2011

Noise, Delays, and Resonance in a Neural Network

Austin Quan
Harvey Mudd College

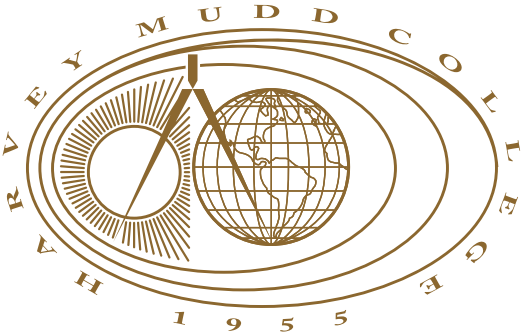

Noise, Delays, and Resonance in a Neural Network

Austin Quan

John Milton, Advisor

Darryl Yong, Reader

May, 2011

HARVEY MUDD
COLLEGE
Department of Mathematics

Abstract

A stochastic-delay differential equation (SDDE) model of a small neural network with recurrent inhibition is presented and analyzed. The model exhibits unexpected transient behavior: oscillations that occur at the boundary of the basins of attraction when the system is bistable. These are known as delay-induced transitory oscillations (DITOs). This behavior is analyzed in the context of stochastic resonance, an unintuitive, though widely researched phenomenon in physical bistable systems where noise can play in constructive role in strengthening an input signal. A method for modeling the dynamics using a probabilistic three-state model is proposed, and supported with numerical evidence. The potential implications of this dynamical phenomenon to nocturnal frontal lobe epilepsy (NFLE) are also discussed.

Acknowledgments

I am ever grateful to John Milton for both giving me an opportunity to work on this project, for putting up with me for this long. A “thank you” is also due Ivan Osorio for funding my work last summer. Finally, a shout out to all my professors, without whose support I would surely have floundered.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A revolution is underway in the field of neuroscience. The study of the nervous system was once the sole domain of experimentalists, but in recent years, there has been an increasing number of scientists who seek to apply more quantitative methods and analysis. The result has been the creation and growth of a new field, known as computational neuroscience, which situates itself between the worlds of experimental biology, nonlinear dynamics, and artificial intelligence.

In this thesis, a stochastic delay differential equation (SDDE) model of a simple two-neuron network is presented and analyzed. This model has previously been analyzed in the context of decision–making (Milton et al., 2010), and belongs to the general class of Hopfield-type models. This particular network is small enough to be mathematically tractable, yet its behavior is much like that of larger networks (Pakdaman et al., 1998b). Our exploration is centered around three topics:

1. *Time delays* are intrinsic to the dynamics of all biological neural networks.

2. *Oscillations* are a consequence of *feedback*, and here we examine two types of periodic behavior.

3. *Noise* is inherent in many biological processes, neural networks included.

We are particularly interested in how these three elements interact to produce different types of resonance in our system.

We place particular focus on a dynamical phenomenon known as delay-induced transitory oscillations (DITOs). These oscillations occur at the

boundary of the basins of attraction of stable fixed points, and have durations significantly longer than the time delay. The analytical mechanism for these transients has been previously examined (Pakdaman et al., 1998a; Gopalsamy and Leung, 1996), but they have not been analyzed in the presence of noise.

Time delayed systems are by definition non-Markovian, which complicates analysis of their stochastic properties. However, here we show that in some cases, the dynamics of these systems can be described using Markovian models. Specifically, we present a three-state probabilistic model that statistically models the structure of our neural network's behavior. This model takes direct influence from similar discrete models used to study bistable systems that exhibit a phenomenon called stochastic resonance. The possible application of these results to the study of "dynamical diseases," specifically nocturnal frontal lobe epilepsy (NFLE), is also discussed.

Chapter 2 introduces the necessary biological background for understanding this model. Chapter 3 details the mathematical concepts involved in this model and its analysis. Chapter 4 presents a delay-differential equation model of two neurons in an inhibitory loop, and performs basic stability analysis. Chapter 5 examines two kinds of stochastic resonance that can occur in the network. Chapter 6 formulates a Markov chain that can be used to model the probabilistic qualities of the system. Chapter 7 discusses the possible biological implications of our work.

Chapter 2

Biological Background: Neural Networks

Neuroscience is primarily concerned with one type of cell, the neuron. What set neurons apart from other cells is their electrical and signalling properties, which can produce a variety of interesting temporal behavior. Furthermore, it allows them to connect to form elaborate networks; the human brain consists of hundreds of billions of neurons, and a single neuron may connect to tens of thousands of others! Unsurprisingly, these networks can have deep and complex dynamics, and significant computational power. Biologically, we see this in the central nervous systems of animals, and theoretically, in the study of artifical neural networks (ANNs), a subfield of artificial intelligence. Here, we provide all that a mathematician needs to know (and no more!) to understand this model.

2.1 Basic Neuroanatomy

The anatomy of a neuron can most simply be viewed as a composition of three parts (Figure 2.1): the soma, dendrites, and the axon. The *soma* is the main cell body, and contains the nucleus. From the soma extend branch-like structures called *dendrites*, which receive chemical signals from other neurons. The *axon* also projects from the soma, and constitutes the majority of the neuron's physical size. Its purpose is to carry electrical impulses, called action potentials, away from the soma. The gap between connected neurons is referred to as the *synapse*.

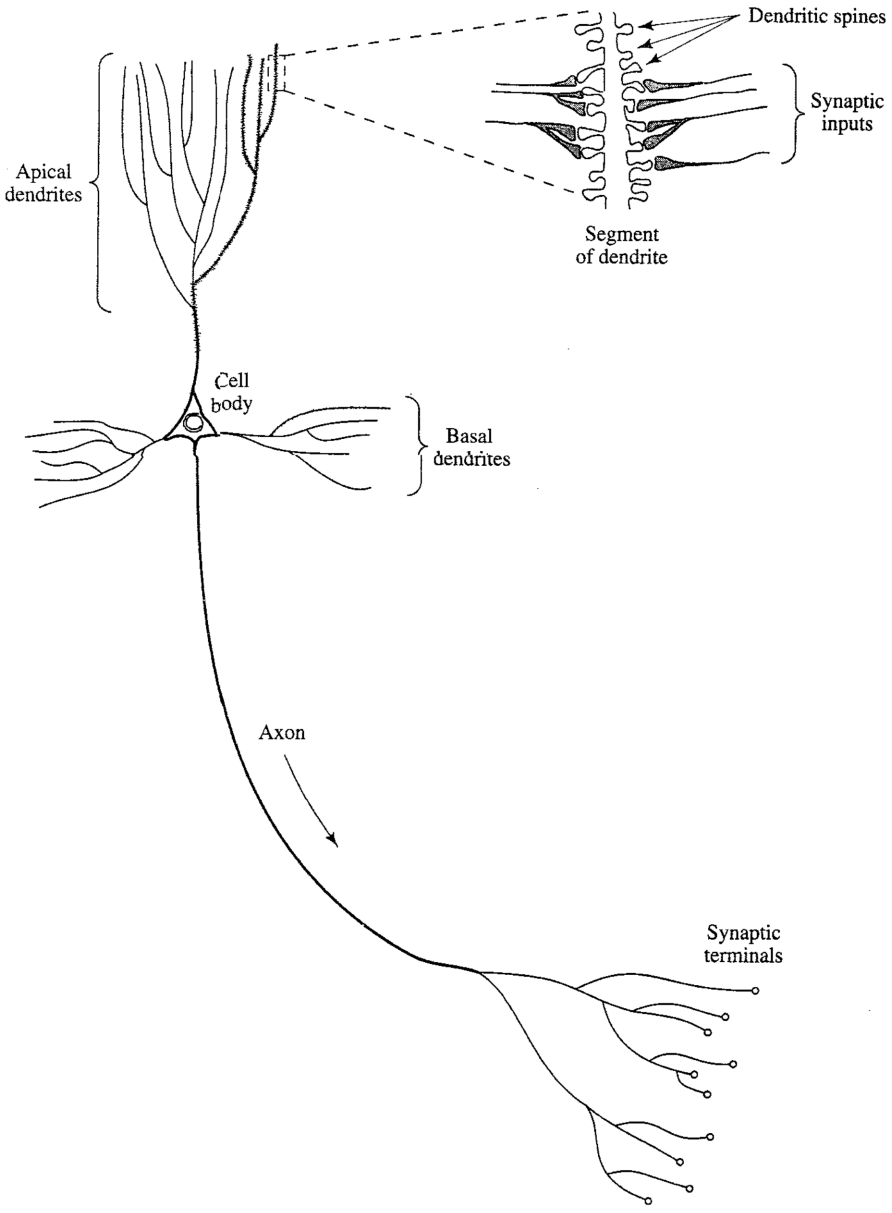

Figure 2.1 A pyramidal neuron. Taken from Haykin (1994).

2.2 Neural Signalling

Signalling between neurons is a complex process, with both electrical and chemical components. A set of chemicals known as *neurotransmitters* are used to pass messages between neurons. Neurotransmitters are released into the synaptic cleft by the *pre-synaptic* neuron. Here, they have the opportunity to bind to receptor sites on the dendrites of the *post-synaptic* neuron, each specific to a neurotransmitter. If a neurotransmitter binds to a receptor, it exerts influence on the behavior of the post-synaptic neuron. This influence is primarily one of two types: excitatory and inhibitory. The binding of an excitatory neurotransmitter causes the neuron to be more likely to spike, or produce an action potential. Equivalently, excitatory neurotransmitters increase the firing rate of the post-synaptic neuron. Likewise, inhibitory neurotransmitters cause the opposite effect: a decrease in neuronal activity. When a neuron fires, an action potential is propagated down the length of its axon. Upon its arrival at the end of the axon, neurotransmitters may be released from *terminal buttons* into various synapses, thus transmitting the signal onward.

An important implication in terms of modelling is that interneural influence is rate-limited. That is, the amount by which one neuron can impact the activity of another is restricted by the number of binding sites on the post-synaptic neuron. Thus, activity in the pre-synaptic neuron will cease to affect the post-synaptic one once it reaches a sufficiently high level.

2.3 Time Delays

Intrinsic to biological neural signalling is the presence a time delay. For example, the fact that the conduction velocity of an action potential is finite introduces a delay between when the pre-synaptic neuron fires, and when the post-synaptic neuron's activity is influenced. Transmission delays are not the only component of the time delay observed in biological nervous systems, there are also cellular and synaptic time delays, among others. For a high-level model, such as the one considered in this document, the cause of the delay is irrelevant; we simply add a delay term to any neural connectivity terms. Many models of neural networks, particularly those used in ANNs, ignore this delay component and treat communication between neurons as instantaneous. As systems of ordinary differential equations (ODE) are much better understood, this is convenient. However, in order to accurately model biological neural dynamics, we must consider delays.

