

Computationally Efficient Cardiac Bioelectricity Models Toward Whole-Heart Simulation

Nathan A. Wedge, Michael S. Branicky, and M. Cenk Çavuşoğlu

Department of Electrical Engineering and Computer Science, Case Western Reserve University, Cleveland, Ohio 44106
{nathan.wedge, mb, cavusoglu}@case.edu

Abstract—This paper studies the characteristics of excitable cell mathematical models, with the goal of developing new insights and techniques in simulating the electrical behavior of the human heart. While very simple models of such behavior can be simulated at real-time or better speeds on powerful computing equipment, the use of realistic cell models or organ-magnitude cell networks make the simulations computationally infeasible. We present an examination of the FitzHugh-Nagumo model and its response to stimulus and, in order to move toward the goal of a full cardiac simulation, we present a method of optimizing single-cell calculations through local interpolation techniques. Additionally, we introduce a separate method of optimizing multi-cell simulations by tracking cellular activations.

Keywords—cardiac bioelectricity modeling, FitzHugh-Nagumo, simulation, optimization.

I. INTRODUCTION AND OVERVIEW

A computationally feasible whole-heart model could be invaluable in the study of human heart pathology and the development of drugs for the treatment of various disorders. The potential value of such a complete and realistic cardiac model is undeniable. Such a model could dramatically expand our presently limited understanding of cardiovascular disease and abnormality while providing a convenient, noninvasive, and inexpensive method of proposing and testing revolutionary drug therapies and other treatment interventions. Present progress (detailing approaches from the electrical, mechanical, and fluidic standpoints) toward this goal is summarized in [1].

Development of techniques for large-scale modeling of systems is a common theme across many fields of research today. Cardiac modeling presents complications not found in many modeling tasks, since excitable cell models are nonlinear in nature, and many ordinary and valuable analysis techniques for differential equations do not apply to nonlinear systems. The simplest cardiac cell models have two variables while realistic models can have dozens, and the total number of cells in the human heart is on the order of 10^{10} (in the tens of billions). Therefore, extraneous details must be abstracted away so that only relevant information is actively simulated.

This paper explores the structure of the FitzHugh-Nagumo excitable cell model in order to characterize which details can be discarded through abstraction and which are important to the model's overall dynamics. Based upon the concept of electrical wave propagation in the human heart, we examine the model and generalize its response to stimuli. Using observations from this analysis, we pose a method of optimizing calculations in the single-cell model with local interpolation

techniques and another method of speeding simulations in multi-cell networks by tracking a list of “active” cells.

II. BACKGROUND

A. FitzHugh-Nagumo Model

Pioneered in the mid-1950s by FitzHugh [2] and based upon research of the giant squid axon by Hodgkin and Huxley, the FitzHugh-Nagumo model is one of the simpler excitable cell models available. It is a reduction of the four-variable Hodgkin-Huxley model, based upon the similarity in speed of the dynamics of the individual variables. Though primarily considered to be qualitative in nature, it is representative of the excitable phenomena of cardiac cells and has been used as a base model for testing of novel computational methods [3]. The FitzHugh-Nagumo model equations relate the cell membrane electrical potential (V), and the gating variable (also called the cell's refractory potential, W):

$$\frac{\partial V}{\partial t} = \frac{1}{\epsilon} \left(V - \frac{V^3}{3} - W \right), \quad (1)$$

$$\frac{\partial W}{\partial t} = \epsilon (V - \gamma W + \beta). \quad (2)$$

The parameters are traditionally interpreted to have positive values. Unless otherwise noted, we applied the following parameter values: $\beta = 0.7$, $\epsilon = 0.2$, and $\gamma = 0.8$.

