J. Neural Eng. **10** (2013) 036005 (12pp) doi:10.1088/1741-2560/10/3/036005

Minimum energy control for *in vitro* neurons

Ali Nabi[1], **Tyler Stigen**[2], **Jeff Moehlis**[1] **and Theoden Netoff**[2]

[1] Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106, USA
[2] Department of Biomedical Engineering, University of Minnesota, Minneapolis, MN 55455, USA

E-mail: nabi@engineering.ucsb.edu

Received 27 December 2012
Accepted for publication 13 March 2013
Published 10 April 2013
Online at stacks.iop.org/JNE/10/036005

Abstract
Objective. To demonstrate the applicability of optimal control theory for designing minimum energy charge-balanced input waveforms for single periodically-firing *in vitro* neurons from brain slices of Long-Evans rats. *Approach*. The method of control uses the phase model of a neuron and does not require prior knowledge of the neuron's biological details. The phase model of a neuron is a one-dimensional model that is characterized by the neuron's phase response curve (PRC), a sensitivity measure of the neuron to a stimulus applied at different points in its firing cycle. The PRC for each neuron is experimentally obtained by measuring the shift in phase due to a short-duration pulse injected into the periodically-firing neuron at various phase values. Based on the measured PRC, continuous-time, charge-balanced, minimum energy control waveforms have been designed to regulate the next firing time of the neuron upon application at the onset of an action potential. *Main result*. The designed waveforms can achieve the inter-spike-interval regulation for *in vitro* neurons with energy levels that are lower than those of conventional monophasic pulsatile inputs of past studies by at least an order of magnitude. They also provide the advantage of being charge-balanced. The energy efficiency of these waveforms is also shown by performing several supporting simulations that compare the performance of the designed waveforms against that of phase shuffled surrogate inputs, variants of the minimum energy waveforms obtained from suboptimal PRCs, as well as pulsatile stimuli that are applied at the point of maximum PRC. It was found that the minimum energy waveforms perform better than all other stimuli both in terms of control and in the amount of energy used. Specifically, it was seen that these charge-balanced waveforms use at least an order of magnitude less energy than conventional monophasic pulsatile stimuli. *Significance.* The significance of this work is that it uses concepts from the theory of optimal control and introduces a novel approach in designing minimum energy charge-balanced input waveforms for neurons that are robust to noise and implementable in electrophysiological experiments.

(Some figures may appear in colour only in the online journal)

1. Introduction

This work investigates the possibility of practical implementation of alternative stimulation protocols to conventional pulsatile methods that are more efficient in terms of performance and the amount of energy used. On a population level, for example, this idea could prove useful in prolonging battery life for neurostimulators that are used to treat various neurological diseases. As a first step, in this work, we look at single *in vitro* neurons and show that by changing conventional pulsatile stimulation methods to continuous-time analytically-derived optimal stimuli, one can greatly reduce the input energy needed for regulating the inter-spike-intervals (ISI) of periodically spiking neurons with intracellular stimulation.

Optimal control theory is employed to design charge-balanced, continuous-time, minimum energy input stimuli

that can change the ISI of a neuron to pre-specified values. The optimal control algorithm uses the neuron's phase response curve (PRC) which is measured experimentally using the so-called direct method (see section 2.4). In previous work [1], following [2], we have applied optimal control theory to phase models of neurons to derive charge-balanced minimum energy ISI regulatory input stimuli. In this work, we extend the theoretical results from [1] and show the applicability of this method in practice by testing the controller on *in vitro* pyramidal neurons in the CA1 region of rat hippocampus. In addition to being minimum energy, charge-balanced, and continuous-time, the designed input stimuli are also low amplitude and of event-based nature. The event-based nature of the input means that it is only applied when a specific event, e.g., an action potential, occurs. This way, the onset of an action potential could be considered as a feedback signal that triggers the input. The controller would then function without any need for continuous feedback from the neuron.

In recent years, various studies have considered different control approaches to neuron systems both on a population level [3–17] and on a single neuron level [1, 2, 18–25]. However, few papers have shown experimental evidence of the applicability of their method in practice. In [3], the authors have employed closed-loop control techniques to test their hypothesis that the persistent synchronized bursting behavior seen in high-density cortical cultures is due to the lack of input from other brain regions. They found that rapid stimulation through multiple electrodes that is fine-tuned and adjusted through closed-loop firing rate feedback reduces the bursting behavior. In [7], a model is developed to describe thalamic deep brain stimulation (DBS) for patients with essential tremor, and the authors present experimental results that support the idea that for high frequency pulsatile stimulation of the ventral intermediate nucleus of the thalamus there is an optimal voltage for maximum tremor suppression. In [12], the authors have applied pallidal and corticopallidal closed-loop stimulation on 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine primates of Parkinson's disease to modulate the pathological oscillatory activity of the basal ganglia-cortical networks. They have suggested that the closed-loop DBS may provide more effective management of advanced Parkinson's disease. In [15], it is demonstrated that a seizure-triggered, feedback transcranial electrical stimulation can reduce spike-and-wave episodes in a rodent model of generalized epilepsy. On single neuron level, the stabilization of the ISI of periodically firing *in vitro* neurons is tested in [21] by designing and implementing a feedback proportional-integral (PI) controller that changes the input current to the neuron based on the history of previous ISIs. Also, in [24], a model-independent control algorithm is considered that regulates an *in vitro* neuron's ISI by inputting an appropriately sized and timed pulse.

The algorithm that we present in this paper utilizes a neuron's PRC to design charge-balanced optimal (i.e., minimum energy) stimulus waveforms that can control the spiking of neurons. The method is tested both computationally and experimentally. In computations, a model of a cortical pyramidal neuron is used [26]. The PRC for the model is obtained using the direct method (see section 2.4). The minimum energy waveforms are tested against conventional pulsatile controls as well as phase shuffled surrogate waveforms and variants of the minimum energy waveforms obtained from suboptimal PRCs. In experiments, the method is tested on *in vitro* CA1 hippocampal neurons. Here, the PRC for each neuron is measured through the direct method, and the minimum energy waveforms are designed and implemented in real-time. The performance of the minimum energy controls in experiments is compared with pulsatile control results of [24], and it is shown that the new charge-balanced waveforms achieve the same level of control, but with at least an order of magnitude less energy.

2. Methods

2.1. GA model

We use a model of a pyramidal neuron introduced by Golomb and Amitai [26], hereafter referred to as the *GA model*, to design and test the performance of the minimum energy control in simulations. The GA model is a conductance-based Hodgkin–Huxley [27] type model with five dimensions that incorporates three potassium currents (I_{Kdr}, I_{KA}, $I_{K,\mathrm{slow}}$), two sodium currents (I_{Na}, I_{NaP}), one leak current (I_L), and an externally applied current (I_{app}). The specifics of this model can be found in [26]. We present a slightly modified version for the voltage equation where a noise term is added and the external current stimulus in the original equations, $I_{\mathrm{app}}(t)$, is split into two parts, I_b and $I_c(t)$, as indicated below:

$$\dot{V} = \frac{1}{c}(-I_{\mathrm{Kdr}}(V,n) - I_{\mathrm{KA}}(V,b) - I_{K,\mathrm{slow}}(V,z) - I_{\mathrm{Na}}(V,h) - I_{\mathrm{NaP}}(V) - I_L(V) + I_b + I_c(t)) + \eta(t), \quad (1)$$

where V is the voltage difference across the neuron's membrane in mV, $c = 1\ \mu\mathrm{F\,cm^{-2}}$ is the membrane capacitance, and I_b is a constant baseline current that induces stable periodic spiking of the neuron in the absence of any control current I_c and any noise $\eta(t)$. For $I_b \approx 0.94\ \mu\mathrm{A\,cm^{-2}}$, the neuron spikes with a period of $T_{\mathrm{s}} = 100$ ms which gives $\omega = \frac{2\pi}{T_{\mathrm{s}}} = 0.0628\ \mathrm{rad\,ms^{-1}}$. The noise process $\eta(t)$ is added to generate variability in the ISI as seen in real neurons. We take $\eta(t)$ to be a zero-mean Gaussian white noise process with standard deviation (std) 0.15, i.e., $\eta(t) = \mathcal{N}(0, 0.15)$, to generate approximately 10% variability in the ISI for the GA model. The value of I_b and the noise characteristics are chosen to approximate the behavior of the *in vitro* neurons under study.

