

Event-based minimum-time control of oscillatory neuron models

Phase randomization, maximal spike rate increase, and desynchronization

Per Danzl · João Hespanha · Jeff Moehlis

Received: 13 August 2009 / Accepted: 22 October 2009 / Published online: 13 November 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

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

Keywords Optimal control · Oscillatory neurons · Phase randomization · Isochrons · Hodgkin–Huxley models

1 Introduction

The motor tremors symptomatic of Parkinson's disease may be linked to the synchronous periodic firing of neurons in the brain (Pare et al. 1990; Nini et al. 1995). This pathological behavior can be mitigated by the use of implanted stimulus electrodes in a treatment called Electrical Deep Brain Stimulation (EDBS) (Benabid et al. 1991). Inspired by the classical ideas of Winfree (2001), we develop a feedback-based desynchronizing control scheme for EDBS which takes advantage of recent developments in nonlinear optimal control theory to account for important biological constraints.

Traditionally, EDBS has utilized relatively high-frequency (100+ Hz) periodic waveforms injected in an open-loop fashion, which after empirically tuning stimulus parameters, can mitigate the synchronous neural activity. The open-loop nature of traditional EDBS is sub-optimal in that the system is constantly stimulating the neurons, regardless of whether they are firing in pathological synchrony or not.

There is much current interest in designing EDBS control systems that use additional feedback electrodes to make the stimulus “demand-controlled” (Tass 1999, 2000). A feedback-based approach is attractive from a clinical perspective in that the biological tissue is only stimulated when necessary, thereby reducing the overall accumulation of negative side effects of electrical stimulation.

Feedback control methods have been designed from a dynamical systems perspective for neurons modeled as two-dimensional oscillators near a Hopf bifurcation (Tass 1999, 2000; Popovych et al. 2006). One effective scheme uses two precisely timed impulses (Tass 2000). Another uses a nonlinear-delay feedback protocol with experimentally tuned stimulation parameters (Popovych et al. 2006). Representing a neuron as a two-dimensional oscillator near a Hopf bifurcation is attractive because it allows one to use the so-called

P. Danzl (✉) · J. Moehlis
Department of Mechanical Engineering, University of California,
Santa Barbara, CA 93106, USA
e-mail: pdanzl@engineering.ucsb.edu

J. Hespanha
Department of Electrical and Computer Engineering, University
of California, Santa Barbara, CA 93106, USA

normal form of the dynamic equation, which is generally much simpler than the underlying model of the particular neuron. However, the phase response curves (a characterization of the neuron's response to stimulus) are symmetric for such normal form models, which is contradictory to experimental observations (Netoff et al. 2005).

In this article, we explore Winfree's idea of using *phaseless sets* to randomize the phase of each neuron in a population, thereby desynchronizing their pathological firing behavior. The central concept is that the phase of an oscillatory dynamical system can be effectively reset using an impulsive stimulus to drive the state of the system to a point where the phase is undefined (Winfree 2001). In the context of oscillatory neurons, this phaseless set is a single point in state space, which is an unstable equilibrium point "inside" the periodic orbit. For more general cases, the phaseless set(s) could also be an unstable periodic orbit (for a planar system), or the stable manifold of an unstable periodic orbit (in higher dimensions). These objects could form the boundary of the basin of attraction of the stable periodic orbit.

There are three primary challenges to the use of electrical stimulus to drive a neuron's state to its phaseless set. First, the Hodgkin–Huxley conductance-based formalism (Hodgkin and Huxley 1952), a general modeling methodology in neuroscience, yields systems of continuous-time ordinary differential equations (ODEs) that tend to be *highly* nonlinear. Many traditional control designs require the controlled system to be either linear or only *weakly* nonlinear. A second challenge is that a control scheme cannot stimulate biological tissue with arbitrarily large signals. The magnitude of the electrical stimulus must be constrained, invalidating the concept of purely impulsive inputs for biological implementation. The third challenge is the fact that the only state that is directly observable is the membrane voltage. A control system cannot measure the dynamic state of the many ion channels that play a critical role in the oscillatory behavior of neural spiking. This poses particular problems for implementing traditional nonlinear feedback control systems, which depend on continuous measurements of state.

The proposed event-based feedback control scheme presented here addresses each of these challenges, and accomplishes the objective of driving a neuron to its phaseless set using a magnitude-constrained stimulus signal in the minimum possible time. Our method combines techniques from modern control engineering, namely Lebesgue sampling (i.e., event-based control) (Åström and Bernhardsson 2003), nonlinear optimal control theory (e.g., Athans and Falb 1966), and a recent computational toolbox for solving the resulting Hamilton–Jacobi–Bellman equations (Mitchell 2007). The method is quite general and can be applied to any conductance-based neuron model that possesses a stable limit cycle.

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

The control design method developed here has several other interesting applications. By changing a simple parameter in the design calculations, we can create a control signal that will optimally increase the spike rate of the neuron. Also, we can visualize the isochrons of the particular neuron model by plotting level sets of an intermediate value function calculated in the course of the stimulus design.

2 Neuron models

This article will present a mathematical framework for controlling neurons which, in absence of stimulation, spike periodically in time. We wish to consider an ODE model that is general, widely known, and has dynamical characteristics typical of spiking neurons. The Hodgkin–Huxley model, in its four-dimensional "space-clamped" form, is a prototypical conductance-based ODE model in the mathematical neuroscience literature (Hodgkin and Huxley 1952). It was originally developed to model the excitable membrane voltage dynamics of the squid giant axon. In a large biologically relevant region of parameter space, the model possesses a strongly attracting stable limit cycle which captures the essential features of a periodically spiking membrane voltage. We will use a two-dimensional reduction of the original four-dimensional model that will enable easier visualization of the mathematical objects used in our optimal stimulus computational method. We note that our proposed methods do not depend on the low dimensionality of the model, and a planar reduction has been chosen only to facilitate the presentation. To illustrate the generality of our method, we will also implement it on the two-dimensional FitzHugh–Nagumo model, another canonical example from the mathematical neuroscience literature (Keener and Sneyd 1998; FitzHugh 1961; Nagumo et al. 1962).