B. Wave Propagation

The important feature of cardiac cellular interaction is the propagation of potential waves through interconnected cells in a complex network. Single cells are surrounded with an insulating membrane (supporting a potential, V) containing selectively permeable ionic channels. The currents through these channels interact with the membrane potential to regulate the activity of the cell. Among other factors, the flow of various ions (sodium, potassium, calcium, etc.) throughout the cardiac tissue is responsible for the propagation of electrical waves through heart tissue. This, in turn, provides the driving force behind the heart's mechanical contraction and its ability to pump blood through the body. To model the diffusion that allows wave propagation to occur in the system, we must alter Equation (1) to include the appropriate term:

$$\frac{\partial V}{\partial t} = \frac{\partial^2 V}{\partial x^2} + \frac{1}{\epsilon} \left(V - \frac{V^3}{3} - W \right). \quad (3)$$

Thus, a sufficient difference in potential between adjacent cells can produce an action potential that flows across a

system. In multiple spatial dimensions, the $\partial^2 V/\partial x^2$ diffusion term becomes $\nabla \cdot (D \cdot \nabla V)$, in which D is the diffusion tensor.

C. Other Models

Past research has developed numerous excitable cell models of varying levels of complexity and with differing specialization. For example, C. Luo and Y. Rudy [4] pioneered a realistic model of ion channel behavior that is composed of over sixty equations. This model is frequently applied in computational simulations that investigate the characteristics of real cardiac cells [5], [6]. The Hodgkin-Huxley model [7] was developed to simulate the activity of a nerve axon (specifically, of the giant squid), but it shares the essential characteristic of excitability with cardiac tissue. Other models developed between the Hodgkin-Huxley (1952) and the Luo-Rudy (1994) model made important contributions to the accuracy of cardiac modeling. For a detailed overview of the evolution of such cardiac cell models, refer to [8].

III. SINGLE-CELL MODEL ANALYSIS

In the process of developing a heart simulation, the preliminary step involves understanding the evolution of the electrical potentials. Therefore, we pay particular attention to the unique characteristics of the cell model.

A. Phase Plane Analysis

Perhaps the most convenient consequence of the two-variable nature of the FitzHugh-Nagumo model is that it can be graphically represented in the phase (V, W) plane. Plotting solutions of the system in the phase plane reveals a well-defined cycle that is robustly attractive across a wide range of initial conditions. The normal course of this cycle, as shown in Fig. 1, represents a single excitation of the modeled cell, and has four distinct states [2]: regenerative, active, absolutely refractory, and relatively refractory. An external electrical stimulus of sufficient magnitude will initiate an excitation cycle, in which the cell progresses through these four states. In the regenerative phase, the cell begins a buildup of potential across its membrane. If the initiating stimulus is of insufficient magnitude, the cell will simply relax back to its equilibrium, rather than experiencing a full excitation cycle. The active phase represents the period in which the cell's membrane potential is temporarily suspended near a peak value. Entering the absolutely refractory phase, the cell's membrane potential drops rapidly toward its equilibrium value. At this point, as shown by the direction field, the gating variable is at a maximum, and the cell is quite resistant to further excitation by any external stimulus. In the final phase, relatively refractory, it is again possible to stimulate the cell into a renewed excitation cycle, though it will display some resistance to such stimulation until it again reaches the equilibrium point.

B. Diffusion Fourier Analysis

Using Fourier analysis, we can interpret the system's frequency response to various diffusive inputs, in a manner

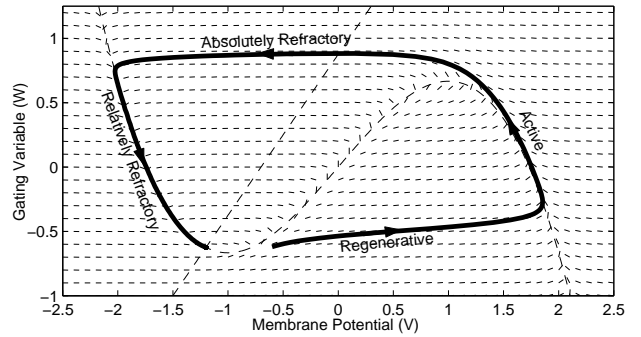


Fig. 1. Phase Plane with Nullclines and Sample Solution Trajectory




Fig. 2. Fourier Plots (Sinusoidal Inputs, $Amplitude = 0.5, 1.0, 1.5$)

analogous to forcing in a linear ordinary differential equation. Since the favorable notion of linearity is not valid in the excitable cell model, the response is not clearly related to either the frequency or the amplitude of the input. Still, we can make a number of useful observations based upon these Fourier plots. We give three such plots that relate input frequency to output frequency components in Fig. 2.