2.2. Experimental preparation

The experiments were performed on brain slices prepared from Long-Evans rats post-natal age 14–21 days old. The rats were deeply anesthetized using isoflurane before extraction of the brain. Once extracted, the brain was bathed in chilled artificial cerebral spinal fluid (aCSF) composed of 124 mM NaCl, 2 mM KCl, 2 mM $MgSO_4$, 1.25 mM NaH_2PO_4, 2 mM $CaCl_2$, 26 mM $NaHCO_3$, and 10 mM D-glucose at pH 7.4, 295 mosM [24]. Transverse slices of the ventral horn

of the hipocampal region were sectioned 350 μm thick on a vibratome (Leica Microsystems, Bannockburn, IL). The slices were placed under the microscope with circulating aCSF and neurons were visualized using differential interference contrast optics (Olympus, Center Valley, PA). Whole cell patch-clamp recordings were performed in the CA1 region of the hipocampus using pyramidal cells. Borosilicate capillary pipettes were pulled to 8 MΩ and filled with intracellular recording fluid composed of 120 mM K-gluconate, 10 mM HEPES, 1 mM EGTA, 20 mM KCl, 2 mM $MgCl_2$, 2 mM Na_2ATP, and 0.25 mM Na_3GTP at pH 7.3, 290 mosM. The neuron's membrane potential was amplified and low-pass filtered at 2.4 kHz (Axon 700B; Molecular Devices, Sunnyvale, CA) and digitized on a real-time Linux computer (NiDAQ 6259; National Instruments, Austin, TX).

2.3. Dynamic clamp

Dynamic clamp is a low latency closed-loop control system that connects a computer or an analog device to one or several (virtual or *in vitro*) neurons. In this setting, a patch-clamp amplifier is connected to a data acquisition card (DAQ) and in turn to a computer through a real-time interface [28]. We use the Real-Time eXperiment Interface (RTXI) software, an open source program for real-time experiments (www.rtxi.org). RTXI can be used with a variety of DAQ cards through Comedi project (www.comedi.org) which runs on the real-time Linux nanokernel (www.rtai.org). RTXI is a modular software with a freely available software repository. We carried out experiments at a rate of 5 kHz, which corresponds to a time step of 0.2 ms. The dynamic clamp is used for measuring PRCs in both the virtual neuron and the *in vitro* neurons. It is also used when the minimum energy waveforms calculated using MATLAB (Mathworks, Natick, MA) are applied to the neurons.

2.4. Estimation of PRC

The so-called direct method [29, 30] is implemented to obtain the PRC in both the simulations and the experiments. In the direct method, a short-duration pulse with charge Q_p is injected into the periodically spiking neuron at various times for which the resulting change in the next spike time is measured. This change, referred to as the spike advance (SA), is recorded as a function of the time at which the stimulus was applied (see figure 1). For the results shown in this paper, the pulsatile input has been applied once every sixth cycle of the neuron's oscillation to allow time for the neuron to settle back on its periodic orbit before the next pulsatile input; this is to minimize interactions between stimuli which could affect the estimate of the PRC. Also the pulse duration is fixed at 1 ms. Once the data points for the SA values are collected, a sixth order polynomial is fitted to the data which is constrained to be zero at both ends of the spiking interval [31, 30]. The reason for constraining the PRC to be zero at both ends is to account for the fact that biological neurons generally show little or no sensitivity to inputs right at their spiking point [31]. It is noted that constraining the fit at both ends leaves only four coefficients for the fit, which is necessary to reasonably capture the features and possible sharp changes in the data.

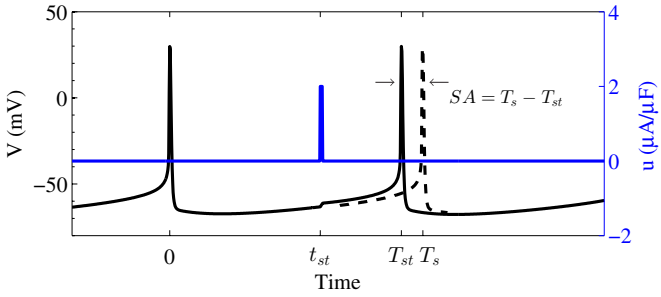

Figure 1. Measuring the PRC with the direct method. Every sixth cycle, a pulse stimulus (blue) is applied at a random time t_{st}, equivalent to a random phase θ_{st}, in the neuron's cycle which changes the next spiking time of the neuron from its natural ISI, T_s to a stimulated ISI, T_{st}. The spike advance, $SA = (T_s - T_{st})$, is measured and recorded.

Assuming linearity, the SA data are normalized by Q_p to emulate the effect of an impulsive input with unit charge. These normalized time data are then scaled according to the relationship $\theta = \frac{2\pi}{T_s}t$ to yield the (infinitesimal) PRC, $Z(\theta)$. This is used to find the minimum energy waveforms as explained in the following section.

In measuring the PRC with the direct method, it is necessary to choose pulse amplitudes that are high enough to produce spike perturbation values beyond those due to the neuron's intrinsic noise. However, if the pulse amplitudes are too high, they can induce instantaneous spikes, which means that the neuron has been perturbed too far off of its periodic orbit to the point where the linearity assumption of the PRC analysis does not hold. This is usually seen when the unperturbed neuron is close to threshold. At these instances, the amount of remaining time to advance the spike, *SA* (ms) may be less than what the pulse stimulus could achieve given the sensitivity of the neuron in that phase: $SA = T_s - T_{st} \geqslant T_s - t_{st}$, with t_{st} being the time at which the stimulus is applied (see figure 1). The effect is that the data points in the last 20%–30% of the period may fall along a straight line of slope -1, that passes through the $(100, 0)$ (ms) point in the spike advance graph (see, for example, figure 4(A)). This line is referred to as the causality line (see [31]) and indicates the upper limit for the spike advance for the neuron. Although some points will inevitably be near or on the line of causality towards the end of the time cycle, a pronounced causality line is an indicator that the stimulation has reached saturation, and thus has induced nonlinearities in the computation of the PRC. In practice, the pulse amplitudes are set by trial and error and on a case by case basis.

To measure the level of induced nonlinearity for the PRCs, a *nonlinearity coefficient*, C_{NL}, is defined as the percentage of the data points that fall in a band neighborhood of the causality line. The width of the band neighborhood is set to be 0.03 which means that if a stimulus causes an advance that is equal to or greater than 97% of the maximum possible advance, then it is considered to have induced nonlinearities in the neuron dynamics. Therefore, any data point that satisfies $(t_{st} - T_{st}) \leqslant 0.03$ is considered to be under the effect of

induced nonlinearity. This way, the more pronounced the causality line, the higher the C_{NL}.

It is noted that for computational models of periodically spiking neurons, the PRC and hence the phase model, can be obtained either by implementing the direct method or by solving the appropriate adjoint equation [32, 33]. For biological neurons however, one can only use the direct method to obtain the PRC experimentally.

2.5. Minimum energy control

The phase model for a neuron under an arbitrary, small external stimulus $u(t)$ is written as [32, 33]

$$\dot{\theta} = \omega + Z(\theta)u(t), \tag{2}$$

where $\theta \in [0, 2\pi)$ is the neuron's phase. By convention, $\theta = 0$ is associated with the onset of a spike for the neuron and the neuron is periodically spiking in the absence of external input with period $T_s = \frac{2\pi}{\omega}$. The natural frequency of the neuron is determined by ω, and $Z(\theta)$ is its PRC.