2.1 Two-dimensional Hodgkin–Huxley model

We consider the two-dimensional reduction of the Hodgkin–Huxley model analyzed by, for example, Moehlis (2006) and Keener and Sneyd (1998). The two states are the membrane voltage V , and a dimensionless gating variable n , which represents the fraction of ion channels that are open in the

membrane. The dynamics are represented by the following ODE system:

$$\dot{V} = \left(I(t) + D\eta(t) + I_b - \bar{g}_{Na}[m_\infty(V)]^3(0.8-n)(V - V_{Na}) - \bar{g}_K n^4(V - V_K) - \bar{g}_L(V - V_L) \right) / C \tag{1}$$

$$\dot{n} = \alpha_n(V)(1 - n) - \beta_n(V)n$$

where

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

In this model, \bar{g}_{Na} , \bar{g}_K , and \bar{g}_L represent the conductances of the sodium, potassium, and leakage channels, respectively. The reversal potentials of these channels are V_{Na} , V_K , and V_L . We represent background noise as $\eta(t)$, a zero-mean white signal, with strength D .

Throughout this article, we will use a baseline current $I_b = 10 \text{ mA}$, for which the system has a stable periodic orbit shown in Fig. 1. In the forthcoming theoretical development, we will consider the noiseless deterministic case when $D = 0$. We will return to the case of nonzero noise after we have computed the optimal control stimulus. The membrane capacitance is $C = 1.0$ in the units of (1), so for notational convenience in the following section, we will incorporate it into our input variable $u(t) = I(t)/C$.

Also, we will use the following shorthand notation for the right-hand side of (1):

$$\begin{aligned} \dot{V} &= f_V(V, n) + u \\ \dot{n} &= f_n(V, n) \end{aligned} \tag{2}$$

2.2 Isochrons and phaseless sets

Our aim is to design a stimulus that will drive the state of the model to regions of state space where the system is very sensitive to noise. By “sensitive to noise” we mean that the (asymptotic) phase can be easily altered by perturbations to the membrane voltage state caused by the noise process [$D\eta(t)$ in (1)]. The *phase* of an oscillator is a measure of the time it has spent evolving since it last passed through the marker event which defines zero phase. For neuron models, the spiking point, which is the blue square shown in Fig. 1, is taken as the marker event defining zero phase. Phase is an important concept for studying periodically spiking neuron models because it also indicates when the neuron will spike next in the absence of stimulus. We seek to randomize the




Fig. 1 Stable periodic orbit for the reduced Hodgkin–Huxley model with $I_b = 10 \text{ mA}$ (red curve). Isochrons, spaced uniformly in phase, are shown as dashed lines. The location of the unstable equilibrium point, where the isochrons converge, is shown by a black X. The spike point (maximum membrane voltage) is shown by the blue square (color online)

phase of a neuron model, but in so doing we will be stimulating the neuron and knocking it off of its periodic orbit. After stimulus, the state will be drawn back in toward the limit cycle, so we must consider an extension of phase that applies to all points in the basin of attraction of the periodic orbit. Following Josic et al. (2006), this extended concept of phase is called *asymptotic phase* and can be defined as follows:

Consider an arbitrary autonomous ODE

$$\dot{x} = f(x), \quad x \in \mathbb{R}^p, \quad p \geq 2 \tag{3}$$

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

$$\lim_{t \rightarrow \infty} |x(t) - \gamma(t + \phi(x))| = 0. \tag{4}$$

State space is foliated with $(p - 1)$ -dimensional hypersurfaces called *isochrons*, which are level sets of asymptotic phase (Guckenheimer 1975). For planar ($p = 2$) systems, the isochrons are simply lines of equal asymptotic phase. A set of isochrons for the reduced Hodgkin–Huxley model are shown as dashed lines in Fig. 1. Points in state space where phase is not defined are called *phaseless sets*.

Intuitively, if the objective is to find states where asymptotic phase is very sensitive to noise, we can simply look for areas where the isochrons become closely bunched together. In Fig. 1, one can see that the isochrons are most dense near the unstable equilibrium point at $(V_s, n_s) = (-59.6, 0.403)$. This unstable equilibrium point does not have a well-defined phase because it is not technically in the basin of attraction of the limit cycle. It is the only phaseless set contained by the basin of attraction of the periodic orbit and is an attractive

target for phase randomization. We will show how to construct a stimulus to reach an arbitrarily small neighborhood of this point, then the noise process will randomly perturb the state onto a nearby isochron, effectively randomizing the asymptotic phase of the neuron.

2.3 Normalized model

We now introduce a simple affine coordinate transformation to simplify the presentation, and to put our dynamical system into a form necessary for the proposed stimulus design method. A critical step in our method uses the Level Set Methods Toolbox (ToolboxLS), a MATLAB-based partial differential equation (PDE) solver (Mitchell 2008). This PDE solver has the requirement that each component should be of the same order of magnitude (Mitchell 2007) and the new state-space coordinates must be centered at the target point, \mathcal{T} . The coordinate transformation is written as:

$$\begin{aligned} x_1 &= \frac{V - V_{\mathcal{T}}}{K_s} \\ x_2 &= n - n_{\mathcal{T}}. \end{aligned} \quad (5)$$

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

Under this transformation, the dynamics become

$$\begin{aligned} \dot{x}_1 &= \frac{1}{K_s} f_V(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) + \frac{1}{K_s} u \\ \dot{x}_2 &= f_n(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) \end{aligned} \quad (6)$$

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

$$\dot{x} = F(x) + \begin{bmatrix} \frac{1}{K_s} \\ 0 \end{bmatrix} u = \begin{bmatrix} F_1(x) \\ F_2(x) \end{bmatrix} + \begin{bmatrix} \frac{1}{K_s} \\ 0 \end{bmatrix} u \quad (7)$$

where

$$F_1(x) = \frac{1}{K_s} f_V(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}) \quad (8)$$

and

$$F_2(x) = f_n(K_s x_1 + V_{\mathcal{T}}, x_2 + n_{\mathcal{T}}). \quad (9)$$

We note that this change of coordinates will only be used to simplify the mathematical presentation of the optimal stimulus design method. When results are presented in Sect. 5, we will revert back to the original variables V and n , and will report stimulus signals as electrical currents, $I(t)$.

3 Event-based control

From the perspective of classical feedback control, this control design problem is challenging in a number of important ways. Conductance-based ODE models of spiking neurons, such as the reduced Hodgkin–Huxley system (1), are complex and highly nonlinear. Furthermore, the magnitude of the stimulus is constrained by biological and hardware implementation limitations. In the case of a single neuron, membrane voltage is the only directly measurable state, and often the measurements are quite noisy. The only thing that can be detected with any certainty is the time at which a voltage spike occurs. Due to such limited observability, classical techniques such as feedback linearization (Khalil 2002) cannot be used.