The most striking contrast between the Fourier plots is in the apparent linearity in frequency of the 0.5 amplitude plot. This is illustrative of an important feature of the system: resistance to small inputs. For any signal of small input amplitude (i.e. 0.5 or less), there is no prevailing excitation triggered, resulting in little response. Instead, the system responds with a potential signal of equal frequency and small amplitude. As a consequence, small diffusive signals may be safely passed over as they are incapable of triggering excitations in the system.

Turning our attention to the other plots, we note that for a given amplitude, there is a cutoff input frequency past which the system displays a qualitatively different response. When driven past this frequency, the system undergoes a single excitation cycle in normal fashion. However, following this, the input holds the gating variable at a near-constant value above equilibrium, preventing any further excitations. A second, higher cutoff frequency also exists, above which the single excitation cycle is also inhibited, resulting in an essentially null response. In the lower frequency regions of the second and third plot (where $Amplitude = 1.0, 1.5$), the system exhibits a periodic excitation caused by the sinusoidal input.

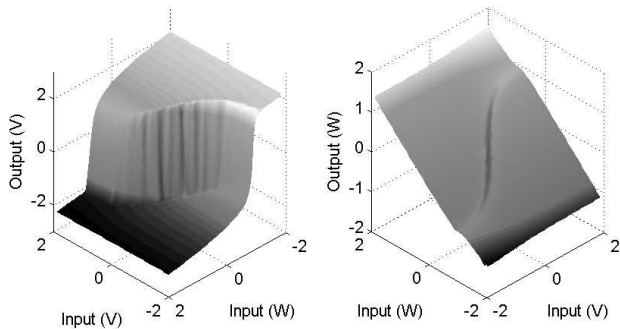


Fig. 3. FitzHugh-Nagumo State Transition Map ($h = 1$)

IV. OPTIMIZING SINGLE-CELL CALCULATIONS

In the autonomous single cell FitzHugh-Nagumo model, variable values at particular times determine the state of the system at all future times. Therefore, we can consider this relationship as a nonlinear mapping between present values and future values, given a timestep. This input/output mapping relationship can be captured graphically (as in Fig. 3) in order to provide a different perspective on the system's time behavior. In particular, these surfaces illustrate the nonlinearities of this cell model. Regions of sharp change (notably, on the diagonal of the input plane) signify value pairs that can evolve distinctly differently if subjected to small perturbations. In contrast, there are also regions that display nearly constant or nearly linear output, in which one variable strongly determines a future output value, and the value of the other variable is only weakly significant. Additionally, the smooth and primarily continuous shapes of the surfaces suggest that the model could be fitted by an interpolated function of the variables.

We can take advantage of this mapping relationship to perform high-accuracy calculations of the system's behavior off-line. During a simulation, these pre-created samples can be interpolated by algorithms such as the nearest neighbor method and locally weighted regression [9]. The nearest neighbor method is a straightforward algorithm that predicts an output by proposing that it is equal to the output of the single nearest sample. Locally weighted regression, by contrast, estimates the output by a weighted average of the outputs of several nearby samples. The sample weights are calculated based on an exponential metric where nearby samples are weighted highest. Hence, both methods can be used to reconstruct a mapping from stored sample values. A sample reconstruction of the V output surface from Fig. 3 is shown in Fig. 4.

Overall, these algorithms can recreate unknown output data well. However, locally weighted regression suffers specifically from large errors near discontinuities. This fallacy is correctable by concentrating samples within problem areas. A primary issue with the nearest neighbor method is its inability to generate smooth surfaces, which results in discontinuous changes in error around the input space. This hinders its accuracy significantly for small sample sizes. Both algorithms have disadvantages, but are highly optimizable through the

