

Massively parallel simulation of cardiac electrical wave propagation on Blue Gene

Jeffrey J. Fox¹, Gregory T. Buzzard², Robert Miller¹, and
Fernando Siso-Nadal¹

¹ Gene Network Sciences,
53 Brown Rd, Ithaca, NY, USA
E-mail: {jeff, robert, siso}@gnsbiotech.com

² Department of Mathematics,
Purdue University,
W. Lafayette, Indiana, USA
E-mail: buzzard@math.purdue.edu

Heart rhythm disorders are a leading contributor to morbidity and mortality in the industrialized world. Treatment and prevention of cardiac rhythm disorders remains difficult because the electrical signal that controls the heart rhythm is determined by complex, multi-scale biological processes. Data-driven computer simulation is a promising tool for facilitating a better understanding of cardiac electrical properties. Conducting detailed, large-scale simulations of the cardiac electrical activity presents several challenges: the heart has a complicated 3D geometry, conduction of the electrical wave is anisotropic, and cardiac tissue is made up of cells with heterogeneous electrical properties. Our group has developed a software platform for conducting massively parallel simulations of wave propagation in cardiac tissue. The code is being used in an on-going study to connect drug-induced modifications of molecular properties of heart cells to changes in tissue properties that might lead to a rhythm disorder. The platform uses a finite difference method for modeling the propagation of electrical waves in cardiac tissue using the cable equation with homogeneous Neumann boundary conditions. We use a novel algorithm which is based on the phase field method for handling the boundary conditions in complicated geometries. To map grid points within the spatial domain to compute nodes, an optimization process is used to balance the trade-offs between load balancing and increased overhead for communication. The performance of the code has been studied by simulating wave propagation in an anatomically realistic model of the left and right ventricles of a rabbit heart on Blue Gene partitions of up to 4,096 processors. Based on these results, the Blue Gene architecture seems particularly suited for cardiac simulation, and offers a promising platform for rapidly exploring cardiac electrical wave dynamics in large spatial domains.

1 Introduction

The primary motivation for developing a detailed 3D model of the electrical activity of the heart is to create a tool that can be used to help identify the underlying electrical mechanisms for dangerous ventricular arrhythmias and to determine the effects of interventions, such as drugs, that may prevent or exacerbate these arrhythmias. High-performance computing systems such as the Blue Gene are a promising resource for conducting the large-scale simulations required for elucidating how molecular modifications to heart cells might lead to dangerous alterations to electrical wave propagation in the heart.

Our group has developed a software platform for conducting massively parallel simulations of wave propagation in cardiac tissue. The motivation for this work is discussed in Sec. 2. The mathematical framework that the platform uses is described in Sec. 3. Next,

several aspects of the code are outlined in Sec. 4, including methods for parallelization. Finally, the performance of the code on up to 4096 processors is studied in Sec. 5.

2 Motivation

The electrical signal that controls the mechanical activity of the heart originates in the sinus node. This signal travels down a specialized conduction system and then propagates through the ventricles, the main pumping chambers of the heart. At the cellular level, the electrical signal is a steep rise and then gradual fall in the cellular membrane potential called the action potential (AP). This electrical activity is manifest on the surface of the body in the form of the electrocardiogram (ECG). The QT interval, the time between the "QRS" complex and the "T" wave in the ECG, roughly corresponds to the period of time between the upstroke of the AP and repolarization.

Heart rhythm disorders, or arrhythmias, refer to a disruption of the normal electrical signal. Catastrophic rhythm disturbances of the heart are a leading cause of death in many developed countries. For example, one type of rhythm disorder called ventricular fibrillation results in about 350K deaths per year in the US alone. In addition, cardiac electrical activity can be unintentionally disrupted by drugs. Unfortunately this problem does not only occur in compounds designed to treat the heart; it can happen in any therapeutic area. Through a number of studies scientists have identified a link between drugs that block certain ion currents (particularly I_{Kr} , the rapid component of the delayed rectifier potassium current), that prolong the QT interval in the ECG, and a deadly type of arrhythmia called Torsades de Pointes. Yet, it remains difficult to determine what specific modifications in the function of these proteins can lead to the initiation of a complex and life-threatening cardiac electrical disturbance. For example, some I_{Kr} inhibitors are widely prescribed and safely used. Thus, the mechanism by which a modification in the functional properties of one or more ion channels might lead to the induction of a dangerous ventricular arrhythmia is still poorly understood.

