$ case shown in Figures 4.7 and 4.8. We see the large-small solution branches are substantially different in magnitude for the center (number 2) oscillator versus the end oscillators (numbers 1 and 3).

Non-uniform coupling: morphing from a line to a ring

By starting with the ring topology, holding two of the coupling links fixed, and varying the third (\tilde{c} in Eqs. (4.20)-(4.22)), we can continuously morph the network from the ring topology to the line topology. In so doing, we can see how the symmetry classes change due to the topology.

The R_3 class splits into two classes, L_3 and L_7 , the R_4 class splits into the L_4 and L_8 classes, the R_5 class splits into the L_5 and L_{10} classes, and the R_6 class splits into the L_6 and L_9 classes. It is interesting to note that the stability properties of the split L classes are inherited from the parent R classes.

4.4.4 N coupled oscillators

These results generalize to N coupled parametrically forced oscillators. Specifically, provided the periodic orbits for the individual oscillators for the uncoupled




Figure 4.18: Example of L_5 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with line topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators.

system are hyperbolic, for every periodic solution which exists for the uncoupled system there will be a nearby periodic orbit for the system with sufficiently small coupling. Furthermore, the periodic orbit will inherit the stability properties from the periodic orbit for the uncoupled system.

Notably, the Implicit Function Theorem argument presented above does not depend on the coupling topology of the system, or any special properties about the coupling strength (for example, all the strengths being equal). That is, regardless of how the oscillators are coupled together, for sufficiently small coupling strengths there will be analogs of the periodic solutions which exist for the uncoupled system. For example, instead of the oscillators having all-to-all coupling, similar results hold for oscillators coupled only to their neighbors. Of course, as the coupling strengths increase away from zero, so that the above arguments no longer hold, the coupling topology will affect the types of states that exist and are stable.

4.5 Discussion

We have discussed periodic solutions which occur for parametrically forced oscillators that are weakly coupled together. A combinatorial approach was presented to enumerate all possible solutions for the uncoupled case. The existence and stability of periodic orbits for the coupled system can be determined by the existence and stability of the individual parametrically forced oscillators when they are uncoupled. Our results follow from an application of the Implicit Function Theorem to an appropriate Poincaré map. The results were confirmed using numerical bifurcation analysis for several simple networks of oscillators with 2:1 excitation resonance. An explicit expression was given for the isotropy subgroup of symmetric solutions, and the generators of those subgroups were identified. We classified the behavior of the




Figure 4.19: Example of L_9 periodic orbit for $b = 0.2, F = 0.5, \omega_f = 2.25, c = 0.03$ with line topology. In the lower panel, the dark line, the light line, and the dashed line correspond to the three different oscillators.

solutions as a function of coupling strength, and classified the symmetry-breaking that occurs for non-zero coupling, which depends on network topology. We also explored the solution bifurcation structure as a function of forcing frequency for our example networks, which could be useful for experimental purposes, for example in MEMS research [104, 20, 60, 82, 111]. These results may be viewed as an analog of general results on the existence of phase-locked solutions for weakly coupled autonomous oscillators, as in [1, 17].




Figure 4.20: $N = 3$ line bifurcation diagram for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter ($\max(|x_1|)$ projection). Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.5, and have symmetry-related solutions tabulated in Table 4.4.




Figure 4.21: $N = 3$ line bifurcation diagram for fixed $b = 0.2$, $F = 0.5$, $c = 0.03$ with ω_f treated as a bifurcation parameter ($\max(|x_2|)$ projection). Solid (resp., dashed) lines indicate stable (resp., unstable) solutions. Stable solutions are labeled by symmetry class per the $0 \lesssim c$ column of Table 4.5, and have symmetry-related solutions tabulated in Table 4.4.

Chapter 5

Reference-tracking Spike Timing Control

This chapter proposes several event-based feedback schemes for controlling spike timing in phase models of neurons. The control schemes presented in this chapter, and also in Chapter 6, are variable shape open-loop waveforms triggered by the detection of a voltage spike. In this chapter, we focus on the objective of driving a neuron to spike in phase with a reference oscillator with the same natural frequency. We propose a biologically-inspired impulsive control system that mimics the signaling behavior of real neurons. We also develop a quasi-impulsive extension that is suitable for experimental implementation. Applied to a pacemaker-driven ensemble, this control scheme can achieve desynchronization using a set of charge-balanced stimuli.

The organization of the chapter is as follows. Introductory material is presented in Section 5.1. In Section 5.2, we discuss the core concept of event-based control, define phase error, and precisely quantify the reference tracking objective. Sections 5.2.3 and 5.2.4 present control strategies based on impulsive and quasi-impulsive inputs, respectively. In Section 5.3, we describe how these methods might be extended to control an ensemble of neurons, including a population driven by a pacemaker. We give concluding remarks in Section 5.4.

5.1 Introduction and background

Much of the motivation for controlling neurons comes from the desire to treat Parkinson's disease, which causes involuntary tremors around 3-6 Hz that typically affect the distal portion of the upper limbs. These tremors have been associated with the synchronization of a cluster of neurons in the thalamus and basal ganglia [77]. For patients with advanced Parkinson's disease who do not respond to drug therapy, electrical deep brain stimulation (EDBS), an FDA-approved therapeutic procedure, may offer relief. In this procedure, a neurosurgeon guides a small electrode into the motor-control region of the brain, and then connects the

electrode to a pacemaker-like device implanted in the chest. As presently implemented, the device sends high-frequency (~ 100 Hz) electrical current stimulation pulses directly into the brain tissue, which has been found to alleviate tremors for some patients [10].

There is much interest in designing electrical DBS control systems that use additional electrodes for feedback, thereby making the stimulus “demand-controlled” [98, 99]. A feedback-based approach is attractive from a clinical perspective in that the biological tissue is only stimulated when necessary, thereby reducing the overall accumulation of negative side effects of electrical stimulation, and also the amount of power required from the implanted battery.

There are, however, challenges to implementing feedback control for neurons. First, the conductance-based Hodgkin-Huxley formalism discussed in Chapter 2 yields systems of continuous-time ordinary differential equations that tend to be *highly* nonlinear. Many traditional feedback control system designs require the controlled system to be either linear or only *weakly* nonlinear. A second challenge is that a control scheme cannot stimulate biological tissue with arbitrarily large signals. The magnitude of the electrical stimulus must be constrained to acceptable ranges for hardware implementation and biological tissue tolerance. A third challenge is the fact that the only state that is directly observable is the neuron’s membrane voltage. A control system cannot measure the dynamic states of the many ion channels that play a critical role in the oscillatory behavior of neural spiking. This poses particular problems for implementing traditional nonlinear feedback control systems, which depend on continuous measurements of state [55]. The control algorithms in this chapter and the next are intended to at least partially overcome these challenges.

We will focus on feedback control of neurons modeled by the Hodgkin-Huxley equations, which display many of the dynamical characteristics of real oscillatory neurons, and constitute the most widely studied model in mathematical neuroscience. The Hodgkin-Huxley model possesses regimes in which the oscillatory spiking can be either advanced or retarded by unipolar stimulation pulses, depending on the time of injection. Also, the Hodgkin-Huxley model exhibits *Class 2 neural excitability*, which means that there is a bistable bifurcation structure featuring a subcritical Hopf bifurcation that leads to a discontinuous relationship between injected baseline current and firing frequency [53]. This model, while not representing human brain neurons, is the prototypical model for neuronal membrane dynamics, and exhibits oscillatory behavior similar to human motor control neurons in the thalamus and basal ganglia regions of the brain. Due to this qualitative similarity, and the fact that the Hodgkin-Huxley model is perhaps the most widely studied and familiar conductance-based neuron model, we choose to consider it as our neuron model in this chapter. Such conductance-based models are amenable to generating control schemes that are experimentally realizable *in vitro*, as shown in recent work [94].

The control algorithms presented in this chapter are designed to achieve charge balance, i.e. the net electrical charge injected into a neuron over one control cycle

is equal to zero. This is important because accumulation of charge leads to irreversible Faradiac reaction products that cause neural tissue damage [64]. There is a trade-off, however. Charge-balanced stimulation has been shown to cause less tissue damage [61, 62], but increases the likelihood of corrosion damage of the stimulating electrode, an issue currently being addressed by materials science and microelectronics research [64]. We focus on charge-balance in the hope that developments in electrode design technology will mitigate the corrosion problem.

As mentioned above, the only observable state is the neuron’s membrane voltage. In practice, background noise often affects the voltage measurement to the extent that the only reliable observation is the detection of voltage spikes, rather than the instantaneous value of the voltage itself. This situation leads us to employ a control methodology known as *event-based control*. The concept of event-based control, sometimes known as Lebesgue sampling, was developed as an improvement to fixed sample-rate feedback control for digital systems [2]. As shown in previous work [29, 30], this methodology finds natural utility in systems involving spiking neurons. In particular, rather than developing a control law based on continuous voltage feedback, we will focus on observing the voltage spikes as *events*. This event-based framework is convenient because we can consider the spike timing objective as an event-timing objective. Much work on event-based control of nonlinear systems has been done, especially with respect to stochastic processes [3]. While we do not consider stochasticity here, the event-based framework is well-suited to overcoming the observability and nonlinearity challenges inherent in spike timing control of oscillatory neurons.

5.2 Control of individual neurons

5.2.1 Event-based control

Measurement of membrane voltage tends to be quite noisy in experiments, making it quite difficult to employ traditional nonlinear feedback control based on continuous measurements of state, e.g. [55]. Neural voltage spikes are, however, reliably detectable. We will use the detection of voltage spikes to enable the controller to know that the neuron has just passed through its zero phase point. The controller has no memory beyond the last spiking event. Based on the time of the detected voltage spike, the controller will stimulate the neuron with an open-loop waveform in order to achieve some control objective. In this chapter, the control objective will be to spike in phase with a reference oscillator. We will develop control schemes that, iteration after iteration, produce the desired dynamical behavior.

5.2.2 Reference-phase tracking

Consider the objective of controlling a neuron to spike (asymptotically) in phase with a reference oscillator with the same natural period as the neuron. The reference

oscillator evolves according to the simple equation

$$\dot{\theta}_r = \omega_r \quad , \quad \theta_r(0) = \theta_{r_0}, \quad (5.1)$$

where $\omega_r = \omega$ is the natural frequency of the reference oscillator, and θ_{r_0} is its initial phase. We will see later in Sections 5.3.1 and 5.3.2 that such a control objective can be useful for controlling an ensemble of synchronized phase neurons to desynchronize by driving each neuron to follow a staggered reference phase trajectory.

The times at which the phase of the reference oscillator crosses zero are the times we want the controlled phase neuron to spike. To achieve this, we must develop a control scheme that, after every event-based open-loop stimulus application, drives the controlled neuron to spike closer (in time) to the zero crossing of the reference oscillator. One can think of the difference in the time at which the controlled neuron spikes to the time that the reference oscillator crosses zero as a *time error*. From a control-theoretic perspective, this is equivalent to the *phase error* at the time of the controlled neuron's spike. When the controlled neuron spikes, we compare its phase to that of its reference oscillator, and construct an open-loop waveform that will actuate the neuron with the goal of correcting all, or a portion of, its phase error by the time the neuron spikes again.

Generally, one would define the phase error as

$$\Delta\theta = \theta - \theta_r.$$

In the scenario presented in this chapter, the phase error is sampled only when the controlled neuron spikes, i.e. $\theta = 0$, so effectively $\Delta\theta = -\theta_r$. However, the phase error, as defined this way, exists on $(-2\pi, 0]$. The topology of the unit phase circle allows us to wrap the phase error to the interval $(-\pi, \pi]$ using the following algorithm (shown here in general form):

$$\Delta\theta = \begin{cases} \theta - \theta_r & , \quad \text{for } |\theta - \theta_r| \leq \pi \\ \theta - \theta_r - \text{sgn}(\theta - \theta_r)2\pi & , \quad \text{for } |\theta - \theta_r| > \pi \end{cases} \quad (5.2)$$

so that the phase error is the shortest distance around the unit phase circle.

When $\Delta\theta < 0$, the controller should speed up the neuron, and when $\Delta\theta > 0$ it should slow down the neuron. This definition of phase error is useful from the perspective of controlling oscillatory neurons. Figure 5.1 illustrates why this definition of phase error is appropriate. The reference oscillator trajectories are shown in dashed lines. The time at which the controlled neuron spikes and triggers the event-based controller is labeled t_0 . The left panel shows a scenario wherein $\Delta\theta < 0$. With this definition of phase error, the controller should seek to drive the controlled neuron to spike at the time labeled t_r , thus speeding it up. The right panel shows a scenario where $\Delta\theta$ is small and slightly positive. If the phase error was not wrapped, for an arbitrarily small error, the controller would try to make the neuron spike almost immediately, which would require a very strong stimulus. Instead, with the phase error wrapping algorithm, the controller's objective is to slow down the neuron using a small stimulus so that the controlled neuron spikes




Figure 5.1: Phase error sign convention examples. Reference trajectories are shown as dashed lines. The controlled neuron’s phase is shown as a solid line. When the controlled neuron’s phase reaches $2\pi = 0 \pmod{2\pi}$, it spikes and triggers assessment of phase error relative to the reference oscillator.

in phase with the reference neuron one period later, as shown by the marker at t_r on the right panel of Figure 5.1.

In this event-based framework, we are interested in how the phase error $\Delta\theta$ changes after a period of control actuation, so we will define $\Delta\theta^+$ to be the phase error at the time of the next spiking event. Thus, we seek a control law that decreases the phase error over one period, i.e. ,

$$\left| \frac{\Delta\theta^+}{\Delta\theta} \right| < 1 \quad \forall \Delta\theta \in (-\pi, \pi), \quad (5.3)$$

excluding $\Delta\theta = 0$ (where $\Delta\theta^+$ should also equal 0).

5.2.3 Impulsive control

In nature, neurons communicate by voltage spikes that are large in magnitude but very short in duration, which naturally limits the production of irreversible Faradaic reaction products that can lead to tissue damage [64]. Signals of this kind are a biological analog of impulses. In fact, dynamical systems researchers in mathematical neuroscience have long used the concept of impulsive coupling to model networks of neurons, see e.g. [14]. Impulsive coupling has been shown to closely correspond to many types of oscillatory biological networks in nature [66, 110]. This has inspired the idea of using impulsive signals (Dirac delta functions) for spike timing control. Impulses are analytically desirable inputs from the perspective of the phase-reduced model, since delta functions turn the calculus into simple algebra. For example, consider the dynamics of generic phase-reduced model over the time interval $[t_I, t_{II}]$ subject to an impulsive input at time \hat{t}

$$\dot{\theta} = \omega + Z(\theta)\tilde{u}\delta(t - \hat{t}), \quad (5.4)$$

where $t_I \leq \hat{t} < t_{II}$, and \tilde{u} is the strength of the impulsive stimulus. The solution is simply

$$\theta(t_{II}) = \theta(t_I) + \omega(t_{II} - t_I) + Z(\theta(t_I) + \omega(\hat{t} - t_I))\tilde{u} \pmod{2\pi}. \quad (5.5)$$

Following [30], we will proceed by using (5.5) as the basic building block of our control scheme. Intuitively, we want to drive the oscillator with impulses timed to occur when its phase corresponds to that of the extremal values of its PRC. For example, if the control objective is to speed up the neuron, the optimal strategy is to stimulate the neuron with a negative impulse, timed to occur when $\theta = \alpha$ (recall $Z(\alpha) = Z_{min} < 0$), followed by a positive impulse, timed to occur when $\theta = \beta$ (recall $Z(\beta) = Z_{max} > 0$).

Since we are not considering noise, we can use (5.5) to predict the phase of the actuated oscillator using simple algebra. The charge-balance constraint is implemented by constraining the control to be in the form of two timed impulses of equal magnitude but opposite sign. Recall that the control objective is to reduce the phase error after each period of actuation, i.e. $|\Delta\theta^+| < |\Delta\theta|$. The following control algorithm, derived using (5.5), gives $|\Delta\theta^+| = K|\Delta\theta|$, where we choose the desired phase error correction factor $K \in [0, 1)$:

$$u(t) = \tilde{u}(\delta(t - t_\alpha) - \delta(t - t_\beta)), \quad (5.6)$$

where

$$\tilde{u} = \frac{(1 - K)\Delta\theta}{Z_{max} - Z_{min}}, \quad t_\alpha = \frac{\alpha}{\omega}, \quad t_\beta = \frac{1}{\omega}(\beta - Z_{min}\tilde{u}). \quad (5.7)$$

We note here that t_α and t_β are the times at which the neuron's phase will equal α and β , respectively (see Figure 2.4). As in the previous sections, t represents the time since the last spiking event t^* and is reset to zero whenever θ crosses the $\theta = 2\pi \pmod{2\pi}$ spike threshold.