Chapter 3

Mathematical Background: Dynamical Systems and Time Delays

The primary tool for studying dynamics is systems ordinary differential equations (ODEs). Here, we'll discuss their extension to delay differential equations (DDEs), and how the presence of a time delay complicates analysis of such systems. A discussion of stochastic resonance, an interesting dynamical phenomenon that has been found to occur in many bistable systems, is also included. As our model, in the most interesting case, exhibits bistability, the stochastic resonance literature provides a useful theoretical framework for our analysis.

3.1 Dynamical Systems, Time Delays, and Noise

Here we cover the major mathematical components of our model. These elements are not just specific to our system, but mathematical models of biological systems in general.

3.1.1 Delay Differential Equations

A typical dynamical system is defined as a system of autonomous first-order ordinary differential equations,

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}),$$

where $f : \mathbb{R}^n \to \mathbb{R}^n$, and is continuous. $\mathbf{x}$ is a vector of *n state variables* whose dynamics are specified by the *state function*, f. It should be noted that this definition specifies a system of first-order differential equations, but given a system with higher-order differential equations, we can construct an equivalent system consisting only of first-order differential equations. Solutions to such systems are specified by an *initial condition*, or a single point in $\mathbb{R}^n$. See Strogatz (1994) and Hale and Kocak (1991) for a complete treatment of systems of ODEs.

Delay differential equations simply extend ODEs to depend on the state of the system in the past, as well as in the present. Hence,

$$\dot{\mathbf{x}}(\mathbf{t}) = f(x(t), x(t-\tau_1), \ldots, x(t-\tau_n)).$$

One immediate consequence of this is that solutions to time delayed systems require that a continuous *initial function* be defined on the interval $[-\max(\tau_i \quad \forall \quad i), 0]$. Furthermore, delay differential equations cannot be solved backwards in time, only forwards. Accordingly, we refer to a DDE system as defining a *semiflow*.

This modification comes at significant cost. Whereas the solution to an ODE system is determined by a single point in $\mathbb{R}^n$, the state of the system at $t = 0$ (the *initial condition*), DDEs depend on a the state at a point $t = -\tau$ in the past. Thus, the phase space for a DDE is an infinite dimensional Banach space, no matter how small the delay. Furthermore, linear stability analysis often is complicated by the fact that the characteristic equations are often transcendental, and typically have an infinite number of roots.

Despite the fact that time delays are ubiquitous in biological systems, many researchers choose to ignore them as it restricts the set of analytical machinery that can be used, and can significantly complicate the dynamics of the system of interest. If the primary interest is in the asymptotics of a system, then this simplification can be acceptable as the existence and local asymptotic stability of equilibria is in some cases the same. On the other hand, if our interest is in transient behavior, delays need to be taken into consideration. That said, the last decade has seen an increased research interest in time-delayed systems given their biological relevance.

Though time delays are analytically inconvenient, their inclusion can lead to a certain richness in the dynamics. A prime example of this is the Ikeda equation,

$$\dot{x}(t) = \sin(x(t-\tau)).$$

In the absence of the delay, this system's behavior is simple: trajectories quickly converge to the nearest stable fixed point which are located at $\pi +$

$2\pi k$ for $k \in \mathbb{Z}$. When we add a delay, the system exhibits a broad range of chaotic behavior depending on our choice of parameters. Figure 3.1 shows several of these results.

3.1.2 Noise

A stochastic process is one for which the future states follow a probability distribution, as opposed to be restricted to one possible outcome. We call such a process *Markovian* if this probability distribution depends on the current state of the process, and not any others. Most standard techniques used in the study of stochastic processes rely upon the process in question possessing the Markov property, or able to be transformed into an equivalent Markov process.

Since the state equations of SDDEs are functions of time-lagged values of the state vector, they are by definition non-Markovian. Thus, the inclusion of noise makes the already difficult problem of dealing with time delays, even more difficult from an analytic perspective. Consequently, the study of SDDEs consists of either analytical work with simple models, or numerical investigations with more complex models.

Given the massive connectivity of neural systems, any given neuron is subject to a large number of inputs. Of these, only some are significant (correlated), while the vast majority are uncorrelated and seemingly random. While some thermal noise is present in the nervous system, the "noise" used in neural network models (including the one used here) is simply a model of this uncorrelated activity.

3.2 Stochastic Resonance

Stochastic resonance is an unintuitive effect that can occur in bistable systems under the influence of noise. In these cases, noise, which is typically viewed as unwanted, can have the effect of strengthening the presence of a weak, periodic input signal in the output of a system. It has been observed in large variety of systems, among them bistable ring lasers, semiconductor devices, chemical reactions, and mechanoreceptors in the tails of crayfish (Gammaitoni et al., 1998). Here we examine how resonance with noise can be observed in two ways.

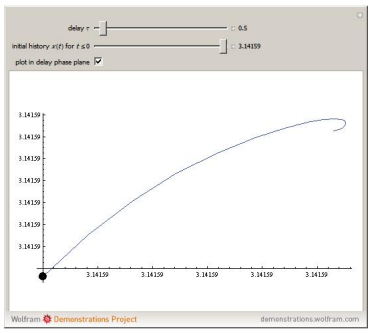

a. $\tau = 0.5$, Initial Function $=$ 3.14159.

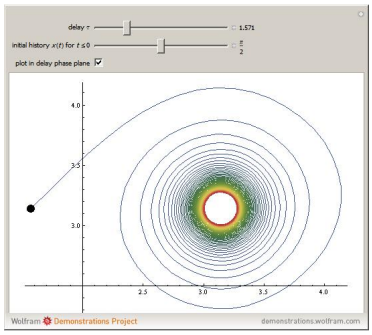

b. $\tau = 1.571$, Initial Function $=$ 1.571.

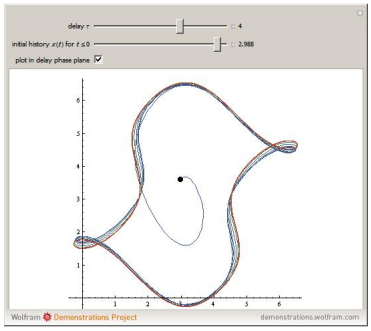

c. $\tau = 4$, Initial Function $= 2.998$.

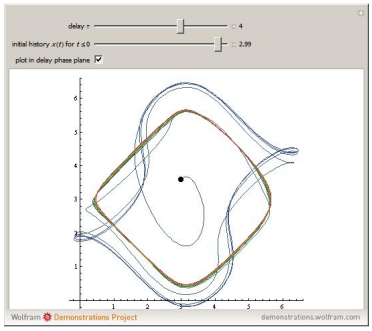

d. $\tau = 4$, Initial Function $= 2.999$.

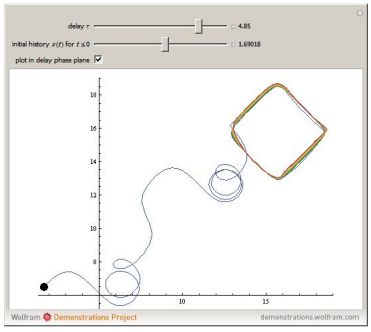

e. $\tau = 4.85$, Initial Function $=$ 1.69018.

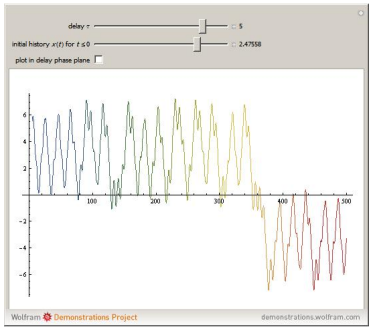

f. $\tau = 5$, Initial Function $=$ 2.47558.

Figure 3.1 A wide variety of behavior can be produced by changing the parameters from the Ikeda equation. Interested readers can explore this on http://demonstrations.wolfram.com/IkedaDelayDifferentialEquation/.

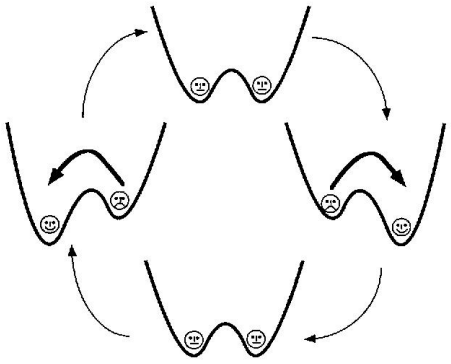

Figure 3.2 Illustration of how a weak periodic signal raises and lowers the potential barrier in a bistable system. From Gammaitoni et al. (1998).

3.2.1 Classical Stochastic Resonance

The mechanism can qualitatively be understood quite easily. Consider an overdamped Brownian particle in a double-well potential. This can be modeled by

$$V(x) = -\frac{a}{2}x^2 + \frac{b}{4}x^4, \tag{3.1}$$

$$\dot{x}(t) = -V'(x) + A_0 \cos(\Omega t + \phi) + \xi(t), \tag{3.2}$$

where $x(t)$ is the position of the particle or output of the system, $V(x)$ is the potential, A_0, a, b, Ω, ϕ are all constants, and $\xi(t)$ is zero–mean unit Gaussian noise. The second term in Equation 3.1 can be interpreted as an input signal into the system, or external periodic forcing.

The behavior of this system is straightforward. The particle rolls down this energy landscape, coming to rest at the bottom of these wells, the lowest possible energy state. These are equivalent to the fixed points of the system. When the system is deterministic (i.e., $\xi(t) = 0$ for all t), the particle will remain at the bottom of whichever well its trajectory led it to first. If $|A_0| > 0$, then the height of the potential barrier between the two wells oscillates, as shown in Figure 3.2.

Now consider the stochastic case. Clearly, there is now some finite probability that the particle will cross the barrier, placing it into the other well. This probability will be influenced by height of barrier at the specific time.

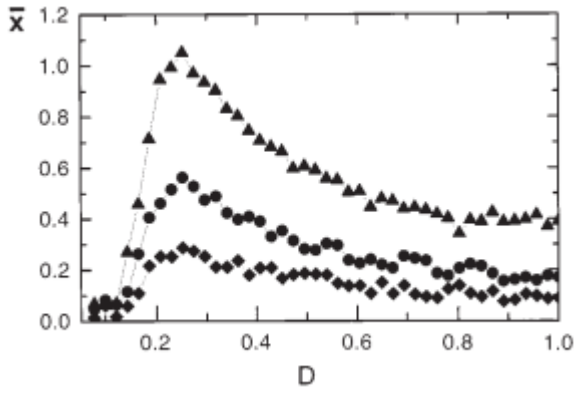

Figure 3.3 Signal versus intensity of the noise. The peak indicates the resonant condition. From Gammaitoni et al. (1998).

Since the periodic forcing is modifying the height of the barrier, noise–induced hopping between basins becomes synchronized with this forcing. Furthermore, there exists some optimal level of noise for which the input signal's presence in the output signal is maximized. This is stochastic resonance.

