

The firing of an excitable neuron in the presence of stochastic trains of strong synaptic inputs

Jonathan Rubin

Department of Mathematics
University of Pittsburgh
Pittsburgh, PA 15260 USA

Krešimir Josić

Department of Mathematics
University of Houston
Houston, TX 77204-3008, USA

November 9, 2005

Abstract

We consider a fast-slow excitable system subject to a stochastic excitatory input train, and show that under general conditions its long term behavior is captured by an irreducible Markov chain. In particular, the firing probability to each input, expected number of response failures between firings, and distribution of slow variable values between firings can be obtained analytically from the distribution of interexcitation intervals. The approach we present immediately generalizes to any pair of input trains, excitatory or inhibitory and synaptic or not, with distinct switching frequencies. We also discuss how the method can be extended to other models, such as integrate-and-fire, that feature a single variable that builds up to a threshold where an instantaneous spike and reset occur. The Markov chain analysis guarantees the existence of a limiting distribution and allows for the identification of different bifurcation events, and thus has clear advantages over direct Monte Carlo simulations. We illustrate this analysis on a model thalamocortical (TC) cell subject to two example distributions of excitatory synaptic inputs, in the cases of constant and rhythmic inhibition. The analysis shows that there is a drastic drop in the likelihood of firing just after inhibitory onset in the case of rhythmic inhibition, relative even to the case of elevated but constant inhibition. This observation provides support for a possible mechanism for the induction of motor symptoms in Parkinson's disease, analyzed in [Rubin and Terman, 2004].

1 Introduction

There has been substantial discussion of the roles of excitatory and inhibitory synaptic inputs in driving or modulating neuronal firing. Computational analysis of this issue generally considers a neuron awash in a sea of synaptic bombardment [Somers et al., 1998, van Vreeswijk and Sompolinsky, 1998, De Schutter, 1999, Tiesinga et al., 2000, Tiesinga, 2005, Chance et al., 2002, Tiesinga and Sejnowski, 2004, Huertas et al., 2005]. In this work, we also investigate the impact of synaptic inputs on the firing of a neuron, but with a focus on the effects of single inputs within stochastic trains. This investigation is motivated by consideration of thalamocortical relay (TC) cells, under the hypothesis that such cells are configured to reliably relay individual excitatory inputs, arising either from strong, isolated synaptic signals or from tightly synchronized sets of synaptic signals, during states of attentive wakefulness, yet are also modulated by inhibitory input streams [Smith and Sherman, 2002]. This viewpoint leads to the question of how the relationship between the activity of a neuron and a stochastic excitatory input train varies under different patterns of inhibitory modulation.

The main goal of this paper is to introduce and illustrate a mathematical approach to the analysis of this relationship. Our approach applies to general excitable systems with separation of time scales, including a single slow variable, and fast onset and offset of inputs. These ideas directly generalize to other neuronal models, such as integrate-and-fire, featuring a slow build-up of potential interrupted by instantaneous spikes and resets. We harness these features to reduce system dynamics to a one-dimensional map on the slow variable. Each iteration of the map corresponds to the time interval from the arrival of one excitatory input to the arrival of the next excitatory input [Othmer and Watanabe, 1994, Xie et al., 1996, Ichinose et al., 1998, Othmer and Xie, 1999, Coombes and Osbaldestin, 2000]. From this map, under the assumption of a bounded excitatory input rate, we derive an irreducible Markov chain, whose bins are indexed by slow variable values and numbers of inputs received since the last firing.

We prove the key result that under rather general conditions, this Markov chain is aperiodic, and hence has a limiting distribution. This limiting distribution can be computed from the distribution of input arrival times. Once obtained, it can be used to deduce much about the firing statistics of the driven cell including the probability that the cell will fire in response to a given excitatory input, the expected number of response failures that the cell will experience between firings, and the distribution of slow variable values attained after any fixed number of unsuccessful inputs arriving between firings. We emphasize that the guaranteed existence of a limiting distribution constitutes a key advantage of the Markov chain framework over direct Monte Carlo simulations for attaining these types of statistics, since there is no guarantee of convergence for the Monte Carlo approach. Moreover, as we illustrate, the limiting distribution for the Markov chain can be computed analytically, eliminating the need for simulations altogether. Finally, the Markov chain analysis allows for the identification of bifurcation events in which variation of model parameters can lead to abrupt changes that affect long-term statistics, although we do not pursue this in detail in this work (see [Doi et al., 1998, Tateno and Jimbo, 2000], which we also comment upon in

the Discussion in Section 9).

We discuss the Markov chain approach in the particular cases of constant inhibition, which may be zero or nonzero, and inhibition that undergoes abrupt switches between two different levels. These choices are motivated by the analysis of TC cell relay reliability in the face of variations in inhibitory basal ganglia outputs that arise in Parkinson's disease (PD) and under deep brain stimulation (DBS), applied to combat the motor symptoms of PD. An important point emerging from experimental results is that parkinsonian changes in the basal ganglia induce rhythmicity in inhibitory basal ganglia outputs [Nini et al., 1995, Magnin et al., 2000, Raz et al., 2000, Brown et al., 2001], while DBS regularizes these outputs, albeit at higher than normal levels [Anderson et al., 2003, Hashimoto et al., 2003]. In recent work, Rubin and Terman provided computational and analytical support for the hypothesis that, given a TC cell that can respond reliably to excitatory inputs under normal conditions, parkinsonian modulations to basal ganglia activity will disrupt reliable relay. On the other hand, regularized basal ganglia activity, even if it leads to unusually high levels of inhibition, can restore reliable TC responses [Rubin and Terman, 2004]. While the computational results presented in [Rubin and Terman, 2004] focused on finite time simulations of two TC models with particular choices of excitatory and inhibitory input strengths, frequencies, and durations, the general framework presented here can be used to analyze how a model TC cell responds to any train of fast excitatory inputs, in the presence of inhibition that stochastically makes abrupt jumps between two levels.

The paper is organized as follows. In Section 2 we review the way in which the dynamics of a fast-slow excitable system under pulsatile drive can be reduced to a map. In Section 3 we use these ideas to construct a Markov chain that captures the relevant aspects of the dynamics of the system. We consider the existence of limiting densities for the constructed Markov chain and their interpretation in Section 4. Further, we discuss the application of these ideas to responses of a population of cells, as well as their extension to populations with heterogeneities and to integrate-and-fire type models, in Section 5. The specific problem of determining the reliability of a reduced model TC cell under various types of inputs is addressed, as an application of the theory, in Section 6 and 7, while an explicit connection to PD and DBS is made in Section 8. The details of the calculations underlying the results of these sections are contained in the Appendices.

In the two examples that we present, one with a uniform and one with a normal distribution of excitatory inputs, a significant decrease in the responsiveness of the model TC cell and a significant increase in its likelihood of being found far from firing threshold result after the onset of the inhibitory phase of a time-varying inhibitory input, relative to the cases of high or low but constant inhibition. These findings are in agreement with [Rubin and Terman, 2004] and, based on the generality of our approach, appear to represent general characteristics of the TC model with a single slow variable.

2 Reduction of the dynamics to maps of the interval

In this section we introduce a general relaxation oscillator subject to synaptic input. Under certain assumptions on the form of the input, the dynamics of the oscillator is accurately captured by a map of an interval.

A general relaxation oscillator

Relaxation oscillators present a class of models of excitable systems. Such models are in general described by equations of the form

$$\begin{aligned} v' &= f(v, w) + I(v, t) \\ w' &= \varepsilon g(v, w), \end{aligned} \tag{2.1}$$

where $0 < \varepsilon \ll 1$. In the neuronal context, the fast variable v models the voltage, while the slow variable w typically models different conductances [Rubin and Terman, 2002]. The input to the oscillator is modeled by the term $I(v, t)$. We will assume that if $I(v, t) = C$ for a constant C in a range of interest, then the v -nullcline, given implicitly by $f(v, w) = -C$, has a three-branched, or N -like shape (see Fig. 1). We will refer to the different branches of this nullcline as the left, right and middle branch, respectively. Under assumptions on the form of f and g that typically hold in practice [Rubin and Terman, 2002], it is well known that if the w -nullcline, given implicitly by $g(v, w) = 0$, intersects the v -nullcline at the middle branch, then (2.1) has oscillatory solutions. On the other hand, if the w -nullcline meets the v -nullcline at the left branch, there are no oscillatory solutions, but the system is excitable. In this case, a kick in the positive v direction can trigger an extended excursion (a spike). In the following we will consider the case when the two nullclines intersect on the left branch of the v -nullcline.

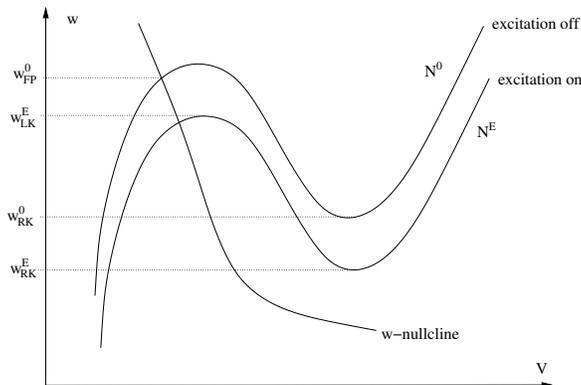


Figure 1: The nullclines of system (2.1) under the assumptions discussed in the text. The upper and lower v -nullclines correspond to $s_{exc} = 0$ and $s_{exc} = 1$ respectively.

In the present work system (2.1) will be used as a model of a neuronal cell, and $I(v, t)$

will model synaptic input, so that

$$I(v, t) = -g_{exc}s_{exc}(t)(v - v_{exc}) - g_{inh}s_{inh}(t)(v - v_{inh}) \quad (2.2)$$

where $v - v_{exc} < 0$ and $v - v_{inh} > 0$ over the relevant range of v -values. The two terms in the sum represent effects of excitatory and inhibitory conductances respectively. We will assume that the synaptic variables $s_{exc}(t)$ and $s_{inh}(t)$ switch between values of 0 (off) and 1 (on) instantaneously and independently [Somers and Kopell, 1993, Othmer and Watanabe, 1994]. This is a reasonable approximation in a situation where the input cells fire action potentials of stereotypical duration, sufficiently widely separated in time to allow for synaptic decay to an extremely small level between inputs, and where the synaptic onset and offset rates are rapid. We need not consider other aspects of the dynamics of the neurons supplying inputs to the model cell, as only the input timing will be relevant.

If the magnitude of I is not too large, then each of the four possible choices for (s_{exc}, s_{inh}) yields a different v -nullcline of system (2.1), each of which has an N -like shape. For simplicity we will first consider the case $s_{inh} = 0$, so that the cell receives only excitatory input. We label the resulting nullclines as \mathcal{N}^i , with $i \in \{E, 0\}$ corresponding to the values of $s_{exc} = 1$ and $s_{exc} = 0$ respectively.

We refer to the left (right) branch of each v -nullcline \mathcal{N}^i as the *silent (active) phase*. Each left branch terminates where it coalesces with the middle branch in a saddle node of equilibria for the v -equation. We denote these bifurcation points, or left knees, by (v_{LK}^i, w_{LK}^i) , and denote the analogous right knees as (v_{RK}^i, w_{RK}^i) , with $i \in \{E, 0\}$ as above. Since $v - v_{exc} < 0$, increasing s_{exc} lowers the v -nullcline in the (v, w) phase plane. See Fig. 1 for the arrangement of the nullclines and the different knees.

By assumption system (2.1) is excitable, with an asymptotically stable critical point (v_{FP}^i, w_{FP}^i) on the left branch of each \mathcal{N}^i . All solutions of (2.1) will therefore approach (v_{FP}^0, w_{FP}^0) in the absence of input. Thus any interesting dynamics is a consequence of synaptic input.

The timing of the synaptic inputs

We denote the times at which the excitatory synaptic inputs occur by t_i , and we assume that each input has the same duration, which we call t^* . We will assume that $t_{i+1} - t_i > t^*$, so that the inputs do not “overlap,” and that the inputs are of equal strength, which yield

$$s_{exc}(t) = \begin{cases} 1 & \text{if } t_i < t < t_i + t^*, \\ 0 & \text{otherwise.} \end{cases}$$

We will comment on trains of inputs with variable amplitudes and durations in Remark 3.2.

Of fundamental interest in the following analysis will be the timing between subsequent excitatory inputs, which gives rise to the distribution of inter-excitatory intervals $T_i = t_{i+1} - t_i$. We will assume that the T_i are independent, identically distributed random variables, with corresponding density $\rho(t)$. Excitation that is T -periodic in time is a special case corresponding to the singular distribution $\rho(t) = \delta(T)$.