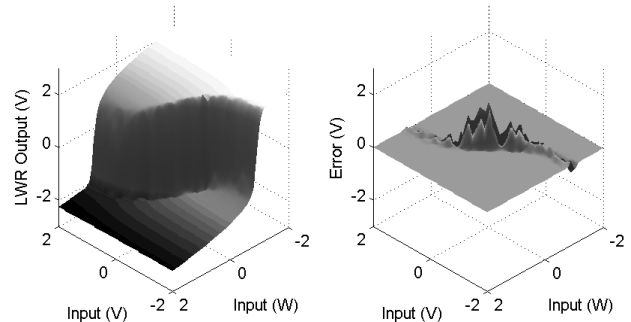


Fig. 4. LWR Reconstruction (8192 Samples) of V Surface and Error Plot

use of clever data structures. Both algorithms are defined for and extendable to higher-dimensional problems, an essential feature for integration into a whole-heart simulation, in which the problem's dimensionality could easily increase tenfold with the addition of a realistic model and spatial dimensions.

V. OPTIMIZING MULTI-CELL SIMULATIONS

Since the FitzHugh-Nagumo system has a stable, attractive equilibrium point, many cells in a simulation will be implicitly inactive, and hence it is wasteful to calculate their unchanging membrane potentials and small propagation currents. It is possible to use a list data structure to track the activation of individual cells in order to avoid these calculations. As a consequence of the mechanism of wave propagation, only direct neighbors of activated cells can be activated in the future. Therefore, two lists that track currently excited cells and the neighbors of these cells will completely describe the subset of a cell network that may be active in the successive timestep. All other cells are assumed to be in equilibrium, so no calculations are performed to update their state. Provided that the computational overhead of list tracking is negligible, this measure can only result in lower computation times. The algorithm for tracking cell activity is as follows:

```

Activity_Track(ActiveList)
  FOR EACH  $cell_i \in ActiveList$ 
    Model.simulate( $cell_i$ );
    FOR EACH  $cell_j \in cell_i.Neighbors$ 
      IF  $cell_j \notin NeighborList \ \& \ cell_j \notin ActiveList$ 
        NeighborList.add.cell( $cell_j$ );
  FOR EACH  $cell_i \in ActiveList$ 
    FOR EACH  $cell_j \in cell_i.Neighbors$ 
      Model.calculate.diffusion( $cell_i, cell_j$ );
  FOR EACH  $cell_i \in NeighborList$ 
    FOR EACH  $cell_j \in cell_i.Neighbors$ 
      IF  $cell_j \in ActiveList$ 
        Model.calculate.diffusion( $cell_i, cell_j$ );
  FOR EACH  $cell_i \in ActiveList$ 
    IF  $cell_i.equilibrium.state() = TRUE$ 
      ActiveList.remove.cell( $cell_i$ );
  FOR EACH  $cell_i \in NeighborList$ 
    IF  $cell_i.equilibrium.state() = FALSE$ 
      ActiveList.add.cell( $cell_i$ );
      NeighborList.remove.cell( $cell_i$ );

```




Fig. 5. Benchmark Wave Types

TABLE I
Activity List Optimization Speedups

	Linear	Grid(1)	Grid(2)	Spiral
Full-Blown Simulation	2.94 s	11.89 s	11.87 s	8.75 s
Activity List Method	0.32 s	4.84 s	6.67 s	9.48 s
Change	89.2%	53.9%	43.8%	-8.3%

In order to test the viability of this algorithm, we created four benchmarking simulation situations (shown in Fig. 5) that are representative of the types of waves that can be produced in a 2D simulation. They are: (1) 120 time units of a 200-by-1 cell line, stimulated by the first cell; (2) 30 time units of a 100-by-100 cell network, stimulated by the center cell; (3) 30 time units of a 100-by-100 cell network, stimulated at two cells on the diagonal; and (4) 60 time units of a 60-by-60 cell network, stimulated to produce a spiral wave pairing. The averaged results of three trials for each situation (on a 2.4 GHz P4 running Matlab 6.5.1 and a fourth order Runge-Kutta solver) are presented in Table I.

From the time data of simulations run with and without the optimization, it is apparent that the improvement in computation time is dependent on the relative level of activity in the system. This conclusion agrees with the basic idea of the optimization: when very few cells are active, there is relatively little to calculate, but when a large portion of cells are active, there is no room to save on calculations. This is apparent from the spiral wave case; most cells are kept active throughout the entire simulation due to the self-reinforcing nature of the wave, so very few cells are skipped by the algorithm. This creates a negligible increase in computation time corresponding to the overhead associated with initializing and maintaining the lists. In the first three cases, the activity lists ensure that cells reaching equilibrium after the passage of a wave are again removed from the set of updated cells. This reduces the computation time roughly proportional to the spatial size of the propagating wave.