This control scheme is, by construction, charge balanced (“bi-phasic charge-balanced with delay” in the terminology of [64]). It corrects the error exactly as we intend over one control period, and any other control waveform would be using energy at a “weaker” region of the PRC, or would violate the charge-balance constraint.

We must remember, however, that the phase-reduced model is a simplified representation of a higher-dimensional conductance-based model, and thus has a shortcoming that must be addressed. The behavior of the phase-reduced model is not necessarily representative of the conductance-based model when the impulses are large enough to drive the oscillator to a phase where the sign of the PRC is different from what it was prior to the impulse. Also, if the oscillator is driven beyond the $\theta = 2\pi = 0 \pmod{2\pi}$ spike threshold, in either direction, the phase-reduced model loses relation to the conductance-based model, since a phase of zero implies a firing and an essential “reset” of the oscillator.

Since we are concerned with asymptotic convergence to a fixed frequency reference trajectory, we can easily avoid these issues by using fractional error correction

Corner Condition	Minimum admissible correction factor K_{min}
$0 < \theta(t_\alpha^+) < \gamma$	$1 + \frac{\alpha(Z_{max} - Z_{min})}{\pi Z_{min}}$
$\theta(t_\alpha^+) < \gamma$	$1 + \frac{(\gamma - \alpha)(Z_{max} - Z_{min})}{\pi Z_{min}}$
$\gamma < \theta(t_\beta^+)$	$1 - \frac{(\beta - \gamma)(Z_{max} - Z_{min})}{\pi Z_{max}}$
$\gamma < \theta(t_\beta^+) < 2\pi$	$1 - \frac{(2\pi - \beta)(Z_{max} - Z_{min})}{\pi Z_{max}}$

Table 5.1: Constraints on the tunable desired contraction factor K . If K is chosen to be greater than the value of the right column, the corner condition in the left column is satisfied.

with

$$0 \leq K_{min} \leq K < 1, \quad (5.8)$$

where K_{min} is constrained by functions of the PRC $Z(\theta)$. Table 5.1 lists the corner conditions for K_{min} . In the table, $\theta(t_\alpha^+)$ refers to the phase of the oscillator immediately after the impulse at $t = t_\alpha$. Likewise $\theta(t_\beta^+)$ refers to the phase immediately after the impulse at $t = t_\beta$.

By using fractional error correction and phase wrapping in our phase error definition, we provide the impulsive event-based control scheme with a way to reduce the total charge delivered over each actuation period while still retaining asymptotic convergence to the specified reference trajectory. We will now extend this concept to a more experimentally relevant context by approximating the impulses by finite magnitude pulses of non-zero duration.

5.2.4 Quasi-impulsive control

Using finite (small) magnitude control pulses is important in the context of stimulating real neurons, since the biological tissue exposed to the electrical stimulus can be damaged by large electrical currents. Also, the phase reduction method that generates the phase models we use makes the assumption that the input acts as a small perturbation. A digital approximation of a Dirac delta function as a rectangular spike with magnitude \tilde{u}/dt , where dt is equal to the sample time, works well for numerical simulation of the phase-reduced nonlinear oscillator model, but is inappropriate for use with the full-dimensional conductance-based model. Such a stimulus can instantaneously jolt the state far off its periodic orbit and yield results that are not closely approximated by the phase reduced model.

To address these issues, we develop a quasi-impulsive control that uses the same control effort as the impulsive control, but extends the duration and confines the magnitude of the impulse to be equal to a threshold C , which is chosen to be greater than or equal to a certain minimum value C_{min} which depends on the PRC, the phase error correction factor K , and the natural frequency of the phase neuron. The

finite duration pulses will be stimulating the neuron at sub-extremal regions of the PRC, so the error correction of this protocol is not exact as in the impulsive case. We will show, however, that when implemented on the full-dimensional Hodgkin-Huxley neuron model, the resulting fractional error correction performance is quite close to the prescribed fractional error correction factor K from the phase-based quasi-impulsive control method.

Analogous to a time-delayed bang-bang control scheme, this method stimulates at magnitudes equal to the threshold constraint $C \geq C_{min}$, using rectangular pulses of opposite sign centered at t_α and t_β with durations such that the integral of each pulse is equal to \tilde{u} . Using a value of fractional error correction $K \geq K_{min}$ satisfying the conditions listed in Table 5.1, we propose the following control scheme:

$$u(t) = \begin{cases} 0 & , \text{ for } 0 \leq t < t_A \\ \text{sgn}(\Delta\theta)C & , \text{ for } t_A \leq t < t_B \\ 0 & , \text{ for } t_B \leq t < t_C \\ -\text{sgn}(\Delta\theta)C & , \text{ for } t_C \leq t < t_D \\ 0 & , \text{ for } t_D \leq t \end{cases} \quad (5.9)$$

where

$$\begin{aligned} t_A &= t_\alpha - \frac{|\tilde{u}|}{2C} \quad , \quad t_B = t_\alpha + \frac{|\tilde{u}|}{2C} \\ t_C &= t_\beta - \frac{|\tilde{u}|}{2C} \quad , \quad t_D = t_\beta + \frac{|\tilde{u}|}{2C} \end{aligned} \quad (5.10)$$

and \tilde{u} is as defined previously in (5.7). The corner conditions that determine the minimum admissible threshold constraint C_{min} are listed on Table 5.2. Together, these constraints ensure that the control signal always stimulates in the right direction and will yield a charge-balanced waveform. In the limit of $C \rightarrow \infty$, this scheme recovers the timing and performance of the purely impulsive control law (5.6).

Theorem *For the phase-reduced neural oscillator model $\dot{\theta} = \omega + Z(\theta)u(t)$ where $u(t)$ is as defined in (5.9), $Z(\theta)$ satisfies the conditions from (2.13), K satisfies the conditions in Table 5.1, and C satisfies the conditions in Table 5.2, the phase error ratio over one period of actuation will be a strict contraction ($|\frac{\Delta\theta^+}{\Delta\theta}| < 1$), implying global monotonic convergence of the oscillator phase $\theta(t)$ to the reference phase $\theta_r(t)$.*

Proof. First, a word on notation. When developing bounds to prove error convergence, underbars \underline{x} and overbars \bar{x} will denote the greatest lower and least upper bounds on the variable x , respectively. The objective of the proof is to show that the error gain, $|\frac{\Delta\theta^+}{\Delta\theta}|$, is strictly less than one for all values of initial error $\Delta\theta \in (-\pi, \pi]$. This implies that the phase error is reduced after each event-driven actuation period. And since the oscillator in absence of input rotates around \mathbb{S}^1 with natural frequency ω , spiking events are persistent in time, which make it impossible for a steady state error to exist.

Condition	Minimum admissible control magnitude C_{min}
$0 < t_A$	$\frac{\omega\pi(1-K)}{2\alpha(Z_{max}-Z_{min})}$
$t_B < t_C$	$\max_{\Delta\theta \in (-\pi, \pi]} \left(\frac{\omega\Delta\theta(1-K)}{(\beta-\alpha)(Z_{max}-Z_{min})-Z_{min}(1-K)\Delta\theta} \right)$
$\theta(t_B) < \gamma$	$\max_{\Delta\theta \in (-\pi, \pi]} \left(\frac{-\omega(1-K)\Delta\theta}{2[(\gamma-\alpha)(Z_{max}-Z_{min})-Z_{min}(1-K)\Delta\theta]} \right)$
$\theta(t_C) > \gamma$	$\frac{\omega\pi(1-K)}{2(\beta-\gamma)(Z_{max}-Z_{min})}$
$\theta(t_D) < 2\pi$	$\max_{\Delta\theta \in (-\pi, \pi]} \left(\frac{-\omega(1-K)\Delta\theta}{2[(2\pi-\beta)(Z_{max}-Z_{min})+Z_{max}(1-K)\Delta\theta]} \right)$

Table 5.2: Constraints on the minimum admissible stimulus constraint $C = C_{min}$. If C is chosen to be greater than the value of the right column, the condition in the left column is satisfied.

If the extension of the impulsive control to the quasi-impulsive case were perfect, we would expect $|\frac{\Delta\theta^+}{\Delta\theta}| = K$. This, however, is the greatest lower bound, since a pulse with nonzero duration implies that the control will be stimulating the neuron at phases where the PRC will be sub-extremal. We will proceed with the proof by developing bounds on the time at which the oscillator will spike (cross the $\theta = 2\pi$ threshold), which we will denote as t^+ , and which will be compared with the time at which the constant frequency reference oscillator spikes to determine the phase error after one period of actuation, $\Delta\theta^+$.

For simplicity, we will develop bounds on t^+ by separately considering the cases $\Delta\theta > 0$ and $\Delta\theta < 0$. When $\Delta\theta = 0$, no control action is taken so that $\Delta\theta^+ = 0$.

Case I: $\Delta\theta > 0$

Intuitively, the control should slow the neuron down when $\Delta\theta > 0$. A control magnitude $C \geq C_{min}$ satisfying the conditions in Table 5.2 guarantees that throughout the duration of the first pulse, the oscillator will have a phase between 0 and γ , the region where $Z(\theta)$ is negative semidefinite. For a positive $\Delta\theta$, the pulse will be positive, so the stimulus can only decrease the velocity of the oscillator below ω . Likewise, admissibility of the control magnitude further guarantees that the oscillator's phase will be between γ and 2π (the region where $Z(\theta)$ is positive semidefinite) during the second pulse which is negative, since $\Delta\theta > 0$. Again this means that the control signal can only decrease the oscillator's velocity below its natural frequency ω .

If there was no control, the neuron would spike again at $t^+ = 2\pi/\omega$, which would result in $\Delta\theta^+ = \Delta\theta$. In view of the argument above, this is, in fact, the lower bound \underline{t}^+ .

Now we step through the dynamics to develop an upper bound for t^+ . We begin

at $\theta(0) = 0$. Then, advancing with zero input until t_A ,

$$\theta(t_A) = \omega t_A = \alpha - \frac{\omega(1-K)\Delta\theta}{2C(Z_{max} - Z_{min})}. \quad (5.11)$$

Now we calculate a lower bound on $\theta(t_B)$. We do this by using Z_{min} as a lower bound on the PRC. Between t_A and t_B , our input is equal to C . We obtain

$$\begin{aligned} \underline{\theta}(t_B) &= \theta(t_A) + (\omega + Z_{min}C)(t_B - t_A) \\ &= \alpha + \frac{(\omega + 2CZ_{min})(1-K)\Delta\theta}{2C(Z_{max} - Z_{min})}. \end{aligned}$$

We then evolve with zero input until t_C :

$$\underline{\theta}(t_C) = \underline{\theta}(t_B) + \omega(t_C - t_B) = \beta - \frac{\omega(1-K)\Delta\theta}{2C(Z_{max} - Z_{min})}. \quad (5.12)$$

The input is then applied again, this time in the negative direction, since we wish to slow the neuron down, and $Z_{max} > 0$. We obtain

$$\begin{aligned} \underline{\theta}(t_D) &= \underline{\theta}(t_C) + (\omega - Z_{max}C)(t_D - t_C) \\ &= \beta + \frac{(\omega - 2CZ_{max})(1-K)\Delta\theta}{2C(Z_{max} - Z_{min})}. \end{aligned}$$

We now solve for the upper bound \bar{t}^+ using the relation

$$\theta(\bar{t}^+) = 2\pi = \underline{\theta}(t_D) + \omega(\bar{t}^+ - t_D), \quad (5.13)$$

giving

$$\bar{t}^+ = t_D + \frac{2\pi - \underline{\theta}(t_D)}{\omega} = \frac{2\pi + (1-K)\Delta\theta}{\omega} \quad (5.14)$$

So for $\Delta\theta > 0$,

$$\frac{2\pi}{\omega} < t^+ \leq \frac{2\pi + (1-K)\Delta\theta}{\omega}. \quad (5.15)$$

In terms of phase, these bounds on t^+ imply $K \leq \frac{\Delta\theta^+}{\Delta\theta} < 1$, as desired.

Case II: $\Delta\theta < 0$

When $\Delta\theta < 0$, the control method seeks to speed up the oscillator. Following the $C \geq C_{min}$ admissibility argument from Case I, but with the signs flipped, we conclude that the control signal cannot slow the oscillator down. Thus we have a simple upper bound: $\bar{t}^+ = 2\pi/\omega$. We can now step through the dynamics in the same manner as Case I, but with $\Delta\theta < 0$, to yield the inequality

$$\frac{2\pi + (1-K)\Delta\theta}{\omega} \leq t^+ < \frac{2\pi}{\omega}. \quad (5.16)$$

Therefore, $K \leq \frac{\Delta\theta^+}{\Delta\theta} < 1$, as claimed.

Thus, for all nonzero values of $\Delta\theta \in (-\pi, \pi]$, the control provides error contraction over one period of actuation, and if $\Delta\theta = 0$, the control takes no action. \square

The solid lines in the plots in Figure 5.2 illustrate the performance of this control algorithm for the phase-reduced model derived from the Hodgkin-Huxley system with the PRC shown previously in Figure 2.4. For this PRC, the minimum admissible values K_{min} and C_{min} are 0.63 and 1.65, respectively. The results shown are for $K = 0.7$ and $C = 1.7$. We see that the gain $\left| \frac{\Delta\theta^+}{\Delta\theta} \right|$ is between 0.7 and 0.8 over the entire interval, quite close to our prescribed K value of 0.7 derived from the optimal impulsive control method (5.6). As discussed in [29], the global stability of the origin of $M : \Delta\theta \mapsto \Delta\theta^+$ determines the global asymptotic stability of the phase error. Here, M is well-behaved, smooth, and is confined to the first and third quadrants (as expected with global monotonic convergence), which results in global asymptotic stability of the phase error.

We now implement the quasi-impulsive control method (5.9) on the full-dimensional neuron model (2.3) using parameters listed in Section 2.2. Our objective is to show that the phase error gain $\left| \frac{\Delta\theta^+}{\Delta\theta} \right|$ is less than one for all initial values of $\Delta\theta$. We will also compare the results to those achieved with the phase-reduced model.

Before outlining our results, we will briefly explain how we implement the control, which was developed for the phase-reduced model, on the full-dimensional model. For a single simulation, we choose an initial error $\Delta\theta$. We initialize the state of the model with phase $\theta(0) = 0$ (the state vector representation of that point on the periodic orbit, $\mathbf{x}(0)$, is known based on information derived during the phase reduction). We then integrate the ODE system (in \mathbf{x} coordinate space) using the electrical stimulus signal $I(t) = cu(t)$, where we recall that c is the constant membrane capacitance (which for the standard Hodgkin-Huxley system is equal to $1.0\mu\text{F}/\text{cm}^2$). The simulation proceeds until a spike is detected (the details of spike detection and phase sampling can be found in [27]). The timing of this spike is compared to the timing of the reference oscillator spike (initialized based on the choice of $\Delta\theta$) to obtain the value of $\Delta\theta^+$.

The results of fifty individual simulations with initial conditions ranging over $\Delta\theta \in (-\pi, \pi]$ are shown as a black line with white circle markers on Figure 5.2. We see that the implementation of the control law based on the phase-reduced model yields very similar results for the full-dimensional system. These results represent a significant improvement over previous work [29]. Here we have monotonic error convergence, whereas previous methods yielded asymptotic error convergence of $\left| \frac{\Delta\theta^+}{\Delta\theta} \right|$, a somewhat weaker control objective.




Figure 5.2: Quasi-impulsive control algorithm performance. The left plot shows the phase error gain. The right plot shows the $\Delta\theta \mapsto \Delta\theta^+$ map. Solid lines are results from the phase-reduced model to be compared with the white circle markers, which are results from the full-dimensional Hodgkin-Huxley system.

5.3 Control of an Ensemble of Neurons

In this section, we describe how the methods for controlling individual neurons described in the previous section might be extended to control an ensemble of neurons, including an ensemble driven by a pacemaker.