Reduction of the dynamics to a map

For the purposes of analysis, we consider that the active phase of the cell is very fast compared to the silent phase. More specifically, we treat the neuron as a three-timescale system. The first timescale governs the evolution of the fast variable (the voltage). The second, intermediate timescale governs the evolution of the slow variable on the right hand branch, while the third, slow timescale governs its evolution on the left hand branch. We measure the duration of inputs on this third, slowest timescale.

Using ideas of singular perturbation theory, these simplifications allow us to reduce the response of an excitable system to a *one dimensional map* on the slow variable w in the singular limit. Similar maps have been introduced previously [Othmer and Watanabe, 1994], in the setting of two rather than three timescales. Since we take the duration of excitation to be $O(1)$ relative to the slowest timescale, it is long compared to the intermediate timescale. Therefore cells that spike are reinjected at w_{RK}^E , and the only accessible points on the left branch of \mathcal{N}^0 lie in the interval $I = [w_{RK}^E, w_{FP}^0)$ (see Fig. 2.) We denote the w coordinates of a point on the trajectory starting at w_0 by $w_0 \cdot t$. We now define a map $M_T : I \rightarrow I$ by $M_T(w_0) = w_0 \cdot T$ where it is assumed that an excitatory input is received at time $t = 0$, at the start of the map cycle, and no other excitatory inputs are received between times 0 and T . In the singular limit, the initial point (v_0, w_0) can be assumed to lie on the left branch of the nullcline \mathcal{N}^0 . Since the evolution on the right branch of \mathcal{N}^E occurs on the intermediate timescale, and the time T exceeds the duration of excitation t^* , the trajectory starting at (v_0, w_0) ends on the left branch of \mathcal{N}^0 at time T . Therefore there is no ambiguity in the definition of the map M_T . We emphasize that when we repeatedly iterate a map M_T , or we compose a sequence of maps M_{T_i} , we assume for consistency that an input arrives at the start of each map cycle and that no other inputs arrive during each cycle.

Fig. 2 (left) shows a schematic depiction of part of a typical trajectory starting at w_0 during an interval of length T following excitation. The trajectory jumps to the right nullcline after excitation is received. The jumps between the branches occur on the fast timescale (triple arrows), while the evolution on the right and left nullcline are intermediate and slow respectively (double and single arrows). The resulting map M_T is shown on the right. Note that this map has two branches. Orbits with initial conditions in the interval $j = (w_{LK}^E, w_{FP}^0)$ will jump to the active phase after excitation is received. On the other hand, orbits with initial conditions in $n = [w_{RK}^E, w_{LK}^E]$ will be blocked from reaching the active phase by the left branch of the \mathcal{N}^E nullcline. In the singular limit the region j is compressed to a single point under M_T , so that $M_T(j) = M_T(w_{RK}^E) = w_{RK}^E \cdot T$. Note that on n , the map M_T has slope less than one, corresponding to the fact that w' decreases as w approaches the fixed point w_{FP}^0 along the left branch of \mathcal{N}^0 .

Periodic excitation

While the case of periodic excitation has been analyzed in much detail elsewhere, we review it here since the analysis shares a common framework with the developments in the subsequent sections.

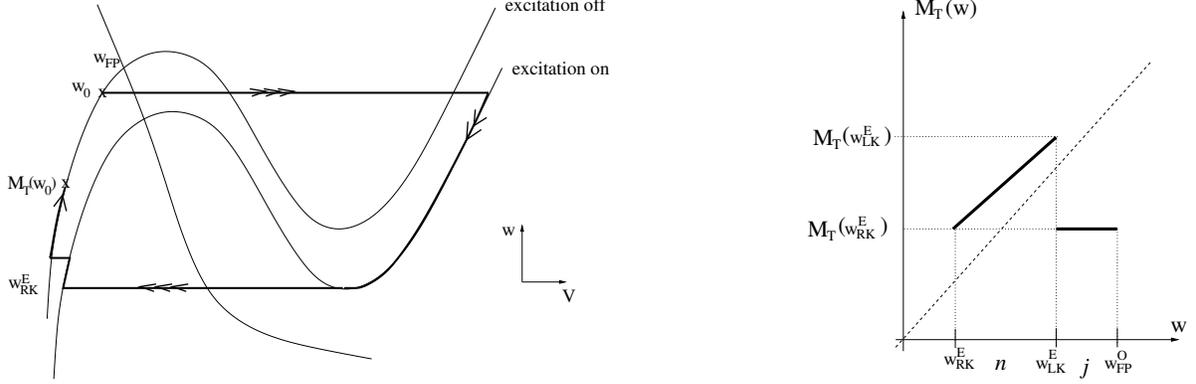


Figure 2: A schematic representation of a typical trajectory starting in region j of the left branch of \mathcal{N}^O , and ending on the same branch after time T (left). $M_T(w_0)$ is defined as the w coordinate of the endpoint of the trajectory starting at the point (v_0, w_0) , namely the point on the left branch of \mathcal{N}^O with $w = w_0$ (right).

As shown in Fig. 3, the fact that $j = (w_{LK}^E, w_{FP}^O)$ is contracted to a single point under M_T implies that all points $w_0 \in I$ get mapped to a periodic orbit of M_T in a finite number of iterations. A simple analysis shows that there exists a natural number N such that a population of cells with initial conditions distributed on I will form N synchronous clusters under N applications of the map M_T . The periodic orbit is obtained from applying N iterations of M_T to $M_T(w_{RK}^E)$ and it consists of the points $\{M_T^i(w_{RK}^E)\}_{i=1}^N$, where $M_T^N(w_{RK}^E) \in j$ (see Fig. 3). Every trajectory is absorbed into this orbit, possibly after an initial transient. This clustering state persists away from the singular limit.

More precisely, consider the following intervals, or bins,

$$j_i = (M_T^{-i}(w_{LK}^E), M_T^{-(i-1)}(w_{LK}^E)] \quad i = 1, 2, \dots \quad (2.3)$$

Note that the i -th iterate of j_i under M_T is contained in j . Therefore, since $M_T^i(j_i) \subset j$, it follows that $M_T^{i+1}(j_i) = M_T(w_{RK}^E)$, or, more generally, $M_T^l(j_i) = M_T^{l-i}(w_{RK}^E)$ for $l > i$.

We can interpret this as follows. A collection of identical oscillators subject to identical input, under the assumption that all oscillators have initial conditions in j_i , will get mapped to the interval j just prior to the i -th input. This collection of cells will respond to the i -th input by firing in unison, and will form a synchronous cluster after $i + 1$ excitatory inputs, since the interval j collapses to a single point after the cells fire.

If the bin j_i contains a fraction q of the initial conditions, then this fraction of cells will fire at the i -th input, and subsequently on every N -th input. Therefore, without knowing the distribution of initial conditions, it is not possible to know what fraction of the cell population will respond to a given input. Consider the two extreme examples: If all cells have initial conditions lying in one bin, the population of cells will respond only to every N -th input. On the other hand if initially every bin j_i contains some fraction of cells, then the population will respond to every input.

The definitions and clustering phenomenon described here carry over identically to the case with inhibition held on at any constant level. We will show in the next section that under general conditions the situation can be quite different when the interexcitation intervals (IEIs) are random and the possibility of sufficiently long IEIs exists.

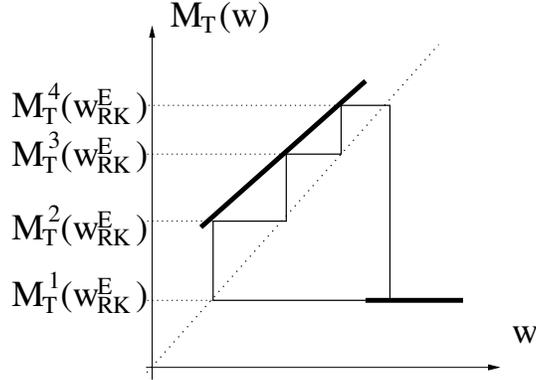


Figure 3: After a finite number of applications of the map M_T , all initial conditions will be mapped to a periodic orbit of period N . In this figure $N = 4$.

Remark 2.1 The map M_T resembles a time T map of the voltage obtained from an integrate-and-fire system. In the present case the map is defined on the slow conductance w , however. If an excitatory input fails to elicit a spike and $g(v, w)$ depends weakly on v , then the input may have little effect on the slow conductance. This is unlike the typical integrate-and-fire model, in which excitatory inputs move the cell closer to threshold.

3 The construction of a Markov chain

We next analyze the long term statistical behavior of a cell that receives excitatory input that is not periodic in time. Our goals are to determine the probability that the cell fires in response to each subsequent input since it last fired and the number of inputs the cell is expected to receive between firings, or, equivalently, the average number of failures before a spike occurs. Our results can be interpreted in the context of a population of cells as well, and this is discussed in Section 5.

A key point in our approach is that, as noted in Section 2, all firing events lead to reinjection at w_{RK}^E . Thus, when a cell fires, its dynamic history before the firing event is rendered irrelevant. This allows us to describe the cell's path through the interval $[w_{RK}^E, w_{FP}^0)$ as a Markov process with a finite number of states. The steps in this process will be demarcated by the arrival times of excitatory inputs. We assume that the IEI time T is a random variable with density ρ . Fundamental in our analysis is the assumption that the support of ρ is contained in $[S, \infty)$, where $S > 0$. As long as the frequency of the cells providing excitatory inputs is bounded, this assumption is not unreasonable. If ρ satisfies this assumption, then

the long term behavior of a population of cells is accurately captured by the asymptotics of a corresponding Markov chain with finitely many states. We show that under certain conditions this Markov chain is aperiodic and irreducible and thus has a limiting distribution. This distribution can be used to describe the firing statistics of the cell.

The states of the Markov chain

We start by again assuming that the cell receives only excitatory input. To simplify the exposition, we first assume that the input is instantaneous, so that $I(v, t) = \sum_{i=-\infty}^{\infty} g_{exc} \delta(t - t_i)(v - v_{exc})$. This assumption will be relaxed subsequently.

In the case of periodic excitation, we considered bins defined by backward iterations of w_{LK}^E , as in equation (2.3). It is now more convenient for notational reasons to consider forward iterations of w_{RK}^E to define the states in the Markov chain. Let N be the smallest number such that $w_{RK}^E \cdot NS > w_{LK}^E$, where $S > 0$ is the lower bound of the support of ρ mentioned above. Therefore, N is the maximal number of excitatory inputs that a cell starting at w_0 can receive before firing. Set

$$I_k = [w_{RK}^E \cdot kS, w_{RK}^E \cdot (k+1)S) \quad \text{if} \quad k+1 < N, \quad (3.4)$$

and define two additional states

$$\begin{aligned} I_{N-1} &= [w_{RK}^E \cdot (N-1)S, w_{LK}^E) \\ I_N &= [w_{LK}^E, w_{FP}^0] = j. \end{aligned} \quad (3.5)$$

The action of the map M_T , for the random variable T , is to transfer cells between bins. The states in the Markov chain generated by M_T are defined as (I_k, l) . A cell is in state (I_k, l) if $w \in I_k$ just prior to receiving its next excitatory input (i.e., after an iteration of M_T), and the cell has received $l-1$ excitatory inputs since it last fired (i.e., the input that is about to arrive, at the start of the next iteration of M_T , will be the l th input received since firing).

To generalize these definitions, we make a slight adjustment to I_k , $k = 2, \dots, N-1$, in the case that the excitatory input has nonzero duration. In such a case, let $S^* = S + t^*$ where t^* is the minimum possible duration of excitation t^* , and replace S by S^* in defining the bins as in (3.4) and (3.5).

Remark 3.1 Note that for $\varepsilon \neq 0$, the actual jump-up threshold lies below w_{LK}^E and is given by the minimal level of w for which an input current of duration t^* pushes (v, w) into the active phase. This value replaces w_{LK}^E in the definitions of I_{N-1} and I_N in (3.5), but we will abuse notation and continue to refer to the threshold as w_{LK}^E in the subsequent analysis. Note further that by this definition, inputs of different amplitudes or durations would give rise to different thresholds. We will return to this point in Remark 3.2 below.

Transition probabilities

To complete the definition of the Markov chain we need to compute the transition probabilities between the different states. The transitions occur at the times at which the cell receives an excitatory input. One way to think about this probability is to imagine a large pool of non-interacting cells, each of which receives a different realization of the input. Assume that cells start out in region I_N just prior to receiving excitatory inputs. After receiving these inputs, the cells spike and get reinjected at w_{RK}^E . The assumption that all cells are reset to the same point is essential in the definition of the Markov chain.

After a time T_1 , just before the arrival of its next input (i.e., its first input since firing), a particular cell will be at $w_{RK}^E \cdot T_1$. The population of cells will be distributed in some interval starting at $w_{RK}^E \cdot S^*$, as shown in Fig. 4. The probability of transitioning from the state (I_N, k) to the state $(I_j, 1)$ equals the fraction of the population of cells in the bin I_j when their next inputs arrive, which is independent of k since the reset after firing is independent of k .