There are several major features which characterize stochastic resonance. First is the resonance peak seen when the signal is plotted as a function of the noise intensity (Figure 3.3). Second, we can look at the switching rate between the basins of attraction of the two stable fixed points. This is known as *Kramer's rate* (Gammaitoni et al., 1998), and is analytically defined as

$$r_k = \frac{\omega_0 \omega_b}{2\pi\gamma} \exp\left(\frac{\Delta V}{D}\right), \tag{3.3}$$

where $\omega_0^2 = V''(x_m)/m$ is the squared angular frequency of the potential in the potential minima at $\pm x_m$, and $\omega_b^2 = |V''(x_b)/m|$ is the squared angular frequency at the top of the barrier, located at x_b. If the noise level is tuned so that the switching rate (the stochastic timescale) is half of the period of the input signal (the deterministic timescale) are the same, we have the resonant condition. This is known as the *timescale matching condition*, and is formally written as

$$2T_k(D) = T_\Omega,$$

where T_k is mean waiting time between subsequent transitions, and T_Ω is the period of the periodic forcing.

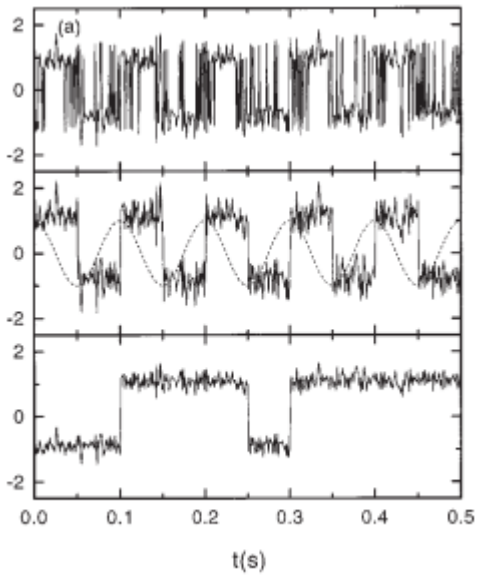

Figure 3.4 Synchronization between the input and output can be seen as the intensity of the noise is varied. The dotted line is the input signal. Taken from Gammaitoni et al. (1998).

Another feature of systems exhibiting stochastic resonance is input–output synchronization. Figure 3.4 shows this. Notable is the phase shift between the input signal and the system output. This can be approximated with Equation 3.4:

$$\bar{\phi} = \arctan\left(\frac{\Omega}{2r_k}\right). \tag{3.4}$$

Finally, residence time distributions provide another way to characterize a resonant system. Figure 3.5 shows the prototypical distribution for these systems.

An important technique implicit in the above discussions is two-state filtering, where each point in the output time series is categorized by which basin it resides in. This reduced data is then used to verify probabilistic binary models used to capture the essential stochastic dynamics of the systems. These models can show the underlying structure of features in the continuous system, such as residence times, switching rates, and resonance.

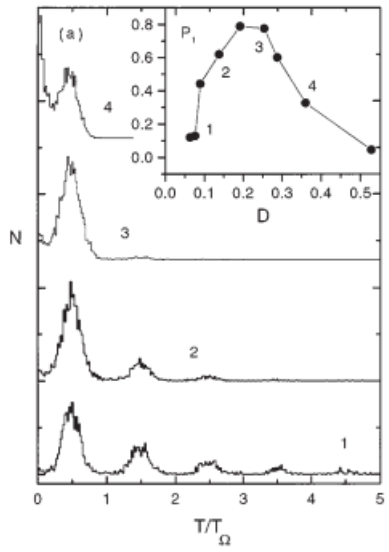

Figure 3.5 Typical residence time distribution of a resonant system as the noise level is increased. The inset shows signal-to-noise versus the intensity of the noise. From Gammaitoni et al. (1998).

3.2.2 Resonance with Delay

A kind of stochastic resonance phenomenon has also been found to occur in systems *without* an explicit periodic term. Time delays are well known to cause oscillations, and noise has been found to interact with these in a resonant sort of way. The mechanism for this is not nearly as well understood as it is for classical stochastic resonance. Here we review two of the major studies on this effect.

Ohira and Sato (1999) were the first to identify this phenomenon. They presented a simple probabilistic model which depended on the state at time τ in the past, though not the present. The model was motivated by a simple first-order DDE with a hyperbolic tangent nonlinearity. By adjusting the parameters for transition probabilities and the time delay, they found that there was some optimal pairing of values such that switches an interval of τ apart were maximized. That is, there was a form of resonance.

These results were extended by Tsimring and Pikovsky in 2001. They proposed a similar two-state model which depended on both the current state, as well as the lagged state. However, there were only two probabilities that defined the system: one for when the current and past state were

the same, and the other for when they were different. Thus, there was no allowance for asymmetric potential wells, although this approach did result in a model that was more mathematically tractable. This model was designed to capture the dynamics of an overdamped particle in a double-well quartic potential as described by the Langevin equation, a prototypical example used in the stochastic resonance literature. They were able to derive analytically values for the transition probabilities by deriving the rate equations for their model, and then equating it with Kramer's rate and solving. This effort was validated by numerical evidence.

Chapter 4

A Noisy Neural Network Model with Time Delays

Here we present a stochastic delay differential equation model of a two-neuron inhibitory network. This model belongs to the class of Hopfield models (Haykin, 1994), which take the general form

$$\frac{dx_i(t)}{dt} = -x_i(t) + \sum_{j=1}^{n} a_{ij} f_j(c_j x_j(t)) + I_i,\ i = 1, 2, \ldots, n.$$

Here, $x_i(t)$ is the firing rate of neuron i; a_{ij} is the connection weight between neuron i and j; f_j is some continuous, monotonic, and bounded function (typically a sigmoid); c_j is some constant; and I_i is the input into neuron i. All constants are positive, with the exception of a_{ij}, whose sign determines whether the influence is excitatory or inhibitory. These types of models gained much attention in the late 1980s and early 1990s, when they were used for their ability to act as content-addressable memories, allowing them to perform tasks such as pattern completion (Haykin, 1994; Olien and Bélair, 1997). These models exhibited multistability, with each equilibrium point corresponding to a specific memory (Haykin, 1994). Hopfield models were the first major model used in artificial neural networks which were dynamic in nature. They are not biologically accurate in that they don't exhibit the primary dynamical feature: spiking. However, they do describe neural dynamics "at least metaphorically" (Pakdaman et al., 1998b), as they contain many of the important features of such systems, namely time delays and rate-limited influence. In this chapter, we formulate our model, and discuss the basic aspects of its stability.

4.1 Model Statement

In this thesis, we consider a neural network consisting of two mutually inhibitory neurons as depicted in Figure 4.1. A time–delayed differential equation model for this network is given by

$$\dot{x} = -x(t) - S_2(y(t-\tau_2)) + I_1(t) + \sigma\xi_1(t), \tag{4.1}$$
$$\dot{y} = -y(t) - S_1(x(t-\tau_1)) + I_2(t) + \sigma\xi_2(t), \tag{4.2}$$

where x, y are the firing rates of the neurons; I_1, I_2 represent external inputs; τ_1, τ_2 are the conduction time delays between the two neurons; ξ_1 and ξ_2 are unit Gaussian noise (and thus σ is their standard deviation); and

$$S_j(u) = \frac{c_j u^2}{\theta_j^2 + u}, \qquad j = 1, 2$$

describe sigmoidal functions representing inhibitory influences. Note that these functions are nonnegative on $u \geq 0$ and increasing. All constants in the model are positive except the delays which are nonnegative. For the remainder of this document, we only consider the case where the delays are symmetric.

Under this formulation, we are making a distinction between correlated and uncorrelated input into these neurons from other neurons. I_1 and I_2 represent correlated inputs, and can be perhaps be interpreted as meaningful or significant. By contrast, we model uncorrelated or random inputs with a stochastic process, specifically a Gaussian one, as we discussed in Chapter 3. As we will see, both play an important role in the model's dynamics.

This specific model was recently applied to context of decision making on the neural level (Milton et al., 2010). Here, fixed points represented two decision outcomes, and the initial function was the information relevant to making one decision or the other. Under this interpretation, DITOs could be seen as a neural corollary to indecision, or "choking", when presented with large amounts of information supporting. There was a brief discussion of the impact of noise on the system, but the study was not comprehensive. They discussed how noise could lead to the "wrong" decision being made, but didn't consider noise-induced transitions, or investigate the stochastic resonance features of the systems dynamics.

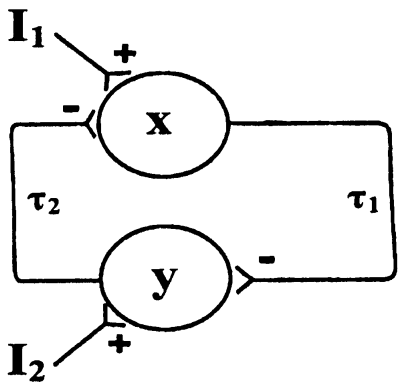

Figure 4.1 Schematic representation of a neural network with mutually inhibitory neurons. "−" indicates an inhibitory connection and "+" indicates an excitatory connection. From Milton et al. (2010).

4.2 Identification of Fixed Points

Stability analysis of this network has been carried out thoroughly in the literature (Olien and Bélair, 1997; Gopalsamy and Leung, 1996; Pakdaman et al., 1998b; Milton et al., 2010), so here we will just touch upon the major points rather than rederive the same results. We consider the case where the inputs into the system, I_1 and I_2, are constant. An equilibrium point (x^*, y^*) of Equation 4.1 must satisfy

$$x^* = I_1 - S_2(y^*) = f_2(y^*), \tag{4.3}$$
$$y^* = I_2 - S_1(x^*) = f_1(x^*). \tag{4.4}$$

Substituting Equation 4.4 in Equation 4.3 yields a single equation for x^*:

$$F(x) \stackrel{def}{=} I_1 - S_2(I_2 - S_1(x^*)) - x^* = 0. \tag{4.5}$$

Now $\lim_{x\to\infty} F(x) < 0$ if S_2 is nonnegative. Thus if $F(0) = I_1 - S_2(I_2) > 0$, then $F(x)$ will have at least one root with $0 < x^* < I_1$. Consideration of Equation 4.4 shows that the root will also have $y^* > 0$ if $I_2 - S_2(I_1) > 0$. Thus, if these two conditions are satisfied there will be at least one equilibrium in the first quadrant. Note that the same argument also holds for any nonlinearities, S_j, which are nonnegative and increasing.