5.3.1 Desynchronizing an ensemble

Consider an ensemble of N identical uncoupled neurons firing in pathological synchrony. One can desynchronize the spike times of the ensemble using N instances of the control methods presented earlier for individual neurons.

Fully desynchronized spike times imply the population's interspike interval is constant and equal to T/N . We will show how the reference-tracking methods from Sections 5.2.3 and 5.2.4 can accomplish this goal using staggered reference phase trajectories.

Reference-phase tracking desynchronization

In a similar spirit, the reference-tracking control scheme from Sections 5.2.3 and 5.2.4 can be extended to the ensemble case by using a set of staggered reference phase trajectories by setting the initial condition of each of the $i = 1, \dots, N$ reference oscillators (5.1) to be $\theta_{r_i}(0) = \frac{2\pi i}{N}$. The i^{th} event-based controller then asymptotically drives its neuron toward the reference trajectory $\theta_{r_i}(t)$, resulting in asymptotic convergence to the desired desynchronized phase trajectories.

5.3.2 Desynchronizing a pacemaker-driven ensemble

In this section, we show how the reference-phase tracking methods presented in Sections 5.2.3 and 5.2.4 can be further extended to address an ensemble of neurons that are being driven towards synchrony by a periodic impulsive stimulus called a pacemaker.

Pacemaker-driven ensemble model

We consider a pacemaker-driven ensemble of N identical phase neurons depicted in Figure 5.3. Such a model is not directly representative of any particular biological network, but is a simple example of pathologically-driven synchronization that can be mitigated with an extension of the control methods proposed in Sections 5.2.3 and 5.2.4.

In the pacemaker-driven ensemble model, each neuron is unidirectionally coupled to a central pacemaker, and has its own event-based controller. We assume each neuron's event-based controller can both observe spikes and stimulate the neuron it controls. The pacemaker fires periodically and is not affected by either the driven neurons or any control stimuli, but its spiking can be measured and known by each controller. We also assume the pacemaker communicates with the neurons by impulsive signals.

Since the neurons do not communicate with each other, we can analyze the dynamics of this system by considering the behavior of a single pacemaker-neuron-controller system. If a controller can drive the pacemaker-driven neuron asymptotically toward a reference phase trajectory with any initial condition $\theta_{r_i}(0) \in [0, 2\pi)$, then a set of such controllers can accomplish the goal of driving each neuron toward a staggered reference phase trajectory for the ensemble.




Figure 5.3: Pacemaker network. Each of the N neurons is driven by control signals u_i and the pacemaker voltage V_P . The controller observes spiking behavior of the individual neuron voltages V_i and the pacemaker voltage V_P .

We define the pacemaker as a simple oscillator

$$\dot{\theta}_P = \omega_P \quad , \quad \theta_P(0) = \theta_{P_0} \quad (5.17)$$

that periodically emits an impulsive signal of fixed (and known) strength K_P when its phase θ_P crosses 0, that is

$$V_P(t) = K_P \delta(\theta_P(t)). \quad (5.18)$$

We will take the natural frequency of the pacemaker to equal that of the phase neurons $\omega_P = \omega$. The pacemaker pulse strength K_P is taken to be positive and of order 1. Since, in this simple model, the dynamics of the pacemaker are fixed and deterministic, we can construct a simple observer that will provide the controller with an (exact) estimate of the phase of the pacemaker. The controller observes the times of two consecutive pacemaker spikes, T_{P_1} and T_{P_2} . Then it is possible to construct the (exact) estimated pacemaker phase trajectory:

$$\theta_P(t) = \frac{2\pi}{T_{P_2} - T_{P_1}}(t - T_{P_2}) \quad \text{mod } 2\pi \quad (5.19)$$

for future times $t \geq T_{P_2}$.

We model neurons as nonlinear phase oscillators that emit a detectable impulsive signal when their phase θ_i crosses 0 according to the following equation

$$\dot{\theta}_i = \omega + Z(\theta_i(t))[V_P(t) + u_i(t)], \quad \theta_i \text{ mod } 2\pi, \quad (5.20)$$

where u_i is the control signal applied to neuron i . For each controller, we define its *event* to be the zero crossing of the phase of its controlled neuron.

Uncontrolled dynamics

For small positive pacemaker strength K_P , the spike times of the uncontrolled system synchronize with the pacemaker. This can be explained as follows. Set $t = 0$, and the first pacemaker spike occurs at time $t = T_{P_1}$. Just before the spike happens, neuron i has phase $\theta_i(T_{P_1}^-)$. Immediately after the spike, the phase has shifted to $\theta_i(T_{P_1}^+) = \theta_i(T_{P_1}^-) + K_P Z(\theta_i(T_{P_1}^-)) \text{ mod } 2\pi$. Now, the pacemaker will fire next at time $T_{P_2} = T_{P_1} + \frac{2\pi}{\omega}$. The neuron's phase right before the next pacemaker spike at $t = T_{P_2}$ is calculated to be

$$\theta_i(T_{P_2}^-) = \theta_i(T_{P_1}^+) + \omega \frac{2\pi}{\omega} \quad \text{mod } 2\pi = \theta_i(T_{P_1}^+). \quad (5.21)$$

In general

$$\theta_i(T_{P_j}^-) = \theta_i(T_{P_{(j-1)}}^+) \quad (5.22)$$

and

$$\theta_i(T_{P_j}^+) = \theta_i(T_{P_j}^-) + K_P Z(\theta_i(T_{P_j}^-)) \quad (5.23)$$

which, by fixing the coupling strength K_P , can be interpreted as a one-dimensional map $M : \theta_i(T_{P_j}^-) \mapsto \theta_i(T_{P_j}^+)$ where

$$M(\theta) = \theta + K_P Z(\theta). \quad (5.24)$$

Phase entrainment is equivalent to attraction to the fixed point at the origin of the map M over the full measure of θ^- state space. In Figure 5.4, we have shown the map M_{HH} for the Hodgkin-Huxley phase oscillator with a rather large pacemaker strength, $K_P = 2$, for easy visualization. The origin of this map is readily verified to be asymptotically attractive on the full measure of its domain, by cobwebbing [97], for example. We note the presence of an unstable fixed point at $\theta^- = \gamma$, which is expected since $Z(\gamma) = 0$.




Figure 5.4: The map M_{HH} for $K_P = 2$.

Impulsive desynchronization

In the same spirit as Section 5.2.3, we propose a reference phase tracking control scheme based on impulsive inputs. When each neuron spikes, it resets its controller's own local clock \tilde{t}_i , and the control algorithm calculates an open loop waveform with the objective of reducing the neurons' phase error relative to a reference trajectory, evaluated when the neuron next spikes. This algorithm, however, is constructed based on the concept of using three impulses instead of two, as in Section 5.2.3. The additional impulse is timed to negate the effect of the pacemaker. The other two impulses are used to accomplish the reference phase tracking while fulfilling the charge balance requirement. In the forthcoming presentation, all times will be relative to \tilde{t}_i , the time since the neuron i last spiked. \tilde{t}_P^+ is the (relative) time when the pacemaker will spike again.

The control signal is composed of three impulses of strength $\bar{u}_{i,\alpha}$, K_P , and $\bar{u}_{i,\beta}$. These values will be defined shortly; they must satisfy

$$\bar{u}_{i,\alpha} - K_P + \bar{u}_{i,\beta} = 0, \quad (5.25)$$

in order to ensure charge-balance over the course of one stimulation interval.

Similar to the reference phase tracking algorithm presented in Section 5.2.3, this algorithm stimulates the phase neuron at the points of the PRC where the $Z(\theta)$ is minimal and maximal, i.e. $\theta_i = \alpha$ and $\theta_i = \beta$. The third impulse is timed to occur at exactly the same time as the pacemaker spike. Under these conditions, we calculate the (relative) time when the phase will be equal to α

$$\tilde{t}_{i,\alpha} = \frac{\alpha}{\omega}. \quad (5.26)$$

At $\tilde{t}_i = \tilde{t}_{i,\alpha}$, the controller applies an impulse with strength

$$\bar{u}_{i,\alpha} = \frac{K_P Z_{max} + (1 - K)\Delta\theta_i}{Z_{max} - Z_{min}}. \quad (5.27)$$

The controller then waits until the time when the controlled neuron's phase will be at the PRC's maximal point $\theta_i = \beta$, which can be computed to be

$$t_{i,\beta} = \frac{\beta}{\omega} - \frac{Z_{min}(K_P Z_{max} + (1 - K)\Delta\theta_i)}{\omega(Z_{max} - Z_{min})}. \quad (5.28)$$

At this time, the controller delivers an impulse with strength

$$\bar{u}_{i,\beta} = \frac{-((1 - K)\Delta\theta_i + K_P Z_{min})}{Z_{max} - Z_{min}}. \quad (5.29)$$

The event-triggered open-loop control waveform can now be expressed as

$$u_i(\tilde{t}_i) = \bar{u}_{i,1}\delta(\tilde{t}_i - t_{i,\alpha}) - K_P\delta(\tilde{t}_i - \tilde{t}_P^+) + \bar{u}_{i,2}\delta(\tilde{t}_i - t_{i,\beta}). \quad (5.30)$$

This charge-balanced impulsive control scheme exactly nullifies the effect of the pacemaker and achieves the reference phase tracking with correction factor K in the same manner as the algorithm of Section 5.2.3.

Quasi-impulsive desynchronization

Similar to the single neuron case examined in Section 5.2.4, we seek to make this control method practical for application to a biological system, using finite magnitude control signals rather than impulses. In the pacemaker system, however, there are three impulses to approximate – the two phase-correcting impulses and the pacemaker-nulling impulse.

Figure 5.5(a) shows simulation results of this control scheme applied to a pacemaker network containing ten phase-reduced Hodgkin-Huxley neurons initialized in a completely synchronized spiking state. Figure 5.5(b) shows results from the same system initialized in the desynchronized state. The phase-entraining effect of the pacemaker makes the system tend toward synchronization in absence of the controller. In these simulations, the pacemaker spike intensity was taken to be $K_P = 2$, the charge magnitude constraint on the control signal was $C = 3$ and the desired

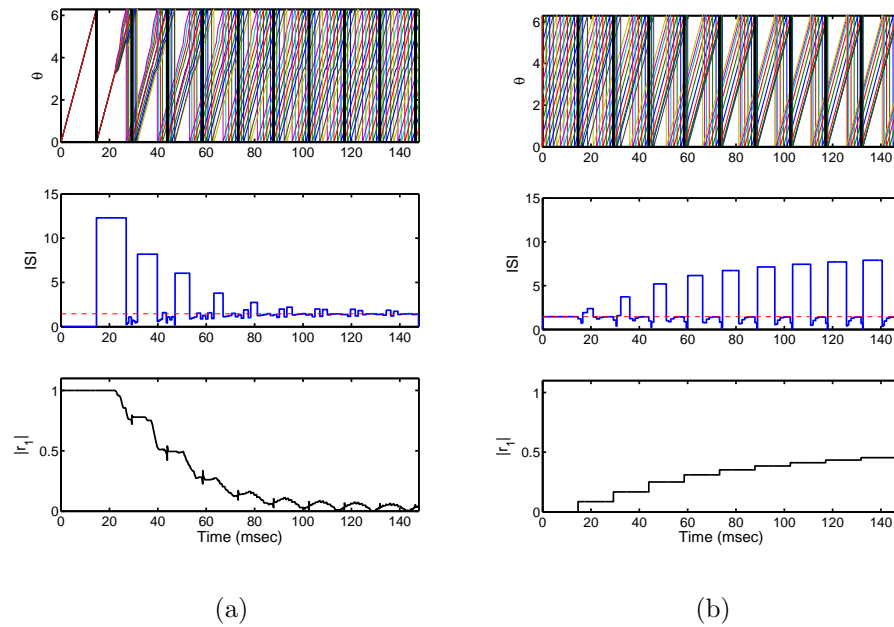


Figure 5.5: (a) Quasi-impulsive anti-pacemaker control scheme applied to a pacemaker-driven ensemble of $N = 10$ phase neurons. The spike timing of the ensemble, as well as the phase distribution, are desynchronized by the action of the controller. (b) When the controller is turned off, the pacemaker drives the ensemble back into spiking synchrony. The top panels show the time evolution of each neuron's phase. The middle panel tracks the most recent value of the ensemble's interspike interval (ISI). The bottom panels show the magnitude of the order parameter $|r_1|$.

fractional error correction was $K = 0.5$. In these plots, synchrony is illustrated in two ways. The middle panels show a measure of spiking synchrony which is the time interval between the last two spikes in the ensemble. Desynchronization of spike times implies that this number is always equal to the natural period of the neuron divided by the number of neurons in the ensemble. In this case, $T/N = 1.46$ msec and is shown by the dashed red line. The other important measure of synchrony is Kuramoto's order parameter [59], as was discussed in Chapter 3, the magnitude of which is a measure of phase synchrony. Recall, phase synchronization implies $|r_1| \rightarrow 1$, while phase desynchronization implies $|r_1| \rightarrow 0$.

5.4 Discussion

We have presented several event-based charge-balanced feedback control algorithms which can be applied to phase models of spiking neurons. This included controlling an individual neuron by using impulsive and quasi-impulsive inputs to make a neuron fire asymptotically in-phase with a reference oscillator. We also described how these methods, particularly the reference-phase tracking methods, might be extended to control an ensemble of neurons, including an ensemble driven by a pacemaker.

We note that the results presented in this chapter on controlling ensembles of neurons require the ability to provide different inputs to each neuron, which would be very difficult to realize experimentally. The challenge remains to develop a feedback control algorithm which desynchronizes a population of *coupled* neurons through a *single*, common input to the whole population, although there are some promising results along these lines [99, 28], as will be discussed in Chapter 6. Such an algorithm might truly deliver on the promise of demand-controlled deep brain stimulation for the treatment of Parkinson's disease.

Chapter 6

Phase-randomizing Control

In this chapter, we present an event-based feedback control method for randomizing the asymptotic phase of oscillatory neurons. Phase randomization is achieved by driving the neuron's state to its phaseless set, where its phase is undefined and is extremely sensitive to background noise. We consider the biologically relevant case of a fixed magnitude constraint on the stimulus signal, and show how the control objective can be accomplished in minimum time. The control synthesis problem is addressed using the minimum-time-optimal Hamilton-Jacobi-Bellman framework, which is quite general and can be applied to any spiking neuron model in the conductance-based Hodgkin-Huxley formalism. We also use this methodology to compute a feedback control protocol for optimal spike rate increase. This framework provides a straightforward means of visualizing isochrons, without actually calculating them in the traditional way. Finally, we present an extension of the phase randomizing control scheme that is applied at the population level, to a network of globally coupled neurons that are firing in synchrony. The applied control signal desynchronizes the population in a demand-controlled way.

This chapter is organized as follows. Section 6.1 provides some background and introductory material. Section 6.2 shows how we normalize the ODE models presented in Chapter 2 for use with the forthcoming control methodology. We outline the control objective, namely driving the neuron to a phase-randomizing target set in minimum time, in Section 6.3. Section 6.4 provides the theoretical basis for the optimal control design and Section 6.5 details the numerical methods used to perform the computations. Results are presented in Section 6.6 including single neuron phase randomization, ensemble phase randomization, and optimal spike rate increase. We also draw a connection to isochron visualization in this section, as well as provide additional results for the FitzHugh-Nagumo model. In Section 6.7, we show how the phase randomizing controller performs on a population of globally coupled neurons. We give concluding remarks in Section 6.8.

6.1 Introduction and background

In this chapter, we explore Winfree’s idea of using *phaseless sets* to randomize the phase of each neuron in a population, thereby desynchronizing their pathological firing behavior. The central concept is that the phase of an oscillatory dynamical system can be effectively reset by using an impulsive stimulus to drive the state of the system to a point where the phase is undefined [109]. In the context of the neuron models considered here, this phaseless set is a single point in state space, which is an unstable equilibrium point “inside” the periodic orbit. For more general cases, the phaseless set(s) could also be an unstable periodic orbit (for a planar system), or the stable manifold of an unstable periodic orbit (in higher dimensions). These objects typically form the boundary of the basin of attraction of the stable periodic orbit.