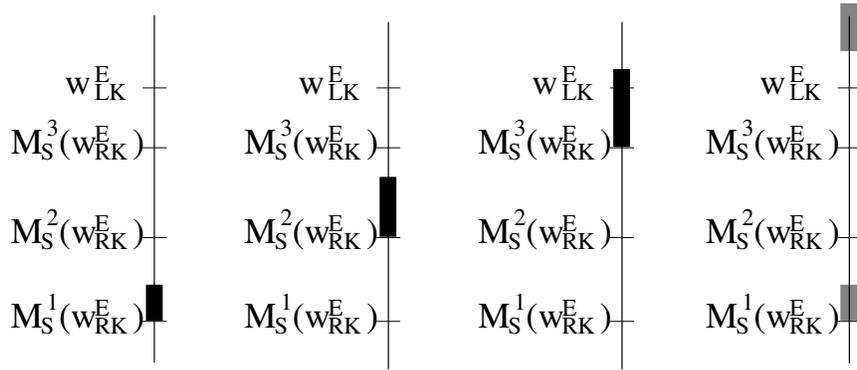


Figure 4: An example of the state of the population of cells that start in I_N and immediately receive inputs and fire. Just before the first input after firing, the cells are at $w_{RK}^E \cdot T$, where T is distributed according to ρ . Therefore all cells lie in an interval bounded below by $M_S^1(w_{RK}^E) = w_{RK}^E \cdot S$ as shown in the leftmost part of the figure. Just before their second inputs the cells are at $w_{RK}^E \cdot (T_1 + T_2)$ where both T_1 and T_2 are distributed according to ρ . After sufficiently many additional inputs arrive, some cells will lie above w_{LK}^E , such that they fire to their next inputs, while others do not, leading the rightmost distribution shown. Note that we have assumed $t^* = 0$; otherwise, each $M_S^i(w_{RK}^E)$ should be replaced by $M_{S^*}^i(w_{RK}^E)$.

If a cell's second input since firing arrives after an additional time T_2 , also chosen from the distribution ρ , then just before this input occurs, the example cell will be at $w_{RK}^E \cdot (T_1 + T_2)$. The fraction of the population starting in bin I_j just prior to their first inputs since firing and ending in bin I_k just prior to their second inputs since firing is the transition probability between the states $(I_j, 1)$ and $(I_k, 2)$. This process can be continued to compute all the transition probabilities. In the example shown in Fig. 4 there are five accessible states, $(I_1, 1)$, $(I_2, 2)$, $(I_3, 3)$, $(I_4, 3)$, and $(I_4, 4)$, so $N = 4$ and the matrix of transition probabilities

has the form

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad (3.6)$$

where each row gives the transition probabilities for cells starting in the same state.

If initial conditions are selected randomly, then some cells may initially lie in transient states that cannot be reached subsequent to reset from w_{RK}^E . We neglect such states in the Markov chain, since including them does not affect the statistics we consider.

The transition probabilities between non-transient states are easiest to compute in the time domain. We first compute the transition probabilities between the states (I_j, l) and $(I_k, l+1)$ for $l \leq j < k \leq N$. Consider independent, identically distributed random variables T_i each with probability density ρ . Let $\sigma_l = \sum_{i=1}^l T_i$ denote the sum of these random variables. The transition probability between the states (I_j, l) and $(I_k, l+1)$ is given by

$$p_{(j,l) \rightarrow (k,l+1)} = P[(I_k, l+1)|(I_j, l)] = P[\sigma_{l+1} \in [kS^*, (k+1)S^*] \mid \sigma_l \in [jS^*, (j+1)S^*]]. \quad (3.7)$$

Since the probability density $\rho^{(l)}$ of the sum σ_l can be computed recursively by the convolution formula

$$\rho^{(l)}(t) = \int \rho^{(l-1)}(u)\rho(t-u)du$$

the conditional probabilities in the expression above can be evaluated. In particular, since

$$P[\sigma_{l+1} \in [z, z + \Delta z] \ \& \ \sigma_l \in [w, w + \Delta w]] \approx \rho^{(l)}(w)\rho(z-w)\Delta w \Delta z$$

it follows that

$$p_{(j,l) \rightarrow (k,l+1)} = \frac{\int_{kS^*}^{(k+1)S^*} \int_{jS^*}^{(j+1)S^*} \rho^{(l)}(w)\rho(z-w)dwdz}{\int_{jS^*}^{(j+1)S^*} \rho^{(l)}(z)dz}. \quad (3.8)$$

Finally, we define the transition probabilities from the states (I_N, l) . By our assumption, a cell in one of these states fires when it receives an excitatory input. Therefore the next state must be of the form $(I_j, 1)$. As discussed above, once a cell fires, it has no memory of the number of excitatory inputs it received prior to this event. Therefore, the transition probability from (I_N, l) to $(I_j, 1)$ is the same for all l . This transition probability can be obtained as

$$p_{(N,l) \rightarrow (j,1)} = P[(I_j, 1)|(I_N, l)] = \int_{jS^*}^{(j+1)S^*} \rho(t)dt.$$

Since no other transitions than the ones described above are possible, this completes the definition of the Markov chain.

Remark 3.2 As mentioned in Remark 3.1, the threshold w_{LK}^E in (3.5) depends on input amplitude and duration. If input amplitudes are not constant, then there will no longer be

a single value N such that a cell fires to its next input if and only if it is in a bin of the form (I_N, j) , no matter how I_N is defined. In such a case, we can, for example, use the threshold defined for the maximal relevant input amplitude to define I_N . Then, the probability distribution of input amplitudes can be used to compute probabilities $p_{(N,l) \rightarrow (N,l+1)}$ and $p_{(N,l) \rightarrow (j,1)}$. For some range of l values, $p_{(N,l) \rightarrow (N,l+1)}$ will be nonzero, and there will be a maximal value l such that no cell requires more than l inputs to fire, and hence (I_N, l) is the last bin in the chain. The case of variable input durations can be handled similarly.

4 Limiting distributions of the Markov chain and their interpretation

We next consider the long term behavior of the Markov chains defined above, and interpret this behavior in terms of the original fast-slow system. For simplicity we assume that the support of ρ is a closed interval. Similar ideas can be used to extend these results to more general cases.

We remind the reader that for a finite state Markov chain with M states and transition matrix A , the probability distribution $\pi = (\pi_1, \dots, \pi_M)$ is a *stationary distribution* if $\sum_i \pi_i A_{i,j} = \pi_j$ for all j . The stationary distribution π is a limiting distribution if

$$\lim_{n \rightarrow \infty} \{A^n\}_{i,j} = \pi_j.$$

A Markov chain is *irreducible* if for any two states i and j , there exists a finite n such that $\{A^n\}_{i,j} > 0$. In other words, there is a nonzero probability of transition between any two states in a finite number of steps. An irreducible finite state Markov chain has a unique stationary distribution [Hoel et al., 1972]. The period d_i of the state i in a Markov chain is the greatest common divisor of the set $\{n | \{A^n\}_{i,i} > 0\}$. For an irreducible chain, all states have the same period d . Such a chain is called *aperiodic* if $d = 1$, and *periodic* if $d > 1$. The following theorem relates these concepts

Theorem 4.1 ([Hoel et al., 1972], p. 73) *For an aperiodic, irreducible Markov chain the stationary distribution is a limiting distribution. If the chain is periodic with period d , then for each pair of states i, j there is an integer r , $0 \leq r < d$, such that $\{A^n\}_{i,j} = 0$ unless $n = md + r$ for some nonnegative integer m , and*

$$\lim_{m \rightarrow \infty} \{A^{md+r}\}_{i,j} = d\pi_j.$$

Conditions for the existence of a limiting distribution

We next show that under very general conditions, the Markov chain constructed in the previous section is irreducible and aperiodic and therefore has a limiting distribution. First notice that by construction there is a nonzero probability of transition from the state $(1, 1)$ to any other state in the chain in a finite number of steps. Similarly, since after a finite number

of steps a cell will spike, the probability of transition from any state to the state (1, 1) in a finite number of steps is nonzero as well. Therefore, the Markov chain under consideration is always irreducible.

We next consider conditions under which the Markov chain is aperiodic. It is sufficient to start with a continuum ensemble of cells at w_{RK}^E . If these cells are subject to different realizations of the input, and after a finite number inputs, a nonzero fraction of cells occupies every bin, then the transition matrix is aperiodic.

To start, note that if ρ is supported on a single point T , so that the input is periodic, then the Markov chain will be periodic. Each point $M_T^j(w_{RK}^E)$ is contained in bin I_j . Therefore the states of the Markov chain are $\{(i, i)\}_{i=1}^N$. In the case depicted in Fig. 3, for example, the transition matrix has the form

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}. \quad (4.9)$$

Consider next what happens as the support of ρ is enlarged to $[T, T + \delta]$ in this example. For small δ , there is no change in the structure of the transition matrix. For some δ_0 sufficiently large however, we have $M_{3(T+\delta_0)}(w_{RK}^E) = w_{LK}^E$, and when $\delta > \delta_0$, as shown in Fig. 4, a fraction of the cells will be in bin (3,3) after 3 inputs, while another fraction will be in bin (4,3). After one more excitatory input the cells in bin (4,3) will fire and transition to (1,1), while the cells in (3,3) will require two more inputs to fire. Therefore, in addition to the states (1,1), (2,2), (3,3), (4,4) that were relevant in the periodic case, the new state (4,3) becomes part of the Markov chain. The transition matrix between these five states for $\delta > \delta_0$ has the form

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 - \varepsilon(\delta) & \varepsilon(\delta) & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (4.10)$$

where $\varepsilon(\delta_0) = 0$, ε increases with δ , and the states are ordered as (1, 1), (2, 2), (3, 3), (4, 3) and (4, 4). It can be checked directly that this transition matrix is aperiodic. This is also a consequence of the following general theorem.

Theorem 4.2 *Suppose that the support of ρ is the interval $[S, U]$, and let $U^* = U + t^*$ where t^* is the duration of excitation. The transition matrix defined in Section 3 is aperiodic, and the corresponding Markov chain has a limiting distribution, if and only if $M_{iU^*}(w_{RK}^E) > w_{LK}^E$ for an integer $0 < i < N$, where N is defined by $M_{(N-1)S^*}(w_{RK}^E) < w_{LK}^E < M_{NS^*}(w_{RK}^E)$, as in Section 3.*

Remark 4.3 In the example discussed above, as the support of ρ widens, the periodic transition matrix in (4.9) is replaced by the aperiodic matrix (4.10). Since the limiting

behavior of periodic and aperiodic Markov chains is very different, the system can be thought of as undergoing a bifurcation as δ is increases past δ_0 .

Proof If $M_{iU^*}(w_{RK}^E) \leq w_{LK}^E$ for all integers $0 < i < N$, then it can be checked directly that the only states that are achievable from an initial condition $M_S^*(w_{RK}^E)$ have the form (i, i) for $0 < i \leq N$. Therefore the transition matrix is of size $N \times N$ and has the form of the matrix (4.9), with ones on the superdiagonal in all rows except that last, which has a one in its first column. Thus, the Markov chain is periodic and has no limiting distribution.

Assume instead that $M_{iU^*}(w_{RK}^E) > w_{LK}^E$ for an integer $0 < i < N$. Consider a continuum of cells that have just spiked and been reset to w_{RK}^E and are receiving independent realizations of the input. First note that just before the i^{th} input following the spike, there will be a nonzero fraction of cells in all bins $(k, i), k \geq i$ that are part of the Markov chain.

The condition $M_{iU^*}(w_{RK}^E) > w_{LK}^E$ and the fact that the support of ρ is an interval imply that some cells will fire on the $i^{th}, (i+1)^{st}, \dots$, and N^{th} inputs after reset. Correspondingly, there will be cells in all bins $(k, 1), k \geq 1$ after the i^{th} input, cells in all bins $(k, 2), k \geq 2$ and $(k, 1), k \geq 1$ after the $(i+1)^{st}$ input, and so on, until all bins of the form (k, j) with $k \geq j$ and $j \in \{1, \dots, N-i+1\}$ are non-empty after N inputs. Similarly, some cells will fire again after input number $2i$, with some other cells firing on each input from $2i+1$ up to $2N$, and after $2N$ inputs, all bins (k, j) with $k \geq j$ and $j \in \{1, \dots, 2N-2i+1\}$ will be non-empty. Continuing this argument inductively shows that all states are necessarily occupied just before the arrival of the $(cN)^{th}$ input, where $c = \lceil \frac{N-1}{N-i} \rceil$, namely the least integer greater than or equal to $(N-1)/(N-i)$, such that cN is bounded above by $(N-1)N$, since $i \leq N-1$. \square

Interpretation of the invariant density

Suppose that the Markov chain that we have derived has a limiting distribution Q . The limiting distribution has two interpretations [Taylor and Karlin, 1998]. Assume that a cell is subjected to excitatory input for a long time. If w is recorded just prior to the moments when the cell receives an excitatory input, then $Q[(I_j, l)]$ is the fraction of recordings that are observed to fall inside (I_j, l) . Consequently, the total mass of the distribution Q that lies in the bins (I_N, j) , namely $\sum_{j=1}^N Q[(I_N, j)]$, is the probability that a cell will respond to the next excitatory input by firing. Similarly, we can think of $Q[(I_j, l)] / \sum_{k=l}^N Q[(I_k, l)]$ as the probability that a cell will be found to have $w \in I_j$ just prior to receiving its l -th excitatory input since it fired last.