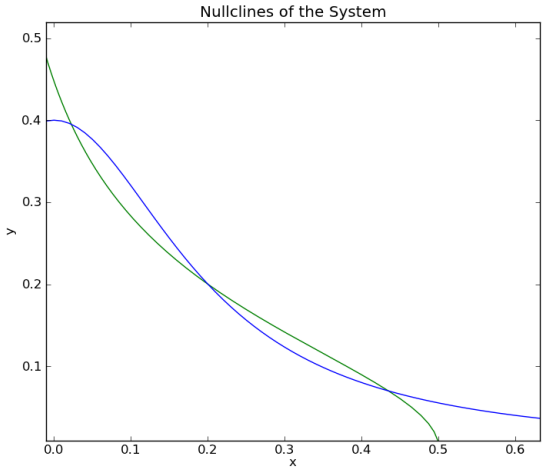

Figure 4.2 Determination of the equilibria for Equations 4.1. Equilibrium points exist at the intersections of the two nullclines. Parameter values: $T_1 = T_2 = 1$, $c_1 = 0.4$, $c_2 = 0.6$, $I_1 = 0.5$, $I_2 = 0.4$, $n_1 = 2$, $n_2 = 2$, $\theta_1 = 0.2$, and $\theta_2 = 0.2$.

With the nonlinearity we have chosen, when the conditions defined above are satisfied, we observe that there can be one, two or three equilibria in the first quadrant. We focus on the situation when there are three fixed points in the first quadrant (Figure 4.2).

It should be noted that when there are three equilibria in the first quadrant it is possible to write Equations 4.1 in the form of two mutually excitatory neural populations (Pakdaman et al., 1998b, a).

4.3 Stability Analysis

A full derivation of the local asymptotic stability of equilibria is given in Milton et al. (2010). Here we'll review the results. In the case where there is a single fixed point, it is locally asymptotically stable. When the system admits three fixed points, then the two "outer" fixed points are asymptotically stable, and the "inner" one is an unstable saddle-type point. It was shown in Olien and Bélair (1997) that the stability of these equilibrium points is unaffected by the delay using the theory of monotone dynamical systems. A monotone dynamical system is one for which, given two trajectories for

a system $\phi_1(x)$ and $\phi_2(x)$, if $\phi_1 > \phi_2$ at $t = 0$, then $\phi_1 > \phi_2$ for all $t > \tau$. Note that a trajectory at time t is defined as a continuous function over the interval $[t - \tau, t]$.

The stable manifold of the saddle-type equilibrium point is a *separatrix*; that is, it separates the x, y plane into regions. Initial conditions which lie to the left of the the separatrix yield solutions which asymptotically approach equilibrium point A and those which lie to the right approach B. These regions are called the *basins of attraction* of the equilibrium points A and B. While the stability of the equilibrium points is unaffected by the presence of the delay, describing the basins of attraction of the equilibrium points becomes considerably more complex. As noted in the previous chapter, when the delays are zero, the basins are subsets of the plane. However, when at least one delay is not zero, the basins are subsets of an infinite dimensional space consisting of the set of all initial functions.

Pakdaman et al. (1998b) have examined the effect of the delay on the separatrix on the restricted set of constant initial functions. They found that even under this restriction, the inclusion of time delays have complex impacts on the location of the separatrix. They also provide proof that there exists a boundary that divides the phase plane such that all trajectories whose (constant-valued) initial functions are on one side converge to the fixed point on that side of the boundary, and the same for trajectories with initial function on the other side of the boundary.

4.4 Potential Function

Much of the theory of stochastic resonance revolves around the existence of a potential function. For a first-order system, a potential function always exists, and is defined by

$$\mathbf{f}(\mathbf{x}) = -\frac{dV}{dx}.$$

For a second-order system, a potential function only exists if the equations specify a conservative vector field. That is, there exists some function $V(x)$ such that

$$\mathbf{f}(\mathbf{x}) = -\Delta V.$$

Equivalently, there is a potential function if and only if

$$\frac{\partial f_1}{\partial y} = \frac{\partial f_2}{\partial x},$$

where $f_1 = \dot{x}$ and $f_2 = \dot{y}$.

Theorem 4.1. *There does not exist a potential function for the system defined in Equations 4.1 when c_i and θ_i are nonnegative.*

Proof. We first differentiate f_1 with respect to y:

$$\begin{aligned}\frac{\partial f_1(x,y)}{\partial y} &= \frac{\partial}{\partial y}(-x(t) - S_2(y(t)) + I_1(t)) \\ &= -\frac{\partial}{\partial y}S_2(y) \\ &= -\frac{\partial}{\partial y}\left(\frac{c_2 y^2}{\theta_2^2 + y^2}\right) \\ &= -\frac{2c_2\theta_2^2 y}{(\theta_2^2 + y^2)^2}.\end{aligned}$$

Now we differentiate f_2 with respect to x:

$$\begin{aligned}\frac{\partial f_2(x,y)}{\partial x} &= \frac{\partial}{\partial x}(-y(t) - S_1(x(t)) + I_2(t)) \\ &= -\frac{\partial}{\partial x}S_1(x) \\ &= -\frac{\partial}{\partial x}\left(\frac{c_1 x^2}{\theta_1^2 + x^2}\right) \\ &= -\frac{2c_1\theta_1^2 x}{(\theta_1^2 + x^2)^2}.\end{aligned}$$

The two quantities will never be equal under the restrictions we've placed on c_i and θ_i, so a potential function does not exist. □

This lack of existence means that we will not be able to analytically build off or use the major analytical approaches in the stochastic resonance literature. This, however, obviously does not rule out a more numerical approach, and thusly one is taken. It should also be noted that some *potential-like* quantity could be defined for the system, such as a Lyapunov function, but this path was not pursued here.

Chapter 5

Resonance, Delays, and Noise

In this chapter we look at two types of oscillations that exhibit a form of stochastic resonance in the presence of noise. First, we'll consider the case where there is some external periodic forcing on the system, as in classical stochastic resonance studies. Then we'll look at a type of periodic behavior that occurs in the absence of any explicit periodic terms, *delay–induced transient oscillations*. These, too, exhibit a kind of resonance with the noise.

5.1 Numerics

The system was numerically solved using either XPPAUT (Ermentrout, 2002) for manual exploration, or Python using Scipy, Numpy, and Pydelay libraries for batch solving (Jones et al., 2001; Flunkert and Schöll, 2009). With XPPAUT, a modified Euler solver was used. Accordingly, a small time step ($dt = 0.001$) was used. Pydelay uses a standard order 4 Runge–Kutta solver, modified (albeit crudely) for stochastic systems. The time step used was also $dt = 0.001$.

Our interest primarily lies in the impact of varying the noise (σ) and time delay (τ) on the system, and so all other parameters were kept constant across all simulations, unless otherwise noted. The values used were $c_1 = 0.4$, $c_2 = 0.6$, $\theta_1 = 0.2$, $\theta_2 = 0.2$, $I_1 = 0.5$, and $I_2 = 0.4$.

5.2 External Periodic Input Signal

We begin with the classical case, where where we have external forcing on the system. We then examine how the system responds both in the absence and presence of a time delay.

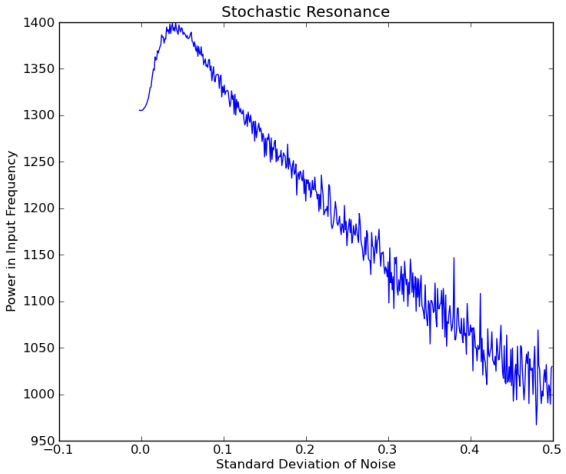

Figure 5.1 The x-axis is the value of σ, and the y-axis if the value of the FFT taken at the frequency of the input signal. We see a resonance peak as the noise intensity (σ) is varied. Here, $A_1 = A_2 = 0.15$, $\omega_1 = \omega_2 = 1.0$, and $\phi_1 = \phi_2 = 0$.

5.2.1 Absence of time delays

Consider our system under external periodic forcing:

$$T_1\dot{x} = -x - S_2(y(t-\tau_2)) + I_1(t) + \xi_1(t) + A_1\sin(\omega_1 t + \phi_1), \tag{5.1}$$
$$T_2\dot{y} = -y - S_1(x(t-\tau_1)) + I_2(t) + \xi_2(t) + A_2\sin(\omega_2 t + \phi_2). \tag{5.2}$$

If we set the time delay to $\tau = 0$, then the system exhibits many of the characteristic features of stochastic resonance. Most significantly, there exists some optimal noise level for which the signal-to-noise ratio is maximized, as shown in Figure 5.1.

5.2.2 Presence of Time Delays

Figure 5.2 shows the impact of increasing the value of τ above 0. Even when the delay is small ($\tau = 1.0$), we see that the presence of the output signal is destroyed. So while our system does exhibit the classical features of stochastic resonance under external periodic forcing when there is no delay, in the more biologically realistic time-delayed case, we do not see this effect.

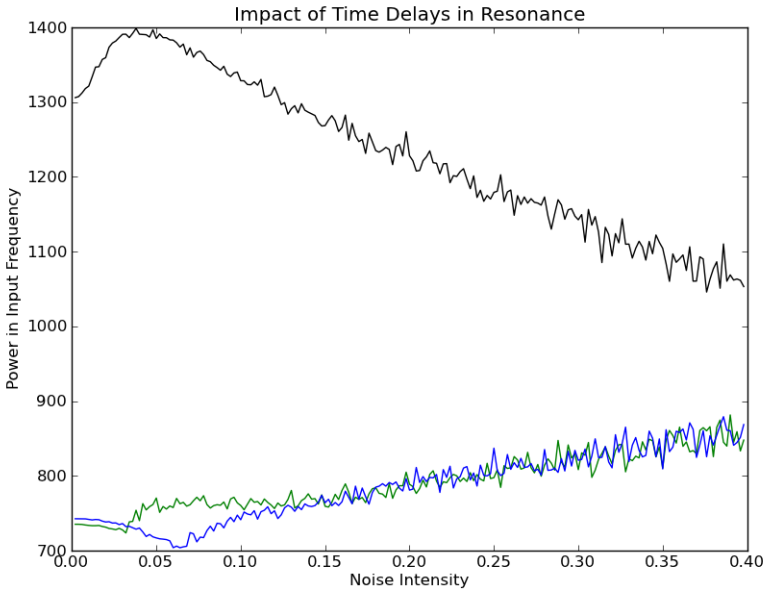

Figure 5.2 The impact of introducing time delays on the strength of the input signal. Black is when $\tau = 0$, green is when $\tau = 1.0$, and blue is when $\tau = 10.0$.