The fact that the only reliable state measurements are the detection of voltage spikes leads us to employ a control methodology known as *event-based control*. The concept of event-based control, sometimes known as Lebesgue sampling, was developed as an improvement to fixed sample-rate feedback control for digital systems (Åström and Bernhardsson 2003). As shown in previous study (Danzl and Moehlis 2007, 2008), this methodology finds natural utility in systems involving spiking neurons. In this study, we design a controller that waits for a voltage spike (which we define to be the *event*) then stimulates the neuron with a pre-computed waveform designed to drive the neuron's state close to the unstable fixed point “inside” the periodic orbit, a point which is surrounded with closely packed isochrons. We will also show that other target sets can be used, which allows the controller to achieve other objectives, like maximally increasing the spike rate.

3.1 Control objectives

From a theoretical standpoint, the objective of our control system is to drive the system's state to a pre-defined target location \mathcal{T} . By setting \mathcal{T} to be the unstable equilibrium point, the controller will cause asymptotic phase randomization. On the other hand, if we set \mathcal{T} to be the spiking point, the controller will cause the neuron's spike rate to increase as much as possible, given a fixed stimulus magnitude constraint. To keep our control stimulus design methodology general, we will not specify \mathcal{T} . In fact, any point in state space can be chosen as a target set; although, it is not clear what the utility of using points besides the unstable equilibrium and the spike point might be. Recall that Eq. 5 translates the system by \mathcal{T} so that the coordinates are centered on the target set. Regardless of what is chosen for \mathcal{T} , in the transformed coordinate system the target is always the origin.

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

$$|u(t)| \leq \mu \tag{10}$$

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

$$\tilde{u} = \frac{u}{\mu}. \tag{11}$$

Equation 7 then becomes

$$\dot{x} = F(x) + b\tilde{u} \tag{12}$$

where

$$b = \begin{bmatrix} \frac{\mu}{K_s} \\ 0 \end{bmatrix}. \tag{13}$$

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

4 Optimal stimulus design

The optimal stimulus design method we present follows an approach similar to that outlined in Mitchell (2007) as applied to the classical double integrator example of Athans and Falb (1966). The optimal control signal is found by first computing an approximation of the minimum-time-to-reach value function, which is a viscosity solution of a Hamilton–Jacobi–Bellman (HJB) PDE. The numerical approximation of the value function is then used to generate the optimal state trajectory and the optimal control signal by forward simulation. We now elaborate on the theoretical details of this method.

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

$$t_{\text{end}}(x, \tilde{u}(t)) = \min\{t : x(t) \in \mathcal{T} | x(0) = x\}. \tag{14}$$

The terminal time can be infinite when the trajectory $x(t)$ never reaches the target set. This can occur if the constraint in (10) is so small that the controller does not have enough control authority to over-ride the system’s natural autonomous dynamics. The terminal time is not known at the outset, and is only found through calculating the optimal stimulus and optimal state trajectories.

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

$$J(x, \tilde{u}(t)) = \int_0^{t_{\text{end}}} g(x(t), \tilde{u}(t))dt + q(x(t_{\text{end}})). \tag{15}$$

Here, $g(x(t), \tilde{u}(t)) \geq 0$ is the time-additive component of the cost function, and $q(x(t_{\text{end}}))$ is the terminal component. We present Eq. 15 in full generality to be consistent with the classical way of presenting this theoretical background. However for our control objective, which corresponds to minimizing the time to reach the target point, these functions take very simple forms:

$$g(x, \tilde{u}) = 1 \tag{16}$$

and

$$q(x(t_{\text{end}})) = 0. \tag{17}$$

Substituting these functions, Eq. 15 simplifies to

$$J(x, \tilde{u}(t)) = t_{\text{end}}(x, \tilde{u}(t)) \tag{18}$$

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

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

$$\mathcal{V}(x) = \inf_{\tilde{u}(t) \in \tilde{\mathcal{U}}} J(x, \tilde{u}) = \inf_{\tilde{u}(t) \in \tilde{\mathcal{U}}} t_{\text{end}}(x, \tilde{u}(t)). \tag{19}$$

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

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

$$0 = \min_{\tilde{u} \in \tilde{\mathcal{U}}} \{1 + \nabla \mathcal{V}(x) \cdot (F(x) + b\tilde{u})\} \tag{20}$$

with the boundary condition

$$\mathcal{V}(x) = 0 \quad \forall x \in \mathcal{T}. \tag{21}$$

This equation represents a critical step in the theoretical foundation of our optimal control design process. We refer the reader to (Bardia and Capuzzo-Dolcetta 1997) for a detailed derivations and proofs of Eq. 20, and an introduction to viscosity solutions of PDEs. If we can find a function $\mathcal{V}(x)$ that satisfies Eq. 20 in the viscosity sense, the rest of the control design is quite straightforward. Unfortunately, equations

of this form can be quite difficult to solve, and no general methods exist.

We can, however, analytically simplify Eq. 20 by realizing that the only the last term, $\nabla\mathcal{V}(x) \cdot b\tilde{u}$, is involved in the minimization. The value of \tilde{u} that minimizes Eq. 20 is typically called the “H-minimal” control, or \tilde{u}_H (Athans and Falb 1966). The vector b has only a single nonzero entry, so we can easily perform the minimization to find

$$\tilde{u}_H = -\text{sgn}(\nabla\mathcal{V}(x) \cdot b) = -\text{sgn}\left(\frac{\partial\mathcal{V}}{\partial x_1}\right), \quad (22)$$

since μ and K_s are strictly positive. We see from Eq. 22 that $\tilde{u}_H \in \{-1, +1\}$ which means the optimal control magnitude will always be equal to its maximum value. In fact, the solutions to all minimum-time optimal control problems are of this “bang–bang” type, as follows from Pontryagin’s Minimum Principle (Pontryagin et al. 1962).