The objective is to find a bounded input stimulus ($|u(t)| \leqslant u_{\max}$) that is charge-balanced and uses minimum energy to steer the neuron (2) from $\theta(0) = 0$ to $\theta(T_{\mathrm{targ}}) = 2\pi$, where T_{targ} is the desired next spiking time for the neuron. The value of $u_{\max}$ is determined by the maximum allowable current that could be injected into the neuron. This is set by considering the tolerance of the tissue and possible hardware limitations. For this study the value of $u_{\max}$ was set to 1 nA/μF. In order to ensure charge-balance for the input, one can define the variable $q(t) : \mathbb{R} \to \mathbb{R}$ as the total accumulated charge in the neuron at time t due to the external input $u(t)$. Then one can write

$$\dot{q} = u(t), \tag{3}$$

with the boundary conditions $q(0) = q(T_{\mathrm{targ}}) = 0$. Minimizing the total input energy is equivalent to minimization of the cost function

$$C = \int_0^{T_{\mathrm{targ}}} u^2(t)\,\mathrm{d}t. \tag{4}$$

Following standard optimal control theory [34], one can write the Hamiltonian for this optimal control problem as

$$\mathcal{H}(\Phi) = u^2(t) + \lambda_1(t)(\omega + Z(\theta)u(t)) + \lambda_2(t)u(t), \tag{5}$$

where $\Phi = [\theta, q, \lambda_1, \lambda_2, u]$, and $\lambda_1(t)$ and $\lambda_2(t)$ are the Lagrange multipliers (or co-states) associated with the θ-dynamics (2) and the q-dynamics (3), respectively. Using this Hamiltonian, the necessary conditions for optimality are written as

$$\dot{\theta} = \frac{\partial \mathcal{H}}{\partial \lambda_1} \quad \Rightarrow \quad \dot{\theta} = \omega + Z(\theta)u(t), \tag{6}$$

$$\dot{\lambda}_1 = -\frac{\partial \mathcal{H}}{\partial \theta} \quad \Rightarrow \quad \dot{\lambda}_1 = -\lambda_1(t)Z'(\theta)u(t), \tag{7}$$

$$\dot{q} = \frac{\partial \mathcal{H}}{\partial \lambda_2} \quad \Rightarrow \quad \dot{q} = u(t), \tag{8}$$

$$\dot{\lambda}_2 = -\frac{\partial \mathcal{H}}{\partial q} \quad \Rightarrow \quad \dot{\lambda}_2 = 0, \tag{9}$$

where prime represents differentiation with respect to θ [34, 35]. From Pontryagin's minimum principle, an optimal control stimulus is one that minimizes the Hamiltonian (5)

$$u^*(t) = \arg \min_{|u(t)| \leqslant u_{\max}} \big(u(t)^2 + \lambda_1^*(t)(\omega + Z(\theta^*)u(t)) + \lambda_2^*(t)u(t)\big),$$

where the search for the control stimulus is constrained to values bounded in magnitude by $u_{\max}$ and the asterisk superscript denotes optimal values. This equation yields

$$\begin{aligned} u^*(t) &= -\tfrac{1}{2}\psi(t), && |\psi(t)| \leqslant 2u_{\max}, \\ u^*(t) &= -\mathrm{sign}(\psi(t))u_{\max}, && |\psi(t)| > 2u_{\max}. \end{aligned} \tag{10}$$

where $\psi(t) = \lambda_1(t)Z(\theta(t)) + \lambda_2$. We note that λ_2 is a constant according to (9).

By substituting (10) in the system of equations (6)–(9) we arrive at a two point boundary value problem (TPBVP) which we solve using a shooting method. The boundary values for this system are $\theta(0) = q(0) = q(T_{\mathrm{targ}}) = 0$ and $\theta(T_{\mathrm{targ}}) = 2\pi$. This formulation can be solved to yield the minimum energy control stimulus for any oscillatory system for which a PRC can be obtained.

2.6. Surrogates of minimum energy control

In order to better evaluate the performance of the minimum energy control waveforms, a set of surrogate waveforms is devised that use the same amount of energy, but will be shown to perform less efficiently. These are phase shuffled surrogates, where the Fourier transform of the PRC is calculated, the phases are randomized and then an inverse Fourier transform is performed. This preserves all the linear statistics, mean, standard deviation and autocorrelation of the data, but the resulting waveform does not have the correct phases [36].

2.7. Variants of minimum energy control

To further evaluate the performance and robustness of the minimum energy waveforms, variants of these waveforms are computed from variations of the PRC. These variant waveforms are shown to achieve less accurate control with comparable energy levels to those of the minimum energy waveforms. The variations of the PRC are found by fitting sixth order polynomials to positive and negative one standard deviations of the PRC. The fits are constrained to zero at both ends of the ISI.

2.8. Single pulse control

The single pulse control method is adopted from [24] to provide another comparison for the results of the minimum energy method. In [24], a model-independent control algorithm is introduced that regulates *in vitro* neurons' ISI by inputting appropriately sized and timed pulses. Here, the pulses' durations have been fixed at 0.2 ms and are always applied at the point of maximum PRC. By applying pulses of different amplitudes and measuring the resulting spike advances, a mapping is found that gives a prediction of

the needed amplitude for a given advance. This mapping is characterized by an inverse sigmoidal function of the form

$$u_{\text{pls}} = C - D \log\left(\frac{B-A}{SA-A} - 1\right), \tag{11}$$

where u_{pls} gives the pulse amplitude (in A) for a given desired spike advance SA (in s). The constant parameters A, B, C and D are found from fitting the sigmoidal function to the raw data.

In this method, the pulses are timed to occur at the point of maximum PRC, where they would be most effective. This yields the lowest pulse amplitude and thus, highest energy efficiency for the family of all single pulse controls of width 0.2 ms. In fact, one can even argue that this results in higher energy efficiency than any other single pulse protocol with pulse durations less than 1 ms. This is because if one decreases the pulse duration, one has to increase the amplitude to deliver the same amount of stimulating charge which, loosely speaking, is needed to achieve the same amount of spike advance. However, this results in an increase in the total input energy. For example, decreasing the pulse duration by a factor of 5 necessitates an increase of the amplitude by a factor of 5, which in turn, amounts to five times more energy according to the fact that energy is proportional to the square of the amplitude multiplied by the time duration: $E \propto u_{\text{pls}}^2 \Delta t$. We have tested and verified this in simulations as can be seen in figures 4(E) and (F).

2.9. Application of the control

A set of seven different target ISI values around the neuron's nominal ISI, $T_{\text{s}} = 100$ ms, were considered: $\text{ISI}_{\mathcal{T}} = [80, 85, 90, 95, 100, 105, 110]$ ms. For the simulations, the calculated minimum energy waveforms were applied to the full five dimensional GA model (1) in the following manner. At the onset of an action potential, a target ISI, $T_{\text{targ}} \in \text{ISI}_{\mathcal{T}}$ is randomly selected and its corresponding minimum energy waveform is applied to (1) as $I_c(t) = c\,u(t)$ for $t \in [0, T_{\text{end}}]$, where T_{end} is the minimum of T_{targ} and the time of next neuron spiking. In other words, the control input is reset to zero at the onset of the next action potential or when it has been applied fully for one cycle. Once reset to zero, the input will remain zero for five cycles of the neuron firing before it comes back on at the onset of the sixth action potential with another randomly selected T_{targ}. The reason for using five interleave cycles is to allow the neuron to return to its original periodic orbit for which the PRC was originally computed. Figure 2 shows an example of how the minimum energy waveforms are applied to the neuron. The surrogate waveforms, the pulse control inputs, and the variants of the minimum energy waveforms are applied in a similar manner. As shown in the figure, first a target ISI (of, in this case, $T_{\text{targ}} = 90$ ms) is set and its corresponding minimum energy input is applied. The measured ISI (for this case, $T_{\text{act}} = 87.2$ ms) is labeled below, indicating that the neuron spiked earlier than desired. Then, the input is set to zero for five cycles before turning back on for another randomly selected target ISI of, in this case, 105 ms, for which $T_{\text{act}} = 108.8$ ms. Note that for the first application of the input shown, the input is reset to zero at the onset of the next spike, whereas for the second application,

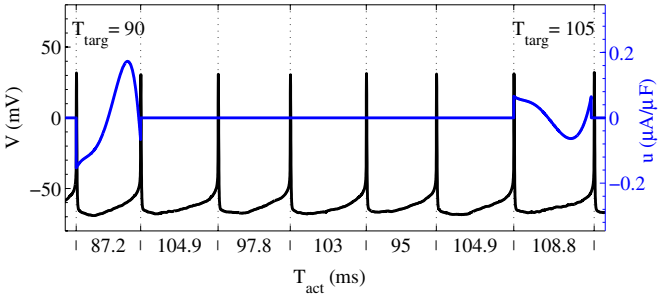

Figure 2. Method of stimulus administration for both the PRC measurements and the control. The figure shows an example for the minimum energy waveforms applied to the noisy GA model. The surrogate and the pulsatile inputs are applied similarly in the simulations. In the experiments, different number of interleave cycles were tested for the minimum energy inputs.

the input has been reset to zero when the full cycle of the input is applied. Also, it is seen that due to the presence of noise, during the middle resetting period the actual ISIs fluctuate around the nominal 100 ms spiking period. In the physiological experiments, the minimum energy waveforms were applied in the same way, except that different interleave cycles were tested. Since the minimum energy waveforms are continuous in time and have low amplitudes, the neurons remain close to the periodic orbit at all times, and so it is not necessary to have five interleave cycles between every two consecutive administrations of the control. This allows for more data recording without compromising the accuracy too much. We have tested zero, three, and five interleave cycles for the control studies in the experiments and have found that with small amplitudes, little control is lost when the interleave cycles are reduced to three.