There are three primary challenges to the use of electrical stimulus to drive a neuron’s state to its phaseless set, which also posed challenges to reference-tracking in Chapter 5. First, the Hodgkin-Huxley conductance-based formalism [50], a general modelling methodology in neuroscience, yields systems of continuous-time ODEs that tend to be *highly* nonlinear. Many traditional control designs require the controlled system to be either linear or only *weakly* nonlinear. A second challenge is that a control scheme cannot stimulate biological tissue with arbitrarily large signals. The magnitude of the electrical stimulus must be constrained, invalidating the concept of purely impulsive inputs for biological implementation. The third challenge is the fact that the only state that is directly observable is the membrane voltage. A control system cannot measure the dynamic state of the many ion channels that play a critical role in the oscillatory behavior of neural spiking. This poses particular problems for implementing traditional nonlinear feedback control systems, which depend on continuous measurements of state.

The proposed event-based feedback control scheme presented here addresses each of these challenges, and accomplishes the objective of driving a neuron to its phaseless set using a magnitude-constrained stimulus signal in the minimum possible time. Our method combines techniques from modern control engineering, namely Lebesgue sampling (i.e. event-based control) [2], nonlinear optimal control theory (e.g. [4]), and a recent computational toolbox for solving the resulting Hamilton-Jacobi-Bellman equations [67]. The method is quite general and can be applied to any conductance-based neuron model that possesses a stable limit cycle.

We develop the event-based minimum-time optimal control algorithm for the case of a single deterministic neuron. We show that when background noise is added to the controlled neuron, the asymptotic phase of the neuron is indeed randomized. The utility of this control protocol is extended by considering a population-level event-based framework. In the case of a globally coupled network with pathological synchronous spiking, we show that this control scheme desynchronizes the network, and the control is active only when synchrony is detected.

The control design method developed here has several other interesting applications. By changing a simple parameter in the design calculations, we can create

a control signal that will optimally increase the spike rate of the neuron. Also, we can visualize the isochrons of the particular neuron model by plotting level sets of an intermediate value function calculated in the course of the stimulus design.

6.2 Normalized model

We now introduce a simple affine coordinate transformation to simplify the presentation, and to put our dynamical system into a form necessary for the proposed stimulus design method. A critical step in our method uses the Level Set Methods Toolbox (ToolboxLS), a MATLAB-based partial differential equation (PDE) solver [68]. This PDE solver has the requirement that each component should be of the same order of magnitude [67] and the new state-space coordinates must be centered at the target point, \mathcal{T} . The coordinate transformation is written as:

$$\begin{aligned} x_1 &= \frac{V-V_{\mathcal{T}}}{K_s} \\ x_2 &= n - n_{\mathcal{T}}. \end{aligned} \tag{6.1}$$

Since the gating variable n lies in the interval $[0, 1]$, we choose the voltage scaling factor, K_s , such that the scaled voltage coordinate $x_1 = \mathcal{O}(1)$. For the reduced Hodgkin-Huxley model, we use $K_s = 80$.

Under this transformation, the dynamics become

$$\begin{aligned} \dot{x}_1 &= \frac{1}{K_s} f_V(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) + \frac{1}{K_s} u \\ \dot{x}_2 &= f_n(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) \end{aligned} \tag{6.2}$$

To simplify notation further, we introduce the vector notation $x = [x_1 \quad x_2]^T$ with dynamics

$$\dot{x} = F(x) + \begin{bmatrix} \frac{1}{K_s} \\ 0 \end{bmatrix} u = \begin{bmatrix} F_1(x) \\ F_2(x) \end{bmatrix} + \begin{bmatrix} \frac{1}{K_s} \\ 0 \end{bmatrix} u \tag{6.3}$$

where

$$F_1(x) = \frac{1}{K_s} f_V(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) \tag{6.4}$$

and

$$F_2(x) = f_n(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}). \tag{6.5}$$

We note that this change of coordinates will only be used to simplify the mathematical presentation of the optimal stimulus design method. When results are presented in Section 6.6, we will revert back to the original variables V and n , and will report stimulus signals as electrical currents, $I(t)$.

In the present work, we design a controller that waits for a voltage spike (which we define to be the *event*) then stimulates the neuron with a pre-computed waveform designed to drive the neuron's state close to the unstable fixed point "inside" the periodic orbit, a point which is surrounded with closely packed isochrons. We will also show that other target sets can be used, which allows the controller to achieve other objectives, like maximally increasing the spike rate.

6.3 Control objectives

From a theoretical standpoint, the objective of our control system is to drive the system's state to a pre-defined target location \mathcal{T} . By setting \mathcal{T} to be the unstable equilibrium point, the controller will cause asymptotic phase randomization. On the other hand, if we set \mathcal{T} to be the spiking point, the controller will cause the neuron's spike rate to increase as much as possible. To keep our control stimulus design methodology general, we will not specify \mathcal{T} . In fact, any point in state space can be chosen as a target set, although it is not clear what the utility of using points besides the unstable equilibrium and the spike point might be. Recall that Equation (6.1) translates the system by \mathcal{T} so that the coordinates are centered on the target set. Regardless of what is chosen for \mathcal{T} , in the transformed coordinate system the target is always the origin.

As mentioned in the preceding section, a challenge inherent in controlling any biological system is that the magnitude of the control signal is constrained. Mathematically, we formalize this constraint as

$$|u(t)| \leq \mu \quad (6.6)$$

where $\mu > 0$ is given, and is determined by the limitations imposed by the biology or by the available hardware for the experimental implementation. To make the forthcoming analytical optimizations simple, we scale the input by the value of the constraint by introducing

$$\tilde{u} = \frac{u}{\mu}. \quad (6.7)$$

Equation (6.3) then becomes

$$\dot{x} = F(x) + b\tilde{u} \quad (6.8)$$

where

$$b = \begin{bmatrix} \frac{\mu}{K_s} \\ 0 \end{bmatrix}. \quad (6.9)$$

By defining an interval $\tilde{\mathcal{U}} = [-1, 1]$, we can restate our objective – find the control signal $\tilde{u}(t) \in \tilde{\mathcal{U}} \times [0, \infty)$ that drives the system from its initial state x_0 to the target set \mathcal{T} (which is the origin $x = 0$ in our translated coordinate framework) in minimum time. In the optimal control literature, this objective is often termed *minimum-time-to-reach*. It should be noted that such a control signal is not guaranteed to exist. As the magnitude constraint, μ , approaches zero it becomes impossible to reach any point except for those on the periodic orbit.

6.4 Optimal stimulus design

The optimal stimulus design method we present follows an approach similar to that outlined in [67] as applied to the classical double integrator example of [4]. The optimal control signal is found by first computing an approximation of the

minimum-time-to-reach value function, which is a viscosity solution of a Hamilton-Jacobi-Bellman (HJB) PDE. The numerical approximation of the value function is then used to generate the optimal state trajectory and the optimal control signal by forward simulation. We now elaborate on the theoretical details of this method.

We begin by defining the terminal time, $t_{\text{end}} \in [0, \infty]$, which is the minimum time at which the state reaches the target set when starting from $x(0) = x$ under the control signal $\tilde{u}(t)$. This is written precisely as

$$t_{\text{end}}(x, \tilde{u}(t)) = \min\{t : x(t) \in \mathcal{T} | x(0) = x\}. \quad (6.10)$$

The terminal time can be infinite when the trajectory $x(t)$ never reaches the target set. This can occur if the constraint in (6.6) is so small that the controller does not have enough control authority to over-ride the system's natural autonomous dynamics. The terminal time is not known at the outset, and is only found through calculating the optimal stimulus and optimal state trajectories.

We now define the *cost* functional, J , of an $(x(t), \tilde{u}(t))$ trajectory starting at $x(0) = x$ (over the time interval $t \in [0, t_{\text{end}}]$) as

$$J(x, \tilde{u}(t)) = \int_0^{t_{\text{end}}} g(x(t), \tilde{u}(t)) dt + q(x(t_{\text{end}})). \quad (6.11)$$

Here, $g(x(t), \tilde{u}(t)) \geq 0$ is the time-additive component of the cost function, and $q(x(t_{\text{end}}))$ is the terminal component. We present Equation (6.11) in full generality to be consistent with the classical way of presenting this theoretical background. However for our control objective, which corresponds to minimizing the time to reach the target point, these functions take very simple forms:

$$g(x, \tilde{u}) = 1 \quad (6.12)$$

and

$$q(x(t_{\text{end}})) = 0. \quad (6.13)$$

Substituting these functions, Equation (6.11) simplifies to

$$J(x, \tilde{u}(t)) = t_{\text{end}}(x, \tilde{u}(t)) \quad (6.14)$$

which is precisely equivalent to our objective of reaching the target set in minimum time. Other functions could be used for $g(x(t), \tilde{u}(t))$ and $q(x(t_{\text{end}}))$ and would imply different control objectives. Consider, for instance, $g(x, \tilde{u}) = 1 + \tilde{u}^2$. The control objective would now be a combination of minimizing t_{end} and minimizing power of the control signal.

We compute the minimum-time-to-reach *value* function, $\mathcal{V}(x)$, from state x at time $t = 0$ to the target set as

$$\mathcal{V}(x) = \inf_{\tilde{u}(t) \in \tilde{\mathcal{U}}} J(x, \tilde{u}) = \inf_{\tilde{u}(t) \in \tilde{\mathcal{U}}} t_{\text{end}}(x, \tilde{u}(t)). \quad (6.15)$$

The value function can take infinite value at points in state space from which the controller cannot drive the state to the target set.

A fact from classical optimal control theory is that, in the minimum-time-to-reach framework, the value function $\mathcal{V}(x)$ is a viscosity solution of the following Hamilton-Jacobi-Bellman equation:

$$0 = \min_{\tilde{u} \in \tilde{\mathcal{U}}} \{1 + \nabla \mathcal{V}(x) \cdot (F(x) + b\tilde{u})\} \quad (6.16)$$

with the boundary condition

$$\mathcal{V}(x) = 0 \quad \forall x \in \mathcal{T}. \quad (6.17)$$

This equation represents a critical step in the theoretical foundation of our optimal control design process. We refer the reader to [5] for a detailed derivations and proofs of Equation (6.16), and an introduction to viscosity solutions of PDEs. If we can find a function $\mathcal{V}(x)$ that satisfies Equation (6.16) in the viscosity sense, the rest of the control design is quite straightforward. Unfortunately, equations of this form can be quite difficult to solve, and no general methods exist.

We can, however, analytically simplify Equation (6.16) by realizing that the only the last term, $\nabla \mathcal{V}(x) \cdot b\tilde{u}$, is involved in the minimization. The value of \tilde{u} that minimizes Equation (6.16) is typically called the ‘‘H-minimal’’ control, or \tilde{u}_H [4]. The vector b has only a single nonzero entry, so we can easily perform the minimization to find

$$\tilde{u}_H = -\text{sgn}(\nabla \mathcal{V}(x) \cdot b) = -\text{sgn}\left(\frac{\partial \mathcal{V}}{\partial x_1}\right), \quad (6.18)$$

since μ and K_s are strictly positive. We see from Equation (6.18) that $\tilde{u}_H \in \{-1, +1\}$ which means the optimal control magnitude will always be equal to its maximum value. In fact, the solutions to all minimum-time optimal control problems are of this ‘‘bang-bang’’ type, as follows from Pontryagin’s Minimum Principle [80].

We substitute (6.18) into (6.16), expand into components, and recall that for a real number y , $-y \text{sgn}(y) = -|y|$. This removes the minimization functional to yield a slightly simpler expression of the Hamilton-Jacobi-Bellman equation:

$$0 = 1 + \nabla \mathcal{V}(x) \cdot F(x) - |\nabla \mathcal{V}(x) \cdot b|. \quad (6.19)$$

The solution $\mathcal{V}(x)$ of Equation (6.19), in the viscosity sense, enables the computation of the optimal state-feedback (and ultimately event-based open-loop) policies \tilde{u}^* and the corresponding state trajectories x^* for any given starting state x . Given $\mathcal{V}(x)$, the state-feedback form of the optimal control policy $\tilde{u}^*(x)$ is simply equal to the H-minimal control evaluated at x .

$$\begin{aligned} \tilde{u}^*(x) &= \arg \min \{1 + \nabla \mathcal{V}(x) \cdot (F(x) + b\tilde{u})\} \\ &= \tilde{u}_H(x) = -\text{sgn}\left(\frac{\partial \mathcal{V}}{\partial x_1}\Big|_x\right). \end{aligned} \quad (6.20)$$

The optimal state trajectories satisfy the system dynamics driven by the optimal state-feedback control law:

$$\begin{aligned}\dot{x}^*(t) &= F(x^*(t)) + b\tilde{u}^*(x^*(t)) \\ &= F(x^*(t)) - b \operatorname{sgn} \left(\frac{\partial \mathcal{V}}{\partial x_1} \Big|_{x^*(t)} \right).\end{aligned}\tag{6.21}$$

We can then compute the open-loop optimal control signal $\tilde{u}^*(x_0, t)$ for all $t \in [0, t_{\text{end}}(x_0)]$ by simulating Equation (6.21) starting from initial position $x(0) = x_0$ until x reaches the target set \mathcal{T} at time $t = t_{\text{end}}(x_0) = \mathcal{V}(x_0)$. The simulation provides the optimal x trajectory, which we use to calculate the optimum control through Equation (6.20). Thus, given any initial condition x_0 we have all the necessary tools to calculate a variable-time-length open-loop control signal $\tilde{u}^*(x_0, t)$.

In our event-based framework, only the spiking state, shown as a blue square in Figure 2.2, is observable and its detection triggers the execution of the precomputed open-loop optimal control signal $\tilde{u}^*(x_{\text{spike}}, t)$

The crucial component of this machinery is an accurate solution or approximation of the cost function $\mathcal{V}(x)$ over the entire domain of x . Nonlinear Hamilton-Jacobi-Bellman PDEs are notoriously difficult to solve, especially in stationary form as Equation (6.19).

6.5 Numerical methods

Our goal is to solve Equation (6.19), with boundary conditions in (6.17), to obtain an approximation to the minimum-time-to-reach value function $\mathcal{V}(x)$. This PDE is both stationary and discontinuous, which renders its solution non-trivial. Computational tools exist for solving time-dependent PDEs, in particular [68] provides a package that uses level set methods to solve HJB equations of the following form (among others):

$$0 = \frac{\partial \phi}{\partial t}(x, t) + H(x, t, \phi, \nabla \phi).\tag{6.22}$$

Following [67, 76], we seek an auxiliary form of our original HJB Equation (6.19) and (6.17) that converts them into a quasi-time-dependent form.

We define a function

$$G(x, \nabla \mathcal{V}(x)) = 1 + \nabla \mathcal{V}(x) \cdot F(x) - |\nabla \mathcal{V}(x) \cdot b|\tag{6.23}$$

which allows us to write our HJB equation (6.19) as

$$\begin{aligned}G(x, \nabla \mathcal{V}(x)) &= 0 \quad \text{on } \mathcal{D} \setminus \partial \mathcal{T} \\ \mathcal{V}(x) &= 0 \quad \text{on } \partial \mathcal{T},\end{aligned}\tag{6.24}$$

where $\partial \mathcal{T}$ is the boundary of our target set \mathcal{T} , and \mathcal{D} is the spatial domain.

We introduce a time-like variable $s \in \mathbb{R}_{\geq 0}$, an auxiliary function $\phi(x, s)$, and perform the following coordinate transformation:

$$\begin{aligned}\mathcal{V}(x) &\leftarrow s \\ \nabla \mathcal{V}(x) &\leftarrow \frac{\nabla \phi(x, s)}{\frac{\partial \phi}{\partial s}(x, s)}\end{aligned}\tag{6.25}$$

which is valid as long as the following technical condition is satisfied [76]:

$$\sum_{i=1}^d p_i \frac{\partial G(x, p)}{\partial p_i} \neq 0 \text{ on } \partial\mathcal{T}. \quad (6.26)$$

This condition essentially states that the boundary conditions must be non-characteristic.