Note that the firing probability $\sum_{j=1}^N Q[(I_N, j)]$ is not the reciprocal of the average number of failures before a spike. However, the average number of failures, call it E_f , can be computed once the $Q[(I_j, l)]$ are known. Clearly,

$$E_f = \sum_{j=1}^{N-1} jF_j, \quad (4.11)$$

where each F_j denotes the probability that exactly j failures occur. Note that j failures occur precisely when $j + 1$ conditions are met. That is, for each $i = 1, \dots, j$, just prior to the i^{th} input since the last firing, the trajectory falls into a bin of the form (I_{k_i}, i) , where $i \leq k_i < N$. Further, the trajectory ends in bin $(I_N, j + 1)$ before the $(j + 1)^{\text{st}}$ input, to which the cell responds by firing. The quantity

$$\frac{\sum_{k=i}^{N-1} Q(I_k, i)}{\sum_{k=i}^N Q(I_k, i)}$$

equals the probability that, just before the i^{th} input arrives, a cell will lie in a bin (I_k, i) with $k < N$. Thus, the failure probabilities are given by

$$F_j = \left(\frac{\sum_{k=1}^{N-1} Q(I_k, 1)}{\sum_{k=1}^N Q(I_k, 1)} \right) \left(\frac{\sum_{k=2}^{N-1} Q(I_k, 2)}{\sum_{k=2}^N Q(I_k, 2)} \right) \cdots \left(\frac{\sum_{k=j}^{N-1} Q(I_k, j)}{\sum_{k=j}^N Q(I_k, j)} \right) \left(\frac{Q(I_N, j + 1)}{\sum_{k=j+1}^N Q(I_k, j + 1)} \right). \quad (4.12)$$

Of course, while some of the $Q(I_k, j)$ may be zero, the exclusion of transient bins in our Markov chain construction ensures that all of the sums in the denominators of (4.12) will be positive for every population of cells.

5 Extensions of the Markov chain approach

The ideas introduced in the previous section are applicable in many other settings. Here we outline extensions to heterogeneous populations of excitable cells subject to identical periodic input and to general excitable systems.

Populations and heterogeneity

It is important to remember that the limiting distribution computed from a Markov chain cannot be used directly to determine what fraction of a population of *identical* cells subject to *identical* stochastic trains of excitatory inputs will respond to a given input. As noted earlier, in this case it is necessary to specify the distribution of initial conditions to compute this fraction. On the other hand, it is possible to draw conclusions about a population of identical cells subject to *different realizations* of the stochastic train of excitatory inputs. In this case, all of the quantities discussed in the previous subsection can be reinterpreted with respect to the firing statistics of a population of cells. For example, the fraction of cells responding with a spike to a given excitatory input equals the probability of a single cell firing to a particular excitatory input, that is $\sum_{j=1}^N Q[(I_N, j)]$.

We can also apply similar ideas to the case of a heterogeneous population of excitable cells all subject to the same periodic input, say of period T . Heterogeneity in intrinsic dynamics may lead to different rates of evolution for different cells in the silent phase, but this disparity can be eliminated by rescaling time separately for each cell, so that all cells evolve at the same unit speed in the silent phase. As a result of this rescaling, the heterogeneity

in the population will become a heterogeneity in firing thresholds. Denote the resulting distribution of thresholds by ϕ , such that for any $t > 0$, $\int_0^t \phi(\tau) d\tau$ gives the fraction of cells with thresholds below t . The distribution ϕ will have support on some interval $[mT, nT]$ for nonnegative integers $m < n$. Thus, for $i \geq 1$, $\delta_i = \int_{(i-1)T}^{iT} \phi(\tau) d\tau$ gives the fraction of cells that will fire in response to the i th input, which is nonzero for $i = m + 1, \dots, n$.

We can define Markov chain states I_j , for $j = 1, \dots, n$, by stating that a cell is in state I_j if it has received $j - 1$ inputs since it last fired. The transition probability $P(j, j + 1)$ from state I_j to state I_{j+1} is 1 if $j \leq m$; if $j = m + 1$, then $P(m + 1, m + 2) = 1 - \delta_{m+1}$ while the probability of transition from state I_m to state I_1 is $P(m + 1, 1) = \delta_{m+1}$; and if $m + 1 < j \leq n - 1$, then

$$P(j, j + 1) = 1 - \gamma_j := \frac{1 - \delta_j - \dots - \delta_{m+1}}{1 - \delta_{j-1} - \dots - \delta_{m+1}} = 1 - \frac{\delta_j}{1 - \delta_{j-1} - \dots - \delta_{m+1}},$$

while $P(j, 1) = \gamma_j = \delta_j / (1 - \delta_{j-1} - \dots - \delta_{m+1})$. Finally, the transition probability from state I_n to state I_1 is 1, and all other transitions have zero probability. If we set $\gamma_{m+1} = \delta_{m+1}$, then the transition matrix takes the form

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 & 0 & \dots & 0 & 0 \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \gamma_{m+1} & 0 & 0 & 0 & \dots & 1 - \gamma_{m+1} & 0 & \dots & 0 & 0 \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \dots & \cdot & \cdot & \dots & \cdot & \cdot \\ \gamma_{n-1} & 0 & 0 & 0 & \dots & 0 & 0 & \dots & 0 & 1 - \gamma_{n-1} \\ 1 & 0 & 0 & 0 & \dots & 0 & 0 & \dots & 0 & 0 \end{bmatrix}.$$

As long as $m < n - 1$, such that the distribution ϕ has support on an interval of length greater than T , not all cells will fire together. Under this condition, there exists $i < n$ such that $\delta_i \neq 0$ and correspondingly $\gamma_i \neq 0$. In this case, the proof of Theorem 4.2 immediately generalizes to show that there exists a sufficiently large number of iterations N such that $\{A^r\}_{i,i} > 0$ for all $r > N$, which implies that A is aperiodic and the Markov chain has a limiting distribution. That is, if the heterogeneity is quite weak, then all cells will always respond to the same inputs in a train, giving an unreliable population response. With a stronger degree of heterogeneity, population responses will disperse, such that in the limit of an infinitely large population, every input will evoke a response from some nonempty subset of the cells, and the statistics of the population response will be given by the limiting distribution of the Markov chain.

General excitable systems and integrate-and-fire type models

An excitable system can be characterized by the existence of a stable rest state and a threshold. In such a system, when a perturbation pushes a trajectory from the rest state across the threshold, the trajectory makes a significant excursion before returning to a small neighborhood of the rest state. Suppose that an excitable system receives transient inputs of a stereotyped form. If we select a time T , then we can define a map $M_T(u)$ having as its domain a set of starting states for the system. To do this, we can simply select each state u in the set as an initial condition, apply the input, and then allow the system to evolve for time T after the onset of the input. The output of the map would be the new state $u(T)$ of the system after this evolution. Like a phase response curve (PRC), the map M_T captures the effect of an input, but is much more general. In particular it is n -dimensional, where n is the number of components in u , and thus in general it captures more detailed information about a system than the PRC.

In theory, bins in n -dimensional space could be defined to implement the Markov chain approach for an n -dimensional map. Indexing these bins efficiently and computing transition probabilities between them could become problematic when n is not small, however. On the other hand, the Markov chain approach discussed in the previous section can be applied directly to systems for which $M_T(u) = u(T)$ lie (approximately) on an interval in the phase space of the system, and which are reset to a fixed value u_{reset} after crossing threshold. The assorted versions of the integrate-and-fire (IF) model (leaky IF, quadratic IF, exponential IF, and so on) in the subthreshold regime satisfy both of these conditions. There is an important difference between the excitable systems considered in detail here and the IF model, however. When the results derived for excitable systems are applied to neuronal models, the Markov chain will be defined using a slowly changing ionic conductance, or an associated activation or inactivation, while in the case of the IF model, it would be defined in terms of the voltage. As a consequence, for the IF case, one would need to take into account the jumps in voltage due to synaptic inputs when defining the states of the Markov chain and computing the transition probabilities. Otherwise, the construction would be equivalent to the one outlined in Sections 3-4.

6 An example: the TC model

A prototypical representative of the class of models to which the analysis outlined in the preceding sections is applicable is a model for a thalamocortical relay (TC) cell relevant for the study of Parkinson's disease (PD) and deep brain stimulation (DBS). The TC model that we consider takes a similar form to the reduced TC model in [Rubin and Terman, 2004], namely

$$\begin{aligned} C_m v' &= -I_L - I_T - g_{exc} s_{exc}(v - v_{exc}) - g_{inh} s_{inh}(v - v_{inh}) \\ w' &= \phi(w_\infty(v) - w) / \tau(v). \end{aligned} \tag{6.13}$$

Here, the leak current $I_L = g_L(v - v_L)$ and the T-type or low-threshold calcium current $I_T = g_T m_\infty(v) w(v - v_{Ca})$, with parameter values $C_m = 1 \mu F/cm^2$, $g_L = 1.5 mS/cm^2$, $v_L = -68 mV$, $g_T = 5 mS/cm^2$, $v_{Ca} = 90 mV$, $g_{exc} = 0.08 mS/cm^2$, $v_{exc} = 0 mV$, $g_{inh} = 0.12 mS/cm^2$, $v_{inh} = -85 mV$, $\phi = 3.5$ and functions $m_\infty(v) = (1 + \exp(-(v + 35)/7.4))^{-1}$, $w_\infty(v) = (1 + \exp((v + 61)/9))^{-1}$, and $\tau(v) = 10 + 400/(1 + \exp((v + 50)/3))$. The levels of s_{exc} and s_{inh} will be determined by stochastic processes discussed below. The assumptions that we make regarding the presence of three timescales in the dynamics are not unreasonable in this model, as shown in [Rubin and Terman, 2004, Stone, 2004].

In this section we show how to analytically compute the transition matrix for this model and compute the stationary distribution for the corresponding Markov chain for two particular IEI distributions, one uniform and one normal, both with inhibition held off for all time and with inhibition held on for all time. We compare these stationary distributions with those found by numerical simulations and find good agreement.

Uniform IEI distribution

Consider system (6.13) with constant inhibition, which may be on or off. Suppose that we fix the duration of each excitatory input at 10 msec. Each IEI interval T , corresponding to the time interval from the offset of one input (10 msec after it began) to the onset of the next input, is chosen from a uniform distribution on $[20, 60]$ msec. The bin intervals I_k are defined as described in Section 3; note that $S^* = S + t^* = 20 + 10 = 30$ msec. We assume that, if an excitatory input causes a cell to fire, then the input is turned off as the firing cell is reset, rather than 10 msec after the reset. This is equivalent to replacing the instantaneous reset with a 10 msec absolute refractory period and is a way to weaken the biologically unrealistic instantaneous reset condition without unduly complicating the analysis. Since $w_{LK}^E \in (w_{RK}^E \cdot (N - 1)S^*, w_{RK}^E \cdot NS^*)$, we have $I_1 = [w_{RK}^E \cdot 20, w_{RK}^E \cdot 50)$, $I_2 = [w_{RK}^E \cdot 50, w_{RK}^E \cdot 80)$, \dots , $I_{N-1} = [w_{RK}^E \cdot 20 + (N - 2)S^*, w_{LK}^E)$, $I_N = [w_{LK}^E, w_{FP}^0)$.