2.10. Measure of efficiency

To show the performance of the minimum energy control protocol, a measure of the input energy is used that is of the form of the cost function written in (4). To find an average input energy for each of the target ISI values $T_{\text{targ},i}$, $(i = 1, \ldots, 7)$, the quantity

$$E_i = \bar{E}_{ij}, \quad \text{where} \quad E_{ij} = \int_0^{T_{\text{end},j}} u_i^2 \, \mathrm{d}t, \tag{12}$$

is calculated, where $\bar{E}_{ij}$ is the average of E_{ij} which is the input energy for u_i, the control corresponding to $T_{\text{targ},i}$, at its jth application.

To characterize the efficiency of the control, the following error measures are defined:

$$e_{\text{rms}} = \sqrt{\frac{\sum_{i=1}^{7} |T_{\text{targ},i} - \overline{T}_{\text{act},i}|^2}{7}}, \tag{13}$$

$$p = \frac{\sum_{i=1}^{7} \sigma_i}{7}, \tag{14}$$

where, $\bar{T}_{\text{act},i}$ and σ_i are the average actual ISI and its standard deviation achieved for the target $T_{\text{targ},i}$, respectively. The quantities e_{rms} and p give an overall performance measure for each recording across all target ISIs of interest. The lower the values of these parameters, the better the performance of the controller.

2.11. Electrophysiology experiments

Neurons in the CA1 region of hippocampus from Long-Evans rat brain slices were stimulated and recorded using whole cell patch-clamp recording techniques. The applied current for achieving periodicity in firing for the neurons was determined using a dynamic clamp and implemented using the RTXI platform. The implementation consists of three parts. In the first part, the neuron was controlled to spike periodically at an average period of 100 ms. This was done by implementing an auxiliary closed-loop PI controller that regulates the baseline current I_b for the neuron. It is important to have stable periodic spiking behavior in the neuron, as the underlying assumption in the formulation developed previously is that the system is on a stable limit cycle. The necessity for having a controller to maintain the periodicity of the neuron comes from the empirical fact that the dynamics of *in vitro* neurons change slowly over time when patched [21]. To compensate for this slow change, a PI controller that makes small perturbations to the DC I_b value at each spike time was used; this was previously designed and tested in [21]. The DC I_b value is typically around 100 pA cm^{-2}, and the amount of perturbation from the PI controller is less than 1 pA cm^{-2} per action potential. The PI controller runs throughout the experiment.

In the second part, the PRC of the neuron was obtained from the direct method. Once the neuron had stabilized around a baseline firing rate, it was stimulated with a 1 ms pulse stimulus every sixth cycle to obtain the PRC. The stimulation phase for the pulse was randomly chosen every time. Although pulses with amplitude 150 pA cm^{-2} were used for most of the experiments, the effect of other pulse amplitudes were also investigated in terms of the amount of induced nonlinearity that is characterized by C_{NL}. The spike advance measurements were stored and transferred to MATLAB, from which the PRC was estimated for each neuron under study.

In the third part, the PRC was used to find the minimum energy waveforms using MATLAB. The optimal input waveforms were then fed back into RTXI where, at the onset of a spike, a target ISI value from $\mathrm{ISI}_{\mathcal{T}}$ was randomly chosen and its corresponding optimal waveform was applied to the neuron for one cycle. The actual ISI of the neuron for that cycle was measured to evaluate the performance of the controller. Since a unit membrane capacitance, c, is assumed, and since $u = I_c/c$, the resulting optimal control inputs $u(t)$ can be viewed as electrical current stimuli. The method was applied to a total of 9 different neurons from four different rats, and a total of 15 recordings were made.

3. Results

3.1. Simulation results

Figure 3 shows the interaction of the minimum energy waveforms and the stabilizing PI controller. As it can be seen, the PI controller's input remains constant through every ISI and only changes at the spike times. For the computational simulations the change in the PI controller input current is recorded to be of order 10 pA cm^{-2} (as opposed to <1 pA cm^{-2} in the electrophysiology experiments) with most changes occurring almost always immediately after the application of the minimum energy waveform for $T_{\mathrm{targ}} = 80$ ms, aiming to restabilize the neuron to its nominal 100 ms firing rate.

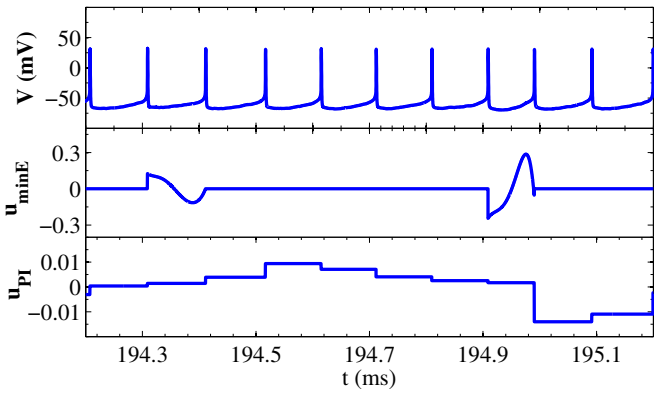

Figure 3. Stimulation of the noisy GA model. Membrane voltage changes (top), application of minimum energy waveforms (middle), and electrical current input change for the stabilizing PI controller per action potential (bottom). The inputs in (B) and (C) are in μA cm^{-2}.

Figure 4 shows results for the spike advance, PRC and its variations, minimum energy waveforms, their surrogates, their positive and negative variants, and amplitude-spike advance mappings for the pulse control methods for the noisy GA model. The minimum energy waveforms, u_{minE}, are obtained from the PRC shown with blue in figure 4(B) and solving equations (6)–(10) for a given target ISI. Similarly, the variants of the minimum energy waveforms, $u_{+\sigma}$ and $u_{-\sigma}$, are computed using equations (6)–(10) with respective variations of the PRC (magenta and cyan curves in figure 4(B)). The target ISI value for each waveform corresponds to the time at which it terminates.

Figure 5 shows results of applying the minimum energy waveforms, the surrogates, the minimum energy waveform variants, and the pulse control inputs to the noisy GA model. The pulse amplitudes used for this simulation are indicated by dashed lines in figures 4(G) and (H). Performance of control is measured by reporting the Pearson R_{cont} value for correlating target ISI against the measured actual ISI, e_{rms} and p values, and the mean and standard deviation statistics for each target ISI. The surrogate waveforms, which use the same amount of energy as the minimum energy waveforms, perform considerably worse, as expected. Also, the variants of the minimum energy waveforms perform less efficiently both in control and the amount of their average energy for achieved ISI. The efficiency of control for the minimum energy waveforms and the pulse control inputs are comparable. However, the minimum energy waveforms use considerably less energy, as shown in figure 5(G). We see that the minimum energy method can reduce the level of the input energy by at least an order of magnitude with respect to the pulsatile method.