We substitute (22) into (20), expand into components, and recall that for a real number y , $-y \text{sgn}(y) = -|y|$. This removes the minimization functional to yield a slightly simpler expression of the Hamilton–Jacobi–Bellman equation:

$$0 = 1 + \nabla\mathcal{V}(x) \cdot F(x) - |\nabla\mathcal{V}(x) \cdot b| \quad (23)$$

The solution $\mathcal{V}(x)$ of Eq. 23, in the viscosity sense, enables the computation of the optimal state-feedback (and ultimately event-based open-loop) policies \tilde{u}^* and the corresponding state trajectories x^* for any given starting state x . Given $\mathcal{V}(x)$, the state-feedback form of the optimal control policy $\tilde{u}^*(x)$ is simply equal to the H-minimal control evaluated at x .

$$\begin{aligned} \tilde{u}^*(x) &= \arg \min\{1 + \nabla\mathcal{V}(x) \cdot (F(x) + b\tilde{u})\} \\ &= \tilde{u}_H(x) = -\text{sgn}\left(\frac{\partial\mathcal{V}}{\partial x_1}\Big|_x\right) \end{aligned} \quad (24)$$

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

$$\begin{aligned} \dot{x}^*(t) &= F(x^*(t)) + b\tilde{u}^*(x^*(t)) \\ &= F(x^*(t)) - b \text{sgn}\left(\frac{\partial\mathcal{V}}{\partial x_1}\Big|_{x^*(t)}\right). \end{aligned} \quad (25)$$

We can then compute the open-loop optimal control signal $\tilde{u}^*(x_0, t)$ for all $t \in [0, t_{\text{end}}(x_0)]$ by simulating Eq. 25 starting from initial position $x(0) = x_0$ until x reaches the target set \mathcal{T} at time $t = t_{\text{end}}(x_0) = \mathcal{V}(x_0)$. The simulation provides the optimal x trajectory, which we use to calculate the optimum control through Eq. 24. Thus, given any initial condition x_0 we have all the necessary tools to calculate a variable-time-length open-loop control signal $\tilde{u}^*(x_0, t)$.

In our event-based framework, only the spiking state, shown as a blue square in Fig. 1, is observable and its detection triggers the execution of the pre-computed open-loop optimal control signal $\tilde{u}^*(x_{\text{spike}}, t)$.

The crucial component of this machinery is an accurate solution or approximation of the cost function $\mathcal{V}(x)$ over the entire domain of x . Nonlinear Hamilton–Jacobi–Bellman PDEs are notoriously difficult to solve, especially in stationary form as Eq. 23.

4.1 Numerical methods

Our goal is to solve Eq. 23, with boundary conditions in (21), to obtain an approximation to the minimum-time-to-reach value function $\mathcal{V}(x)$. This PDE is both stationary and discontinuous, which renders its solution nontrivial. Computational tools exist for solving time-dependent PDEs, in particular Mitchell (2008) provides a package to solve HJB equations of the following form (among others):

$$0 = \frac{\partial\phi}{\partial t}(x, t) + H(x, t, \phi, \nabla\phi) \quad (26)$$

using level set methods. Following Mitchell (2007) and Osher (1993), we seek an auxiliary form of our original HJB Equations 23 and 21 that converts them into a quasi-time-dependent form.

We define a function

$$G(x, \nabla\mathcal{V}(x)) = 1 + \nabla\mathcal{V}(x) \cdot F(x) - |\nabla\mathcal{V}(x) \cdot b| \quad (27)$$

which allows us to write our HJB equation (23) as

$$\begin{aligned} G(x, \nabla\mathcal{V}(x)) &= 0 \quad \text{on } \mathcal{D} \setminus \partial\mathcal{T} \\ \mathcal{V}(x) &= 0 \quad \text{on } \partial\mathcal{T} \end{aligned} \quad (28)$$

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

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

$$\begin{aligned} \mathcal{V}(x) &\leftarrow s \\ \nabla\mathcal{V}(x) &\leftarrow \begin{aligned} \nabla\phi(x, s) \\ \frac{\partial\phi}{\partial s}(x, s) \end{aligned} \end{aligned} \quad (29)$$

which is valid as long as the following technical condition is satisfied (Osher 1993):

$$\sum_{i=1}^d p_i \frac{\partial G(x, p)}{\partial p_i} \neq 0 \quad \text{on } \partial\mathcal{T}. \quad (30)$$

This condition essentially states that the boundary conditions must be noncharacteristic.

This allows us to cast the Hamilton–Jacobi–Bellman Eq. 28 into the form of Eq. 26 which is compatible with ToolboxLS (Mitchell 2007). The resulting auxiliary PDE is:

$$0 = \frac{\partial\phi}{\partial s}(x, s) + (\nabla\phi(x, s) \cdot F(x) - |\nabla\phi(x, s) \cdot b|). \quad (31)$$

with boundary conditions

$$\begin{aligned} \phi(x, 0) &= 0 \quad \text{for } x \in \partial\mathcal{T} \\ \phi(x, 0) &< 0 \quad \text{for } x \in \mathcal{T} \setminus \partial\mathcal{T} \\ \phi(x, 0) &> 0 \quad \text{for } x \in \mathcal{D} \setminus \mathcal{T}. \end{aligned} \tag{32}$$

Our new objective is to solve (31) for $\phi(x, s)$. Then we can extract the cost function

$$\mathcal{V}(x) = \{s \mid \phi(x, s) = 0\}. \tag{33}$$

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

An approximation of $\phi(x, s)$ is solved by creating a uniform grid over a rectangular region of state space enclosing the periodic orbit. ToolboxLS solves (31) using a Lax–Friedrichs scheme for the spatial dimension and a Runge–Kutta time stepping scheme. The value function $\mathcal{V}(x)$ is extracted, using Eq. 33, by interpolation in the quasi-time variable s .

5 Results

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

5.1 Phase randomization: single neuron results

We seek a phase-randomizing control scheme, so we choose the target set to be the unstable equilibrium point inside the periodic orbit, shown as the black X in Fig. 1. A particular challenge of the reduced Hodgkin–Huxley model is that this unstable fixed point is near the periodic orbit. This situation requires the radius of the target set approximating circle ε to be quite small, which in turn drives the need for a fine grid on which to compute the solution of the Hamilton–Jacobi–Bellman PDE. The following set of calculations uses $\varepsilon = 0.005$ and $\mu = 10 \text{ mA}$ on a 151×151 grid, uniformly spaced in the scaled x coordinates.

Figure 2 shows uniformly spaced level sets of the resulting approximation of the value function surface over the domain of interest. Recall that the value of this surface $\mathcal{V}(x)$ at a point x is the minimum time it will take the system to reach the target set under optimal control. At the spike point, the value is 6.88 which indicates that when the event-based controller is activated, it will drive the neuron to the target set in approximately 6.88 ms.