VI. DISCUSSION AND CONCLUSIONS

The computation times for full-blown simulation of our benchmark problems are predictable based on the number of cells involved in the simulation and the timescale. Assuming the overhead of the simulation is small (i.e. the simulation is reasonably large), these benchmarks require roughly $1.5 \mu\text{s}/\text{cell} \cdot \text{timestep}$. It is interesting to note that a simple extension of this idea leads to a predicted requirement of roughly four hours to simulate one second of the FitzHugh-Nagumo dynamics at the magnitude scale of the human heart

using the full-blown approach. Further accounting for the added complexity of a realistic cell model quickly makes the simulation a dauntingly unrealistic task. To make these high-level simulations feasible for use in heart pathology and drug research, it would be necessary to improve the computational cost by several orders of magnitude. Clearly, our optimization technique is only a precursory step toward such a goal.

In this paper, we outlined some of the basic characteristics of the FitzHugh-Nagumo excitable cell model and extended the analysis to the periodicity conditions and properties of the system. By examining the time evolution of variables, we introduced the idea of using nearest neighbor and locally weighted regression methods to capture the input/output relationship of the model. Additionally, we showed that an activity-list-based optimization technique demonstrates marked benefits in speeding calculations for simulations that involve some degree of inactivity (i.e. cells at equilibrium). The extendibility of this technique to realistic, quantitative models has yet to be verified, but we expect that it will have greater benefits as the overhead of list tracking becomes insignificant compared to the problem magnitude. Furthermore, we hope to find novel ways to apply locally weighted regression to a simplification of the multi-cell simulation environment. Considering the computational complexity of problems in this field and the present level of development, there is wide latitude for innovation toward full-scale simulations.

ACKNOWLEDGMENT

We would like to thank Dr. Niels Otani for his invaluable guidance and motivation. This research was supported in part by the National Science Foundation under grants 0208919 (Branicky) and CISE IIS-0222743 (Çavuşoğlu).

REFERENCES

- [1] P. Hunter, A. Pullan, and B. Smaill, "Modeling total heart function," *Annual Review of Biomedical Engineering*, vol. 5, pp. 147–177, 2003.
- [2] R. FitzHugh, "Impulses and physiological states in theoretical models of nerve membrane," *Biophysical Journal*, vol. 1, pp. 445–466, 1961.
- [3] J. Rogers and A. McCulloch, "A collocation–Galerkin finite element model of cardiac action potential propagation," *IEEE Transactions on Biomedical Engineering*, vol. 41, pp. 743–757, 1994.
- [4] C. Luo and Y. Rudy, "A dynamic model of the cardiac ventricular action potential: Simulations of ionic currents and concentration changes," *Circulation Research*, vol. 74, pp. 1071–1096, 1994.
- [5] Y. Wang, R. Kumar, M. Wagner, R. Wilders, D. Golod, W. Goolsby, and R. Joyner, "Electrical interactions between a real ventricular cell and an anisotropic two-dimensional sheet of model cells," *American Journal of Physiology – Heart and Circulatory Physiology*, vol. 278, pp. 452–460, 2000.
- [6] N. Otani, "Computer modeling in cardiac electrophysiology," *Journal of Computational Physics*, vol. 161, pp. 21–34, 2000.
- [7] A. Hodgkin and A. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve membranes," *Journal of Physiology*, vol. 117, pp. 500–544, 1952.
- [8] D. Noble and Y. Rudy, "Models of cardiac ventricular action potentials: Iterative interaction between experiment and simulation," *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, vol. 359, pp. 1127–1142, 2001.
- [9] S. Schaal, C. Atkeson, and S. Vijayakumar, "Scalable techniques from nonparametric statistics for real time robot learning," *Applied Intelligence – Special Issue on Scalable Robotic Applications of Neural Networks*, vol. 17, pp. 49–60, 2002.