For the default parameters of system (6.13) and the excitation characteristics described here, with inhibition held at $s_{inh} = 0$, we find $N = 3$, while with inhibition held at $s_{inh} = 1$, we find $N = 5$. With $s_{inh} = 0$, the states of the Markov chain are the bins $(1, 1)$, $(2, 1)$, $(2, 2)$, $(3, 2)$, $(3, 3)$. The transition matrix for this case, call it P^0 , is computed analytically in Appendix A. The unit dominant eigenvector of $(P^0)^T$ gives the limiting distribution

$$v^0 = [.3404 \ .1135 \ .0922 \ .3617 \ .0922]^T.$$

As discussed in Section 4, the values of v^0 represent the likelihood that a cell is found in a given bin, if bin memberships are recorded just prior to the onset of an excitatory input. For comparison, we simulated a single cell, modeled by (6.13) with the technical modification mentioned in Remark 6.1 below, over 70 sec, after an initial transient of 10 sec. This simulation yielded the vector

$$v_{num}^0 = [.3276 \ .1289 \ .0878 \ .3629 \ .0878]^T$$

of proportions of inputs during which the cell belonged to each relevant bin, which agree

nicely with the analytically computed expectations v^0 . Over the entire simulation, the cell never failed to fire to three consecutive inputs. Grouping our analytical values indicates that, just before onset of excitation, on 34.04% of the observations we expect $w \in [w_{RK} \cdot 20, w_{RK} \cdot 50)$, on 20.57% of the observations we expect $w \in [w_{RK} \cdot 50, w_{RK} \cdot 75.5)$, and on 45.39% of the observations we expect $w > w_{RK} \cdot 75.5$. In particular, this implies that a cell will fire to roughly 45% of its inputs, for these choices of parameters. Further, from equations (4.11)-(4.12) with Q values given by entries of v^0 , it is expected that a successful input will be followed by 1.20 inputs that fail to elicit a spike, before the next successful input occurs; this is in good agreement with direct simulation results showing a mean of 1.19 unsuccessful inputs between successful ones. As discussed in Section 5, these results can also be interpreted in terms of a population of identical cells, as long as the cells receive different realizations of the input.

With $s_{inh} = 1$, the states of the Markov chain are the bins (1,1), (2,1), (2,2), (3,2), (3,3), (4,2), (4,3), (4,4), (5,2), (5,3), (5,4) and (5,5). The transition matrix, P^I , is computed analytically in Appendix B. $(P^I)^T$ has the unit dominant eigenvector v^I , which is shown below together with the bin proportions listed in vector v_{num}^I obtained from direct numerical simulation for 80 sec, with a transient consisting of the first 10 sec removed from consideration:

$$\begin{aligned} v^I &= [.2290 \ .0763 \ .0859 \ .1622 \ .0215 \ .0569 \ .0624 \ .0005 \ .0004 \ .2211 \ .0833 \ .0005]^T \\ v_{num}^I &= [.2272 \ .0784 \ .0726 \ .1927 \ .0194 \ .0388 \ .0669 \ .0022 \ .0007 \ .2171 \ .0820 \ .0002]^T. \end{aligned}$$

Grouping the bins (I_5, j) reveals that approximately a cell will fire in response to about 30% of its inputs. This is a higher failure rate than in the case of $s_{inh} = 0$, which is expected because inhibition changes the locations of w_{LK}^E and w_{RK}^E , such that the time to evolve from w_{RK}^E to w_{LK}^E is longer with inhibition on than with inhibition off. Similarly, based on v^I , the case of $s_{inh} = 1$ has a substantially higher expected number of failed inputs between spike-inducing inputs, namely $E_f = 1(.0013) + 2(.7240) + 3(.2731) + 4(.0016) = 2.28$ from equations (4.11)-(4.12), than the 1.20 expected in the case of $s_{inh} = 0$; direct numerical simulation gives a similar estimate of $E_f = 2.35$ in this case.

Remark 6.1 There is an additional complication that arises with a nonzero duration of excitation. If an excitatory input arrives but fails to make a cell fire, the cell jumps from the left branch of \mathcal{N}^0 to the left branch of \mathcal{N}^E , and evolves on this branch of \mathcal{N}^E while the input is on. The rates of flow on the left branches of \mathcal{N}^0 and \mathcal{N}^E may differ, however. This can be easily accounted for in selecting the boundaries for all bins I_k with $k < N - 1$. However, this complicates selection of the upper boundary for I_{N-1} , which relates to a left knee location. To get around this, we adjust the flow in our simulations of (6.13) in this example such that the rate of change is the same on the left branches of \mathcal{N}^0 and \mathcal{N}^E . In theory, the way that we do this introduces the possibility that solutions may escape from the left branch of \mathcal{N}^0 in the vicinity of w_{LK}^0 and fire without receiving input. However, in practice this does not occur because w_{LK}^0 is sufficiently large, relative to the IELs, that it is never reached.

Normal, or other, IEI distributions

If IEIs are taken from a non-uniform distribution, we can still use equation (3.8) to obtain the transition matrix elements analytically, albeit with numerical evaluation of the integrals that arise, which are analogous to those given in Appendices A and B for the uniform case. Once the transition matrix is obtained, the limiting distribution for the Markov chain is computed as the unit dominant eigenvector of its transpose, as previously.

For example, consider IEIs of the form $T = t_{ref} + X$, where t_{ref} is a fixed constant and X is selected from a truncated normal distribution. In particular, suppose that $t_{ref} = 20$ msec, that \tilde{X} is selected from a normal distribution with mean 20 msec and standard deviation 10 msec, and that

$$X = \begin{cases} 0, & \tilde{X} < 0, \\ \tilde{X}/M_X, & 0 \leq \tilde{X} \leq 40, \\ 40, & 40 < \tilde{X}, \end{cases}$$

where M_X is a constant correction factor such that $\int_{\mathbf{R}} X = 1$. This is a reasonable choice that keeps IEIs within the bounds present in the example in the previous subsection. For this example, with other simulation parameters fixed as in the previous subsection and $s_{inh} = 0$, we obtained the transition matrix

$$P_n^0 = \begin{bmatrix} 0 & 0 & .1474 & .8526 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ .8576 & .1424 & 0 & 0 & 0 \\ .8576 & .1424 & 0 & 0 & 0 \end{bmatrix}$$

for bins (1, 1), (2, 1), (2, 2), (3, 2), (3, 3). The dominant eigenvector v_n^0 of $(P_n^0)^T$ matches closely with a vector $(v_n^0)_{num}$ obtained from direct counting of bin membership in the last 90 sec of a 100 sec numerical simulation:

$$\begin{aligned} v_n^0 &= [.4073 \quad .0670 \quad .0594 \quad .4108 \quad .0594]^T \\ (v_n^0)_{num} &= [.3883 \quad .0812 \quad .0625 \quad .4070 \quad .0610]^T. \end{aligned}$$

Based on these results, we expect that a cell will respond to approximately 47% of the inputs, with an average of 1.15 failures between successful spikes from equations (4.11)-(4.12); this agrees nicely with the estimate 1.13 that we obtained from direct numerical simulations.

For completeness, we conclude with the results of an analogous calculation done with

$s_{inh} = 1$. This yields

$$P_n^I = \begin{bmatrix} 0 & 0 & .2548 & .7249 & 0 & .0203 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & .7352 & 0 & .2642 & .0006 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .1074 & 0 & 0 & .5612 & .3314 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & .1448 & .8552 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ .8576 & .1424 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ .8576 & .1424 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ .8576 & .1424 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

for bins $(1, 1), (2, 1), (2, 2), (3, 2), (3, 3), (4, 2), (4, 3), (5, 2), (5, 3), (5, 4)$. The corresponding dominant eigenvector v_n^I and an example vector $(v_n^I)_{num}$ of state occupancy probabilities from direct simulations are

$$\begin{aligned} v_n^I &= [.2638 \ .0438 \ .0672 \ .2234 \ .0072 \ .0169 \ .0701 \ .2133 \ .0942]^T \\ (v_n^I)_{num} &= [.2587 \ .0505 \ .0697 \ .2132 \ .0121 \ .0263 \ .0637 \ .2329 \ .0728]^T, \end{aligned}$$

where we have omitted the $(5, 2)$ entry since the probability of occupancy there is less than 0.5×10^{-5} . These results imply that a cell with $s_{inh} = 1$ will respond to approximately 31% of excitatory inputs and is thus less reliable than a cell with $s_{inh} = 0$. The average number of failures between spikes is 2.27, from equations (4.11)-(4.12), which is similar to the 2.24 obtained from direct simulations and exceeds that found with $s_{inh} = 0$.

7 Inhibitory and excitatory inputs

We next consider the more complex case in which the postsynaptic cell receives a stochastic excitatory input train while subjected to modulation by an inhibitory input. The corresponding analysis illustrates how switches between epochs with different bin transition characteristics can be handled naturally within the Markov chain framework through products of transformation matrices. Moreover, the example that we present will be used in the next section to provide insights about a possible mechanism for the therapeutic effectiveness of DBS.

Transition matrices

In this section we assume that for a system of the form (2.1) with input given by (2.2), the function $s_{inh}(t)$ turns on and off abruptly, with a relatively long period between transitions. Previously, we have discussed how to derive transition matrices for the case of constant inhibition of any level. Thus, if we can also derive transition matrices encoding the probabilities of passing between various bins during the onset and offset of inhibition, then we

can form a transition matrix for the time from one inhibitory offset to the next, or for other corresponding times, by matrix multiplication.

One complication in this derivation is that the bin boundaries differ between the cases of zero and nonzero inhibition. Even with the assumption that the rate of evolution in the silent phase is independent of input level, differences in bin boundaries remain due to differences in right knee positions, leading to different starting points in the silent phase, and differences in left knee positions, leading to different cut-offs for firing. A second difficulty is that, due to the stochasticity of the input trains, a change in inhibition level may occur at any time relative to the start of the particular IEI during which it happens, or even at a time when excitation is on.

To handle these issues, we define the transition matrix $P^{0 \rightarrow I}$ for the onset of inhibition such that its (i, j) entry encodes the probability that a cell is in bin i of the inhibition-off partition at the moment preceding the arrival of the last excitatory input before inhibition turns on, call it time t_{last} , and is in bin j of the inhibition-on partition at the moment just before the arrival of the first excitatory input after inhibition turns on. There are several advantages to basing the cycle length only on excitatory input times, rather than the time when inhibition turns on. Most importantly, in this formulation, it does not matter whether inhibition turns on while the excitation is still on (that is, for $t \in (t_{last}, t_{last} + d)$, where d is the duration of excitation) or after it is off, if we continue to assume that the rate of silent phase evolution is input-independent. Let $w_{RK}^{(i,j)}$ ($w_{LK}^{(i,j)}$) denote the right (left) knee of nullcline $\mathcal{N}^{(i,j)}$, where $i \in \{E, 0\}$ indicates whether excitation is on or off and $j \in \{I, 0\}$ indicates whether inhibition is on or off. The decision as to whether the excitation that arrives at time t_{last} causes the cell to fire or not is based on the knee positions without inhibition, and if a firing occurs, instantaneous reset to $w_{RK}^{(E,0)}$ follows; we neglect the probability zero case of excitation and inhibition turning on at precisely the same moment, and therefore if a cell fires, then inhibition will always turn on after the cell is reset.

One additional complication may arise due to the fact that $w_{RK}^{(E,0)} < w_{RK}^{(E,I)}$. That is, at the onset of inhibition, w may lie below the lower boundary of the inhibition-on partition. To account for this possibility requires the inclusion of additional bins in this partition, with a corresponding adjustment of the matrix P^I to allow for matrix multiplication.

Analogously to the case of $P^{0 \rightarrow I}$, the (i, j) entry of the transition matrix $P^{I \rightarrow 0}$ for the offset of inhibition encodes the probability that a cell is in bin i of the inhibition-on partition at the moment preceding the arrival of the last excitatory input before inhibition turns off and is in bin j of the inhibition-off partition at the moment just before the arrival of the first excitatory input after inhibition turns off. The complication in this case is that post-inhibitory rebound (PIR) may occur if, when the offset of inhibition occurs, either excitation is off and $w > w_{LK}^{(0,0)}$ or excitation is still on and $w > w_{LK}^{(E,0)}$. Rebound leads to reinjection into the silent phase followed by evolution there. However, the duration of this evolution, from the time of reinjection to the time that the next excitatory input arrives, depends on when the inhibition turned off, relative to the time of the previous input, which complicates calculations.

Given that $P^{0 \rightarrow I}$, $P^{I \rightarrow 0}$ can be computed, the appropriate transition matrix for the case

of excitatory and inhibitory input trains is obtained by multiplication of the transposes of the separately computed transition matrices $P^0, P^I, P^{0 \rightarrow I}$, and $P^{I \rightarrow 0}$. Specifically, using $[M]^T$ to denote the transpose of matrix M and $(M)^n$ to denote M to the n^{th} power, the matrix

$$[P^{I \rightarrow 0}]^T ([P^I]^T)^m [P^{0 \rightarrow I}]^T ([P^0]^T)^n$$

is the transition matrix to consider after the offset of inhibition, just before the arrival of the first subsequent excitatory input. The (i, j) entry of this matrix gives the probability that a cell that starts in bin j of the inhibition-off partition at this moment ends up in bin i of the inhibition-off partition at the analogous moment after the next offset of inhibition, assuming that n excitatory inputs arrive before inhibition turns on and m excitatory inputs arrive after inhibition turns on. Of course, when the IEIs and the durations of inhibitory on and off periods are selected randomly from distributions, the probabilities that different numbers of excitatory inputs arrive are also random; in Appendix C, these are calculated for the example of uniform distributions.