Fig. 2 Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple circle). The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, n) using a control signal bounded by $\mu = 10$ (color online)

The plot of the state-feedback form of the optimal control law is shown in Fig. 3a. As mentioned in Sect. 4, the optimal control law takes only two values, $\{-\mu, +\mu\}$, with the sign determined by simply checking the sign of the value function’s x_1 -directional derivative, per (22). The optimal trajectory, shown as a blue dashed line in Fig. 3a, was computed by simulating the dynamical system beginning at the spike point shown as the blue square. At each time step, the optimal control signal (which depends only on the current location in state-space) is applied to the system. Under influence of the optimal control signal, the system initially follows the upper lobe of its periodic orbit, then begins to deviate to the left as it approaches the lower left knee of the orbit, at which point it cuts across the orbit to reach the target set in approximately 6.88 ms.

Figure 3b shows the time profile of the optimal trajectory, as well as the optimal control signal, which was simply recorded point-by-point during the simulation. This optimal control signal will be the open-loop waveform that is used by the event-based controller to randomize the phase of the neuron.

5.2 Phase randomization: ensemble results

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




Fig. 3 **a** State-feedback optimal control law (grayscale background) with the periodic orbit (red line), optimal state trajectory (blue dashed line) and the target set (purple circle). Gray regions signify maximum control in the negative direction ($I = -10$ mA), and white regions signify maximum control in the positive direction ($I = +10$ mA). **b** Optimal voltage (black line) and gating variable (green line) trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, starting from the spike point, in minimal time ($t_{\text{end}} \approx 6.88$ ms) (color online)

$$\theta = \frac{2\pi(t_{\text{next}} \bmod T)}{T} \quad (34)$$

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

Figure 4a shows a histogram representing the distribution of phases caused by the noise process $D\eta(t)$ without any control input. Note the tight unimodal shape of this distribution is to be expected given the attractive strength of the periodic orbit and that D is small.

The simulations are repeated, but now we add the optimal control signal as input to each neuron, in addition to the background noise. The resulting phase histogram is shown in Fig. 4b. The control scheme greatly widens the phase distribution of the ensemble after a single application. If the phase

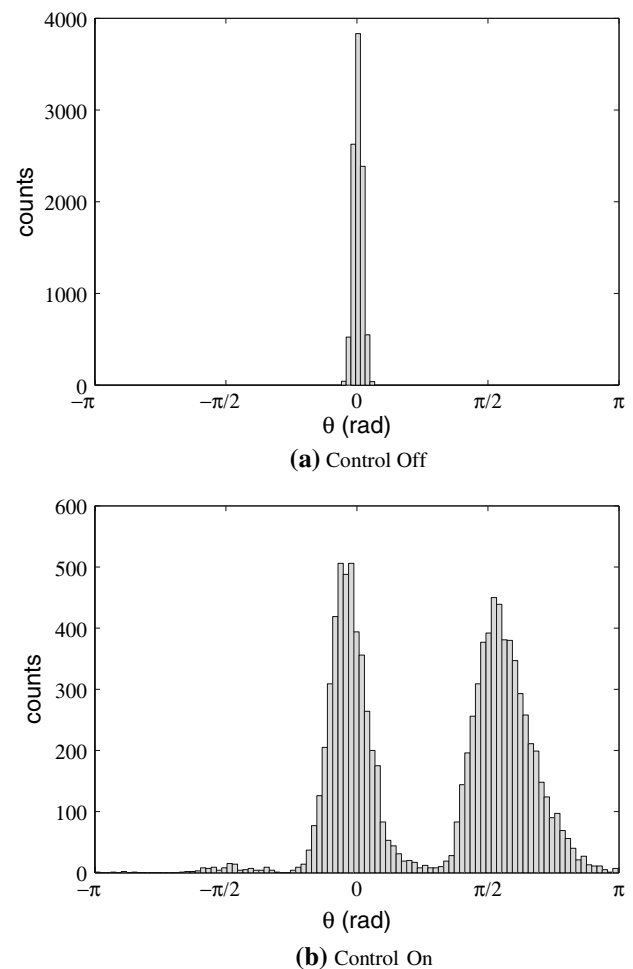


Fig. 4 Ensemble asymptotic phase results for 10^5 uncoupled neurons, each starting from the spiking point. **a** Neurons are exposed only to unit-magnitude noise, **b** neurons are also exposed to the phase randomizing control signal; **a** control off, **b** control on

randomization was perfect, we would expect a uniform distribution, but in Fig. 4b we see a nonuniform bimodal distribution. This arises from the fact that we have approximated the phaseless set, which is a single point, by a small ε -ball. Due to the geometry of the ODE's vector field around the target set, optimal trajectories tend to approach this ε -ball non-uniformly, which for this particular case results in a bimodal distribution. The precise shape of the distribution is governed by the underlying ODE, the value of ε , the magnitude of the noise D , and the magnitude constraint μ .

Complete phase randomization after a single application of the optimal control signal is not necessary. As long as the phase distribution is significantly spread out, we can simply apply the control signal multiple times. We will illustrate this in the context of a globally coupled network in Sect. 6.




Fig. 5 Level sets of the numerical approximation of the value function V for optimal spike-rate-increasing control. Superimposed are the periodic orbit (red line) and the target set (purple ellipse). Again, we consider $\mu = 10$ (color online)

5.3 Spike rate control

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

Figure 5 shows the level sets of the value function for this example. The state-feedback control law and optimal state trajectories are shown in Fig. 6a and b. The optimal control signal, shown in Fig. 6b, is interesting in that it is not simply using the maximum magnitude control in a single direction, which would be equivalent to increasing the baseline current. To optimally increase the firing rate of this neuron, the controller switches between $+\mu$ and $-\mu$. We see that the optimal state trajectory, the blue dashed line on Fig. 6a, stays quite close to the unperturbed periodic orbit, so the control is not radically altering the dynamics, just speeding up the spiking from its uncontrolled natural period of 11.84 ms to a controlled period of 7.74 ms.