3 Mathematical Framework

A standard model¹ for the dynamics of the membrane potential, V (measured in mV), in cardiac tissue is given by the continuous cable equation:

$$\frac{\partial V}{\partial t} = \nabla \cdot D \nabla V - \frac{1}{C_{mem}} \sum I_{ion} \quad (1)$$

C_{mem} is the membrane capacitance (measured in pF/cm^2), which is usually taken to be 1, I_{ion} is the sum of all ionic currents (measured in pA/cm^2), ∇ is the 3D gradient operator, and D is the effective diffusion coefficient (measured in cm^2/s), which is in general a 3x3 matrix to allow for anisotropic coupling between cells. The equations for the ionic currents are ordinary differential equations (ODEs) and contain the details of the AP models. For isolated tissue, the current flow must be 0 in directions normal to the tissue boundary. This zero-flux condition is represented by the condition $n \cdot (D \nabla V) = 0$, where n is the unit vector normal to the heart surface, which implies that in the absence of the I_{ion} term, voltage is conserved. Fenton, et al² provide a method of deriving D from knowledge of

the fiber direction at each point in the heart. The phase field method³ can be used to deal with irregular geometries characteristic of cardiac anatomy. In this method ϕ is a smooth function so that the region of interest (the ventricle), R , is given by $R = \phi > 0$ so that $\phi = 0$ outside of R and $\phi = 1$ on a large portion of R . The region $0 < \phi < 1$ is called the transition layer and should be thought of as a thin layer on the inside of the boundary of R . Then the modified equation is

$$\phi \frac{\partial V}{\partial t} = \nabla \cdot (\phi D \nabla V) - \phi \frac{1}{C_{mem}} \sum I_{ion}. \quad (2)$$

It can be shown that voltage is conserved in the limit as the transition layer width tends to 0. However, there are still some delicate issues in terms of implementation. In order to recover exactly the voltage conservation property of the cable equation with homogeneous Neumann conditions, we developed a finite-difference algorithm⁴ by first discretizing the phase-field modified cable equation on a regular grid, then taking a limit as the width of the phase field tends to 0. This is discussed further in the next section.

4 Overview of the Software

The platform uses a finite difference method for modeling the propagation of electrical waves in cardiac tissue using the cable equation (a reaction-diffusion type partial differential equation, or PDE) with homogeneous Neumann boundary conditions. We use a novel algorithm that is based on a limiting case of the phase field method for handling the boundary conditions in complicated geometries. To map grid points within the spatial domain to compute nodes, an optimization process is used to balance the trade-offs between load balancing and increased overhead for communication. The code is written in C++, and uses several standard libraries, including SUNDIALS, MPI, BOOST, BLAS/LAPACK, and SUPERLU.

4.1 Numerical Methods

Our software uses a novel finite difference scheme for evaluating $\nabla \cdot (\phi D \nabla V)$ on a regular grid in 2 or 3D⁴. The scheme is obtained as a limiting case of the phase field method using a sharp transition layer given by $\phi = 1$ in R and $\phi = 0$ outside of R . As a result, this scheme gives exact voltage conservation over the cells in R . This scheme is motivated by the formula for summation by parts that is the analog of integration by parts⁵. We assume that the problem lies on a regular grid containing a bounded region of interest, R , with grid points $p = (x_1, x_2, x_3)$ (the 2D case will be obtained simply by ignoring the third dimension). The grid space in the i^{th} coordinate is Δx_i , and e_i represents the vector that is Δx_i in the i^{th} coordinate and 0 in the other coordinates. Given a function f on the grid

and using $V_i(p)$ to denote the change in $V(p)$ during a single time step, define:

$$\Delta_i f(p) = \frac{f(p + e_i) - f(p)}{\Delta x_i} \quad (3)$$

$$\bar{\Delta}_i f(p) = \frac{f(p + e_i) - f(p - e_i)}{2\Delta x_i} \quad (4)$$

$$\Delta_j^i f(p) = \begin{cases} \Delta_i f(p) & \text{if } i = j \\ \frac{1}{2}(\bar{\Delta}_j f(p) + \bar{\Delta}_j f(p + e_i)) & \text{if } i \neq j \end{cases} \quad (5)$$

$$\Delta_{ij} f(p) = \begin{cases} \Delta_i \Delta_i f(p - e_i) & \text{if } i = j \\ \bar{\Delta}_i \bar{\Delta}_j f(p) & \text{if } i \neq j. \end{cases} \quad (6)$$