This allows us to cast the Hamilton-Jacobi-Bellman Equation (6.24) into the form of Equation (6.22) which is compatible with ToolboxLS [67]. The resulting auxiliary PDE is

$$0 = \frac{\partial \phi}{\partial s}(x, s) + (\nabla \phi(x, s) \cdot F(x) - |\nabla \phi(x, s) \cdot b|), \quad (6.27)$$

with boundary conditions

$$\begin{aligned} \phi(x, 0) &= 0 & \text{for } x \in \partial\mathcal{T} \\ \phi(x, 0) &< 0 & \text{for } x \in \mathcal{T} \setminus \partial\mathcal{T} \\ \phi(x, 0) &> 0 & \text{for } x \in \mathcal{D} \setminus \mathcal{T}. \end{aligned} \quad (6.28)$$

Our new objective is to solve (6.27) for $\phi(x, s)$. Then we can extract the cost function

$$\mathcal{V}(x) = \{s \mid \phi(x, s) = 0\}. \quad (6.29)$$

Solving HJB PDEs of this type using ToolboxLS requires a target set of non-zero area, and our target is always the origin, so we approximate the target by a small ε -ball (disk) around the origin.

An approximation of $\phi(x, s)$ is solved by creating a uniform grid over a rectangular region of state space enclosing the periodic orbit. ToolboxLS solves (6.27) using a Lax-Friedrichs scheme for the spatial dimension and a Runge-Kutta time stepping scheme. The value function $\mathcal{V}(x)$ is extracted, using Equation (6.29), by interpolation in the quasi-time variable s .

6.6 Results

Numerical computations were implemented in the scaled and translated x -coordinate system using ToolboxLS. However, for clarity of presentation, all following results will be presented in the original coordinates, and control signals will be plotted as un-scaled control currents in mA (recall $I(t) = u(t)/C = \mu \tilde{u}(t)/C$).

6.6.1 Phase randomization: single neuron results

We seek a phase-randomizing control scheme, so we choose the target set to be the unstable equilibrium point inside the periodic orbit, shown as the black X in Figure 2.2. A particular challenge of the reduced Hodgkin-Huxley model is that this unstable fixed point is near the periodic orbit. This situation requires the radius of

the target set approximating circle ε to be quite small, which in turn drives the need for a fine grid on which to compute the solution of the Hamilton-Jacobi-Bellman PDE. The following set of calculations uses $\varepsilon = 0.005$ and $\mu = 10$ mA on a 151×151 grid, uniformly-spaced in the scaled x -coordinates.

Figure 6.1 shows uniformly-spaced level sets of the resulting approximation of the value function surface over the domain of interest. Recall that the value of this surface $\mathcal{V}(x)$ at a point x is the minimum time it will take the system to reach the target set under optimal control. At the spike point, the value is 6.88 which indicates that when the event-based controller is activated, it will drive the neuron to the target set in approximately 6.88 msec.






Figure 6.1: Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple circle). The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, n) using a control signal bounded by $\mu = 10$.

The plot of the state-feedback form of the optimal control law is shown in Figure 6.2(a). As mentioned in Section 6.4, the optimal control law takes only two values, $\{-\mu, +\mu\}$, with the sign determined by simply checking the sign of the value function's x_1 -directional derivative, per (6.18). The optimal trajectory, shown as a blue dashed line in Figure 6.2(a), was computed by simulating the dynamical system beginning at the spike point shown as the blue square. At each time step, the optimal control signal (which depends only on the current location in state-space) is applied to the system. Under influence of the optimal control signal, the system initially follows the upper lobe of its periodic orbit, then begins to deviate to the left as it approaches the lower left knee of the orbit, at which point it cuts across the orbit to reach the target set in approximately 6.88 msec.

Figure 6.2(b) shows the time profile of the optimal trajectory, as well as the optimal control signal, which was simply recorded point-by-point during the simulation. This optimal control signal will be the open loop waveform that is used by the event-based controller to randomize the phase of the neuron.




(a)




(b)

Figure 6.2: (a) State-feedback optimal control law (grayscale background) with the periodic orbit (red line), optimal state trajectory (blue dashed line) and the target set (purple circle). Gray regions signify maximum control in the negative direction ($I = -10$ mA), and white regions signify maximum control in the positive direction ($I = +10$ mA). (b) Optimal voltage (black line) and gating variable (green line) trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, starting from the spike point, in minimal time ($t_{\text{end}} \approx 6.88$ msec).



(a) Control Off



(b) Control On

Figure 6.3: Ensemble asymptotic phase results for 10^5 uncoupled neurons, each starting from the spiking point. In (a), neurons are exposed only to unit-magnitude noise. In (b), the neurons are also exposed to the phase randomizing control signal.

6.6.2 Phase randomization: ensemble results

In the previous section, we showed that the optimal control signal does, in fact, drive the system to the target set. We now examine the effectiveness of this control scheme as a phase randomizer. Consider an ensemble of 10^5 identical uncoupled neurons, now each under the influence of independent identically distributed noise $D\eta(t)$, where we set $D = 1$ mA. Each neuron is simulated, starting at the spiking state, and the time of the next spike is recorded as t_{next} . This spike time is converted to phase by the following equation:

$$\theta = \frac{2\pi(t_{\text{next}} \bmod T)}{T}, \quad (6.30)$$

where the mod operator handles the case where the neuron is driven so close to the unstable equilibrium point that it slowly spins around and away from it for longer than the natural period before returning to the stable periodic orbit.

Figure 6.3(a) shows a histogram representing the distribution of phases caused by the noise process $D\eta(t)$ without any control input. Note the tight unimodal shape of this distribution is to be expected given the attractive strength of the periodic orbit and that D is small.

The simulations are repeated, but now we add the optimal control signal as input to each neuron, in addition to the background noise. The resulting phase histogram is shown in Figure 6.3(b). The control scheme greatly widens the phase distribution of the ensemble after a single application. If the phase randomization was perfect, we would expect a uniform distribution, but in Figure 6.3(b) we see a non-uniform bimodal distribution. This arises from the fact that we have approximated the phaseless set, which is a single point, by a small ε -ball. Due to the geometry of the ODE's vector field around the target set, optimal trajectories tend to approach this ε -ball non-uniformly, which for this particular case results in a bimodal distribution. The precise shape of the distribution is governed by the underlying ODE, the value of ε , the magnitude of the noise D , and the magnitude constraint μ .

Complete phase randomization after a single application of the optimal control signal is not necessary. As long as the phase distribution is significantly spread out, we can simply apply the control signal multiple times. We will illustrate this in the context of a globally-coupled network in Section 6.7.

6.6.3 Spike rate control

By applying the optimal minimum-time control design method using the spike point as a target, the firing rate of the neuron can be increased. Targeting the spike point is equivalent to asking how to make the neuron spike again as soon as possible given the constraint on the control magnitude. Here, we use $\mu = 10$ and $\varepsilon = 0.01$.

Figure 6.4 shows the level sets of the value function for this example. The state feedback control law and optimal state trajectories are shown in Figures 6.5(a) and

6.5(b). The optimal control signal, shown in Figure 6.5(b), is interesting in that it is not simply using the maximum magnitude control in a single direction, which would be equivalent to increasing the baseline current. To optimally increase the firing rate of this neuron, the controller switches between $+\mu$ and $-\mu$. We see that the optimal state trajectory, the blue dashed line on Figure 6.5(a), stays quite close to the unperturbed periodic orbit, so the control is not radically altering the dynamics, just speeding up the spiking from its uncontrolled natural period of 11.84 msec to 7.74 msec.





Figure 6.4: Level sets of the numerical approximation of the value function \mathcal{V} for optimal spike-rate-increasing control. Superimposed are the periodic orbit (red line) and the target set (purple ellipse). Again, we consider $\mu = 10$.


6.6.4 Connection to isochrons

The same numerical techniques used to calculate the optimal spike rate increasing control stimulus can be used to visualize an approximation of the model's isochrons in a straightforward manner. If we were to set the stimulus magnitude constraint $\mu = 0$ and repeat the optimal spike rate increase stimulus design procedure discussed in the previous section, we would be exactly calculating isochrons. To understand why this is the case, consider the meaning of the value function. For the control objective of driving the system to the spiking point, the value function is the time it takes, under optimal stimulus, for the system to be driven from position x to the spiking point. Now, if the control magnitude bounds are zero, we have eliminated the controller's ability to affect the dynamics of the system. So the value function in the case of $\mu = 0$ can be interpreted as the time it takes, under no stimulus, for the system to be driven from position x to the spiking point. Level sets of this value function are, therefore, curves connecting points in phase space that will reach the spiking point at the same time, in the absence of stimulus, i.e. the isochrons.

There are numerical difficulties with setting $\mu = 0$ so we create approximations to isochrons by calculating level sets of the value function for the case of small μ .



(a)



(b)

Figure 6.5: (a) State-feedback optimal control law for optimal spike rate increase. Notation in this figure is the same as Figure 6.2. (b) Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom). Here, the objective is to reach an ε -ball centered at the spike point in minimal time starting from the spike point. This control protocol decreases the interspike interval neuron from 11.84 msec to 7.74 msec, resulting the fastest spiking rate possible with a stimulus constrained by 10 mA.

Figure 6.6 shows the isochron approximation calculated using $\mu = 1$. This compares well with the isochrons shown on Figure 2.2 that were calculated accurately using backwards integration techniques as in [22, 58].




Figure 6.6: Level sets approximating isochrons for the reduced Hodgkin-Huxley model generated using parameters $\varepsilon = 0.01$, $\mu = 1$.

6.6.5 FitzHugh-Nagumo model

The proposed method is not dependent on the details of the particular model, such as how close the unstable fixed point is to the periodic orbit. To illustrate the broad applicability of the optimal stimulus design method, we select another classical neural model, the FitzHugh-Nagumo model [37, 71] in the dimensionless form outlined in Section 2.2.3

We take the time scale separation variable $\delta = 0.01$, and set the parameter $a = 0.6$ to position the unstable fixed point near the center of the region enclosed by the periodic orbit, which is encircled by the purple ε -radius target circle in Figure 6.8(a). Its worth noting that this neuron model has very different time, input, and voltage scaling than the Hodgkin-Huxley model. Since the unstable equilibrium point is far inside the periodic orbit, this model would be particularly difficult to control by using a timed impulse to instantaneously kick the state from the limit cycle to unstable equilibrium, as proposed in [109]. By employing the control synthesis method presented above, we can precisely calculate the optimal control signal to drive the FitzHugh-Nagumo model from its spiking point to an ε disk centered at the unstable fixed point. We take this target radius to be $\varepsilon = 0.01$, compute the value function, and plot its level sets in Figure 6.7. Figures 6.8(a) and 6.8(b) show the optimal trajectories and the optimal control stimulus when the control magnitude constraint is $\mu = 0.2$.




Figure 6.7: Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple line) for the FitzHugh-Nagumo model. The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, w) using a control signal bounded by $\mu = 0.2$.

6.7 Extension: globally coupled networks

We now apply our control scheme in a more biologically relevant context. We consider a network of N neurons with all-to-all electrotonic coupling in a noisy environment.

By applying the optimal phase randomizing control stimulus, we seek to desynchronize the spike times of the network. Following [54], we write dynamical equations for a network of reduced Hodgkin-Huxley neurons obeying (2.7) as


$$\dot{x}_i = F(x_i) + e_1(D\eta_i + I(t))/C + \frac{k}{N}M_1 \sum_{j=1}^N (x_j - x_i) \quad (6.31)$$

where $x_i = [V_i, n_i]^T$ is the state vector of the i^{th} neuron, η_i is uncorrelated white noise input to the i^{th} neuron, and D is the magnitude of the noise. The vector $e_1 = [1, 0]^T$ and the matrix


$$M_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

account for the fact that the neurons are stimulated and coupled only through the voltage component of the state. The variable k is the coupling strength, which we take to be small. Here, $I(t)$ is the control stimulus which is applied to all neurons simultaneously. Again we use an event-based feedback framework, but in the network context we need a population-level event definition. To this end, we introduce the concept of a network observable - a quantity from which we can infer the level of synchrony in the network. We use the average voltage as our network observable, written simply as

$$\bar{V}(t) = \frac{1}{N} \sum_{i=1}^N V_i(t). \quad (6.32)$$



(a)



(b)

Figure 6.8: (a) State-feedback optimal control law, optimal state trajectory, and the target set for the FitzHugh-Nagumo model with $\mu = 0.2$. Notation in this figure is the same as in Figure 6.2. (b) Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, in minimal time ($t_{\text{end}} \approx 0.5$).

Were these neurons modeled as phase oscillators, we could simply use Kuramoto’s order parameter, as in [27]. In this context, however, phase is unknown - only voltages are measurable.

We envision a simplified representation of electrical deep brain stimulation with two electrodes. One electrode is an input that stimulates each neuron with the same control signal $I(t)$. The other electrode is an output that measures the average voltage $\bar{V}(t)$ of the population (a simplification of local field potential where we neglect spatial weighting). We define an *event* as the value of the network observable crossing some threshold, indicating that most of the neurons are spiking at the same time. This threshold is a tunable parameter, generally greater than zero for a network of reduced Hodgkin-Huxley neurons.

In the uncontrolled system with uniform initial distribution, the neurons become synchronized due to coupling, even in the presence of moderate noise. One can view this as a simple example of pathological neural synchronization, linked to the symptoms of neural disorders such as Parkinson’s disease [10] and epilepsy [91].

To demonstrate the effectiveness of the proposed control scheme, we simulate the network represented by Equation (6.31) for $N = 100$ neurons, each starting synchronized at the spike point. This corresponds to the pathological spike synchronization described above. The synchronous spiking is detected by a super-threshold value of $\bar{V}(t)$, our network observable. Super-threshold average voltage constitutes an event, which triggers the controller to stimulate all the neurons with the optimal control waveform calculated in Section 6.6.1. Recall that in the absence of coupling, this stimulus causes the neurons’ states to approach the phase randomizing unstable fixed point. If the coupling is relatively weak, the controller is still able to drive the neurons close enough to the unstable equilibrium for the phase randomization to be effective. The neurons’ phases are randomized by the stimulus, thus diminishing the value of the network observable $\bar{V}(t)$. The controller then shuts off until the coupling draws the network back toward synchronous spiking ($\bar{V}(t) \geq V_{\text{thresh}}$) at which point the optimal stimulus is again triggered. We note that if the system does not develop synchronized spiking again, the controller will continue to remain off, thus achieving demand-controlled desynchronization similar to [81].

Figure 6.9(a) shows results for a population of 100 neurons globally coupled with uniform coupling strength $k = 0.01$ and each under the influence of independently drawn white noise with magnitude $D = 1.0$. We have kept the control magnitude constraint $\mu = 10$ mA and used the exact optimal stimulus waveform calculated in Section 6.6.1.

In this simulation, we begin with the controller switched off, to illustrate the synchronous spiking of the network. We turn the controller on at $t = 20$ msec, but recall that this is an event-based controller, so it remains inactive until the population level observable, the average voltage, crosses the threshold V_{thresh} . The first population-level spiking event detected by the controller occurs at $t \approx 23$ msec, at which point the open-loop optimal stimulus waveform is triggered. After a single application, there is still a significant level of spiking synchrony, closely related to the bimodal distribution caused by the randomization described in Section 6.6.2.

The controller is triggered again at $t \approx 37$ msec, and again the open-loop stimulus is activated. This time, the spike times of the population are randomized to a much greater degree, which can be easily visualized by constructing a raster plot of the spike times across the population, shown as Figure 6.9(b).

We note that the resulting distribution of spike times is not uniform; there are loose vertical “bands” of spiking activity in the raster plot after the control is turned on. This is due to two factors – the coupling, and the approximation of the target set by a disk rather than a point (which is related to the change in distribution due to a single application of the controller, as discussed above in Section 6.6.2). However, the pathology is due primarily to synchronized spiking, shown as the tightly synchronized vertical stripes in the raster plot before the controller is activated. The event-based controller destroys this type of synchronization, which may provide therapeutic utility for pathological conditions, such as Parkinson’s disease, that involve neuronal synchronization.


6.8 Discussion

In this chapter, we have developed an event-based control system that can drive a neuron’s state to reach a target set in minimum time under stimulus magnitude constraints. By choosing the unstable equilibrium as the target set, we have shown that this control system can effectively randomize the phase of an oscillatory neuron. If computed using the spike point as the target set, this control protocol can maximally increase the firing rate for a given stimulus magnitude constraint.


We have illustrated this method using two-dimensional neuron models so that the resulting value function can be easily visualized as level sets in the plane, but we remark here that the method is not limited to planar systems. While the computational time for the Hamilton-Jacobi-Bellman PDE grows with dimension, it is still quite feasible to consider models with higher dimensions, such as the four-dimensional Hodgkin-Huxley model. Furthermore, by targeting the spiking point and setting the magnitude constraint to be small (ideally zero), the numerical viscosity solutions of the Hamilton-Jacobi-Bellmann PDE can be used to visualize isochrons – a useful tool for characterizing the phase space of spiking neuron ODE models.