3.2. Experimental results

In figure 6, the measured spike advances, the PRC, and the calculated minimum energy waveforms for an *in vitro* neuron are shown. The correlation factor for the spike advance fit is

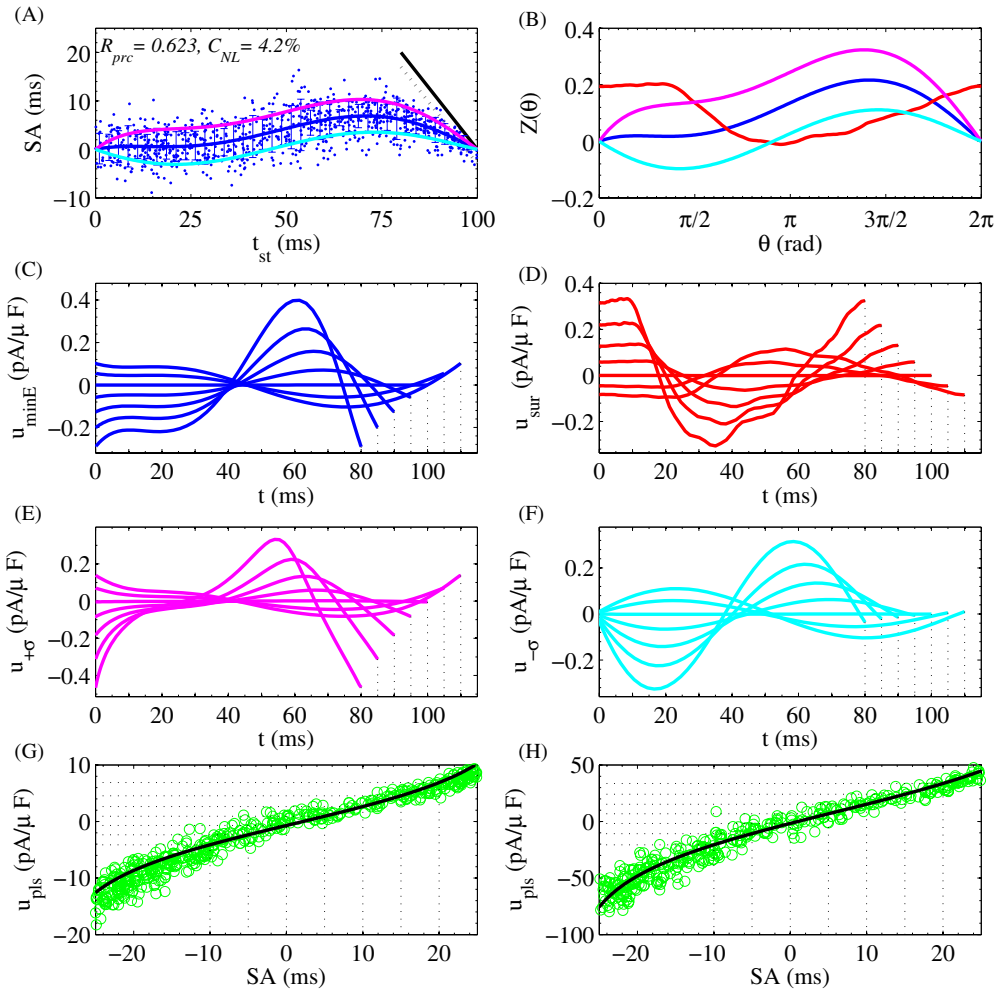

Figure 4. Results for the noisy GA model. (A) Spike advance data and fit for the PRC measurement using a 2 μA cm^{-2}, 1 ms pulse stimulating at randomly selected times every sixth cycle of the neuron's firing. Error bars indicate the standard deviation of the data at different stimulation times. The magenta and cyan curves are, respectively, sixth order fits to plus and minus one std points of the PRC data, constrained to zero at both ends. The diagonal black and dashed gray lines on the right indicate the line of causality and the boundary of its 3% band neighborhood used to calculate C_{NL}. (B) The PRC in terms of phase (blue), its surrogate (red), and its positive (magenta) and negative (cyan) stds. (C)–(F) Charge-balanced minimum energy stimuli computed from each PRC of corresponding color in (B). The target ISI value for each waveform corresponds to the time at which it terminates. (G) and (H) Pulse amplitude data and (inverse sigmoidal) fit for the pulse method with 1 ms and 0.2 ms pulses, respectively. Pulse amplitudes of target ISIs of interest are marked with dashed lines. From (A) and (B), it is seen that the PRC is maximum at 70% of the full cycle.

found to be $R_{\mathrm{prc}} = 0.350$. One can see the formation of the causality line in figure 6(A). For the example shown in this figure, $C_{\mathrm{NL}} = 11.0\%$. We recall that to ensure reliable PRCs, one must choose pulse amplitudes that are small enough to avoid large nonlinear responses, while at the same time being large enough that they can produce a notable variability in the spike time of the neuron. The PRC for this neuron is shown in figure 6(B).

When the minimum energy waveforms, shown in figure 6(C), are applied to the *in vitro* neuron, the results shown in figure 7 are obtained. We note that in the results presented in this figure, we have eliminated any T_{act} value less than 30 ms and greater than 200 ms, which amount to less than 1% of the data, as outliers. The first subplot in this figure summarizes the target ISI versus actual ISIs achieved in each case, while the others are the histograms for each of the target ISI values of interest. The statistics of the actual ISIs achieved by the neuron are reported in each case in the form of a mean and standard deviation: $T_{\mathrm{act}} = \mathrm{mean} \pm \mathrm{std}$. It is seen that for the case with $T_{\mathrm{targ}} = 100$ ms, the mean is 100.7 ms and the standard deviation is 8.6 ms. Considering the minimum energy waveform for this case, which is practically zero according to figure 6(C), we conclude that the standard deviation is a result of the intrinsic noise in the neuron. The standard deviation in all of the other cases is seen to be of the same order as in this case, which is indicative of consistency across the different cases. By looking at the mean values, however, one observes that in general, as the target ISI is moved away from the nominal (unstimulated) ISI of $T_{\mathrm{s}} = 100$ ms, the error in the mean values tends to increase. Arguably, one reason for this is that by targeting ISIs that are further away from T_{s}, the magnitude of the minimum energy waveforms increases, which in turn pushes the phase model of the neuron toward the edge of the range for which it is valid. A total of 15 recordings were made, as summarized in table 1.