5.3 Delay–Induced Transient Oscillations

Our system exhibits another kind of oscillatory behavior when we have a sufficiently large time delay, delay-induced transient oscillations (DITOs). Most interestingly, these oscillations occur in the absence of an explicitly periodic component. This state is characterized by the following traits:

1. Period of oscillation is twice the delay (2τ).
2. Direction of oscillation is approximately parallel to the separatrix.
3. Duration of oscillatory state is many times the length of the delay.

Pakdaman et al. (1998a) examined the underlying analytical cause for DITOs, by considering a simplification of this system. Using this, they were able to clearly identify the mechanism for the oscillation. Furthermore, by using geometric arguments to construct a one-dimensional map to describe the temporal evolution of these oscillations, they demonstrated that the duration of the transient scaled exponentially with the length of the time delay.

This state can induced in the deterministic case by placing the initial function close to the unstable fixed point. Figure 5.3 shows such behavior occurring. Under the influence of some stochastic process, these oscillations may become *noise-induced*. For a given set of parameter values, the length of this state can vary greatly, though it always exists in some form whenever a trajectory crosses the separatrix.

5.3.1 Resonance

These oscillations also exhibit a form of stochastic resonance. As Figure 5.5 shows, there is some value for which the presence of the DITO frequency is maximized. Ohira and Sato classify this type of phenomenon as "stochastic resonance-like" (1999).

5.3.2 Identification

Given the periodic nature of DITOs, it is natural to characterize them in the frequency domain. These transient states can identified by taking the discrete Fourier transform of sequential intervals ("windows") of a time series produced by numerically solving our equations, and then isolating the power in the DITO frequency, $f = \frac{1}{2\tau}$. Collecting all of these, we can

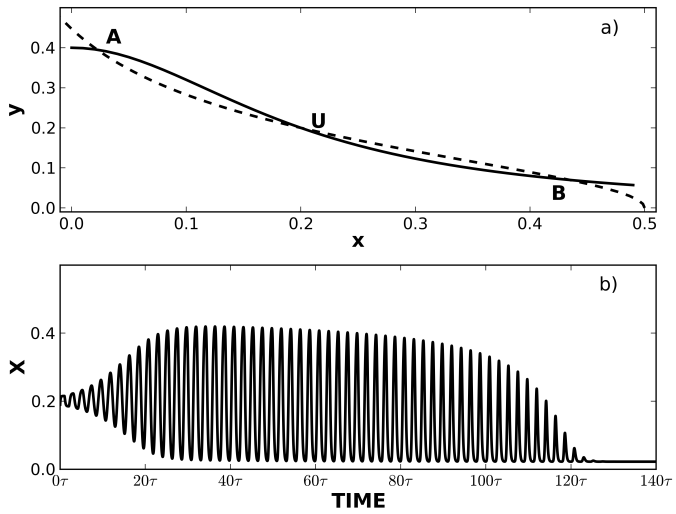

Figure 5.3 A Delay–Induced Transient Oscillation (DITO).

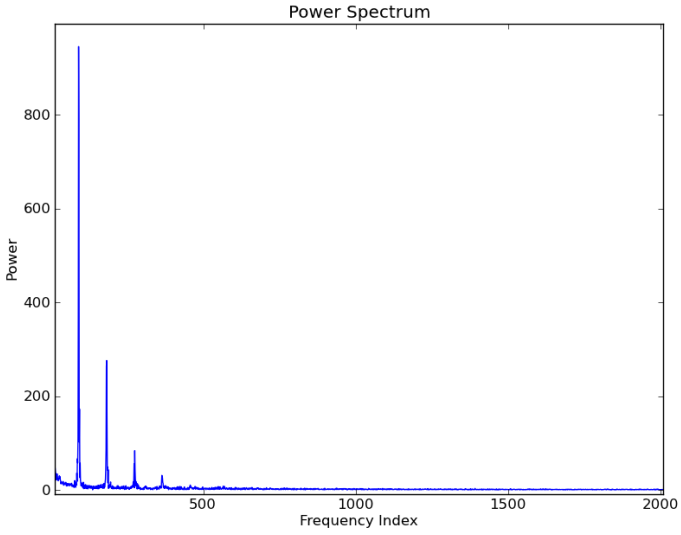

Figure 5.4 Power spectrum of a time series containing a DITO.

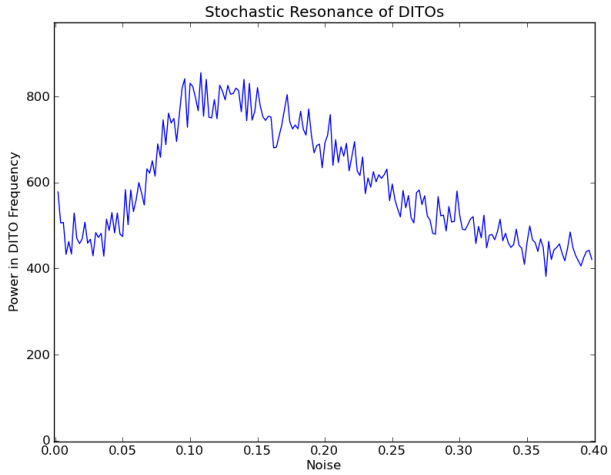

Figure 5.5 A plot of signal-to-noise as the noise level is increased. The peak of this value identifies the resonant condition.

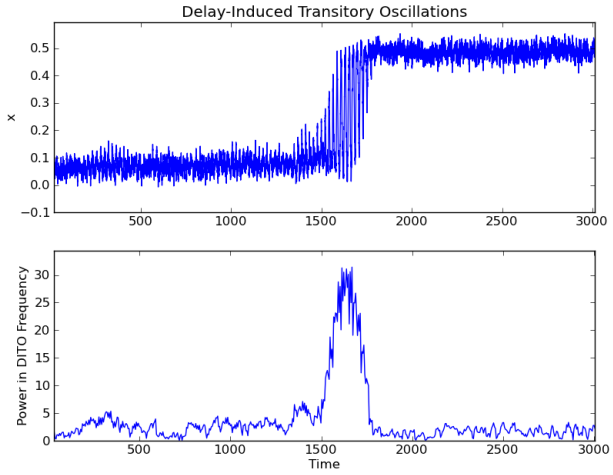

Figure 5.6 A time series from the system (top) with temporal evolution of the DITO frequency (bottom). This particular instance involves the transition of the trajectory from one stable fixed point to the other.

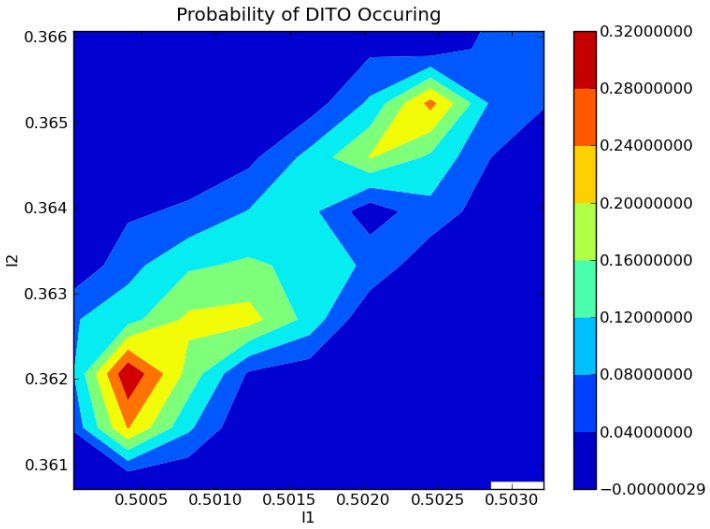

Figure 5.7 This shows the dependence of the probability of a DITO of a certain length occuring (> 100 time steps) as the inputs, I_1 and I_2, are varied. This was achieved by running 1000 trials at each point on a grid covering the parameter range shown. Here, $\tau = 8.0$ and $\sigma = 0.05$.

produce a new time series which shows when these oscillatory states arise. The results of such a sequence of processing is shown in Figure 5.6.

As we would expect, these states have a clear signature in the frequency domain, which makes detecting their presence straightforward. We can automate this process by setting some threshold above which we consider the system oscillating. This threshold needs to be set and confirmed manually, and only holds for nearby parameter values. Using this method, we can determine many of the statistical features of this state, namely the distribution of oscillation durations. For example, Figure 5.7 shows the probability of a DITO of duration greater than 100 time units occuring, as a function of the values of I_1 and I_2.

5.4 Summary

We have shown that our model displays resonance of two kinds: with periodic external forcing, and in the presence of delay. An interesting thing to note is that these cannot both exist simultaneously. In the external forc-

ing case, increasing the delay eliminates the effects of stochastic resonance. Similarly, when we have a delay present, and thus DITOs, adding forcing destroys this oscillating state.

As shown, collecting data about DITOs is a reasonably simple process using standard time series techniques in the frequency domain. While this is sufficient to recover many of the statistical properties of the state, it lacks the ability to describe the structure to these dynamics, the way binary models can for stochastically resonant systems. Furthermore, it requires some oscillation threshold to be set and confirmed manually, allowing automated analysis to only be carried out on parameters local to the ones which such a threshold was specifically determined for. The development of a model and estimation method to remedy these issues is the topic of Chapter 6. Unsurprisingly, its construction follows the ones described in Chapter 3.

Chapter 6

A Discrete Model

Here we present a discrete model to describe the dynamics of our system, similar to the binary ones used in stochastic resonance studies. In particular, it allows us to recover useful qualities about the system such as residence time distributions, time until first oscillation, and expected duration of transient oscillations. The models are stochastic, and are defined by probability of transitions between states. Although the other models mentioned in Chapter 3 were able to derive values for these probabilities analytically, the systems they considered were considerably simpler. Thus, the primary method used here to estimate transition probabilities is empirical. However, repeating this process for different parameter values reveals the dependence of the transition probabilities on parameters, and therefore the statistical properties of the dynamics of the system on the same parameters. The results derived from the method are not comprehensive, but more serve to demonstrate the usefulness of this as a method of analysis. Limits on the applicability of this model to different parameter ranges are discussed.

6.1 Formulation

We propose a three-state probabilistic model for representing the dynamics of our network. The model is schematically shown in Figure 6.1. The S_0 and S_2 represent the states seen in many of the previously referenced models; that is, a trajectory "firmly" in the basin of attraction for one of the stable fixed points. This will be clarified momentarily. The middle state, S_1, represents a trajectory that is oscillating. We'll refer to the stable fixed points as r_0 and r_2, and the unstable one as r_1.

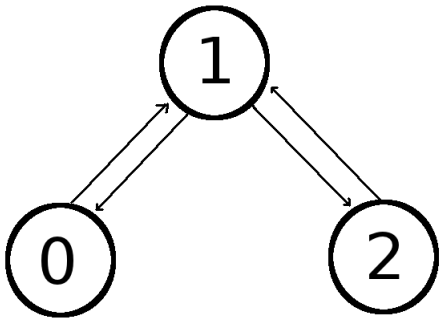

Figure 6.1 A diagram of the three-state model.