We have also shown that the control system can be extended to a network of globally coupled neurons. In this extended form, the event-based controller is triggered by a network observable. In the case of weak coupling, we have shown that this controller can successfully mitigate pathologically synchronized spiking.

An important concern with this design method is that it depends strongly on having a good model for the neuron, which in practice, can be challenging. We hope the model identification procedure outlined in Section 2.6 could be a useful starting point for the development of an accurate model for the real neurons of interest.



(a)



(b)

Figure 6.9: (a) Network results for $N = 100$ globally coupled Hodgkin-Huxley neurons, $k = 0.01$, $D = 1.0$, $\mu = 10$. Top panel shows voltage traces for each neuron. Middle panel shows the control signal applied to the population. Controller is activated at $t = 20$ msec indicated by the vertical dashed line. Bottom panel is the population-level observable $\bar{V}(t)$ with the threshold $V_{\text{thresh}} = 0$ shown as a red dashed line. (b) Raster plot of spike times for the network showing desynchronization.

Chapter 7

Master-slave Control: Pairwise Spike Timing

We propose an event-based feedback control scheme to stabilize the spike time difference between the controlled (slave) neuron and a reference (master) neuron. The natural frequencies of the two neurons need not be identical, but must satisfy a given inequality. The methods proposed in this chapter are different from those in Chapter 5 in that it is not necessary for the controller to know the entire phase response curve. Instead, the experimentalist chooses a fixed phase, usually the maximal point, for a single pulse stimulation. The control scheme then determines what magnitude the pulse should be, based on the last spike times of the master and slave neurons.

This chapter is organized as follows. Introductory material is presented in Section 7.1. Master and slave neuron models are discussed in Section 7.2, along with the concept of a spike advance curve, which is an alternative to a phase response curve. Our control objective is specified in Section 7.3. The algorithms developed to achieve the desired spike timing are presented in Section 7.4. We walk through a numerical example of the algorithms in Section 7.4.3. Finally, we give concluding remarks in Section 7.5.

7.1 Introduction and background

During the course of recent experimental collaborations [94], we have been introduced to the problem of pairwise spike timing control. This problem arises from a natural extension of single neuron spike timing control. The experimental setup is as follows. There are two periodically spiking neurons, a master which is not controllable, and a slave which is controllable (but which could have a different natural frequency).

Our goal is to realize a specified time interval between the spikes of the master and the slave neuron. A challenge, based on certain experimental considerations, is that the controller does not know the phase response curve. The experimental

implementation is only able to measure a spike advance curve that relates the spike time advance as a function of stimulus magnitude for a given monophasic current pulse applied at a given phase after spiking. Essentially, the experimentalist chooses a phase at which maximum spike advance is achieved for a fixed magnitude pulse. Then, by varying the magnitude of the pulse, the so-called *spike advance curve* can be measured for the particular neuron.

Knowing only the natural frequencies of the two neurons and the spike advance curve for the slave neuron, we propose a nested pair of event-based algorithms that drive the two neurons to spike with a specified interspike interval drawn from an prescribed admissible interval.

7.2 Master-slave neuron models

Master neuron

The master neuron is a simple clock oscillator that is deterministic and is not affected by the slave neuron. Experimentally, we envision the master neuron as a phase oscillator spiking periodically in an isolated apparatus. The master neuron's dynamics satisfy the following equations:

$$\dot{\theta}_m = \omega_m, \quad \theta_m(0) = \theta_{m,0}. \quad (7.1)$$

In the framework of this problem, the master neuron's spikes ($\theta = 0$ -crossings) are detectable by the controller. The natural period of the master neuron is simply

$$T_m = \frac{2\pi}{\omega_m}. \quad (7.2)$$

Slave neuron

The slave neuron is a nonlinear phase oscillator with a input sensitivity function $Z(\theta_s, I)$ that produces the spike advance function when combined with the stimulus protocol. The function $Z(\theta_s, I)$ will not be used explicitly, and is only included to provide a connection to the methods discussed in Chapter 5 that involve phase response curves. The slave neuron evolves along trajectories satisfying

$$\dot{\theta}_s = \omega_s + Z(\theta_s, I)I/c, \quad \theta_s(0) = 0. \quad (7.3)$$

In this study, it is assumed that the slave neuron is always on a strongly attracting periodic orbit, so that the phase model holds exactly. We note here that the master neuron is not coupled to the slave neuron, and neither neuron affects the other's evolution. In the absence of input, the slave's natural period is

$$T_s = \frac{2\pi}{\omega_s}. \quad (7.4)$$

The slave neuron's spikes are detectable by the controller, and constitute the *event* that activates the event-based control algorithm used to stabilize a given interspike time interval between the master and the slave neuron.

Spike advance curves

The spike advance curve is an experimentally measurable characteristic of the slave neuron. It is essentially a convolution of the input sensitivity function $Z(\theta_s, I)$ and the (fixed) shape of the stimulus pulse. The stimulus protocol administers a monophasic current pulse timed for the optimal phase θ_c , which is usually at or near the point at which the neuron's phase response curve is maximal (meaning the neuron is most sensitive to input). For example, if the controlled neuron were a Hodgkin-Huxley neuron with baseline current $I_b = 10$ mA, then $\theta_c = \beta$ as shown in Figure 2.4. Even though the entire phase response curve may not be known, it is experimentally feasible to assess which phase is most sensitive to stimulation. This way of representing the system is quite robust, because if the chosen value of θ_c is actually not quite the true maximal point, the algorithm will still work, only it will take slightly more control energy than if the true maximal phase was used.




Figure 7.1: Typical spike advance curve $f(I)$ invertible over (I_{\min}, I_{\max}) . This function characterizes the slave neuron's response to a stimulus of magnitude I applied at the maximal phase θ_c .

Figure 7.1 shows a caricature of a typical spike advance function. We will restrict the domain of our spike advance curve to ensure the resulting function is invertible. Currently, we are using only a small linear region [94]. The input current magnitude is denoted by I and the resulting spike (time) advance by Δs . We will need to truncate the spike advance curve, since it is only invertible over a limited (interior) region. With some abuse of notation, we will denote this truncated spike time advance curve by

$$f(I) : (I_{\min}, I_{\max}) \mapsto (\Delta s_{\min}, \Delta s_{\max}). \quad (7.5)$$

Here, $|\Delta s_{\min}|$ corresponds to the maximum possible spike delay and Δs_{\max} corresponds to the maximum possible spike advance. These values are critical because they determine the maximum possible discrepancy between the period of the master and the slave, which must satisfy

$$T_s - \Delta s_{\max} \leq T_m \leq T_s + |\Delta s_{\min}|. \quad (7.6)$$

Our algorithm requires that T_m and T_s satisfy Eq. (7.6). The $f(I)$ function was truncated to be a one-to-one function, so there exists an inverse function $f^{-1}(\Delta s)$ which allows us to calculate which value of input stimulus magnitude in the interval (I_{\min}, I_{\max}) that will result in a desired spike advance Δs . This inverse function will be an important component of the proposed event-based algorithm.

7.3 Control objective

Like the other control strategies presented in this dissertation, the master-slave algorithm is an event-based control scheme. The controller is triggered on spikes of the slave neuron, and stimulates using an open-loop pulsatile waveform computed at the time of the event (slave neuron spike). When the slave neuron spikes, we will estimate the phase of the master neuron based on its last spike time and known period. In theory, this estimation is exact because we are not considering noise. We find the next several master spike times and add the desired interspike interval, Δt_D (in units of time). This forms the set of target spike times we want our slave neuron to spike at. The challenge here is the fact that we cannot achieve any desired advance, only advances within the interval $(\Delta s_{\min}, \Delta s_{\max})$.

Conceptually, we break the control scheme into two stages. The first objective to drive the slave neuron to spike at the same time as the master, plus the desired spike time offset. Once this objective is achieved, the second stage of the control simply applies the necessary advance or delay each period to effectively adjust the firing period of the slave to equal the firing period of the master. Unless the two natural periods are equal, this will require a fixed stimulus to the slave neuron each period. It is from this second stage of control that the constraints on the natural periods, depending on the maximum advance and delay, is derived as in (7.6). We note here that this two-stage framework is only used for conceptual explanation—the algorithms presented here treat every spiking event identically and the different conceptual staged are handled implicitly by the internal mathematics of the algorithms.

7.4 Algorithms

7.4.1 Pre-processing

Before the event-based control is activated, some quantities that will be used in the online algorithm can be pre-computed. The first is the positive integer i_m^* , which represents the maximum number of slave control iterations necessary to hit a desired spike time:

$$i_m^* = \left\lceil \frac{T_m}{\Delta s_{\max} + |\Delta s_{\min}|} \right\rceil \quad (7.7)$$

where $\lceil \cdot \rceil$ is the round-up operator. The rationale behind this notion is that each target spike time will be T_m time units away from the next, and each time the

control is activated (at $t = t_0$, the spike of the slave neuron), the controller can adjust the timing of the next spike to be anywhere in the time interval $(t_0 + T_s - \Delta s_{\max}, t_0 + T_s + |\Delta s_{\min}|)$. So in n control periods, it can adjust the timing of the n^{th} slave spike to be anywhere in the interval $(t_0 + nT_s - n\Delta s_{\max}, t_0 + nT_s + n|\Delta s_{\min}|)$. For sufficiently large $n \leq i_m^*$, the following inequality holds:

$$n(\Delta s_{\max} + |\Delta s_{\min}|) \geq T_m. \quad (7.8)$$

This means that the control “window” is (eventually) larger than the period of the target spikes - a situation that guarantees success in achieving a single target spike time within this time interval. As stated before, after a single target spike time is achieved, we then take the difference of the interspike intervals $T_s - T_m$ as our desired spike advance.

7.4.2 Event-based algorithm

At each event (slave spike), the controller needs to determine whether to drive the neuron to achieve the maximum spike time advance, maximum delay, or some intermediate spike advance. This decision task is accomplished by two nested algorithms: **event_control**, which is the high-level driver; and **spike_advance**, which is called to calculate the desired spike time advance each control period. We assume that our implementation-level control system tracks current time in the variable t , and the most recent spike time of the master neuron in the variable t_{last} .

Algorithm 1 **event_control**(t, t_{last})

- 1: $t_0 = t$
 - 2: $\theta_m(t_0) = \frac{2\pi(t_0 - t_{\text{last}})}{T_m}$
 - 3: $\Delta s = \mathbf{spike_advance}(t_0, \theta_m(t_0))$
 - 4: $I = f_s^{-1}(\Delta s)$
 - 5: apply pulse with amplitude I at time $t_0 + \frac{\theta_c}{2\pi}T_s$
-

Explanation–Algorithm 1: **event_control**

The event-based control system is triggered when the slave neuron spikes. Recall that the controller already has pre-computed the variables discussed in Section 7.4.1. At the instant the controller detects a spike from the slave neuron, we sample the current time, t , and the time of the last recorded master neuron spike, t_{last} . What follows is a line-by-line explanation of this algorithm.

Line 1: Assign the current value of t to the variable t_0 . This will serve as a time offset for future calculations.

Line 2: Estimate the phase of the master neuron based on its last recorded spike time.

Line 3: Call the **spike_advance** algorithm to determine the optimal slave spike time advance to command.

Algorithm 2 spike_advance($t_0, \theta_m(t_0)$)

```
1: initialize  $i = 1, \text{term} = 0$ 
2:  $\mathcal{I}_t = \left\{ \begin{array}{l} t_t = t_0 + [2\pi - \theta_m(t_0)] T_m / 2\pi + \Delta t_D + k T_m, \\ k \in \mathbb{Z}_{\geq 0} \mid t_0 < t_t < t_0 + i_m^* T_s \end{array} \right\}$ 
3:  $\mathcal{I}_C = (t_0 + T_s - \Delta s_{\max}, t_0 + T_s + |\Delta s_{\min}|)$ 
4: if  $\mathcal{I}_t \cap \mathcal{I}_C \neq \emptyset$  then
5:    $\Delta s = T_s - (\min \{\mathcal{I}_t \cap \mathcal{I}_C\} - t_0)$ 
6: else
7:    $i++$ 
8:   while  $\text{term} == 0$  do
9:      $\mathcal{I}_A = (t_0 + i T_s - i \Delta s_{\max}, t_0 + i T_s)$ 
10:     $\mathcal{I}_D = (t_0 + i T_s, t_0 + i T_s + i |\Delta s_{\min}|)$ 
11:    if  $\mathcal{I}_t \cap \mathcal{I}_A \neq \emptyset$  then
12:       $\Delta s = \Delta s_{\max}, \text{term} = 1$ 
13:    else if  $\mathcal{I}_t \cap \mathcal{I}_D \neq \emptyset$  then
14:       $\Delta s = \Delta s_{\min}, \text{term} = 1$ 
15:    else
16:       $i++$ 
17:    end if
18:  end while
19: end if
20: return  $\Delta s$ 
```

Line 4: Calculate the current amplitude needed for the stimulus to cause the desired spike advance.

Line 5: Apply the stimulus pulse when the controller expects the slave neuron's phase to be at the optimum stimulus point, θ_c .

Explanation–Algorithm 2: spike_advance

This algorithm is called from the **event_control** algorithm, and is used to compute the desired spike advance. This function would be trivial if the controller could command any arbitrary spike advance. Since the choice of spike advance is restricted to the range $(\Delta s_{\min}, \Delta s_{\max})$, we must take care in deciding what to do if we cannot reach the first desired spike time in one iteration of control. The dynamic information passed to this function includes the current slave spike time offset, t_0 , and the estimated phase of the master, $\theta_m(t_0)$. We will now describe each line of this algorithm.

Line 1: Initialize an integer counter, i , and a boolean flag, **term**.

Line 2: Calculate a set of target spike times, \mathcal{T}_t . These are based on when the master neuron will spike next, and what the desired time interval is between the master and slave spikes, Δt_D . The set is truncated using our estimate of the maximum possible control periods, i_m^* , necessary to reach any T_m -periodic spike train, as calculated in Equation (7.7).

Line 3: Calculate the first interval of control, \mathcal{I}_C . This represents the amount the controller can change the next slave spike time by applying up to the maximum advance or delay in this control period.

Line 4: Check to see if a target spike time lies within the first interval of control. If not, this will return \emptyset (the *null* or *empty* set).

Line 5: If the intersection of the set of target spike times, \mathcal{T}_t , and the first interval of control, \mathcal{I}_C , is non-empty, choose the first target spike time in the intersection set and compute the necessary spike advance to apply this control period. At this point, the Line 4 **if** statement is done—proceed to Line 20.

Line 6: If the intersection of \mathcal{T}_t and \mathcal{I}_C is empty, we cannot achieve our control objective in one control period. We must now enter into an iterative portion of the algorithm that will decide whether to apply the maximum spike advance Δs_{\max} or maximum spike delay Δs_{\min} .

Line 7: Increment the counter, i . This counter tracks how many (hypothetical) control periods must be used to result in enough cumulative spike advance or delay for a slave spike to coincide with a master spike.

Line 8: This **while** loop runs until we have found how many periods of successive maximal advances or delays is necessary to achieve our objective—codified by setting the terminal flag, **term**, to 1 (true). This **while** loop is guaranteed to terminate within i_m^* iterations.

Line 9: Compute the i^{th} advance control time interval. This represents the interval of control that can be achieved in i control periods by applying the maximum advance, Δs_{\max} , each time.

$T_s = 1$	$\theta_c = \frac{3\pi}{2}$
$T_m = 1.2$	$\theta_s(0) = 0^-$
$\Delta s_{\max} = 0.1$	$\theta_m(0) = \frac{3\pi}{2}$
$\Delta s_{\min} = -0.3$	$\Delta t_D = \frac{T_m}{3} = 0.4$
$f(x) = x$	$f^{-1}(x) = x$

Table 7.1: Values used for the numerical example in Figure 7.2.

Line 10: Compute the i^{th} delay control interval. This represents the interval of control that can be achieved in i control periods by applying the maximum delay, Δs_{\min} , each time.

Line 11: Check to see if any target spike times lie within this i^{th} interval of maximally advancing control, \mathcal{I}_A .