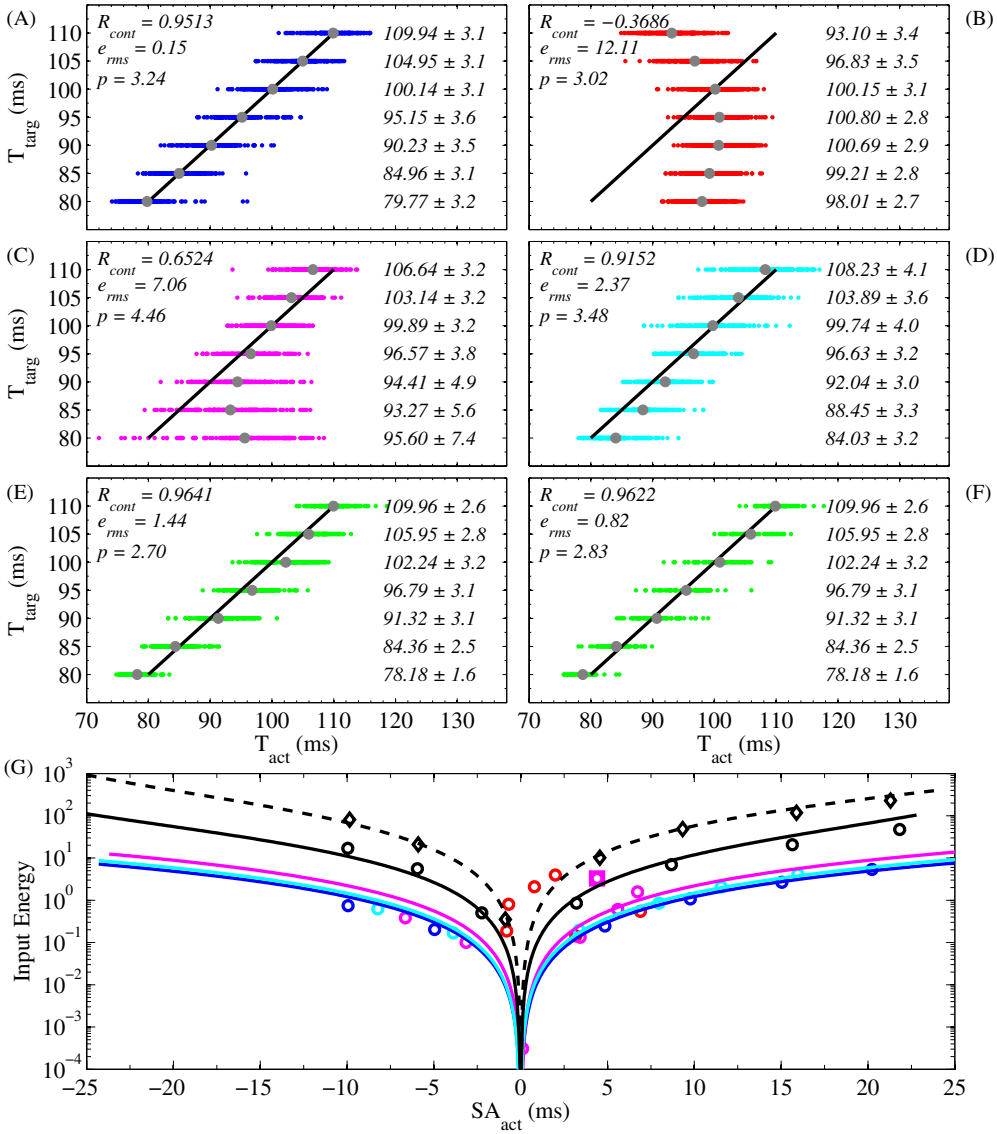

Figure 5. Results for the noisy GA model stimulated with (A) the minimum energy waveforms, (B) the surrogate waveforms, (C) and (D) the respective positive and negative variants of the minimum energy waveforms, (E) and (F) respective 1 ms and 0.2 ms single pulses of select amplitudes applied at 70% of the full cycle. The mean values of the actual ISIs are shown with gray markers along the straight unit slope black line. The statistics of control efficiency for each case are reported as mean ± std. (G) Input energy as a function of actual SA achieved for the minimum energy waveforms (blue), surrogates (red), minimum energy positive (magenta) and negative (cyan) variants, the pulse method with 1 ms pulses (black), and the pulse method with 0.2 ms pulses (dashed). The energy curve for the surrogate inputs are not shown, due to very low performance. Markers show the energy associated with the mean actual ISI values from (A) to (F). In fitting the magenta data, the point plotted with a square was considered an outlier and was excluded from the fit. The energy curves have been shifted horizontally by a bias value to yield a minimum at zero for ease of comparison. The biases were 0.7, 1.3, 0.7, −2.2, and −1.0 ms for the blue, magenta, cyan, solid black, and dashed curves, respectively.

4. Discussion

For a periodically firing neuron, the optimal control method presented here uses the neuron's phase model and PRC to compute minimum energy waveforms for specific target ISIs that are charge-balanced, continuous-time, low frequency and low amplitude. The minimum energy waveform for each T_{targ} is computed independently. It is seen that these waveforms resemble the shape of the PRC in their respective time spans. The inputs that are designed to decrease the ISI (i.e., those with $T_{\text{targ}} < 100$ ms) reach their maximum positive value at a time around 70% of their respective value of T_{targ}. This matches the location at which the PRC is maximally positive, hence increasing the dynamics of the neuron due to (1). Charge-balance is achieved by these waveforms assuming negative values in the first part of their time span where the PRC is close to zero and hence least effective. A similar argument could also be made for those waveforms that are designed to increase the ISI (i.e., those with $T_{\text{targ}} > 100$ ms). They start positive at first, but take negative values as the time progresses and the PRC becomes increasingly positive. It is also noted that the accuracy of the controller could be slightly affected by the target ISI. This is because the phase model is only valid for small inputs. Therefore, by demanding larger changes in the

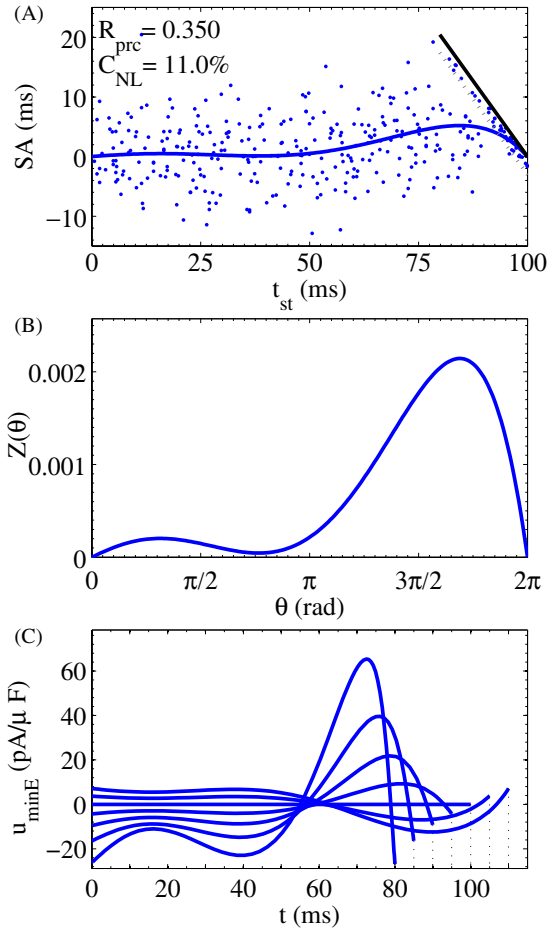

Figure 6. (A) Spike advance, (B) PRC and (C) minimum energy control inputs for an *in vitro* neuron. For the PRC measurements, the neuron was stimulated with a 150 pA cm^{-2}, 1 ms pulse every sixth cycle.

Table 1. Results for a total of 15 different recordings from 9 different *in vitro* neurons. PlsAmp gives the pulse amplitude in pA cm^{-2} that was used to obtain the PRC. ILC indicates the number of interleave cycles that was allowed between every two consecutive control applications.

	Cell/Rec.	PlsAmp	R_{prc}	$\%C_{NL}$	R_{cont}	e_{rms}	p	ILC
1	1/1	200	0.559	21.1	0.160	9.9	7.3	0
2	1/2	100	0.352	13.2	0.594	4.9	9.1	0
3	1/3	150	0.325	17.8	0.717	2.4	11.3	0
4	2/1	200	0.573	17.4	0.697	5.0	6.5	0
5	2/2	100	0.312	7.6	0.666	3.4	8.9	0
6	2/3	150	0.404	15.5	0.650	8.0	12.0	0
7	3/1	200	0.467	22.3	0.478	3.6	12.0	0
8	3/2	100	0.330	13.4	0.644	1.9	10.5	0
9	4/1	150	0.412	7.8	0.559	5.7	7.3	5
10	5/1	150	0.645	37.8	0.423	4.0	14.9	5
11	5/2	100	0.419	25.6	0.214	8.3	16.9	3
12	6/1	150	0.350	11.0	0.820	2.0	8.0	3
13	7/1	150	0.368	24.1	0.404	6.2	12.7	0
14	8/1	150	0.373	7.8	0.773	2.5	6.7	3
15	9/1	150	0.418	11.1	0.599	3.4	10.8	3

neuron's ISI, thus requiring larger stimuli, the phase model is pushed beyond the local region for which it is valid, resulting in less accurate results.

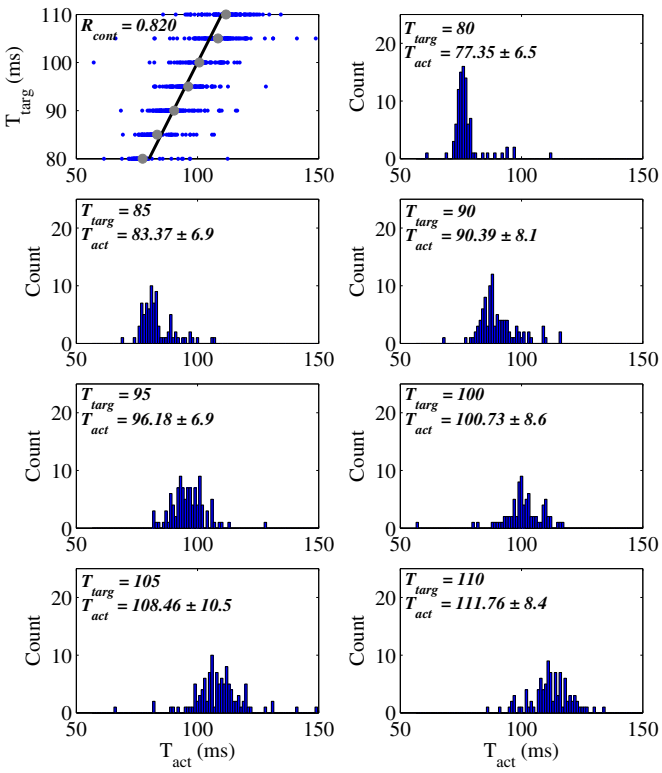

Figure 7. Results for an *in vitro* neuron. The upper left subplot summarizes the control outcome, where the data points are the actual ISI values T_{act}, achieved for each of the target ISI values T_{targ}. The mean values of the actual ISIs are shown with gray markers around the straight unit slope black line. Other subplots show histograms and statistical results for control to each of the target ISIs. The actual ISIs achieved for each case are reported in the form: T_{act} = mean ± std.