In the general discussion given here, let us assume that m, n are bounded above by $M, N < \infty$, respectively. Given this, we would like to claim that the expected proportion of trials for which a cell will belong to each inhibition-off bin are given by the entries of the unit dominant eigenvector of the matrix

$$\sum_{m=1}^M \sum_{n=1}^N c_{m,n} [P^{I \rightarrow 0}]^T ([P^I]^T)^m [P^{0 \rightarrow I}]^T ([P^0]^T)^n, \quad (7.14)$$

where each coefficient $c_{m,n}$ denotes the probability of occurrence of the corresponding exponent pair (m, n) . Substantiating this claim necessitates justifying whether this eigenvector exists and truly represents a limiting distribution, which will be true if the matrix in (7.14) is irreducible and aperiodic, as discussed in Section 4. Recall that the proof of Theorem 4.2 gives an upper bound c for the number of inputs after which occupancy of all states is guaranteed, under the assumptions of the theorem, for constant inhibition. Now, let c_I, c_0 denote the respective upper bounds for matrices P^I, P^0 , based on the IEI distribution. A sufficient pair of conditions to guarantee the existence of a limiting distribution for (7.14) are that if $p(m) > 0$, then $m \geq c_I$, while if $p(n) > 0$, then $n \geq c_0$. If these conditions are violated, then of course the states of the system may still tend to some limiting distribution, particularly if $p(m), p(n)$ are substantial for some values above the respective bounds, but this must be verified numerically.

Similarly, the unit dominant eigenvector of the matrix

$$\sum_{m=1}^M \sum_{n=1}^N c_{m,n} [P^{0 \rightarrow I}]^T ([P^0]^T)^n [P^{I \rightarrow 0}]^T ([P^I]^T)^m, \quad (7.15)$$

if it exists and represents a limiting distribution, gives the expected proportions of trials for which a cell will belong to each inhibition-on bin at the moment of arrival of the first

excitatory input after inhibition turns on. The dominant eigenvectors of the matrices

$$\sum_{m=1}^M \sum_{n=1}^N c_{m,n} ([P^0]^T)^n [P^{I \rightarrow 0}]^T ([P^I]^T)^m [P^{0 \rightarrow I}]^T$$

$$\sum_{m=1}^M \sum_{n=1}^N c_{m,n} ([P^I]^T)^m [P^{0 \rightarrow I}]^T ([P^0]^T)^n [P^{I \rightarrow 0}]^T$$

have similar interpretations.

Remark 7.1 In fact, justifying the existence of the limiting distribution for the matrix in (7.15) requires an additional technical adjustment, relative to that for matrix (7.14). The added complication arises because the fixed point of $\mathcal{N}^{(0,I)}$ is higher than that of $\mathcal{N}^{(0,0)}$, which may lead to a violation of irreducibility. This is simply a technical point and can be handled by replacing $([P^I]^T)^m$ by the product of a sequence of non-identical matrices that incorporate successively larger numbers of bins, assuming m is not too small.

TC cell example revisited: Uniform distributions

Consider again the TC model (6.13). We first compute the transition matrices analytically under the assumptions that IEs are selected from a uniform distribution on [20,60] msec, excitatory duration is fixed at 10 msec, and both durations of inhibitory inputs and time intervals from inhibitory offset to onset are selected from a uniform distribution on [125,175] msec. For simplicity, we assume that after rebound, a cell will only reach I_1 before the next excitation arrives; since the possibility of rebounding and reaching I_2 before this excitation, and of rebound in general, are relatively rare for our parameter values, this should not be a major source of error. To maintain tractability, we also neglect the fact that $w_{RK}^{(E,I)} > w_{RK}^{(E,0)}$, based on the fact that the right knees are fairly close for $g_{inh} = 0.12$; this will introduce some error into our results.

We computed the matrices from (7.14) analytically, using (3.8) as in the constant inhibition case discussed in Appendices A and B. In this example, the coefficients $c_{m,n} = p(m)p(n)$ in (7.14) can also be computed analytically, and we discuss some details of this calculation in Appendix C. After the $p(m)$ are computed, the desired eigenvector of matrix (7.14) can be obtained. Since rebound can throw off the count of inputs having occurred since the last firing event, we focus on the dominant eigenvector $v^{0 \rightarrow I}$ corresponding to inhibitory onset, which we compare to its counterpart, $v_{num}^{0 \rightarrow I}$, generated numerically from 301 inhibitory onsets during the last 90 sec of a 100 sec simulation:

$$v^{0 \rightarrow I} = [.3530 \ .0801 \ .1598 \ .2558 \ .0333 \ .0677 \ .0363 \ 0 \ 0 \ .0140 \ 0 \ 0]^T$$

$$v_{num}^{0 \rightarrow I} = [.3089 \ .1030 \ .0897 \ .2724 \ .0199 \ .0432 \ .0698 \ 0 \ 0 \ .0664 \ .0066 \ 0]^T.$$

The bins here are exactly those listed in Section 6 for the uniform IEI distribution with $s_{inh} = 1$, namely (1,1), (2,1), (2,2), (3,2), (3,3), (4,2), (4,3), (4,4), (5,2), (5,3), (5,4), and (5,5). Note that the zero entries here (signifying values less than 0.5×10^{-5}) correspond to the fact that these distributions are only relevant just after each onset of inhibition, and in the absence of inhibition, cells rarely reach the levels of the upper bins attained with inhibition,

which are included in this list. These bins become reoccupied as subsequent inputs arrive during the inhibition-on phase.

The agreement between $v^{0 \rightarrow I}$ and $v_{num}^{0 \rightarrow I}$ is fairly good, although not as good as those obtained with constant inhibition in Section 6, presumably due to rebound and w_{RK} issues. Note that a cell will be expected to fire well under 10% of the time to the first input that arrives just after the onset of inhibition. This low number fits the prediction of phase plane analysis, which implies that the onset of inhibition interferes with responsiveness to the subsequent excitatory input [Rubin and Terman, 2004]. Moreover, with the cell lying in the bottom bin over 30% of the time, which is a significant increase over the percentage there in the $s_{inh} = 1$ case, the compromise of responsiveness by rhythmic inhibition will endure beyond the first excitation after inhibitory onset.

TC cell example revisited: Normal distributions

We also calculated the transition matrices and dominant eigenvector, for comparison to direct numerical simulation, when IEIs were selected from the truncated normal distribution described earlier and inhibition on and off durations were selected from a similar distribution. Specifically, we assumed that these durations were of the form $t_{ref}^I + Y$, $t_{ref}^I + Z$, where Y was obtained by choosing \tilde{Y} from a normal distribution with mean 25 msec and standard deviation 10 msec and setting

$$Y = \begin{cases} 0, & \tilde{Y} < 0, \\ \tilde{Y}/M_Y, & 0 \leq \tilde{Y} \leq 50, \\ 50, & 50 < \tilde{Y}, \end{cases}$$

with normalization constant M_Y . The random variable Z was defined similarly but independently. To compute the relevant dominant eigenvector, as done above in the uniform distribution case, we need coefficients $c_{m,n}$, which we computed as $c_{m,n} = p(m)p(n)$ from the numerically obtained probabilities $p(1) = .0014, p(2) = .1625, p(3) = .6763, p(4) = .1625, p(5) = .0014$ with $p(j) = 0$ for $j \geq 6$. In this case, we actually used transition probabilities obtained from long-time simulations to compute the transition matrices $P_n^{I \rightarrow 0}, P_n^{0 \rightarrow I}$, which we do not display here, although these could have been computed from (3.8) as well. The dominant eigenvector $v_n^{0 \rightarrow I}$, corresponding to bin occupancy at the arrival time of the first excitatory input after inhibitory onset, and an example of the bin occupancy proportions $(v_n^{0 \rightarrow I})_{num}$ obtained from the last 90 sec of a 100 sec simulation are

$$v_n^{0 \rightarrow I} = [.3354 .0463 .1233 .3685 .0083 .0394 .0350 .0439 0]^T,$$

$$(v_n^{0 \rightarrow I})_{num} = [.3984 .0549 .1071 .3132 .0082 .0302 .0384 .0467 .0027]^T$$

for bins (1, 1), (2, 1), (2, 2), (3, 2), (3, 3), (4, 2), (4, 3), (5, 3), (5, 4). These results show that a cell can be expected to fire reliably to the first excitatory input after onset of inhibition less than 5% of the time and will lie in the bottom bin over 30% of the time, as also seen in the uniform distribution example.

8 Explicit connection to parkinsonian reliability failure

In PD, the inhibitory output from the basal ganglia may become rhythmic. DBS eliminates this rhythmicity, leading to inhibition from the basal ganglia that is elevated but, when summed over all inputs to a cell, is roughly constant, possibly with fast oscillations around a high constant level. A possible mechanism for the induction of motor symptoms in PD and for their amelioration by DBS, analyzed in [Rubin and Terman, 2004], is that rhythmic basal ganglia inhibitory outputs periodically compromise TC response reliability, while the regularization of these outputs by DBS restores responsiveness. The drastic drop in the number of TC cells expected to fire and the accumulation of cells in bins far from firing threshold found just after inhibitory onset in the case of rhythmic inhibition in our examples, relative even to the case of elevated but constant inhibition, offers a strong demonstration of the feasibility of this idea. The approach that we have taken to obtain these results is based on limiting distributions, and hence eliminates the possibility of spurious outcomes due to transient effects in simulations.

If we adjust parameters to an extreme case, so that TC cells fire in response to every excitatory input when $s_{inh} = 0$ and when $s_{inh} = 1$, then we still find some failures when inhibition is made to turn on and off rhythmically. In this case, we can read off from $v^{0 \rightarrow I}$ the probabilities with which a cell will experience each possible number of failures due to inhibitory onset, as seen in the examples discussed above. Even if no failures occur when inhibition is held on, there may be multiple failures following inhibitory onset if the difference $w_{RK}^{(E,I)} - w_{RK}^{(E,0)}$ is large.

Inhibitory offset may lead to response failure as well, though PIR. Cells that rebound when inhibition turns off are reset to $w_{RK}^{(E,0)}$. While the assumption of reliability in the inhibition-off case implies that

$$w_{RK}^{(E,0)} \cdot S^* > w_{LK}^{(E,0)}, \quad (8.16)$$

the time from the offset of inhibition until the arrival of the next excitatory input may be less than S^* , since inhibition may turn off at any moment in the IEI. However, under assumption (8.16), a cell will clearly experience at most one failure after rebound, since it lies above $w_{RK}^{(E,0)}$ after its first failure. Thus, while PIR may lead to bursts that do not reflect inputs to TC cells as well as a small number of response failures, the onset of inhibition represents the more significant obstacle to TC reliability under rhythmic inhibition, assuming perfect reliability in the constant inhibition case.

9 Discussion

We have considered a fast-slow excitable system subject to a stochastic excitatory input train, and we have shown how to derive an irreducible Markov chain that can be used to analytically compute the system's firing probability to each input, expected number of response failures between firings, and distribution of slow variable values between firings, in the infinite-time limit. We have illustrated this analysis on a model TC cell subject to a

uniform or to a truncated normal distribution of excitatory synaptic inputs, in the cases of constant inhibition and of inhibition that switches abruptly between two levels. The analysis immediately generalizes to any pair of input trains, excitatory or inhibitory and synaptic or not, with distinct switching frequencies, as well as to other models, such as integrate-and-fire, that feature a single variable that builds up to a threshold where an instantaneous spike and reset, possibly followed by a refractory period, occur.

In the TC cell case, our results generalize earlier findings suggesting how the modulation of inhibition by outputs of the basal ganglia can compromise TC responsiveness to excitatory inputs, with possible relevance to PD and DBS [Rubin and Terman, 2004]. In fact, basal ganglia output areas in rat show abrupt firing rate fluctuations on the time scale of seconds-to-minutes even in non-parkinsonian states [Ruskin et al., 1999], and the ideas that we have introduced could be used to consider how different fluctuation characteristics affect TC reliability. TC cells are also inhibited by thalamic reticular (RE) cells, which in turn are targets of excitatory corticothalamic inputs. Cortical oscillations, particularly abrupt transitions between cortical up and down states [Steriade et al., 1993a, Cowan and Wilson, 1994], could naturally lead to jumps in the levels of inhibition from RE cells to TC cells, providing an alternative source for the type of inhibition that we consider. In this context, our results provide a way to quantify the expected extent of the transient loss of thalamic relay reliability during the transitions between up and down states of different depths, as well as the likelihood that transitions from down to up will be signaled by PIR thalamic bursts [Steriade et al., 1993b]. Huertas et al. have also considered the relay properties of TC cells, specifically those in the dorsal lateral geniculate nucleus, in the presence of RE inhibition [Huertas et al., 2005], using simulations of an integrate-and-fire-or-burst model with an oscillatory driving input based on retinal ganglion cell activity. In their simulations, as found here and in previous work such as [Rubin and Terman, 2004], rhythmic inhibition to TC cells led to TC bursts and a failure of TC cells to respond to excitatory signals, although TC-RE interactions in their model gave rhythmic TC bursting phase-locked to the driving stimulus, which would not be present for the types of synaptic drive we consider.