Line 12: If the intersection of the set of target spike times, and the i^{th} interval of maximally advancing of control is non-empty, this means that by continuing to apply to maximum spike advance, Δs_{\max} , we can eventually achieve our control objective. So set $\Delta s = \Delta s_{\max}$, and switch the boolean termination flag, **term**, to 1 (this will get us out of the **while** loop).

Line 13: Check to see if any target spike times lie within this i^{th} interval of maximally delaying control, \mathcal{I}_D .

Line 14: If the intersection of the set of target spike times, and the i^{th} interval of maximally delaying of control is non-empty, this means that by continuing to apply to maximum spike delay, Δs_{\min} , we can eventually achieve our control objective. So set $\Delta s = \Delta s_{\min}$, and switch the boolean termination flag, **term**, to 1 (this will get us out of the **while** loop).

Lines 15-16: If both interval intersection checks return the null set, increment the counter i by one and return to Line 8.

Line 20: The algorithm will return the desired spike advance Δs to the **event_control** algorithm that called it.

7.4.3 Numerical example

We will now present a simple numerical example to illustrate the inner details of these algorithms. Figure 7.2 is provided to assist in understanding this example.

Setup

The example will use numerical values listed in Table 7.1. First, check the relationship between the period of the two neurons.

$$\begin{array}{rcccl} T_s - \Delta s_{\max} & \leq & T_m & \leq & T_w + |\Delta s_{\min}| \\ 0.9 & \leq & 1.2 & \leq & 1.3 \end{array}$$




Figure 7.2: Example master-slave control scenario. Slave phase, θ_s , shown as solid blue line with spike events as blue ‘x’ markers. Master phase, θ_m , shown as red dashed line with red ‘+’ markers. Target slave spike times (master spike times offset by Δt_D) shown by green ‘o’ markers. The thin black dotted line is the optimal control phase, $\theta_c = 3\pi/2$, which is where we apply the stimulus pulse. In this example, the slave spikes converge to the desired phase offset (relative to the master spikes) after two periods of control. Then the controller simply adjusts the period of the slave each time to match that of the master.

The inequality is satisfied, so we may proceed. We will now compute i_m^* , the maximum number of control iterations necessary to achieve the desired spike timing difference

$$i_m^* = \left\lceil \frac{1.2}{0.1 + 0.3} \right\rceil = 3.$$

We have assumed that we start time t exactly when the slave neuron spikes, so $\theta_s(0) = 0^-$. By 0^- , we mean that the neuron’s phase is just crossing zero from the other side as time begins.

So the first event is triggered immediately. The **event_control** algorithm activates. The controller samples t and sets $t_0 = 0$, per Line 1.

The phase of the master neuron is known from the initial condition listed in Table 7.1 and no time has elapsed, so $\theta_m(t_0) = \theta_m(0) = 3\pi/2$. Now the **event_control** algorithm calls the **spike_advance** algorithm with arguments $t_0 = 0$ and $\theta_m(0) = 3\pi/2$.

After initializing $i = 1$ and term = 0, the **spike_advance** algorithm calculates the set of target spike times for this control period

$$\mathcal{T}_t = \{0.7, 1.9\}.$$

The first control interval is calculated in Line 3

$$\mathcal{I}_C = (0.9, 1.3).$$

We find $\mathcal{T}_t \cap \mathcal{I}_C = \emptyset$ in Line 4, which means the controller cannot hit the target spike time during this period of control. So we proceed to Lines 6 and 7 and increment

our counter, $i = 2$. We now enter into the **while** loop at Line 8. The first task is to compute the two maximal control intervals:

$$\mathcal{I}_A = (1.8, 2)$$

and

$$\mathcal{I}_D = (2, 2.6).$$

Line 11 checks and finds $\mathcal{T}_t \cap \mathcal{I}_A = 1.9 \neq \emptyset$, so we proceed to Line 12, which sets $\Delta s = \Delta s_{\max} = 0.1$, as well as switching the boolean flag, **term** = 1, which ends the **while** loop. The **spike_advance** algorithm returns $\Delta s = 0.1$ to the main **event_control** algorithm.

The **event_control** algorithm implements the desired spike advance by applying a pulse of magnitude $I = f^{-1}(\Delta s) = 0.1$ at time $t = 0.75$. This stimulus will cause the slave neuron to spike next at $t = t_0 + T_s + \Delta s = 0.9$. When the slave neuron spikes at $t = 0.9$, **event_control** is activated again. The algorithm samples t and sets $t_0 = 0.9$. It also reads the last recorded master spike, which happened at $t = 0.3$, which is used to calculate $\theta_m(t_0) = \pi$ in Line 2. Again, the algorithm calls the **spike_advance** algorithm with input data $t_0 = 0.9$ and $\theta_m(t_0) = \pi$. At Line 2, the target spike time set is computed

$$\mathcal{T}_t = \{1.9, 3.1\}.$$

The first control interval is calculated

$$\mathcal{I}_C = (1.8, 2.2).$$

Now at Line 4, we find that the intersection is non-empty

$$\mathcal{T}_t \cap \mathcal{I}_C = 1.9$$

so we proceed to Line 5 to calculate the exact spike advance needed to hit the target spike time

$$\Delta s = T_s - (\min \{\mathcal{T}_t \cap \mathcal{I}_C\} - t_0) = 1 - (1.9 - 0.9) = 0.$$

This means that, since we chose these simple numbers for the example, the slave neuron's next spike will coincide exactly with the master neuron if we use no stimulus whatsoever. Had the numbers been slightly different, we could have made a slight advance or delay here. However, this saves a step, since at the next event, we will have the slave neuron synchronized to its target trajectory. This is not the end of the example, though. Because $T_m = T_s + 0.2$, we need to apply a spike advance of -0.2 every period to maintain the synchrony to the target spike train. We will illustrate how the algorithms provide this functionality.

The next slave spike will occur at $t = 1.9$. The algorithm **event_control** is triggered, sets $t_0 = 1.9$, and reads $t_{\text{last}} = 1.5$. Next, we calculate

$$\theta_m(t_0) = \frac{2\pi}{3}.$$

The `spike_event` algorithm is called and calculates

$$\mathcal{T}_t = \{3.1, 4.3\}.$$

Also,

$$\mathcal{I}_C = (2.8, 3.2).$$

We see that $\mathcal{T}_t \cap \mathcal{I}_C = 3.1 \neq \emptyset$, so Line 5 calculates

$$\Delta s = 1 - (\min \{3.1\} - 1.9) = -0.2$$

and `spike_advance` returns $\Delta s = -0.2$ as expected. This is the steady state case and the two neurons will continue firing in this correct pattern, with the controller commanding a $\Delta s = -0.2$ each time, so that the next spike will occur at $t = 3.1$, the preceding spike will occur at $t = 4.3$, and so on.

7.5 Discussion

The control methodology presented in this chapter is a result of ongoing collaborations with the Netoff laboratory at the University of Minnesota. Many important aspects of this problem are still not understood. The issue of noise and stochasticity is particularly important in the context of experimental neuroscience. Real neurons do not follow the simple deterministic dynamics used in the construction of the algorithms presented here. Initial studies have, however, yielded encouraging results, some of which were presented in [94]. A full investigation is the subject of a forthcoming study.

We note here that this framework may be fairly straightforward to extend to more complicated variants of this problem. Multiple instances of this control system could be implemented to any number of slave neurons, and by setting each slave i 's spike time offset $\Delta t_{D,i} = 0$, we can control synchrony in the ensemble, or by setting $\Delta t_{D,i} = (i - 1)/T_m$, we can desynchronize the ensemble. Once the algorithm has been rigorously tested in the pairwise context, it will be applied to larger ensembles.

Chapter 8

Conclusion

Oscillatory neurons play important roles in dynamical diseases of the nervous system, such as Parkinson's disease. In this dissertation, we have studied the dynamical behavior of such neurons, both individually and in ensemble. We considered conductance-based ODE models, as well as several types of reduced-order models, such as nonlinear phase oscillator models. Modeling neurons as phase oscillators allows one to explore the complex behavior arising from the nonlinear response to stimulus as a function of phase. We presented a set of event-based feedback control schemes, based on reduced model dynamics, that exploit the nonlinearities to achieve spike timing control.

In Chapter 2, we considered a variety of issues related to the modeling of oscillatory neurons. Beginning with the conductance-based Hodgkin-Huxley formalism, we introduced several reduced models including a planar representation of the Hodgkin-Huxley model, the closely related FitzHugh-Nagumo model, and single dimension phase-reduced models. Also included in this chapter were tools to estimate model parameters from experimental data.

Chapter 3 presented results from a study of ensemble response to independent random stimuli consisting of impulses with interspike intervals drawn from a Poisson distribution. This model of background noise has been shown to be a reasonable representation of the actual neural background activity in mammalian biology [89]. We provided a theoretical framework, based on the Kramers-Moyal expansion, for calculating the population-level partial phase synchronization due to this background noise. It was shown by simulation that the theoretical predictions are closely matched by numerical results for both full-dimensional conductance-based models and phase-reduced models. We developed a phase sampling algorithm to enable such a comparison. Taking advantage of the computational efficiency of phase models, we performed a parametric study of partial phase synchronization of an ensemble of neurons using a large range of mean spike rates and noise magnitudes. The resulting phase distributions were shown to respond quite differently to input than a uniform distribution. We concluded Chapter 3 with a discussion of how phase synchrony is related to spiking synchrony, with the important result being the two concepts are not equivalent. Simple examples were developed wherein

phase synchrony can be made arbitrarily high, as measured by the magnitude of the Kuramoto order parameter, but yet the ensemble can have no spike synchronization whatsoever. We also showed how an ensemble can have complete spike synchronization in the absence of phase synchronization, except at the instantaneous spike times (a measure zero set in time).

We considered oscillators that are parametrically forced in Chapter 4. In particular, we studied a nonlinear oscillator model with strong 2:1 resonance. Results were presented for the single oscillator case, then extended to populations with weak coupling. The Implicit Function Theorem was used to prove the persistence of solutions from the uncoupled case. We categorized the solutions based on symmetry considerations and detail bifurcation diagrams that track the size of solutions as a function of forcing frequency. Chaotic trajectories were seen in the case of strong negative coupling. The effect of network topography is illustrated by considering a network of three oscillators that morphs from a line to a ring.

We switched our focus to feedback control starting with Chapter 5, which considers spike timing control in the context of reference phase tracking. Event-based control algorithms were developed, first for individual neurons. We characterized stability with respect to a reference-tracking objective and showed how biologically inspired impulsive control schemes can achieve this objective. Impulses, however, are not physically realizable, so we proposed a quasi-impulsive algorithm that uses biphasic pulses of finite magnitude and nonzero duration. Conditions for global asymptotic stability were presented, and depend on the shape of the neuron's underlying phase response curve. Extensions of the impulsive and quasi-impulsive control schemes were presented that are applicable to simple networks, such as an uncoupled ensemble and a pacemaker-driven ensemble.

Chapter 6 presented an event-based phase-randomizing control scheme based on the Hamilton-Jacobi-Bellman framework for minimum-time optimal control. The state space target was chosen based on the Winfree's idea [109] of driving the oscillator to an unstable equilibrium point that is a phaseless set. When the state is near this point, its asymptotic phase is extremely sensitive to noise. The Hamilton-Jacobi-Bellman framework enables the design of a controller that drives the neuron to this target in minimum time using a constrained magnitude stimulus. The resulting control scheme is an experimentally-realizable technique to randomize the asymptotic phase of an oscillatory neuron under realistic laboratory conditions. Through the development of the control synthesis procedure, various extensions were investigated including optimal spike rate increase and an interesting connection to isochron visualization using intermediate numerical results. A population level extension of the control scheme was presented that was shown to effectively randomize the phases of a network of globally coupled neurons initialized in synchrony.

Chapter 7 presents ongoing research on a pairwise spike timing control protocol. This problem was encountered during the course of experimental collaboration and has resulted in the development of a nested pair of event-based control algorithms that enable the stabilization of a desired interspike time interval between an un-

controlled master neuron and a controlled slave neuron. The algorithms presented in this chapter are distinct from those presented in Chapter 5 because here the two neurons do not necessarily have the same natural frequency. If the two neurons have commensurate frequencies, the master-slave algorithm can drive the interspike interval between the two neurons to a desired value chosen from an admissible set.

The studies presented in this dissertation are but a subset of the wide breadth of research opportunities in these areas. Building on the ensemble results of Chapter 3 and the parametric response results of Chapter 4, it would be interesting to consider the response of a neuronal ensemble to common sinusoidal stimulus. Sinusoids and similar periodic waveforms are among those employed by tradition EDBS, and have been studied by many in the mathematical neuroscience community. However, the phase sampling tools presented in Section 3.4 could provide valuable insight into how higher dimensional conductance based neurons compare to the phase reduced populations that are often studied. Also, in this dissertation, we have focused primarily on synchronization from the perspective of its pathological role in dynamical disease. It has been suggested that spike timing could also play an important role in cognition and higher brain function [51], and it would be interesting to explore the possible effects of partial phase synchrony in this context. Coupling and network heterogeneity are two other compelling issues worth investigating in the context of neuronal phase distribution and its physiological consequences.

From a controls perspective, there are many important considerations to be addressed by future research. The effects of stochasticity in terms of process and measurement noise would be natural extensions worth exploring, in relation to the feedback control schemes presented in Chapters 5, 6, and 7. Robustness, though non-trivial to define in the nonlinear context, would be an interesting characteristic to investigate for these control schemes. The natural variation in model parameters seems like an ideal place to start along these lines. In terms of experimental application to physiological systems, it will be important to develop control schemes that take network structure and heterogeneity into account. A tractable first step may be to consider a Hamilton-Jacobi-Bellman style minimum-time control scheme, like the one presented in Chapter 6, for small numbers of coupled heterogeneous neurons. There are no theoretical barriers, although the complexity of computations grows exponentially with state dimension. Results for two or three coupled neurons are certainly within reach, and may provide valuable insight to the related problem of cardiac arrhythmia.

The algorithms presented in this dissertation rely on precise individual neuron measurements and the ability to stimulate each neuron independently. In reality, it may only be possible to measure and stimulate a small fraction of the neurons involved in a particular pathology. This raises important questions about controllability and observability that may be difficult to assess given the nonlinearity of the underlying neuronal dynamics and the complexity of the network topology. Indeed, a unified theoretical framework for controlling such complex neuronal systems is still a largely untouched research objective. Such theoretical developments, however, are certainly worth pursuing, as they could ultimately lead to a control

synthesis procedure. Designing optimal feedback controllers for EDBS systems that use micro-electrode arrays to stimulate and measure many neurons simultaneously may be an important application for the emerging field of network control engineering. The motivation for this line of research is the opportunity to make a positive impact on the health and well-being of the many patients living with dynamical diseases of the nervous system, including Parkinson's disease and epilepsy.

Bibliography

- [1] P. Ashwin and J. Swift. The dynamics of n weakly coupled identical oscillators. *Journal of Nonlinear Science*, 2:69–108, 1992.
- [2] K. J. Åström and B. Bernhardsson. Systems with Lebesgue sampling. In A. Rantzer and C. I. Byrnes, editors, *Directions in Mathematical Systems Theory and Optimization*, volume XIII. Springer, 2003.
- [3] K. J. Åström and B.M. Bernhardsson. Comparison of Riemann and Lebesgue sampling for first order stochastic systems. In *Proceedings of the 41st IEEE Conference on Decision and Control*, pages 2011–2016, Las Vegas, NV, 2002.
- [4] M. Athans and P. L. Falb. *Optimal Control: An Introduction to the Theory and its Applications*. McGraw-Hill New York, 1966.
- [5] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhauser, 1997.
- [6] J. Bélair, L. Glass, U. an der Heiden, and J. Milton. Dynamical disease: identification, temporal aspects and treatment strategies of human illness. *Chaos*, 5:1, 1995.
- [7] J. Bélair, L. Glass, U. an der Heiden, and J. Milton. *Dynamical Disease: Mathematical Analysis of Human Illness*. AIP Press Woodbury, NY, 1995.
- [8] I. Bena, C. Van den Broeck, R. Kawai, M. Copelli, and K. Lindenberg. Collective behavior of parametric oscillators. *Physical Review E*, 65:036611, 2002.
- [9] I. Bena and C. Van den Broeck. Coupled parametric oscillators. *Europhysics Letters*, 48:498–504, 1999.
- [10] A. L. Benabid, P. Pollak, C. Gervason, D. Hoffmann, D. M. Gao, M. Hommel, J. E. Perret, and J. De Rougemont. Long-term suppression of tremor by chronic stimulation of the ventral intermediate thalamic nucleus. *The Lancet*, 337:403–406, 1991.
- [11] M. Bennett, M. F. Schatz, H. Rockwood, and K. Wiesenfeld. Huygen’s clocks. *Proceedings of the Royal Society A*, 458:563–579, 2002.