In view of figure 5(G), it should be noted that although the surrogate inputs are designed to have the same energy content as the minimum energy waveforms, i.e., same value of C in (4), when applied to the model, their average input energy measure for the actual ISI that they achieve is different from that of the minimum energy waveforms, i.e., different values for E_i in (12). A similar argument can also be made for the variants of the minimum energy waveforms. If one compares the energy content of the variants to that of the actual minimum energy waveforms using (4), one finds that the variant waveforms' energy content is slightly less than that of the minimum energy ones. However, when applied to the noisy model, their average input energy measure for a given actual ISI is slightly higher. This, together with their lower performance in control, illustrates the robustness of the designed minimum energy waveforms. The performance of the variants of the minimum energy waveforms are also slightly lower than that of the pulsatile methods. However, considering the amount of difference in their respective energy levels, one may favor using the optimal control methodology, even with rough PRC estimates, over conventional pulsatile methods in regulating the ISI for a neuron.

Results in table 1 suggest that in finding the PRC, there is a direct relationship between the pulse amplitude and the nonlinearity coefficient C_{NL}: the higher the pulse amplitude,

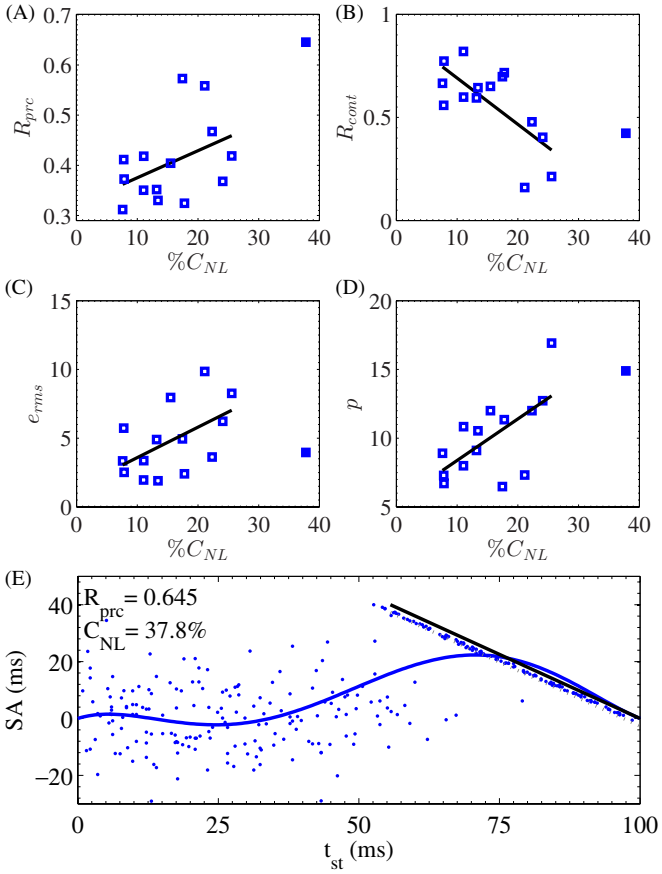

Figure 8. Analysis of the experimental results for *in vitro* neurons. (A) and (B): Respective correlation coefficients for the PRC and the control versus the nonlinearity coefficient. (C) Root mean square error and (D) average stds of the actual ISIs for the data in each recording versus C_{NL}. The fitted lines are least squares fits excluding the solid marker as an outlier. It is seen that better results are achieved when C_{NL} values are less than 20%. (E) The PRC of an overstimulated neuron corresponding to the solid marker in A–D. In this experiment, the PRC measurement has been done with pulse amplitudes of 200 pA cm^{-2}.

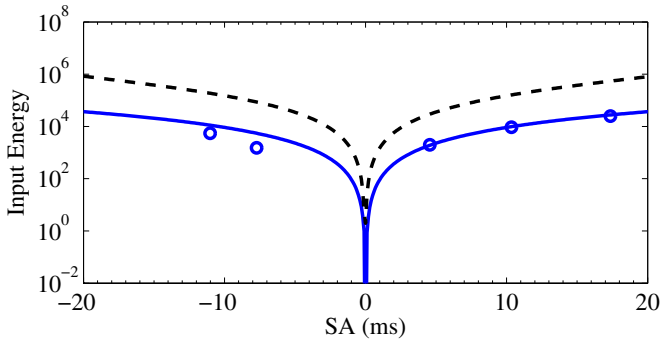

Figure 9. Input energy comparison for the method presented in this paper (markers with solid fit) with that of [24] (dashed). The *x*-axis presents the spike time advance (T_s–T_{st} in figure 1). The data points for the recording of figure 7, the corresponding least squares quadratic fit, and the energy curve form [24] have been shifted horizontally to yield a minimum at zero for ease of comparison. We see that the method presented in this paper can reduce the input energy by at least an order of magnitude.

the higher the value of C_{NL}. The reason for this is that larger amplitude pulses are more likely to induce immediate spikes in the last 20%−30% of the phase resulting in higher C_{NL} values.

It is also seen from table 1 that the PRC correlation coefficient (R_{prc}), the control correlation coefficient (R_{cont}), the root mean square error between the target ISIs and actual ISIs (e_{rms}), and the average std (p) for each neuron show a direct relationship with C_{NL}, as seen in figure 8. When C_{NL} is high, i.e., when the pulse amplitude is high (e.g., 200 pA cm^{-2}), R_{prc} values are also higher (see figure 8(A)). This is because the PRC fit function is a least squares fit to the data and, by design, approaches zero at 2π. This produces a higher R_{prc} value as the fit becomes close to the dense causality line resulting from overstimulation. A PRC that is affected by nonlinearities is not reliable and oftentimes results in inefficient control characterized by lower R_{cont}, larger e_{rms}, and larger p values. This is better illustrated in figures 8(B), (C) and (D). Higher R_{cont} values are obtained for C_{NL} values less than 20%. These values of C_{NL} correspond to R_{prc} values that are mostly less than 0.5. Also, the e_{rms} and p are smaller when C_{NL} is smaller. The data point with $C_{NL} = 37.8\%$ in this figure (shown with solid marker) indicates overstimulation in measuring the PRC, shown in figure 8(E). Although this recording gives a high value of $R_{prc} = 0.645$, the resulting $R_{cont} = 0.423$ value is not very high, proving the unreliability of the PRC. We did not observe a notable effect from changing the number of interleave cycles in the last column of the table. The results shown in figures 6 and 7 are for the 12th row in the table.

To evaluate the performance of the minimum energy control method on *in vitro* neurons, we compare the input energy obtained for this method with that of the pulse method presented in [24]. The pulse method in [24] was implemented on the same type of *in vitro* neurons, prepared in exactly the same manner. The result of this comparison is shown in figure 9. We note that implementing both methods on the same biological neuron may not give accurate results as the dynamics of these cells change slowly over time when patched. Therefore, for a direct comparison of energy and control, these experiments were done *in-silico* where experimental conditions will not confound the findings.

Even though the timing of the pulse control method is carefully chosen to yield minimum pulse amplitudes, we find that the (continuous-time) minimum energy control method presented here achieves the same task with an order of magnitude less energy. This is a significant achievement, especially since the designed minimum energy waveforms, in contrast to the pulse control inputs, are charge-balanced. This matches the simulation results in figure 5(D). It is important to point out that for the CA1 hippocampal neurons that mostly have Type I PRCs, the minimum energy protocol uses significantly less energy if the charge-balance constraint is lifted [1].