	S_0	S_1	S_2
S_0	p_0	$1 - p_0$	0
S_1	p_{10}	p_1	p_{12}
S_2	0	$1 - p_2$	p_2

Table 6.1 Transition probability matrix.

The transition probabilities between states determine the dynamics of the model. These can be represented as a 3×3 matrix (Table 6.1), which we'll denote P. Under this formulation, p_i represents the probability of remaining in state S_i in the next time step.

6.2 Justification

Implicit in our choice of a Markov chain model is the assumption that the system it is approximating is Markovian. This is certainly an issue as our system is non-Markovian by definition. However, by analyzing the time series produced by numerically solving our model, we can see that this dependence on the past is still finite, and thus we can rescale our time when transforming it to the discrete model in such a way that it satisfies the Markovian property.

Typically in time series analysis, the autocorrelation of a data set is used to determine the correlation of given values with time-lagged values of the series. The partial autocorrelation goes a step further and removes the linear effects of earlier dependencies (smaller lags) from the calculation of cor-

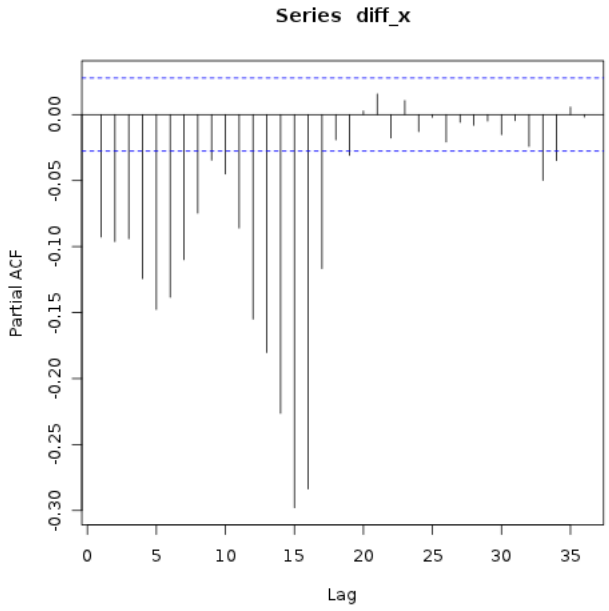

Figure 6.2 The partial autocorrelation function of a time series from our system. Here $\tau = 8.0$, and the standard deviation of our noise is 0.10. Data is downsampled to one data point per time unit, and the first difference is taken to make the series stationary. We see that there is moderate correlation up to about 16 to 20 time steps in the past, after which correlation becomes statistically insignificant.

relation of later values. More intuitively, it removes "echoes" of correlation at smaller lag values at larger lags, thus painting a more accurate picture of dependence on past states.

Figure 6.2 shows the partial autocorrelation function of one such time series, where $\tau = 8.0$. We see that there is significant negative correlation with time-lagged values up to a point, after which the correlation drops to statistically significant. The point of no correlation seems to roughly fall between 2.0 and 2.5 times the value of τ. It then follows that if we rescale our time so that one step of our discrete model correspond to an interval of at least this length, then all dependencies will be contained within this chunk.

We can test this hypothesis by downsampling our data so that the space between data points is the length of this interval (2.5τ). As Figure 6.3

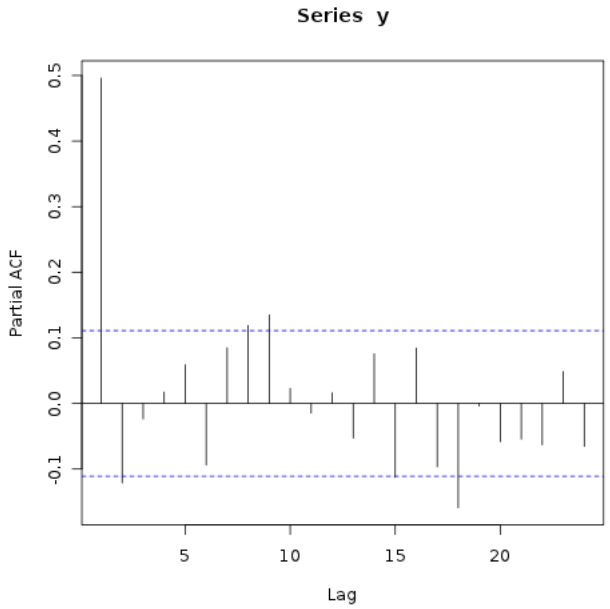

Figure 6.3 The partial autocorrelation function of a time series from our system (same as Figure 6.2) after downsampling to one data point per 20 time steps (2.5τ). We only see significant correlation at a lag of 1, indicating that the system behaves like a random walk at this time scale.

shows, when we take the partial autocorrelation of this series, there is no significant correlation other than at lag 1. This means that our time series, statistically speaking, has the same correlation properties of a random walk: it is only dependent on its previous state. Thus, if we let each time step of our discrete model represent an interval of length 2.5τ, then the Markovian property is satisfied.

6.3 Three-State Filtering

The next issue becomes how we would transform a time series from solving our system to our discrete model. We need a method to filter a time series into three states, the way two-state filters are used in the stochastic resonance literature to reduce the system down to a binary one. More specifically, we need to determine how to classify intervals of length 2.5τ as belonging to one of these states. We propose the following simple method:

1. Divide the data into subsequent intervals of length 2.5τ.
2. If all of the points in an interval are within the approximated basin of attraction of one of the stable fixed points, classify it as belonging to the corresponding stable state, either S_0 or S_2.
3. Otherwise classify the interval as belonging to S_1.

As mentioned in Chapter 4, the delay impacts the location of the separatrix in a complex way, and so we approximate it. This is done by systematically solving the system for different constant intial functions, and using a binary search type algorithm to determine within an error tolerance of 0.001, the location of the boundary at various y values. These points were then fit with a cubic polynomial, and we use this fit as an approximation of our boundary. It should be emphasized, however, that the boundary that we approximate only determines the basins of attraction for constant initial functions, not the actual trajectories we are working with.

Figure 6.4 shows the results of this method of filtering. As shown, at various parameter values, the time which is classified as S_1 corresponds quite reasonably to existence of an oscillatory state. Notable also is that the transitions between states are relatively clean. This is to say that each state that the system assumes remains relatively stable for reasonably long intervals, and that transitions between them are usually paired with some noticeable qualitative change in the behavior of the solution. This serves as preliminary evidence of the validity of our method.

6.4 Estimation of Transition Probabilities

For the discrete models mentioned from the stochastic resonance literature, the authors were all able to analytically derive at least some approximation of transition probabilites. This was done either by way of equating Kramer's rate with the theoretical switching rate, or by choosing probabilities proportional to the size of the potential barrier. We cannot use either approach for our system. For the first case, we can't use Kramer's rate as our system has three states, while models of stochastic resonance use two states. Additionally, there is no potential function for this system, as it is not conservative as shown in Chapter 4, thus we cannot use that to pick our value for transition probabilities. Both of these models were used to describe a system that is simpler than our system, so we shouldn't expect that those analytical tools carry over. Thus, the primary method employed in estimation is an empirical one.

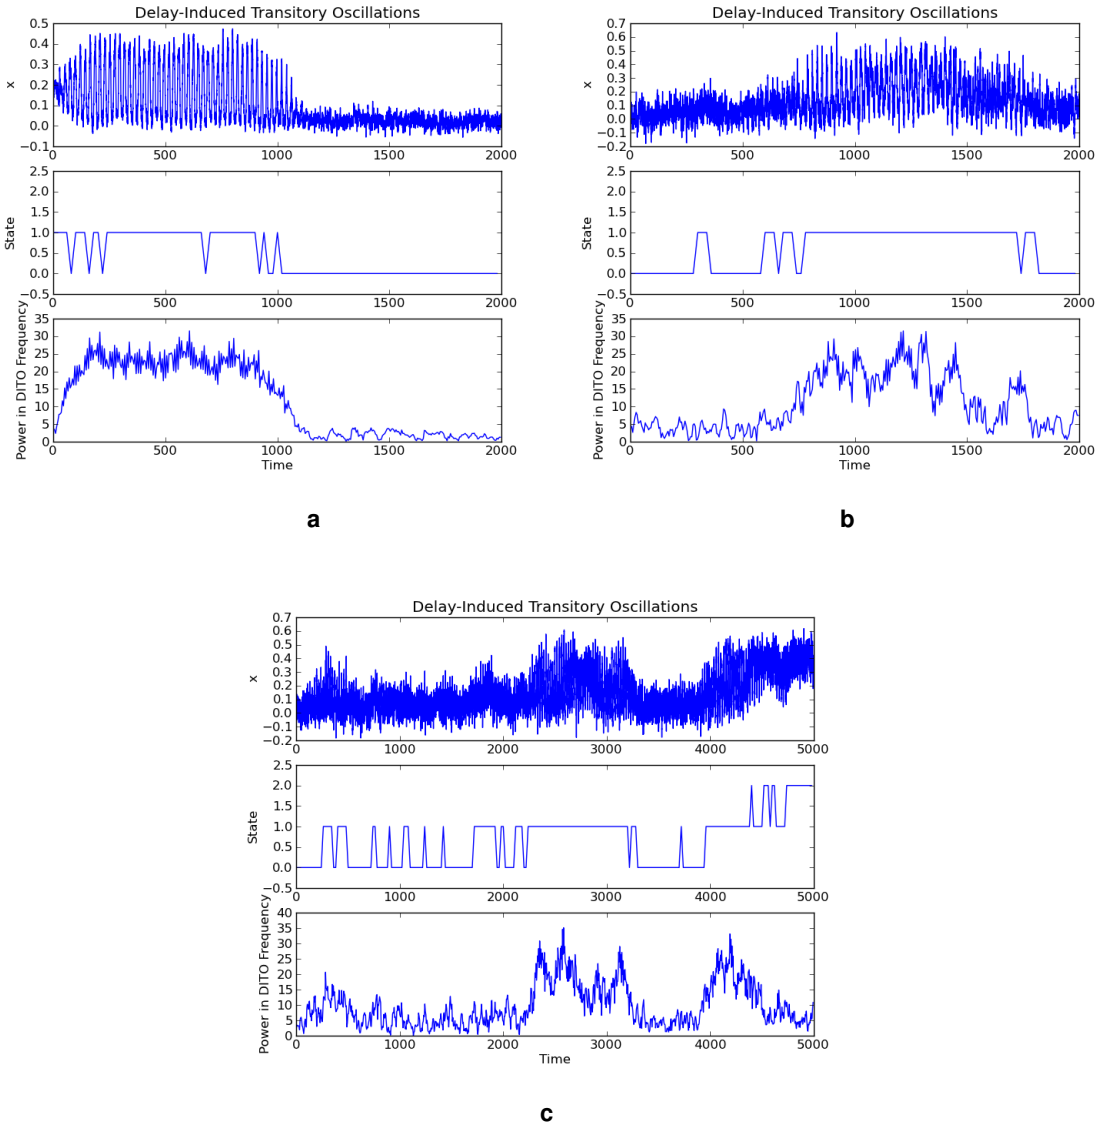

Figure 6.4 An example of using three-state filtering of a time series containing a DITO. Intervals classified as belonging to S_1 correlate highly with oscillations in the output.