Given this definition, we propose that the discrete analog of the spatial part of the modified equation cable equation is:

$$\begin{aligned} \phi(p)V_i(p) = & \sum_{i,j=1}^3 \left[\frac{\phi(p)}{2} (\Delta_i D_{ij}(p) \Delta_j^i V(p) + \Delta_i D_{ij}(p - e_i) \Delta_j^i V(p - e_i)) \right. \\ & + \frac{D_{ij}(p)}{2} (\Delta_i \phi(p) \Delta_j^i V(p) + \Delta_i \phi(p - e_i) \Delta_j^i V(p - e_i)) \\ & \left. + \phi(p) D_{ij}(p) \Delta_{ij} V(p) \right]. \end{aligned} \quad (7)$$

By taking a limit as the phase field width tends to 0, we obtain what we call the sharp interface limit. In this case, either $\phi(p) = 0$ or $\phi(p) = 1$. When $\phi(p) = 1$, the expression above defines $V_i(p)$ as a weighted sum of the values of V at the neighbors of p . When $\phi(p) = 0$, the expression is interpreted by multiplying both sides by $\phi(p)$ and requiring that the remaining sum on the right be 0. In this case, the first and third lines drop out due to $\phi(p) = 0$, so we are left with an equation determined by setting the remaining part of the sum to 0. This gives a system of linear equations corresponding to the Neumann boundary conditions. Solving this system, which need be done only once for a given (nonmoving) boundary, produces a scheme that gives exact voltage conservation (up to roundoff error). This allows simulation of complex geometries without the need to include a transition layer that may introduce numerical artifacts. Employing this algorithm has several additional advantages. First, the implementation of this method is straightforward. Second, it enables the use of a simple finite difference scheme on a regular cubic grid without the need for a finite element representation or for additional computation of boundary values. This ensures that the complexity of the numerical scheme does not increase substantially and also facilitates parallelization of the code. Moreover, the outer layer of the ventricle, known as the epicardial layer, may be on the order of 1 mm. Since cell heterogeneity plays an important role in arrhythmias⁶, an accurate representation of tissue requires reducing or eliminating the transition layer. Additionally, from a computational point of view, a wider transition layer is very costly. In a model of rabbit ventricles, with a grid spacing of 0.25 mm, the phase field method with a transition layer of 4 grid points requires 742,790 copies of the ionic model, versus 470,196 for the method described here, a reduction of nearly 40 percent. The bulk of simulation time is devoted to the ionic terms, so this corresponds roughly to the same reduction in simulation time. Third, an important benefit of choosing this method over the use of finite elements is that the 3D code is essentially the same for all geometries considered, so that only the input file containing the anatomic structure and

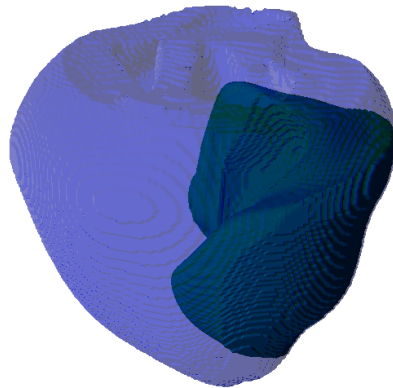


Figure 1. Visualization of the spreading activation wave front from a simulation of wave propagation in a model of the rabbit left and right ventricles using the Fox et al AP model.

conductivity tensor changes from one geometry to the next, thereby avoiding the complex and time-consuming task of 3D mesh generation. Each grid point in the 3D model can be assigned a different ionic model, thus allowing the specification of the locations of epicardial, endocardial, and M cells. This flexibility also allows the effect of variations in the distribution of these cell types on global wave propagation to be studied.