5.4 Connection to isochrons

The same numerical techniques used to calculate the optimal spike rate increasing control stimulus can be used to visualize an approximation of the model's isochrons in a straightforward manner. If we were to set the stimulus magnitude constraint $\mu = 0$ and repeat the optimal spike rate increase stimulus design procedure discussed in the previous section, we would be exactly calculating isochrons. To understand why this is the case, consider the meaning of the value function. For the control objective of driving the system to the

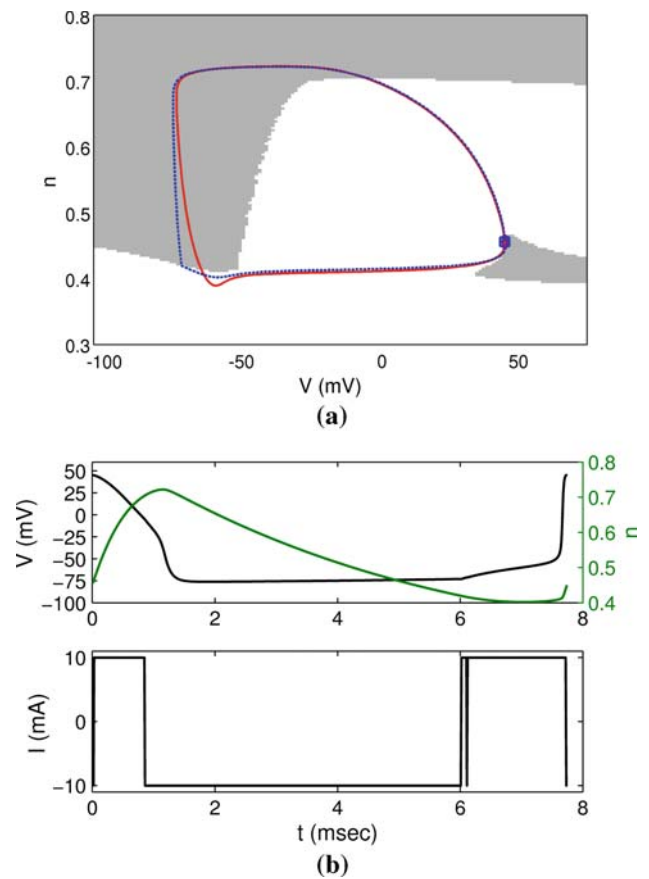


Fig. 6 **a** State-feedback optimal control law for optimal spike rate increase. Notation in this figure is the same as Fig. 3. **b** Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom). Here, the objective is to reach an ϵ -ball centered at the spike point in minimal time starting from the spike point. This control protocol decreases the interspike interval neuron from 11.84 to 7.74 ms, resulting the fastest spiking rate possible with a stimulus constrained by 10 mA (color online)

spiking point, the value function is the time it takes, under optimal stimulus, for the system to be driven from position x to the spiking point. Now, if the control magnitude bounds are zero, we have eliminated the controller's ability to affect the dynamics of the system. So the value function in the case of $\mu = 0$ can be interpreted as the time it takes, under no stimulus, for the system to evolve from position x to the spiking point under its natural dynamics. Level sets of this value function are, therefore, curves connecting points in phase space that will reach the spiking point at the same time, in absence of stimulus, i.e., the isochrons.

There are numerical difficulties with setting $\mu = 0$ so we create approximations to isochrons by calculating level sets of the value function for the case of small μ . Figure 7 shows the isochron approximation calculated using $\mu = 1$. This compares well with the isochrons shown in Fig. 1 that were calculated accurately using backward integration techniques as by Campbell et al. (1989) and Josic et al. (2006).




Fig. 7 Level sets approximating isochrons for the reduced Hodgkin-Huxley model generated using parameters $\varepsilon = 0.01, \mu = 1$ (color online)




Fig. 8 Level sets of the numerical approximation of the value function \mathcal{V} with the periodic orbit (red line) and the target set (purple line) for the FitzHugh–Nagumo model. The value of \mathcal{V} represents the minimum time to reach the phaseless set, starting at a point (V, w) using a control signal bounded by $\mu = 0.2$ (color online)

5.5 FitzHugh–Nagumo model

The proposed method is not dependent on the details of the particular model, such as how close the unstable fixed point is to the periodic orbit. To illustrate the broad applicability of the optimal stimulus design method, we select another classical neural model, the FitzHugh–Nagumo model (FitzHugh 1961; Nagumo et al. 1962) in a dimensionless form taken from Keener and Sneyd (1998):

$$\begin{aligned} \delta \dot{V} &= V(V + a)(1 - V) - w + u \\ \dot{w} &= V - 0.5w. \end{aligned} \tag{35}$$

We take the time scale separation variable $\delta = 0.01$, and set the parameter $a = 0.6$ to position the unstable fixed point near the center of the region enclosed by the periodic orbit, which is encircled by the purple ε -radius target circle in Fig. 9a. Its worth noting that this neuron model has very different time, input, and voltage scaling than the Hodgkin–Huxley model. Since the unstable equilibrium point is far inside the periodic orbit, this model would be particularly difficult to control using a timed impulse to instantaneously kick the state from the limit cycle to unstable equilibrium, as proposed by Winfree (2001). By employing the control synthesis method presented above, we can precisely calculate the optimal control signal to drive the FitzHugh–Nagumo model from its spiking point to an ε disk centered at the unstable fixed point. We take this target radius to be $\varepsilon = 0.01$, compute the value function, and plot its level sets in Fig. 8. Figure 9 shows the optimal trajectories and the optimal control stimulus when the control magnitude constraint $\mu = 0.2$.

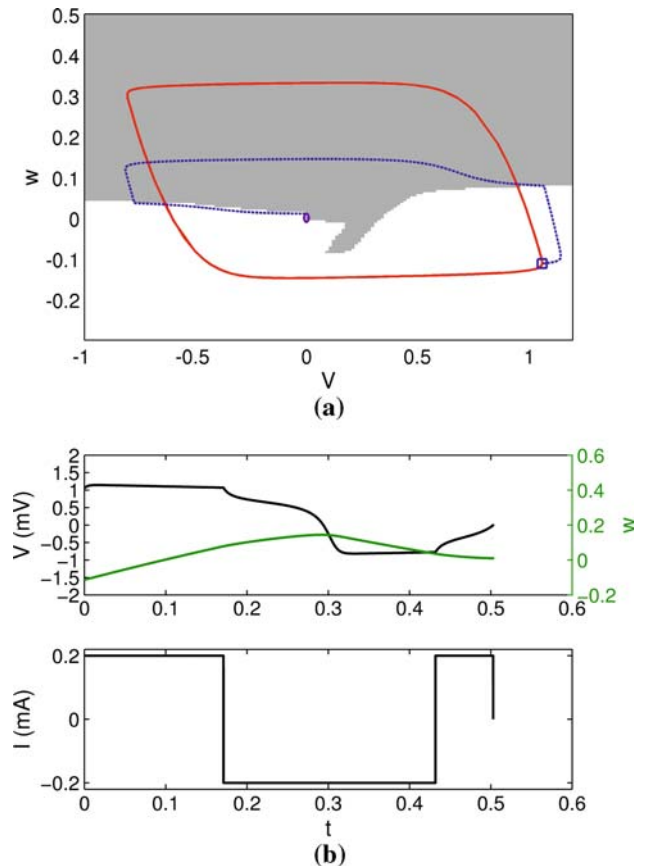


Fig. 9 a State-feedback optimal control law, optimal state trajectory, and the target set for the FitzHugh–Nagumo model with $\mu = 0.2$. Notation in this figure is the same as in Fig. 3. **b** Optimal voltage and gating variable trajectories (top) and the optimal control signal (bottom) for reaching the target set, an ε -ball centered at the unstable fixed point, in minimal time ($t_{\text{end}} \approx 0.5$ ms) (color online)