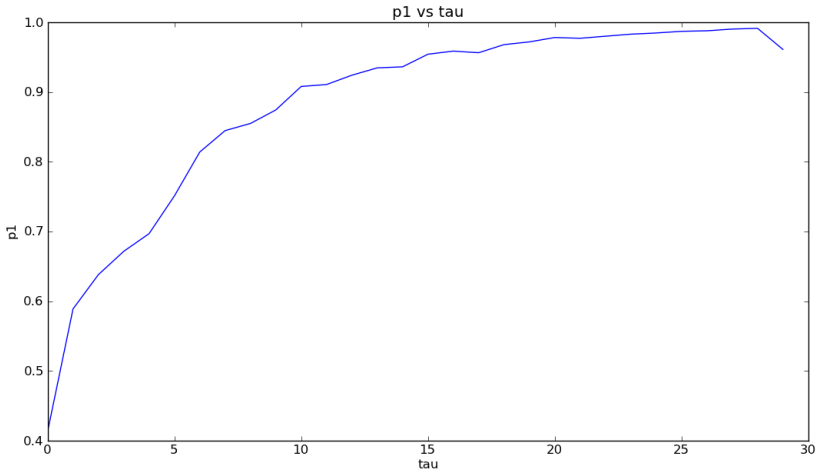

Figure 6.5 The estimated values of p_1 at various values of τ, with σ held constant at 0.08. Each estimate was found by averaging the estimates of 500 trials, 10,000 time units in length.

The process is as one would expect. We set the initial function to one of the two stable fixed points, and run the solver for a relatively long period ($t_{final} = 10,000$ time steps), so that the system has time for its dynamics to play out and to reach an equilibrium if there is one. This is repeated many times, and then repeated the same number of times but placing the initial function at the other fixed point. Taking the time series from all of these trials and three-state filtering them then allows us to estimate the transition probabilities between states. The resulting estimates can then be used to produce a transition probability matrix.

6.5 Results: Parameter Dependence

Figures 6.5 and 6.6 show the results of estimating p_1 (continuation of oscillating state) as the delay and noise level are varied. In both cases, we see a very clear, sigmoid-like trend. We can also see the threshold level of noise needed for there to be a nonneglible chance of reaching the separatrix.

These plots also make what values of τ and σ are valid for this model. We can see that at large values of either, nearly all intervals get classified as oscillating. While this does expose a limitation of this model and method

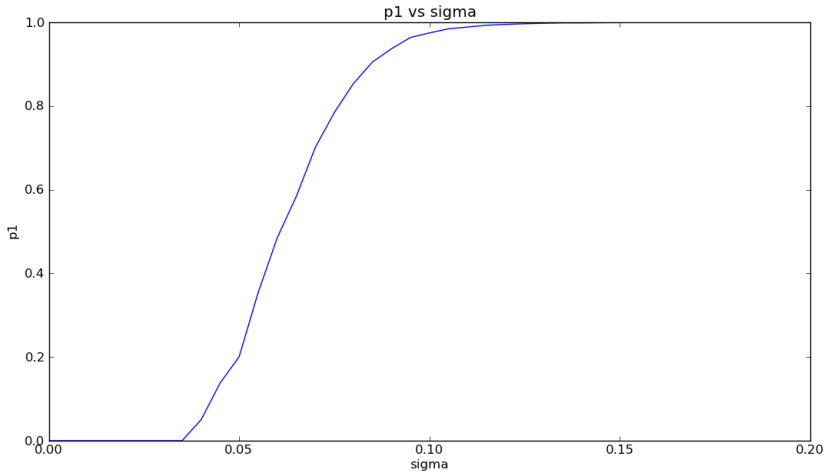

Figure 6.6 The estimated values of p_1 at different noise levels. $\tau = 8.0$ for all trials. Each point estimated was the result of 500 trials, each 10000 time units long.

of estimation, it should also be noted that these parameter values are not of much interest, biologically speaking. For τ, a large value is not realistic, and for σ, the system looks roughly like white noise at these extreme values.

6.6 Computation of Features

Having now established some of the general relationships of parameter values on the transition probabilities, we can now use those to estimate the probabilities for parameter value combinations that we do not explicitly estimate using the method from Section 6.4.

We can then by extension recover the statistical properties of the dynamics of our network in these cases quite easily. For a given state S_i with probability p_i of remaining in the same state in the next time step, we can calculate the expected duration of this state:

$$E[T_i] = \tau \sum_{j=1}^{\infty} j p_i (1 - p_i)^{j-1} \tag{6.1}$$

$$= \frac{\tau}{p_i}. \tag{6.2}$$

For the stable states, this defines the average residence time in that state, as well as the expected first passage time. For state S_1, this value represents the expected number of oscillations for any given transient.

6.7 Validation

Here we compare the results of the discrete model to data from the real model. Figure 6.7 compares the distribution of oscillation times from a time series (using the Fourier method outlined in Chapter 5) with the distribution obtained from simulating our discrete model for an equivalent amount of time. The two data sets are distributed similarly, but generally speaking, the estimation for p_1 tends to be low, but not unreasonably so. In the case indicated in the figure, p_1 was estimated to be 0.943, but a value of 0.97 would fit the data "best," as will be discussed shortly.

This isn't too much of a surprise given the simplicity of the method, and the likely cause is quick switches between states that we can observe when we three-state filter a time series. That said, the distribution of oscillation times does follow the type of distribution we are using quite well, our estimates are simply a bit shy of the best value. This "best" value can be calculated by solving

$$\mu = \frac{1}{p},$$

where μ is the mean oscillation duration of your time series.

We also compare the residence time distribution of one of the states. Figure 6.8 shows the results of this effort. We have here pretty reasonable agreement as well. It should be noted, however, that this comparison is somewhat circular, as the definition for leaving a basin, and leaving a stable state in the discrete model are the same.

6.8 Summary and Discussion

Here we have shown that the dynamics of our non-Markovian system can be approximated by a three-state Markov chain. The transition probabilities for this model can be estimated empirically from time series data using a simple filtering scheme, and dependence on parameters can be determined. These estimates fit the data reasonably well, though they tended to be on the low side. This, however, is too be expected given the simplicity of the estimation method. More importantly, the residence time distributions

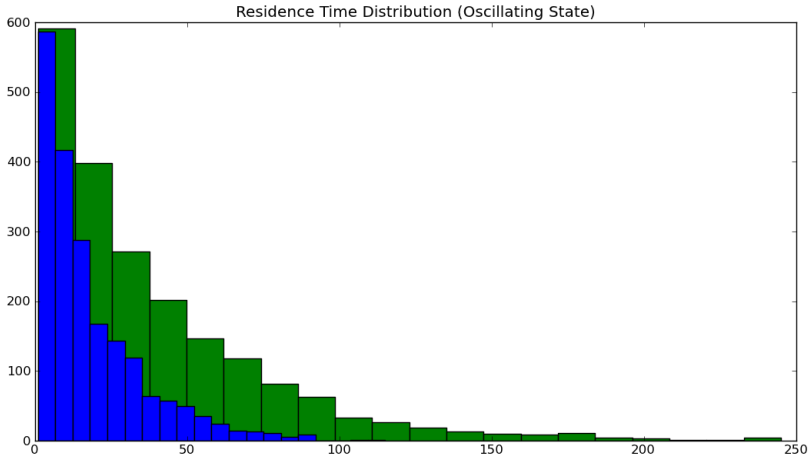

Figure 6.7 Comparison of residence time distribution/oscillating distribution for the original system (green) versus the discrete system (blue). $\tau = 10$ and $\sigma = 0.08$. Each system was run for the equivalent of 10,000 time steps over 500 trials.

follow quite closely to the type of distribution we are using to model it. If a better method for estimation, or an analytical derivation of these values could be found, a very good fit could likely be achieved.

While this model does show reasonable agreement with data classified using Fourier characterization, it has two potential flaws. First, there is no indication of a resonant condition that we found in Chapter 5. This would be a desirable quality in such a model. Second, the model itself does exhibit oscillations or periodic switching, rather we assign a state as oscillating. This simplification, though it fits our intuitive understanding of the dynamics, pulls it further away from similar models in the stochastic resonance literature, and thus makes it more difficult to build on previous results. Finally, its lack of analytically determined transition probabilities fundamentally limits the range of parameters that it can explain.

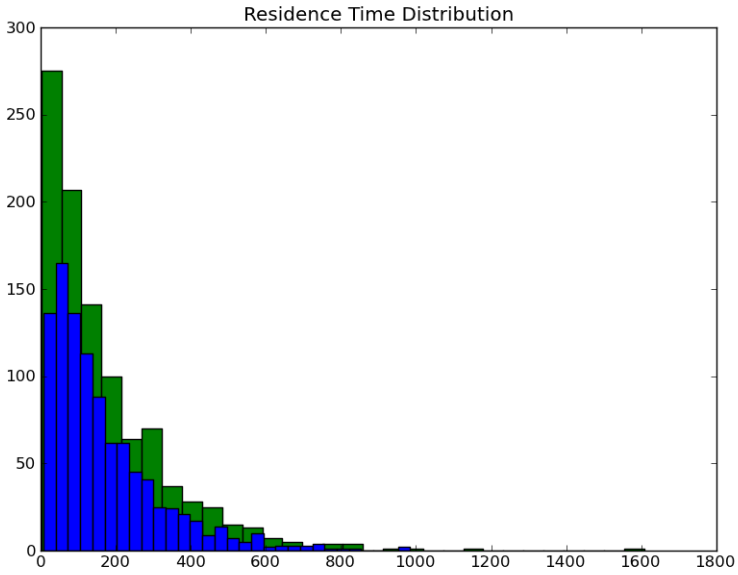

Figure 6.8 Residence time distribution for state S_0 in original system (blue) versus the discrete model (green). Here $\tau = 10$ and $\sigma = 0.10$. Each system was run for the equivalent of 10,000 time steps over 500 trials.

Chapter 7

Biological Significance

To this point, we have only considered our model as a mathematical object, ascribing not particular interpretation. However, given that it is a model of a biological system, our work clearly has some biological implications. Here we present one potential application of our work to a specific kind of epilepsy.

7.1 Nocturnal Frontal Lobe Epilepsy

Epilepsy is a neurological disorder that affects millions of individuals worldwide. An epileptic seizure is a paroxysmal event in which there is widespread synchronization of neural firing observed in the brain. Figure 7.1 shows an EEG (encephaloelectrogram) of an individual before and during a seizure. The mechanisms that cause a seizure to occur are at present poorly understood, though an very active area of research interest.