We have proved that under general conditions, the Markov chain that we have derived is aperiodic and hence has a limiting distribution, which contrasts with Monte Carlo simulations, in which convergence properties cannot be forecast. Moreover, as we have also demonstrated, this distribution can be computed analytically, using equation (3.8), which eliminates any issues concerning transient effects seen in simulations. Our work is related to two earlier studies in which a Markov operator was derived to track transitions between oscillator phases, relative to sinusoidal forcing, at successive jump-up [Doi et al., 1998] or threshold-crossing [Tateno and Jimbo, 2000] times, in the presence of noise. Neither of these works, however, used a Markov chain to track transitions linked to successive input arrivals but not to firing events. Moreover, neither arrived at analytically computable transition probabilities or proved the existence of a limiting distribution, as we have done. A nice feature of the previous papers was the numerical tracking of sub-dominant eigenvalues to identify bifurcations, defined in a stochastic sense, relating to changes in mode-locking. The approach that we have presented could also be used to study bifurcations. For example, as

mentioned in Remark 4.3, a change in the range of IEs can cause the transition probability between a pair of bins to switch from 0 to nonzero, which may abruptly change the existence, or the nature, of the corresponding limiting distribution.

As a related, alternative approach, one could try to analyze the long term behavior of a cell by studying the random dynamical system $w_{n+1} = M_T(w_n)$, where the interinput interval T is a random variable with a known density [Lasota and Mackey, 1994]. If it could be found, then the limiting density for the variable w could be used to compute the statistics of interest for the cell. However, finding this limiting density is, in general, quite difficult. The Markov chain approach that we have presented can in fact be viewed as a discretization of this procedure. As a result, we recover a “coarse grained” version of the full limiting density for w , which is sufficient to determine the statistics of interest.

In a series of papers [Othmer and Watanabe, 1994, Xie et al., 1996, Othmer and Xie, 1999], Othmer and collaborators studied the application of step function forcing to a piecewise linear excitable system similar to those which we consider. In their work, which focused on the existence of mode-locked (harmonic or subharmonic) solutions and on chaos, they used a map and defined bins as we do, but their bins were based on positions of the knees of nullclines and their projections to other nullclines rather than the flow along branches of nullclines, and they did not consider a Markov chain for transitions between bins. Moreover, their analysis was restricted to periodic forcing, whereas ours accommodates, and indeed is particularly well suited for, stochasticity in input timing. LoFaro and Kopell [LoFaro and Kopell, 1999] also used 1-d maps to study a forced excitable system, but in their work the excitable system was a neuron mutually coupled via inhibitory synapses to an oscillatory cell and the map was a singular Poincaré map with each iteration corresponding to the time between jumps to the active phase. Similarly, Keener et al. [Keener et al., 1981] and subsequent authors have used firing time maps to study mode-locking in integrate-and-fire and related models with periodic stimuli. Alternatively, other works have used 1-d maps based on interinput intervals to study mode-locked, quasiperiodic, and chaotic responses of excitable systems to periodic, purely excitatory inputs [Coombes and Osbaldestin, 2000, Ichinose et al., 1998].

Clearly, the transformation of some combination of excitatory and inhibitory synaptic inputs into postsynaptic neuronal responses is a fundamental operation present within any nontrivial nervous system. As a result, various manifestations of this transformation have been studied, computationally and experimentally, by many researchers. In our work, as in [Smith and Sherman, 2002], we consider the excitatory input stream as a drive to the postsynaptic cell and the inhibitory input as a modulation that alters the cell’s responsiveness to its drive. We neglect such intriguing effects as long-term synaptic scaling [Desai et al., 2002], changes in intrinsic excitability [Aizenman et al., 2003], and changes in the balance of excitation and inhibition [Somers et al., 1998], which could become relevant in the asymptotic limit, as well as the effect of attention [Tiesinga, 2005]. Moreover, we assume that successful thalamic relay consists of single spike responses to an input train, neglecting the idea that, by virtue of their stereotyped form and reliability, thalamic bursts may serve a relay function [Person and Perkel, 2005, Babadi, 2005]. Other authors have considered how variations in intrinsic properties of postsynaptic cells determine the input characteristics that

induce these cells to spike most reliably [Fellous et al., 2001, Schreiber et al., 2004] and how neuronal processing varies under more general changes in input spike patterns than the ones that we have considered here (e.g., [Hunter and Milton, 2003, Tiesinga and Toups, 2005]). While these issues are beyond the scope of this work, our approach can accommodate input trains that vary stochastically in a variety of ways (see Remark 3.2), and hence it may be well suited for the future exploration of such issues.

10 Appendices

Appendix A: TC cell with $s_{inh} = 0$

With $s_{inh} = 0$, a cell that has just spiked is guaranteed to spike again after at most $N = 3$ subsequent inputs. Although the intervals I_k are defined in terms of intervals of the slow variable w , clearly there is also a particular time interval associated with each I_k . In the particular model (6.13), with $s_{inh} = 0$, these time intervals are $[20, 50)$, $[50, 75.5)$, $[75.5, \infty)$, where $w_{RK}^E \cdot 75.5 \approx w_{LK}^E$. In practice, we used a simulation protocol, rather than determination of the left knee of \mathcal{N}^E directly from (6.13), to compute the value 75.5. That is, we found 75.5 to be the minimum value of t such that, given an initial condition with $w = w_{RK}^E \cdot t$ and with v at the corresponding position on the left branch of \mathcal{N}^0 , the model cell would fire in response to an excitatory input of duration 10 msec. This adjustment represents the modification to w_{LK}^E discussed in Remark 3.1 and is therefore consistent with the rest of the analysis that we present.

Using these time intervals allows us to compute the transition probabilities between bins (I_k, j) for $j, k = 1, \dots, 3$, with $j \leq k$. Let $p_{(k_1, j_1) \rightarrow (k_2, j_2)} = P[(I_{k_2}, j_2) | (I_{k_1}, j_1)]$. To compute these probabilities, we start with the fate of a cell just after firing, computing $p_{(3, j) \rightarrow (k, 1)}$ for each relevant (j, k) . First, note that since all cells fire from a bin of the form (I_3, j) for some $j \leq 3$, and all firing cells are reset together to w_{RK}^E regardless of where they fired from, $p_{(3, j) \rightarrow (k, 1)}$ is independent of j . In this example, with $g_{inh} = 0.12$, we have $p_{(3, j) \rightarrow (1, 1)} = P[T_1 \in [20, 50)]$, where T_1 denotes the time from reset to the onset of the next excitatory input. Further, $P[T_1 \in [20, 50)] = 3/4$, since T_1 is selected from a uniform distribution on $[20, 60]$. Similarly, $p_{(3, j) \rightarrow (2, 1)} = 1/4$ and $p_{(3, j) \rightarrow (k, 1)} = 0$ for all $k > 2$.

Next, we seek values for $p_{(1, 1) \rightarrow (k, 2)}$ for each $k > 1$ and $p_{(2, 1) \rightarrow (k, 2)}$ for each $k > 2$. We have

$$p_{(1, 1) \rightarrow (2, 2)} = P[T_1 + T_2 + 10 \in [50, 75.5) | T_1 \in [20, 50)],$$

since the duration of excitation is 10 msec. This can be computed as the ratio of two areas, either geometrically or by integration, leading to the result that $p_{(1, 1) \rightarrow (2, 2)} = 2601/9600$. Since $N = 3$ for $s_{inh} = 0$, it follows that $p_{(1, 1) \rightarrow (3, 2)} = 6999/9600$, and that $p_{(2, 1) \rightarrow (3, 2)} = p_{(2, 2) \rightarrow (3, 3)} = 1$. This completes the calculation for $s_{inh} = 0$. The corresponding transition

matrix for $s_{inh} = 0$ is thus

$$P^0 = \begin{bmatrix} 0 & 0 & 2601/9600 & 6999/9600 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 3/4 & 1/4 & 0 & 0 & 0 \\ 3/4 & 1/4 & 0 & 0 & 0 \end{bmatrix}$$

where the states of the Markov chain are the bins $(1, 1), (2, 1), (2, 2), (3, 2), (3, 3)$.

Appendix B: TC cell with $s_{inh} = 1$

The case $s_{inh} = 1$ requires more analytical computations than the case $s_{inh} = 0$, since $N = 5$ for $s_{inh} = 1$. In this case, $p_{(5,j) \rightarrow (k,1)}$ are identical to $p_{(3,j) \rightarrow (k,1)}$, computed for $s_{inh} = 0$ above, and are nonzero only for $k = 1, 2$. The bins for $s_{inh} = 1$ correspond to times $[20,50), [50,80), [80,110), [110,128), [128, \infty)$. Similarly to the previous case,

$$p_{(1,1) \rightarrow (2,2)} = P[T_1 + T_2 + 10 \in [50, 80) | T_1 \in [20, 50)] = 3/8,$$

with $p_{(1,1) \rightarrow (3,2)} = 1/2$ and $p_{(1,1) \rightarrow (4,2)} = 1/8$ by analogous calculations. Along the same lines,

$$p_{(2,1) \rightarrow (3,2)} = P[T_1 + T_2 + 10 \in [80, 110) | T_1 \in [50, 60)] = 5/8,$$

while $p_{(2,1) \rightarrow (4,2)} = 37/100$ and $p_{(2,1) \rightarrow (5,2)} = 1/200$. As we proceed further, certain probability calculations become more involved, because bin membership may be achieved by more than one path. For example, we see already that a cell may reach bin $(I_3, 2)$ from $(I_1, 1)$ or from $(I_2, 1)$. Thus,

$$\begin{aligned} p_{(3,2) \rightarrow (4,3)} &= P[T_1 + T_2 + T_3 + 20 \in [110, 128) | T_1 \in [20, 50) \text{ and } T_1 + T_2 + 10 \in [80, 110)] \\ &+ P[T_1 + T_2 + T_3 + 20 \in [110, 128) | T_1 \in [50, 60) \text{ and } T_1 + T_2 + 10 \in [80, 110)] \\ &= 2123/14250. \end{aligned}$$

The full set of calculations reveal that the transition matrix P^I for $s_{inh} = 1$ has the form

$$\begin{bmatrix} 0 & 0 & \frac{3}{8} & \frac{1}{2} & 0 & \frac{1}{8} & 0 & \cdot & \cdot & \cdot & 0 & 0 \\ 0 & 0 & 0 & \frac{5}{8} & 0 & \frac{37}{100} & 0 & 0 & \frac{1}{200} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{4} & 0 & \frac{6011}{13500} & 0 & 0 & \frac{2057}{6750} & 0 & 0 \\ 0 & \cdot & \cdot & \cdot & \cdot & 0 & \frac{2123}{14250} & 0 & 0 & \frac{12127}{14250} & 0 & 0 \\ 0 & \cdot & \cdot & \cdot & \cdot & \cdot & 0 & \frac{243}{10000} & 0 & 0 & \frac{9757}{10000} & 0 \\ 0 & \cdot & 0 & 1 & 0 & 0 \\ 0 & \cdot & 0 & 1 & 0 \\ 0 & \cdot & 0 & 1 \\ \frac{3}{4} & \frac{1}{4} & 0 & \cdot & 0 \\ \frac{3}{4} & \frac{1}{4} & 0 & \cdot & 0 \\ \frac{3}{4} & \frac{1}{4} & 0 & \cdot & 0 \\ \frac{3}{4} & \frac{1}{4} & 0 & \cdot & 0 \end{bmatrix}$$

where the states of the Markov chain are the bins (1,1), (2,1), (2,2), (3,2), (3,3), (4,2), (4,3), (4,4), (5,2), (5,3), (5,4) and (5,5).