6 Extension: globally coupled networks

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

By applying the optimal phase randomizing control stimulus, we seek to desynchronize the spike times of the network. Following Keener and Sneyd (1998), we write dynamical equations for a network of Hodgkin–Huxley neurons obeying (1) as

$$\dot{x}_i = F(x_i) + e_1(D\eta_i + I(t))/C + \frac{k}{N}M_1 \sum_{j=1}^N (x_j - x_i) \quad (36)$$

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

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

account for the fact that the neurons are stimulated and coupled only through the voltage component of the state. The variable k is the coupling strength, which we take to be small.

Here, $I(t)$ is the control stimulus which is applied to all neurons simultaneously. Again we use an event-based feedback framework, but in the network context we need a population-level event definition. To this end, we introduce the concept of a network observable—a quantity from which we can infer the level of synchrony in the network. We use the average voltage as our network observable, written simply as

$$\bar{V}(t) = \frac{1}{N} \sum_{i=1}^N V_i(t). \quad (37)$$

Were these neurons modeled as phase oscillators, we could simply use Kuramoto’s order parameter, as used by Danzli et al. (2008). In this context, however, phase is unknown—only voltages are measurable.

We envision a simplified representation of electrical deep brain stimulation with two electrodes. One electrode is an input that stimulates each neuron with the same control signal $I(t)$. The other electrode is an output that measures the average voltage $\bar{V}(t)$ of the population (a simplification of local field potential where we neglect spatial weighting).

We define an *event* as the value of the network observable crossing some threshold, indicating that most of the neurons are spiking at the same time. This threshold is a tunable parameter, generally greater than zero for a network of two-dimensional Hodgkin–Huxley neurons.

In the uncontrolled system with uniform initial distribution, the neurons become synchronized due to coupling, even in the presence of moderate noise. One can view this as a simple example of pathological neural synchronization, linked to

the symptoms of neural disorders such as Parkinson’s disease (Benabid et al. 1991) and epilepsy (Steriade 2003).

To demonstrate the effectiveness of the proposed control scheme, we simulate the network represented by Eq. 36 for $N = 100$ neurons, each starting synchronized at the spike point. This corresponds to the pathological spike synchronization described above. The synchronous spiking is detected by a super-threshold value of $\bar{V}(t)$, our network observable. Super-threshold average voltage constitutes an event, which triggers the controller to stimulate all the neurons with the optimal control waveform calculated in Sect. 5.1. Recall that in the absence of coupling, this stimulus causes the neurons’ states to approach the phase randomizing unstable fixed point. If the coupling is relatively weak, the controller is still able to drive the neurons close enough to the unstable equilibrium for the phase randomization to be effective. The neurons’ phases are randomized by the stimulus, thus diminishing the value of the network observable $\bar{V}(t)$. The controller then shuts off until the coupling draws the network back toward synchronous spiking ($\bar{V}(t) \geq V_{\text{thresh}}$) at which point the optimal stimulus is again triggered. We note that if the system does not develop synchronized spiking again, the controller will continue to remain off, thus achieving demand-controlled desynchronization similar to Popovych et al. (2006).

Figure 10a shows results for a population of 100 neurons globally coupled with uniform coupling strength $k = 0.01$ and each under the influence of independently drawn white noise with magnitude $D = 1.0$. We have kept the control magnitude constraint $\mu = 10$ mA and used the exact optimal stimulus waveform calculated in Sect. 5.1.

In this simulation, we begin with the controller switched off, to illustrate the synchronous spiking of the network. We turn the controller on at $t = 20$ ms, but recall that this is an event-based controller, so it remains inactive until the population level observable, the average voltage, crosses the threshold V_{thresh} . The first population-level spiking event detected by the controller occurs at $t \approx 23$ ms, at which point the open-loop optimal stimulus waveform is triggered. After a single application, there is still a significant level of spiking synchrony, closely related to the bimodal distribution caused by the randomization described in Sect. 5.2. The controller is triggered again at $t \approx 37$ ms, and again the open-loop stimulus is activated. This time, the spike times of the population are randomized to a much greater degree, which can be easily visualized by constructing a raster plot of the spike times across the population, shown in Fig. 10b.

We note that the resulting distribution of spike times is not uniform; there are loose vertical “bands” of spiking activity in the raster plot after the control is turned on. This is due to two factors—the coupling and the approximation of the target set by a disk rather than a point (which is related to the change in distribution due to a single application of the controller, as

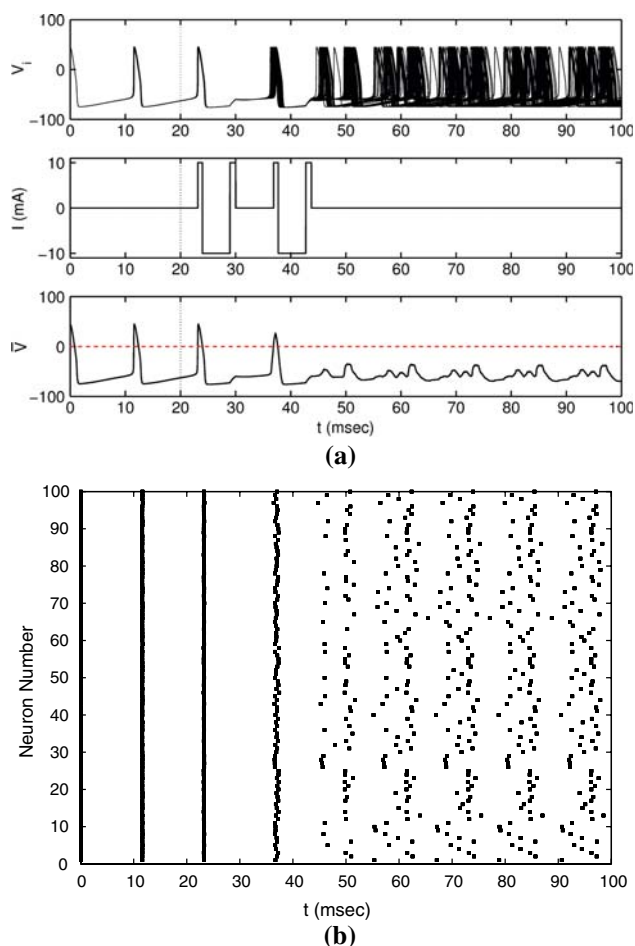


Fig. 10 **a** Network results for $N = 100$ globally coupled Hodgkin-Huxley neurons, $k = 0.01$, $D = 1.0$, $\mu = 10$. *Top panel* shows voltage traces for each neuron. *Middle panel* shows the control signal applied to the population. Controller is activated at $t = 20$ ms indicated by the vertical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

discussed in Sect. 5.2). However, the pathology is primarily due to synchronized spiking, shown as the tightly synchronized vertical stripes in the raster plot before the controller is activated. The event-based controller destroys this type of synchronization, which may provide therapeutic utility for pathological synchronization.