The characteristic feature of nocturnal frontal lobe epilepsy (NFLE) is that seizures occur predominantly during sleep. Autosomal dominant NFLE is associated with mutations in the α-4 subunit of neuronal acetylcholine receptors (Mann and Mody, 2008). However, even in the same individual, seizures are highly variable with respect to their timing, semiology and electroencephalographic features. These observations emphasize that the NFLE genetic mutation is a necessary, but not sufficient, condition for seizure occurrence in NFLE. In other words, there must be additional factors that taken together with the receptor defect determine when the seizure occurs.

The identification of these additional factors may make it possible to develop methods based on seizure anticipation, such as delivery of well-

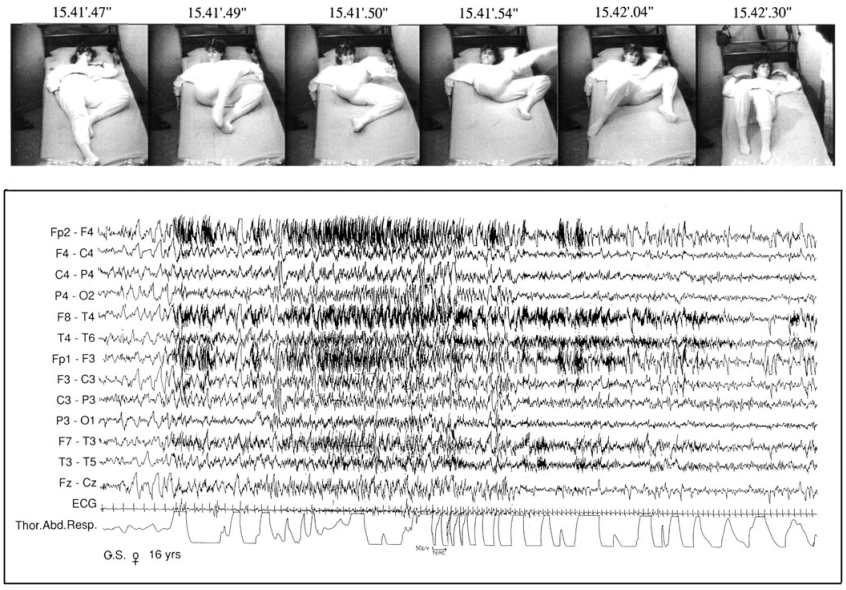

Figure 7.1 Electrical features of a seizure, as measured by an EEG. From Provini et al. (1999).

timed electrical stimuli (Osorio et al., 1998), to treat NFLE. In this way investigations of NFLE can provide clues for the development of seizure-prediction techniques in other types of epilepsy as well.

7.2 A Possible Mechanism

In many ways seizures have many of the same qualitative characteristics as the DITOs seen in our model:

1. Spontaneous synchronization of activity
2. Duration significantly longer than the neural time scale
3. Low probability events (long duration ones)

We then conjecture that perhaps the delay-induced oscillations we see in our simple model are equivalent to the oscillation events observed during a seizure. We could then interpret the stable fixed points of our systems as stages of sleep. Shifts between sleep stages could then be seen as the emergence of a new fixed point, leading to the bistable state, and the eventual

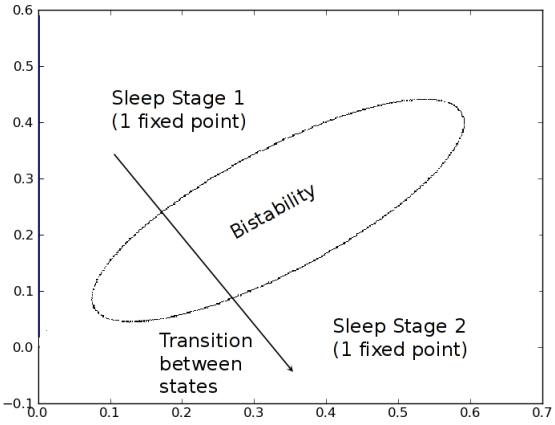

Figure 7.2 A possible explanation for NFLE. During sleep, the values of I_1 and I_2 change, causing the temporary appearance of bistability in the system. It is during this bistability that we see DITOs, and, we hypothesize, seizures. Altered functioning to these neurons can change the path through this bistable zone, potentially increasing the possibility of a long-lived oscillatory event.

noise-induced transition to the other stable fixed point as the previous one disappears. Perhaps, then, predisposition to seizures during sleep can be seen as equivalent to parameters in our model that are tuned in such a way that long–lived oscillations become more probable. Figure 7.2 visualizes this possibility. Pushing it even further, we could even consider seizure events as a kind of stochastic resonance. These ideas are the topic of Milton et al. (2011).

While these are interesting ideas, much work would need to be done to validate this hypothesis. First, we'd need to increase the complexity of our model, both in terms of the type of neuron model used, but also the size of the network. We'd need to see that some of the behavior, namely DITOs, that we see in our simple model, also is exhibited in the complexified ones. The choice of network should also be based on actual neurobiology, and the model used should be one capable of spiking. This effort would need to be paired with the acquisition of some form of neurological data, to verify the results of our models. In any case, this presents an interesting possbility, and one which should be pursued.

Chapter 8

Conclusion

In this thesis, we have explored the interplay of noise, time delays, and resonance in a simple two-neuron network model. This system was shown to exhibit two kinds of resonance. The first was classical stochastic resonance under external periodic forcing, and our system's behavior exhibited the canonical features of this phenomenon. The second was a relationship between the noise level, and delay-induced transitory oscillations. We also developed a three-state probabilistic model for describing the structure of the dynamics of this systems, and particular, the transient oscillations. It succeeded in being able to capture aspects of the non–Markovian behavior of this system with a Markovian model. Although we were not able to determine the parameters of this model analytically, we outline a method for estimating them numerically. These probabilities demonstrate simple and consistent dependencies on the systems parameters, which allowed prediction of the statistical aspects of the systems behavior at untested parameter values. The possible relevance of DITOs to mechanisms of seizure genesis in NFLE was also discussed.

An obvious direction for future work is seeing if the behavior observed at this level scales with model and network complexity. A spiking neural network model of a network from the neurobiological literature would lead to further understanding of the relevance of simple models to actual nervous system functioning. The three-state system proposed also could be further developed, perhaps deriving analytical approximations of transition probabilities.

Bibliography

Ermentrout, Bard. 2002. *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to XPPAUT for Researchers and Students*. Philadelphia : Society for Industrial and Applied Mathematics.

Flunkert, Valentin, and Eckehard Schöll. 2009. pydelay—a Python tool for solving delay differential equations. arXiv:0911.1633 [nlin.CD].

Gammaitoni, Luca, Peter Hänggi, Peter Jung, and Fabio Marchesoni. 1998. Stochastic resonance. *Rev Mod Phys* 70(1):223–287. doi:10.1103/RevModPhys.70.223.

Gopalsamy, Gopal, and Issic Leung. 1996. Delay induced periodicity in a neural netlet of excitation and inhibition. *Physica D: Nonlinear Phenomena* 89(3-4):395–426. doi:10.1016/0167-2789(95)00203-0. URL http://www.sciencedirect.com/science/article/B6TVK-3VS8HXG-1N/2/1c338520068122581625a60e35680ba3.

Hale, Jack, and Huseyin Kocak. 1991. *Dynamics and Bifurcations*. New York: Springer.

Haykin, Simon. 1994. *Neural Networks: A Comprehensive Foundation*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1st ed.

Jones, Eric, Travis Oliphant, and Pearu Peterson. 2001. SciPy: Open source scientific tools for Python. URL http://www.scipy.org/.

Mann, Edward, and Istvan Mody. 2008. The multifaceted role of inhibition in epilepsy: Seizure-genesis through excessive GABAergic inhibition in autosomal dominant nocturnal frontal lobe epilepsy. *Current Opinion in Neurology* 21(2):155–160. doi:10.1097/WCO.0b013e3282f52f5f.

Milton, John, Paulami Naik, Clarence Chan, and Sue Ann Campbell. 2010. Indecision in neural decision making models. *Mathematical Modeling of Natural Phenomena* 5.

Milton, John, Austin Quan, and Ivan Osorio. 2011. *Epilepsy: The intersection of neurosciences, biology, mathematics, engineering, and physics*, chap. Nocturnal Frontal Lobe Epilepsy: Metastability in a Dynamic Disease?, 501–510. Boca Raton: CRC Press.

Ohira, Toru, and Yuzuru Sato. 1999. Resonance with noise and delay. *Phys Rev Lett* 82(14):2811–2815. doi:10.1103/PhysRevLett.82.2811.

Olien, Leonard, and Jacques Bélair. 1997. Bifurcations, stability, and monotonicity properties of a delayed neural network model. *Physica D: Nonlinear Phenomena* 102(3-4):349–363. doi:10.1016/S0167-2789(96)00215-1. URL http://www.sciencedirect.com/science/article/B6TVK-3SPFDKG-Y/2/89956648708079d0977b491daed13bc7.

Osorio, Ivan, Mark G. Frei, and Steven B. Wilkinson. 1998. Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset. *Epilepsia* 39(6):615–627. doi:10.1111/j.1528-1157.1998.tb01430.x.

Pakdaman, Khashayar, C. Grotta-Ragazzo, and C.P. Malta. 1998a. Transient regime duration in continuous-time neural networks with delay. *Phys Rev E* 58(3):3623–3627. doi:10.1103/PhysRevE.58.3623.

Pakdaman, Khashayar, C. Grotta-Ragazzo, C.P. Malta, Ovide Arino, and Jean-Francois Vibert. 1998b. Effect of delay on the boundary of the basin of attraction in a system of two neurons. *Neural Networks* 11(3):509–519. doi:10.1016/S0893-6080(97)00112-3. URL http://www.sciencedirect.com/science/article/B6T08-3T3TKV8-C/2/4f392c2edfee66b6cd4150fc9d22400c.

Provini, Federica, Giuseppe Plazzi, Paolo Tinuper, Stefano Vandi, Elio Lugaresi, and Pasquale Montagna. 1999. Nocturnal frontal lobe epilepsy. *Brain* 122(6):1017–1031. doi:10.1093/brain/122.6.1017. URL http://brain.oxfordjournals.org/content/122/6/1017.abstract. http://brain.oxfordjournals.org/content/122/6/1017.full.pdf+html.

Strogatz, Steven H. 1994. *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering (Studies in Nonlinearity)*. Studies in nonlinearity, Perseus Books Group, 1st ed.

Tsimring, Lev, and Arkady Pikovsky. 2001. Noise-induced dynamics in bistable systems with delay. *Physical Review Letters* 87(25):250,602. doi: 10.1103/PhysRevLett.87.250602.