- [12] A. Beuter and K. Vasilakos. Tremor: Is Parkinsons disease a dynamical disease? *Chaos*, 5(1):35–42, 1995.
- [13] A. Boose, C. Jentgens, S. Spieker, and J. Dichgans. Variations on tremor parameters. *Chaos*, 5(1):52–56, 1995.
- [14] P. Bressloff. Resonantlike synchronization and bursting in a model of pulse-coupled neurons with active dendrites. *Journal of Computational Neuroscience*, 6(3):237–249, 1999.
- [15] P. C. Bressloff and S. Coombes. Synchronization of synaptically coupled neural oscillators.
- [16] Y. Bromberg, M. C. Cross, and R. Lifshitz. Response of discrete nonlinear systems with many degrees of freedom. *Physical Review E*, 73:016214, 2006.
- [17] E. Brown, P. Holmes, and J. Moehlis. Globally coupled oscillator networks. In E. Kaplan, J. Marsden, and K. R. Sreenivasan, editors, *Perspectives and Problems in Nonlinear Science: A Celebratory Volume in Honor of Larry Sirovich*, pages 183–215. Springer-Verlag, New York, 2003.
- [18] E. Brown, J. Moehlis, and P. Holmes. On the phase reduction and response dynamics of neural oscillator populations. *Neural Computation*, 16:673–715, 2004.
- [19] E. Brown, J. Moehlis, P. Holmes, E. Clayton, J. Rajkowski, and G. Aston-Jones. The influence of spike rate and stimulus duration on noradrenergic neurons. *Journal of Computational Neuroscience*, 17:13–29, 2004.
- [20] E. Buks and M. L. Roukes. Electrically tunable collective response in a coupled micromechanical array. *Journal of Microelectromechanical Systems*, 11:802–807, 2002.
- [21] R.J. Butera, J. Rinzel, and J.C. Smith. Models of respiratory rhythm generation in the pre-Botzinger complex. I. Bursting pacemaker neurons. *Journal of Neurophysiology*, 82(1):382–397, 1999.
- [22] A. Campbell, A. Gonzalez, D. L. Gonzalez, O. Piro, and H. A. Larrondo. Isochrones and the dynamics of kicked oscillators. *Physica A*, 155(3):565 – 584, 1989.
- [23] W. T. Coffey, Y. P. Kalmykov, and J. T. Waldron. *The Langevin Equation: With Applications to Stochastic Problems in Physics, Chemistry, and Electrical Engineering*. World Scientific, Singapore, 2004.
- [24] A.H. Cohen, S. Rossignol, and S. Grillner. *Neural Control of Rhythmic Movements in Vertebrates*. Wiley, 1988.

- [25] J. A. Connor, D. Walter, and R. McKown. Neural repetitive firing: Modifications of the Hodgkin-Huxley axon suggested by experimental results from crustacean axons. *Biophysical Journal*, 18:81–102, 1977.
- [26] H. Daido. Generic scaling at the onset of macroscopic mutual entrainment in limit cycles with uniform all-to-all coupling. *Physical Review Letters*, 73:760–763, 1994.
- [27] P. Danzl, R. Hansen, G. Bonnet, and J. Moehlis. Partial phase synchronization of neural populations due to random Poisson inputs. *Journal of Computational Neuroscience*, 25(1):141–157, 2008.
- [28] P. Danzl, J. Hespanha, and J. Moehlis. Event-based minimum-time control of oscillatory neuron models. *Biological Cybernetics*, 101(5):387–399, 2009.
- [29] P. Danzl and J. Moehlis. Event-based feedback control of nonlinear oscillators using phase response curves. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 5806–5811, New Orleans, LA, 2007.
- [30] P. Danzl and J. Moehlis. Spike timing control of oscillatory neuron models using impulsive and quasi-impulsive charge-balanced inputs. In *Proceedings of the 2008 American Control Conference*, pages 171–176, Seattle, WA, 2008.
- [31] P. Danzl and J. Moehlis. Weakly coupled parametrically forced oscillator networks: existence, stability, and symmetry of solutions. *Nonlinear Dynamics*, 59(4):661–680, 2010.
- [32] J.E. Dennis and R.B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, 1996.
- [33] E. Doedel, A. Champneys, T. Fairgrieve, Y. Kuznetsov, B. Sandstede, and X. Wang. *AUTO 97: Continuation and bifurcation software for ordinary differential equations*. <http://indy.cs.concordia.ca/auto/>, 1997.
- [34] R.J. Elble and W.C. Koller. *Tremor*. Johns Hopkins University Press, 1990.
- [35] G. B. Ermentrout. Type I membranes, phase resetting curves, and synchrony. *Neural Computation*, 8:979–1001, 1996.
- [36] G. B. Ermentrout. *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to XPPAUT for Researchers and Students*. Philadelphia, 2002.
- [37] R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961.
- [38] R. W. Friedrich and M. Stopfer. Recent dynamics in olfactory population coding. *Current Opinion in Neurobiology*, 11:468–474, 2001.

- [39] R. F. Galán, Nicolas Fourcaud-Trocmé, G. B. Ermentrout, and N. N. Urban. Correlation-induced synchronization of oscillations in olfactory bulb neurons. *Journal of Neuroscience*, 26(14):3646–3655, 2006.
- [40] Wulfram Gerstner. Time structure of the activity in neural network models. *Physical Review E*, 51(1):738–758, Jan 1995.
- [41] D. Goldobin and A. Pikovsky. Collective modes in parametrically excited oscillator arrays. *Europhysics Letters*, 59:193–198, 2002.
- [42] D. S. Goldobin and A. Pikovsky. Synchronization and desynchronization of self-sustained oscillators by common noise. *Physical Review E*, 71:045201, 2005.
- [43] D. S. Goldobin and A. Pikovsky. Synchronization of self-sustained oscillators by common white noise. *Physica A*, 351:126–132, 2005.
- [44] C. M. Gray, P. König, A. K. Engel, and W. Singer. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337, 1989.
- [45] J. Guckenheimer. Isochrons and phaseless sets. *Journal of Mathematical Biology*, 1:259–273, 1975.
- [46] J. Guckenheimer and P. J. Holmes. *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1983.
- [47] B. S. Gutkin, G. B. Ermentrout, and A. D. Reyes. Phase response curves give the responses of neurons to transient inputs. *Journal of Neurophysiology*, 94:1623–1635, 2005.
- [48] D. Hansel, G. Mato, and C. Meunier. Synchrony in excitatory neural networks. *Neural Computation*, 7(2):307–337, 1995.
- [49] A.V.M. Herz, T. Gollisch, C.K. Machens, and D. Jaeger. Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *Science*, 314(5796):80–85, 2006.
- [50] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117:500–544, 1952.
- [51] F. Hoppensteadt and E. Izhikevich. *Weakly Connected Neural Networks*. Springer-Verlag, New York,, 1997.
- [52] C. Huygens. Instructions concerning the use of pendulum-watches for finding the longitude at sea. *Philosophical Transactions of the Royal Society of London*, 4:937, 1669.

- [53] E. M. Izhikevich. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. Springer, New York, 2007.
- [54] J. Keener and J. Sneyd. *Mathematical Physiology*. Springer, New York, 1998.
- [55] H. K. Khalil. *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 2002.
- [56] W. Klimesch. Memory processes, brain oscillations and EEG synchronization. *International Journal of Psychophysiology*, 24:61–100, 1996.
- [57] N. Kopell. We got rhythm: Dynamical systems of the nervous system. *Notices of the AMS*, 47(1):6–16, 2000.
- [58] J. Koscic, E. T. Shea-Brown, and J. Moehlis. Isochron. *Scholarpedia*, page 2509, 2006.
- [59] Y. Kuramoto. *Chemical Oscillations, Waves, and Turbulence*. Springer, Berlin, 1984.
- [60] R. Lifshitz and M. C. Cross. Response of parametrically driven nonlinear coupled oscillators with application to micromechanical and nanomechanical resonator arrays. *Physical Review B*, 67:134302, 2003.
- [61] J.C. Lilly, G.M. Austin, and W.W. Chambers. Threshold movements produced by excitation of cerebral cortex and efferent fibers with some parametric regions of rectangular current pulses (cats and monkeys). *Journal of Neurophysiology*, 15(4):319–341, 1952.
- [62] J.C. Lilly, J.R. Hughes, E.C. Alvord Jr, and T.W. Galkin. Brief, noninjurious electric waveform for stimulation of the brain. *Science*, 121(3144):468–469, 1955.
- [63] Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1506, 1995.
- [64] D.R. Merrill, M. Bikson, and J.G.R. Jefferys. Electrical stimulation of excitable tissue: design of efficacious and safe protocols. *Journal of Neuroscience Methods*, 141(2):171–198, 2005.
- [65] D. C. Michaels, E. P. Matyas, and J. Jalife. Mechanisms of sinoatrial pacemaker synchronization - a new hypothesis. *Circulation Research*, 61:704–714, 1987.
- [66] R.E. Mirollo and S.H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [67] I. M. Mitchell. A toolbox of level set methods. Technical Report UBC CS TR-2007-11, 2007.

- [68] I. M. Mitchell. The flexible, extensible and efficient toolbox of level set methods. *Journal of Scientific Computing*, 35(2):300–329, 2008.
- [69] J. Moehlis. Canards for a reduction of the Hodgkin-Huxley equations. *Journal of Mathematical Biology*, 52(2):141–153, 2006.
- [70] J. Moehlis. On the dynamics of coupled parametrically forced oscillators. *Proceedings of 2008 ASME Dynamic Systems and Control Conference*, 2008. DSCC2008-2189.
- [71] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962.
- [72] H. Nakao, K. Arai, K. Nagai, Y. Tsubo, and Y. Kuramoto. Synchrony of limit cycle oscillators induced by random external impulses. *Physical Review E*, 72:026220, 2005.
- [73] A. H. Nayfeh and D. T. Mook. *Nonlinear Oscillations*. Wiley-Interscience, New York, 1979.
- [74] K. M. Newell, F. Gao, and R. L. Sprague. The dynamical structure of tremor in tardive dyskinesia. *Chaos*, 5:43, 1995.
- [75] K. Okuda. Variety and generality of clustering in globally coupled oscillators. *Physica D*, 63:424–436, 1993.
- [76] S. Osher. A level set formulation for the solution of the Dirichlet problem for Hamilton–Jacobi equations. *SIAM Journal on Mathematical Analysis*, 24:1145, 1993.
- [77] D. Pare, R. Curro’Dossi, and M. Steriade. Neuronal basis of the Parkinsonian resting tremor: a hypothesis and its implications for treatment. *Neuroscience*, 35:217–226, 1990.
- [78] D. Paydarfar and D.M. Buerkel. Dysrhythmias of the respiratory oscillator. *Chaos*, 5(1):18, 1995.
- [79] A. Pikovsky, M. Rosenblum, and J. Kurths. *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press, Cambridge, 2001.
- [80] L. S. Pontryagin, K. N. Trirogoff, and L. W. Neustadt. *The Mathematical Theory of Optimal Processes*. Wiley New York, 1962.
- [81] O. V. Popovych, C. Hauptmann, and P. A. Tass. Control of neuronal synchrony by nonlinear delayed feedback. *Biological Cybernetics*, 95(1), 2006.

- [82] J. F. Rhoads, S. W. Shaw, K. L. Turner, J. Moehlis, B. E. DeMartini, and W. Zhang. Generalized parametric resonance in electrostatically actuated microelectromechanical oscillators. *Journal of Sound and Vibration*, 296:797–829, 2006.
- [83] J. Rinzel and R. N. Miller. Numerical calculations of stable and unstable periodic solutions to the Hodgkin-Huxley equations. *Mathematical Biosciences*, 49:27–59, 1980.
- [84] J. Ritt. Evaluation of entrainment of a nonlinear neural oscillator to white noise. *Physical Review E*, 68(4):041915, Oct 2003.
- [85] R. Rose and J. Hindmarsh. The assembly of ionic currents in a thalamic neuron I. The three-dimensional model. *Proceedings of the Royal Society B*, 237:267–288, 1989.
- [86] M. Schaus. Neural oscillator identification via phase-locking behavior. Master’s thesis, University of California, Santa Barbara, 2005.
- [87] S.J. Schiff and T. Sauer. Kalman filter control of a model of spatiotemporal cortical dynamics. *Journal of Neural Engineering*, 5:1–8, 2008.
- [88] W. Singer. Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology*, 55:349–374, 1993.
- [89] W.R. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience*, 13(1):334–350, 1993.
- [90] G. S. Stent. A physiological mechanism for Hebb’s postulate of learning. *Proceedings of the National Academy of Sciences*, 70:997–1001, 1973.
- [91] M. Steriade. *Neuronal Substrates of Sleep and Epilepsy*. Cambridge University Press, 2003.
- [92] M. Steriade, E. G. Jones, and R. R. Llinás. *Thalamic Oscillations and Signaling*. Wiley New York, 1990.
- [93] C. F. Stevens and A. M. Zador. Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience*, 1:210–217, 1998.
- [94] T. Stigen, P. Danzl, J. Moehlis, and T. Netoff. Linear Control of Neuronal Spike Timing Using Phase Response Curves. In *Proceedings of the 31st Annual International Conference of the IEEE EMBS*, pages 1541–1544, Minneapolis, MN, 2009.
- [95] S. Strogatz. *Sync: The Emerging Science of Spontaneous Order*. Hyperion, New York, 2003.

- [96] S. H. Strogatz. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D*, 143:1–20, 2000.
- [97] S.H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Perseus Books, 2001.
- [98] P. A. Tass. *Phase Resetting in Medicine and Biology*. Springer, New York, 1999.
- [99] P. A. Tass. Effective desynchronization by means of double-pulse phase resetting. *Europhysics Letters*, 53:15–21, 2000.
- [100] J. Teramae and D. Tanaka. Robustness of the noise-induced phase synchronization in a general class of limit cycle oscillators. *Physical Review Letters*, 93(20):204103, 2004.
- [101] E. G. Tezak, A. H. Nayfeh, and D. T. Mook. Parametrically excited non-linear multidegree-of-freedom systems with repeated natural frequencies. *Journal of Sound and Vibration*, 85:459–472, 1982.
- [102] R. D. Traub and R. Miles. Pyramidal cell-to-inhibitory cell spike transduction explicable by active dendritic conductances in inhibitory cell. *Journal of Computational Neuroscience*, 2(4):291–298, 1995.
- [103] R. D. Traub, R. Miles, and R. K. S. Wong. Model of the origin of rhythmic population oscillations in the hippocampal slice. *Science*, 243:1319–1325, 1989.
- [104] K. L. Turner, S. A. Miller, P. G. Hartwell, N. C. MacDonald, S. H. Strogatz, and S. G. Adams. Five parametric resonances in a microelectromechanical system. *Nature*, 396:149–152, 1998.
- [105] W. M. Usrey and R. C. Reid. Synchronous activity in the visual system. *Annual Review of Physiology*, 61:435–456, 1999.
- [106] H. U. Voss, J. Timmer, and J. Kurths. Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos*, 14(6):1905–1933, 2004.
- [107] S. Wiggins. *Introduction to Applied Nonlinear Dynamical Systems and Chaos, Second Edition*. Springer, New York, 2003.
- [108] A. Winfree. Patterns of phase compromise in biological cycles. *Journal of Mathematical Biology*, 1:73–95, 1974.
- [109] A. Winfree. *The Geometry of Biological Time, Second Edition*. Springer, New York, 2001.

- [110] A. T. Winfree. Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology*, 16(1):15, 1967.
- [111] W. Zhang, R. Baskaran, and K. Turner. Tuning the dynamic behavior of parametric resonance in a micromechanical oscillator. *Applied Physics Letters*, 82:130–132, 2003.