To integrate the discretized cable equation, we use one of two methods. The simplest method, and that used in the scaling study below, is operator splitting. For this method we solve the ionic term for a fixed time interval using CVODE and then solve the diffusion term using the Euler method. More recently we have investigated the method of lines, which performs significantly better (2 to 10 times speed up depending on the problem). For this method we produce a coupled system of ODEs, then integrate with the stiff solver in the SUNDIALS package (from the CVODE subset of <http://acts.nersc.gov/sundials/> using backwards differentiation). Each grid point corresponds to a separate set of equations given by the ionic terms plus the terms from spatial coupling as described above. We use Krylov subspace preconditioning with a diagonal approximation to the Jacobian. To define the fiber directions at boundary points, we start with fiber directions at points in R , then determine the fiber direction at p in the boundary, $B(R)$, by taking an average of the directions of nearest neighbors. In the simplified case described above, we then project this onto the nearest coordinate axis (or select one axis in the case of a tie). Fig. 1 shows an example of wave propagation in a model of the rabbit ventricle using the sharp interface method.

4.2 Parallelization

We use MPI for parallelization of the code. Parallelization of the 3D simulation requires almost exclusively nearest neighbor communication between compute nodes. Specifically, each node is assigned a portion of the 3D spatial domain, and the nodes pass information

about the boundary of the domain to nodes that are responsible for neighboring domains. In order to create simulations to match experimental canine ventricular wedge preparations (these preparations are typically $2 \times 1.5 \times 1$ to $3 \times 2 \times 1.5$ cm³ in size), 192,000 to 576,000 grid points will be required. On the other hand, a typical human heart has dimensions on the order of 10 cm, so directly simulating wave propagation in the human ventricle, a primary goal of electrophysiological modeling, requires 10 to 100 million grid points. Since the bulk of computation time is used on evaluation of the ionic terms and local integration and since communication between grid points is nearest neighbor in form and minimal in size, this problem is well suited to a parallel implementation using a large number of processors. Nevertheless, there are many interesting problems about how best to distribute the cells to the processors to maximize speed; improving load balancing for a large number of processors is an ongoing project. Load balancing is a particularly important issue to consider when using spatial domains that are anatomically realistic. In this case, because the problem is defined on a regular grid, only a portion of the grid nodes correspond to active tissue; the rest are empty space. For example, an anatomical model of a rabbit ventricular wedge uses a grid that is $119 \times 109 \times 11$, so there are 142,681 nodes, but only 42,457 are active. All 142,681 are assigned to tasks, so some tasks get more active nodes than others. To map nodes to tasks, an optimization process is used to balance the trade-offs between load balancing and increased overhead for communication.

To limit the complexity of the message-passing scheme, we consider the tissue as embedded in a larger rectangular domain (a fictitious domain in the language of numerical PDE) and subdivide this rectangular domain into a product grid. This results in a collection of disjoint boxes which together cover the fictitious domain. E.g., if the fictitious domain has dimensions 110 by 150 by 120, the product domain might consist of boxes of size 10 by 5 by 6. A single processor may be assigned more than one box in the resulting decomposition. In the current implementation of the algorithm, the size of the boxes in the product grid is the same for all the boxes and is chosen using an optimization algorithm in order to distribute the number of active nodes nearly evenly among all the processors while at the same time limiting the amount of message passing that is required among processors. In particular, the boxes are assigned to processors in such a way that nearest neighbor boxes are assigned to nearest neighbor processors in order to take advantage of the torus topology of the communications network. In most cases of interest, this algorithm distributes the active nodes to within 1 percent of the optimal distribution. Moreover, given the data to determine the active nodes and the size of the network, the optimization algorithm can be run in advance of the simulation to produce the correct division of labor. Further improvements in this algorithm are possible.

5 Parallel Performance

We have carried out a number of simulations on the Blue Gene/L system to illustrate the parallel performance of the cardiac simulation application. We simulated wave propagation in a 3D model of the rabbit ventricular anatomy⁸ using the Fox et al AP model⁷. We conducted two different comparisons. First, we compared performance between two different sized spatial domains using the same number of compute nodes (32 nodes, 64 processors). We compared the time to simulate 10 msec of wave propagation in the full rabbit ventricle versus a wedge of the ventricle. The results are shown in the table below. The

Tissue	Total grid pts.	Active	% diff. from opt.	Run time (sec)	% in int.
wedge	142,000	42,000	13	80	67
ventricle	2,846,000	470,000	3	659	93

improved performance of the code on the rabbit ventricle is likely explained by looking at the column ”% diff. from opt.”. This quantity measures the percent difference between the optimal number of spatial grid points assigned to each compute node and the actual number of grid points assigned to the ”most loaded” compute node. Ideally, each compute node would be responsible for an identical number of spatial grid points. However, the task assignment process must strike a balance between achieving optimal division of labor and minimizing communication time. Most likely, the simple explanation as to why the full ventricle has a more homogeneous load is the fact that it has more nodes than the wedge so percent differences will be smaller. In addition, it may be that specific geometrical features play a role.