7 Summary and discussion

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

increase the firing rate for a given stimulus magnitude constraint.

We have illustrated this method using two-dimensional neuron models so that the resulting value function can be easily visualized as level sets in the plane, but we remark here that the method is not limited to planar systems. While the computational time for the Hamilton–Jacobi–Bellman PDE grows with dimension, it is still quite feasible to consider models with higher dimensions.

Furthermore, by targeting the spiking point and setting the magnitude constraint to be small (ideally zero), the numerical viscosity solutions of the Hamilton–Jacobi–Bellman PDE can be used to visualize isochrons—a useful tool for characterizing the phase space of spiking neuron ODE models.

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

An important concern with this design method is that it depends strongly on having a good model for the neuron, which in practice, can be challenging. In future study, we seek to combine the optimal stimulus design procedure outlined here with ongoing research in model identification. We also plan to examine the influence of network size and coupling strength on the effectiveness of the phase desynchronizing control scheme presented here.

Acknowledgments We thank G. Orosz, P. Atzberger, and E. Shea-Brown for insightful discussions regarding this study. We also recognize I. Mitchell for his exceptional Level Set Methods Toolbox, which provided several numerical routines that were important for generating these results. P. Danzl is supported by the National Science Foundation through the Integrative Graduate Education and Research Traineeship program. J. Hespanha is supported by the National Science Foundation under grants ECCS-0725485 and ECCS-0835847. J. Moehlis is supported by the National Science Foundation under Grant NSF-0547606.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Åström KJ, Bernhardsson B (2003) Systems with Lebesgue sampling. In: Rantzer A, Byrnes CI (eds) Directions in mathematical systems theory and optimization, vol XIII. Springer, Berlin
- Athans M, Falb PL (1966) Optimal control: an introduction to the theory and its applications. McGraw-Hill, New York
- Bardi M, Capuzzo-Dolcetta I (1997) Optimal control and viscosity solutions of Hamilton–Jacobi–Bellman equations. Birkhauser, Boston
- Benabid AL, Pollak P, Gervason C, Hoffmann D, Gao DM, Hommel M, Perret JE, De Rougemont J (1991) Long-term suppression of

- tremor by chronic stimulation of the ventral intermediate thalamic nucleus. *Lancet* 337:403–406
- Campbell A, Gonzalez A, Gonzalez DL, Piro O, Larrondo HA (1989) Isochrones and the dynamics of kicked oscillators. *Physica A* 155(3):565–584
- Danzl P, Moehlis J (2007) Event-based feedback control of nonlinear oscillators using phase response curves. In: *Proceedings of the 46th IEEE conference on decision and control*, pp 5806–5811, New Orleans, LA
- Danzl P, Moehlis J (2008) Spike timing control of oscillatory neuron models using impulsive and quasi-impulsive charge-balanced inputs. In: *Proceedings of the 2008 American control conference*, pp 171–176, Seattle, WA
- Danzl P, Hansen R, Bonnet G, Moehlis J (2008) Partial phase synchronization of neural populations due to random Poisson inputs. *J Comput Neurosci* 25(1):141–157
- FitzHugh R (1961) Impulses and physiological states in theoretical models of nerve membrane. *Biophys J* 1(6):445–466
- Guckenheimer J (1975) Isochrons and phaseless sets. *J Math Biol* 1:259–273
- Hodgkin AL, Huxley AF (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 117:500–544
- Josic K, Shea-Brown ET, Moehlis J (2006) Isochron. *Scholarpedia* 1(8):1361. www.scholarpedia.org
- Keener J, Sneyd J (1998) *Mathematical physiology*. Springer, New York
- Khalil HK (2002) *Nonlinear systems*. Prentice Hall, Upper Saddle River
- Mitchell IM (2007) A toolbox of level set methods. Technical Report UBC CS TR-2007-11
- Mitchell IM (2008) The flexible, extensible and efficient toolbox of level set methods. *J Sci Comput* 35(2):300–329
- Moehlis J (2006) Canards for a reduction of the Hodgkin–Huxley equations. *J Math Biol* 52(2):141–153
- Nagumo J, Arimoto S, Yoshizawa S (1962) An active pulse transmission line simulating nerve axon. *Proc IRE* 50(10):2061–2070
- Netoff TL, Acker CD, Bettencourt JC, White JA (2005) Beyond two-cell networks: experimental measurement of neuronal responses to multiple synaptic inputs. *J Comput Neurosci* 18(3):287–295
- Nini A, Feingold A, Sloviter H, 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(4):1800–1805
- Osher S (1993) A level set formulation for the solution of the Dirichlet problem for Hamilton–Jacobi equations. *SIAM J Math Anal* 24:1145
- Pare D, Curro’Dossi R, Steriade M (1990) Neuronal basis of the Parkinsonian resting tremor: a hypothesis and its implications for treatment. *Neuroscience* 35:217–226
- Pontryagin LS, Trirogoff KN, Neustadt LW (1962) *The mathematical theory of optimal processes*. Wiley, New York
- Popovych OV, Hauptmann C, Tass PA (2006) Control of neuronal synchrony by nonlinear delayed feedback. *Biol Cybern* 95(1):69–85
- Steriade M (2003) *Neuronal substrates of sleep and epilepsy*. Cambridge University Press, New York
- Tass PA (1999) *Phase resetting in medicine and biology*. Springer, New York
- Tass PA (2000) Effective desynchronization by means of double-pulse phase resetting. *Europhys Lett* 53:15–21
- Winfree A (2001) *The geometry of biological time*, 2nd edn. Springer, New York