It is worth mentioning that if one can obtain an estimate of the oscillatory dynamics model for *in vitro* cells, for example if one uses estimation techniques such as Kalman filtering to characterize the ion channels for the cell, then one can numerically solve the appropriate adjoint equation to find the PRC for the estimated model [32, 33]. Of course, in doing this,

one has to tolerate the potential errors in the estimation of the model.

In our approach to stimulus waveform design, maximum stimulus amplitude was taken into account as a safety constraint [37]. Other safety factors could also be added. One advantage of a smooth and low amplitude waveform is that the charge densities are much lower and control can be achieved much easier within safety bounds. However, we did not take into account physical properties of the electrode and the electrode–tissue interface, which should be taken into consideration for an implanted device [38]. One also needs to account for the location of stimulation and potential correlations between the frequency and the amplitude of the input waveform to ensure an effective input (see [7]). Adding more realism to the consideration of the stimulus waveform is a direction in which we would like to continue this research.

Acknowledgments

The authors would like to thank Oscar Miranda for his help and consultation with the experimental setup. AN and JM were supported by the National Science Foundation grant NSF-1000678. TN and TS were supported by the National Science Foundation CAREER award 0954797.

References

[1] Nabi A and Moehlis J 2009 Charge-balanced optimal inputs for phase models of spiking neurons *ASME 2009 Dynamic Systems and Control Conf. (Hollywood, CA, 12–14 Oct.)*

[2] Moehlis J, Shea-Brown E and Rabitz H 2006 Optimal inputs for phase models of spiking neurons *ASME J. Comput. Nonlinear Dyn.* **1** 358–67

[3] Wagenaar D, Madhavan R, Pine J and Potter S 2005 Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation *J. Neurosci.* **25** 680–8

[4] Popovych O V, Hauptmann C and Tass P A 2006 Control of neuronal synchrony by nonlinear delayed feedback *Biol. Cybern.* **95** 69–85

[5] Feng X J, Shea-Brown E, Greenwald B, Kosut R and Rabitz H 2007 Optimal deep brain stimulation of the subthalamic nucleus—a computational study *J. Comput. Neurosci.* **23** 265–82

[6] Feng X J, Greenwald B, Rabitz H, Shea-Brown E and Kosut R 2007 Toward closed-loop optimization of deep brain stimulation for Parkinson's disease: concepts and lessons from a computational model *J. Neural Eng.* **4** L14–21

[7] Cooper S E, Kuncel A M, Wolgamuth B R, Rezai A and Grill W M 2008 A model predicting optimal parameters for deep brain stimulation in essential tremor *J. Clin. Neurophysiol.* **25** 265–73

[8] Danzl P, Hespanha J and Moehlis J 2009 Event-based minimum-time control of oscillatory neuron models *Biol. Cybern.* **101** 387–99

[9] Schöll E, Hiller G, Hövel P and Dahlem M A 2009 Time-delayed feedback in neurosystems *Phil. Trans. R. Soc.* A **367** 1079–96

[10] Nabi A and Moehlis J 2010 Nonlinear hybrid control of phase models for coupled oscillators *ACC'10: 2010 American Control Conf. (Baltimore, MD, 30 June–2 July)* pp 922–3

[11] Foutz T J and McIntyre C C 2010 Evaluation of novel stimulus waveforms for deep brain stimulation *J. Neural Eng.* **7** 066008

[12] Rosin B, Slovik M, Mitelman R, Rivlin-Etzion M, Haber S, Israel Z, Vaadia E and Bergman H 2011 Closed-loop deep brain stimulation is superior in ameliorating Parkinsonism *Neuron* **72** 370–84

[13] Wilson C J, Beverlin II B and Netoff T 2011 Chaotic desynchronization as the therapeutic mechanism of deep brain stimulation *Front. Syst. Neurosci.* **5** 50

[14] Nabi A and Moehlis J 2011 Single input optimal control for globally coupled neuron networks *J. Neural Eng.* **8** 065008

[15] Berényi A, Belluscio M, Mao D and Buzsáki G 2012 Closed-loop control of epilepsy by transcranial electrical stimulation *Science* **337** 735–7

[16] Nabi A, Mirzadeh M, Gibou F and Moehlis J 2013 Minimum energy desynchronizing control for coupled neurons *J. Comput. Neurosci.* **34** 259–71

[17] Krook-Magnuson E, Armstrong C, Oijala M and Soltesz I 2013 On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy *Nature Commun.* **4** 1376

[18] Schiff S J and Sauer T 2008 Kalman filter control of a model of spatiotemporal cortical dynamics *J. Neural Eng.* **5** 1–8

[19] Schiff S J 2010 Towards model-based control of Parkinson's disease *Phil. Trans. R. Soc.* A **368** 2269–308

[20] Danzl P, Nabi A and Moehlis J 2010 Charge-balanced spike timing control for phase models of spiking neurons *Discrete Contin. Dyn. Syst. Ser.* A **28** 1413–35

[21] Miranda-Domnguez O, Gonia J and Netoff T 2010 Firing rate control of a neuron using a linear proportional-integral controller *J. Neural Eng.* **7** 066004

[22] Nabi A and Moehlis J 2012 Time optimal control of spiking neurons *J. Math. Biol.* **64** 981–1004

[23] Dasanayake I and Li J-S 2011 Optimal design of minimum-power stimuli for phase models of neuron oscillators *Phys. Rev.* E **83** 061916

[24] Stigen T, Danzl P, Moehlis J and Netoff T 2011 Controlling spike timing and synchrony in oscillatory neurons *J. Neurophysiol.* **105** 2074–82

[25] Nabi A, Mirzadeh M, Gibou F and Moehlis J 2012 Minimum energy spike randomization for neurons *ACC'12: 2012 American Control Conf. (Montreal, Canada)* pp 4751–6

[26] Golomb D and Amitai Y 1997 Propagating neuronal discharges in neocortical slices: computational and experimental study *J. Neurophysiol.* **78** 1199–211

[27] Hodgkin A L and Huxley A F 1952 A quantitative description of membrane current and its application to conduction and excitation in nerve *J. Physiol.* **117** 500–44

[28] Dorval A D, Christini D J and White J A 2001 Real-time Linux dynamic clamp: a fast and flexible way to construct virtual ion channels in living cells *Ann. Biomed. Eng.* **29** 897–907

[29] Glass L and Mackey M C 1988 *From Clocks to Chaos: The Rhythms of Life* (Princeton, NJ: Princeton University Press)

[30] Netoff T, Schwemmer M and Lewis T 2012 Experimentally estimating phase response curves of neurons: theoretical and practical issues *Phase Response Curves in Neuroscience (Springer Series in Computational Neuroscience* vol 6*)* ed N Schultheiss, A Prinz and R Butera (New York: Springer) pp 95–129

[31] Netoff T I, Banks M I, Dorval A D, Acker C D, Haas J S, Kopell N and White J A 2005 Synchronization in hybrid neuronal networks of the hippocampal formation *J. Neurophysiol.* **93** 1197–208

[32] Ermentrout G B 1996 Type I membranes, phase resetting curves and synchrony *Neural Comput.* **8** 979–1001

[33] Brown E, Moehlis J and Holmes P 2004 On the phase reduction and response dynamics of neural oscillator populations *Neural Comput.* **16** 673–715

[34] Kirk D E 1970 *Optimal Control Theory: An Introduction* (New York: Dover)
[35] Lenhart S and Workman J T 2007 *Optimal Control Applied to Biological Models* (London: Chapman and Hall)
[36] Theiler J, Eubank S, Longtin A, Galdraikian B and Farmer J D 1992 Testing for nonlinearity in time series: the method of surrogate data *Physica* D **58** 77–94
[37] Agnew W and McCreery D 1990 Considerations for safety with chronically implanted nerve electrodes *Epilepsia* **31** S27–32
[38] Merrill D R, Bikson M and Jefferys J G R 2005 Electrical stimulation of excitable tissue: design of efficacious and safe protocols *J. Neurosci. Methods* **141** 171–98