The second comparison is a strong scaling performance example. We compared simulation times using the same anatomy as a function of number of processors for 10 msec of wave propagation. The table below shows results for the full rabbit ventricle. Fig. 2 shows

Nproc	Run time (sec)	Ideal (sec)	% in int.	Speed up	Ideal speed up
32	1203	1203	93	32	32
64	620	602	93	62.08	64
128	316	301	88	121.92	128
256	166	150	85	232	256
512	87	75	82	442.56	512
1024	49	38	77	785.6	1024
4096	18	9	54	2138.56	4096

how the time to solution as a function of processor number compares to an ideal scaling. The performance of the code is quite good for the partition sizes studied here.

6 Concluding Remarks

To conclude, large scale resources like Blue Gene allow researchers to explore qualitatively different problems. For the example application described above, conducting a dose-response simulation of the effect of a compound on a 3D canine ventricle at a variety of pacing frequencies would require many months of simulation time on a standard computer cluster, but with Blue Gene one could conduct several such simulations in a day. Thus, using the software described in this study on Blue Gene is a promising platform for conducting multi-scale biomedical simulations.

Acknowledgments

This work was supported by NIH grants R44HL077938 and R01HL075515. We are grateful to IBM for providing computer time for the parallel performance studies.

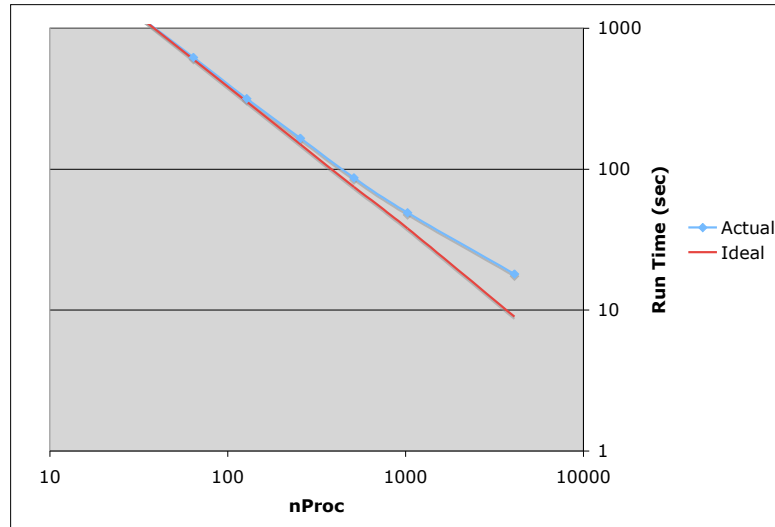


Figure 2. Run time as a function of processor number. The diamond points are from scaling runs on the Blue Gene/L system, and the solid line is the ideal run time assuming perfect scaling.

References

1. *Special issue on cardiac modeling*, Chaos **8**, (1998).
2. F. H. Fenton, E. M. Cherry, A. Karma, W. J. Rappel, *Modeling wave propagation in realistic heart geometries using the phase-field method*, Chaos **15**, 013502 (2005).
3. A. Karma, W. J. Rappel, *Numerical simulation of three-dimensional dendritic growth*, Phys Rev Lett **77**, 450–453 (1996).
4. G. T. Buzzard, J. J. Fox, F. Siso-Nadal, *Sharp interface and voltage conservation in the phase field method: application to cardiac electrophysiology*, SIAM J. Sci Comp. **in press**, (2007).
5. R. L. Graham, D. E. Knuth, O. Patashnik, *Concrete Mathematics*, (Addison-Wesley, New York, 1989).
6. C. Antzelevitch, W. Shimizu, *Cellular mechanisms underlying the long QT syndrome*, Curr. Opin. Cardiol. **17**, 43–51 (2002).
7. J. J. Fox, J. L. McHarg, R. F. Gilmour, Jr., *Ionic mechanism of electrical alternans*, Amer. J. Physiol. **282**, H516–H530 (2002).
8. F. J. Vetter, A. D. McCulloch, *Three-dimensional analysis of regional cardiac function: a model of rabbit ventricular anatomy*, Progress Biophys. Mol. Bio. **69**, 157–183 (1998).