Appendix C: Coefficients for matrix (7.14) in the case of time-dependent inhibition and excitation

To calculate the coefficients $c_{m,n} = p(m)p(n)$ analytically, we compute $p(1)$, which is the probability that exactly one excitatory input arrives during an epoch of constant inhibition, and then we compute $p(m)$ for each $m = 2, \dots, 6$ as the probability of at most m inputs arriving minus the probability of at most $m - 1$ inputs arriving. We stop at $m = 6$ since at most 6 excitatory inputs can arrive during 175 msec, given the IEIs and excitation duration that we consider. Let $t = 0$ denote the time of inhibitory offset, which will last for time $T_I \in [125, 175]$. Define T_{-1} as the time from the last excitatory onset before $t = 0$ to the first excitatory onset after $t = 0$. Let T_1 denote the time of this first excitatory onset. Given that each excitation endures for 10 msec, $T_{-1} \in [30, 70]$, while $T_1 \in [0, T_{-1}]$; see Fig. 5. Finally, let T_2 denote the IEI following the end of the excitatory input that occurs at time T_1 . Since $T_1 < T_I$ must hold, the value of $p(1)$ is the probability that $T_1 + T_2 + 10 < T_I$. Thus,

$$p(1) = \frac{\int_{125}^{175} \int_{30}^{70} \int_0^{T_{-1}} \int_{\min(T_I - T_1 - 10, 60)}^{60} dT_2 dT_1 dT_{-1} dT_I}{\int_{125}^{175} \int_{30}^{70} \int_0^{T_{-1}} \int_{20}^{60} dT_2 dT_1 dT_{-1} dT_I} = \frac{27}{51200}.$$

For $m > 1$, each quantity $p(m)$ is given as the ratio of two multiple integrals as well, with one additional nested integral in each, relative to those used to calculate $p(m - 1)$. Since these integrals can be evaluated exactly, we obtain exact fractional representations for each, but to save space, we simply give decimal approximations here: $p(2) \approx .1985$, $p(3) \approx .6075$, $p(4) \approx .1875$, $p(5) \approx .0060$, $p(6) \approx 4.731 \times 10^{-6}$.

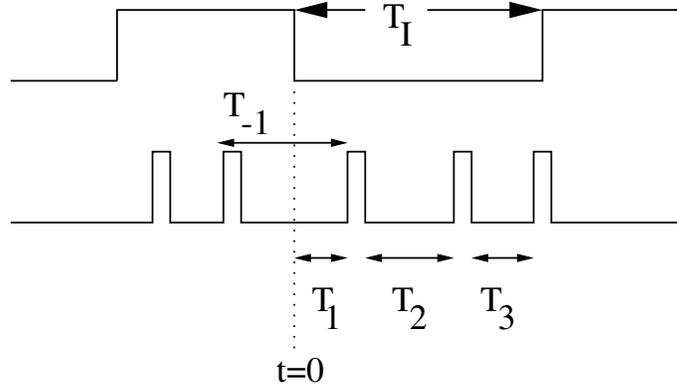


Figure 5: The notation for the calculation of the probabilities $p(m)$. Note that the actual number of excitatory inputs arriving during the interval of constant inhibition will be between one and six.

Acknowledgment. The authors received partial support for this work from National Science Foundation grants DMS-0414023 (JR), DMS-0244529 (KJ), and ATM-0417867 (KJ).

References

- [Aizenman et al., 2003] Aizenman, C. D., Akerman, C. J., Jensen, K. R., and Cline, H. T. (2003). Visually driven regulation of intrinsic neuronal excitability improves stimulus detection in vivo. *Neuron*, **39**(5):831–842.
- [Anderson et al., 2003] Anderson, M., Postpuna, N., and Ruffo, M. (2003). Effects of high-frequency stimulation in the internal globus pallidus on the activity of thalamic neurons in the awake monkey. *J. Neurophysiol.*, **89**:1150–1160.
- [Babadi, 2005] Babadi, B. (2005). Bursting as an effective relay mode in a minimal thalamic model. *J. Comput. Neurosci.*, **18**:229–243.
- [Brown et al., 2001] Brown, P., Oliviero, A., Mazzone, P., Insola, A., Tonali, P., and Lazzaro, V. D. (2001). Dopamine dependency of oscillations between subthalamic nucleus and pallidum in Parkinson’s disease. *J. Neurosci.*, **21**:1033–1038.
- [Chance et al., 2002] Chance, F., Abbott, L., and Reyes, A. (2002). Gain modulation from background synaptic input. *Neuron*, **35**:773–782.
- [Coombes and Osbaldestin, 2000] Coombes, S. and Osbaldestin, A. (2000). Period-adding bifurcations and chaos in a periodically stimulated excitable neural relaxation oscillator. *Phys. Rev. E*, **62**:4057–66.

- [Cowan and Wilson, 1994] Cowan, R. L. and Wilson, C. J. (1994). Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *J Neurophysiol*, 71(1):17–32.
- [De Schutter, 1999] De Schutter, E. (1999). Using realistic models to study synaptic integration in cerebellar Purkinje cells. *Rev. Neurosci.*, **10**:233–245.
- [Desai et al., 2002] Desai, N. S., Cudmore, R. H., Nelson, S. B., and Turrigiano, G. G. (2002). Critical periods for experience-dependent synaptic scaling in visual cortex. *Nat Neurosci*, **5**(8):783–789.
- [Doi et al., 1998] Doi, S., Inoue, J., and Kumagai, S. (1998). Spectral analysis of stochastic phase lockings and stochastic bifurcations in the sinusoidally forced van der Pol oscillator with additive noise. *J Stat Phys*, 90(5/6):1107–1127.
- [Fellous et al., 2001] Fellous, J. M., Houweling, A. R., Modi, R. H., Rao, R. P., Tiesinga, P. H., and Sejnowski, T. J. (2001). Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *J Neurophysiol*, 85(4):1782–1787.
- [Hashimoto et al., 2003] Hashimoto, T., Elder, C., Okun, M., Patrick, S., and Vitek, J. (2003). Stimulation of the subthalamic nucleus changes the firing pattern of pallidal neurons. *J. Neurosci.*, **23**:1916–1923.
- [Hoel et al., 1972] Hoel, P. G., Port, S. C., and Stone, C. J. (1972). *Introduction to stochastic processes*. Houghton Mifflin Co., Boston, Mass. The Houghton Mifflin Series in Statistics.
- [Huertas et al., 2005] Huertas, M., Groff, J., and Smith, G. (2005). Feedback inhibition and throughput properties of an integrate-and-fire-or-burst network model of retinogeniculate transmission. *J. Comp. Neurosci.*, **19**:147–180.
- [Hunter and Milton, 2003] Hunter, J. D. and Milton, J. G. (2003). Amplitude and frequency dependence of spike timing: implications for dynamic regulation. *J Neurophysiol*, 90(1):387–394.
- [Ichinose et al., 1998] Ichinose, N., Aihara, K., and Judd, K. (1998). Extending the concept of isochrons from oscillatory to excitable systems for modeling excitable neurons. *International Journal of Bifurcations and Chaos*, 8(12):2375–2385.
- [Keener et al., 1981] Keener, J. P., Hoppensteadt, F. C., and Rinzel, J. (1981). Integrate-and-fire models of nerve membrane response to oscillatory input. *SIAM J. Appl. Math.*, **41**(3):503–517.
- [Lasota and Mackey, 1994] Lasota, A. and Mackey, M. C. (1994). *Chaos, fractals, and noise: Stochastic aspects of dynamics*, volume 97 of *Applied Mathematical Sciences*. Springer-Verlag, New York, second edition.

- [LoFaro and Kopell, 1999] LoFaro, T. and Kopell, N. (1999). Timing regulation in a network reduced from voltage-gated equations to a one-dimensional map. *J. Math. Biol.*, **38**(6):479–533.
- [Magnin et al., 2000] Magnin, M., Morel, A., and Jeanmonod, D. (2000). Single-unit analysis of the pallidum, thalamus, and subthalamic nucleus in parkinsonian patients. *Neuroscience*, **96**:549–564.
- [Nini et al., 1995] Nini, A., Feingold, A., Sloviter, H., and Bergman, H. (1995). Neurons in the globus pallidus do not show correlated activity in the normal monkey, but phase-locked oscillations appear in the MPTP model of parkinsonism. *J. Neurophysiol.*, **74**:1800–1805.
- [Othmer and Watanabe, 1994] Othmer, H. and Watanabe, M. (1994). Resonance in excitable systems under step-function forcing. I. Harmonic solutions. *Adv. Math. Sci. Appl.*, **4**:399–441.
- [Othmer and Xie, 1999] Othmer, H. and Xie, M. (1999). Subharmonic resonance and chaos in forced excitable systems. *J. Math. Biol.*, **39**:139–171.
- [Person and Perkel, 2005] Person, A. L. and Perkel, D. J. (2005). Unitary IPSPs drive precise thalamic spiking in a circuit required for learning. *Neuron*, **46**(1):129–140.
- [Raz et al., 2000] Raz, A., Vaadia, E., and Bergman, H. (2000). Firing patterns and correlations of spontaneous discharge of pallidal neurons in the normal and tremulous 1-methyl-4-phenyl-1,2,3,6 tetrahydropyridine vervet model of parkinsonism. *J. Neurosci.*, **20**:8559–8571.
- [Rubin and Terman, 2002] Rubin, J. E. and Terman, D. (2002). Geometric singular perturbation analysis of neuronal dynamics. In *Handbook of dynamical systems, Vol. 2*, pages 93–146. North-Holland, Amsterdam.
- [Rubin and Terman, 2004] Rubin, J. E. and Terman, D. (2004). High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model. *J. Comput. Neurosci.*, **16**:211–235.
- [Ruskin et al., 1999] Ruskin, D. N., Bergstrom, D. A., Kaneoke, Y., Patel, B. N., Twery, M. J., and Walters, J. R. (1999). Multisecond oscillations in firing rate in the basal ganglia: robust modulation by dopamine receptor activation and anesthesia. *J. Neurophysiol.*, **81**(5):2046–2055.
- [Schreiber et al., 2004] Schreiber, S., Fellous, J.-M., Tiesinga, P., and Sejnowski, T. J. (2004). Influence of ionic conductances on spike timing reliability of cortical neurons for suprathreshold rhythmic inputs. *J. Neurophysiol.*, **91**(1):194–205.
- [Smith and Sherman, 2002] Smith, G. and Sherman, S. (2002). Detectability of excitatory versus inhibitory drive in an integrate-and-fire-or-burst thalamocortical relay neuron model. *J. Neurosci.*, **22**:10242–10250.

- [Somers and Kopell, 1993] Somers, D. and Kopell, N. (1993). Rapid synchronization through fast threshold modulation. *Biol. Cybern.*, 68:393–407.
- [Somers et al., 1998] Somers, D., Todorov, E., Siapas, A., Toth, L., Kim, D., and Sur, M. (1998). A local circuit approach to understanding integration of long-range inputs in primary visual cortex. *Cerebral Cortex*, 8:204–217.
- [Steriade et al., 1993a] Steriade, M., Nunez, A., and Amzica, F. (1993a). A novel slow (≈ 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J Neurosci*, 13(8):3252–3265.
- [Steriade et al., 1993b] Steriade, M., Nunez, A., and Amzica, F. (1993b). Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *J Neurosci*, 13(8):3266–3283.
- [Stone, 2004] Stone, M. (2004). Varied stimulation of a relaxation oscillator. Master’s thesis, University of Houston.
- [Tateno and Jimbo, 2000] Tateno, T. and Jimbo, Y. (2000). Stochastic mode-locking for a noisy integrate-and-fire oscillator. *Phys. Lett. A*, 271:227–236.
- [Taylor and Karlin, 1998] Taylor, H. M. and Karlin, S. (1998). *An introduction to stochastic modeling*. Academic Press Inc., San Diego, CA, third edition.
- [Tiesinga, 2005] Tiesinga, P. (2005). Stimulus competition by inhibitory interference. *Neural Comput.*, 17:2421–2453.
- [Tiesinga et al., 2000] Tiesinga, P., Jose, J., and Sejnowski, T. (2000). Comparison of current-driven and conductance-driven neocortical model neurons with Hodgkin-Huxley voltage-gated channels. *Phys. Rev. E*, 62:8413–8419.
- [Tiesinga and Sejnowski, 2004] Tiesinga, P. and Sejnowski, T. (2004). Rapid temporal modulation of synchrony by competition in cortical interneuron networks. *Neural Comput.*, 16:251–275.
- [Tiesinga and Toups, 2005] Tiesinga, P. H. E. and Toups, J. V. (2005). The possible role of spike patterns in cortical information processing. *J Comput Neurosci*, 18(3):275–286.
- [van Vreeswijk and Sompolinsky, 1998] van Vreeswijk, C. and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.*, 10:1321–1371.
- [Xie et al., 1996] Xie, M., Othmer, H., and Watanabe, M. (1996). Resonance in excitable systems under step-function forcing. II. Subharmonic solutions and persistence. *Physica D*, 98:75–110.