Estimation of Activation Times in Cardiac Tissue Using Graph Based Methods

Mikael Wallman[1,2], Nic Smith[3], and Blanca Rodriguez[1]

[1] Oxford University Computing Laboratory, Oxford, United Kingdom
[2] Fraunhofer-Chalmers Centre, Chalmers Science Park, Gothenburg, Sweden
[3] Division of Imaging Sciences, King's College London, United Kingdom

Abstract. The bidomain and monodomain equations are well established as the standard set of equations for the simulation of cardiac electrophysiological behaviour. However, the computational cost of detailed bidomain/monodomain simulations limits their applicability to scenarios in which results are needed in real time (e.g. clinical scenarios). In this study, we present a graph based method which relies on point to point path finding to estimate activation times in cardiac tissue with minimal computational costs. Activation times are compared to bidomain simulation results for heterogeneous tissue slabs and an anatomically-based rabbit ventricular model. Differences in activation times between our proposed graph based method and bidomain results are less than 10% of the total activation time and computational performance is orders of magnitude faster with the graph based method. These results suggest that the graph based method could provide a viable alternative to the bidomain formalism for the fast estimation of activation times when the need for fast performance justifies limited loss of accuracy.

1 Introduction

During the last two decades, the bidomain equations, and the closely related monodomain equation, have emerged as a gold standard for simulating cardiac electrophysiology [1]. Although able to provide a sophisticated representation of cellular mechanisms and intercellular interactions, solving the resulting PDEs is very computationally expensive. Additionally, in many situations, e.g. when determining the activation times in a steady state setting, the level of complexity provided by the bidomain formalism is higher than necessary, making bidomain simulations inefficient solutions to the problem at hand.

For such situations, several alternative ways of describing cardiac propagation exist, ranging from early models relying on Huygens principle [2], via cellular automata models [3], to models explicitly derived from the bidomain equations, such as the eikonal equations [4]. Common between all of these simplified models is that finding the order of activation requires a sequence of computations for all the spatial nodes. Thus, in each step of the algorithm, all data needed to find the wavefront are computed. While this is desirable in many situations, for some applications, even focussing the computational efforts on locating the wavefront is inefficient.

D.N. Metaxas and L. Axel (Eds.): FIMH 2011, LNCS 6666, pp. 71–79, 2011.

Instead of computing the wavefront, we might determine the path the current takes from a point of initial activation to any point in the tissue, representing the fastest route between the two points. The analytic calculation of these paths in a continuous setting remains an open problem for most situations. However, some progress has been made towards closed form solutions [5], and several algorithms exist for solving the corresponding problem on graphs. Cardiac tissue can thus be described as a connected graph, allowing activation times to be approximated in a very efficient manner.

This paper presents a novel approach to estimating activation times in cardiac tissue using a graph-based method. Results are compared to finite element solutions to the bidomain equations in an attempt to characterise method accuracy. The comparison is performed in three different models: (1) a succession of 7 cardiac tissue slabs of decreasing resolution to evaluate methods convergence; (2) a cardiac tissue slab incorporating fiber rotation and a central region of slow conduction; (3) an anatomically-based rabbit ventricular model.

2 Methods

2.1 The Bidomain Model

The bidomain equations describe the cardiac tissue as two continuous and completely interpenetrating domains. At each point in space, two electrical potentials exist, one intra-cellular (ϕ_i), and one extra-cellular (ϕ_e). This also implies a transmembrane potential, v, at each point in space, defined as

$$v = \phi_i - \phi_e \ . \tag{1}$$

With the transmembrane potential defined is this way, the bidomain equations can be written as

$$\begin{aligned} &\beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(\eta, V_m)) - \nabla \cdot (\sigma_i \nabla (V_m + \phi_e)) = I_{S_i}, \\ &\nabla \cdot ((\sigma_i + \sigma_e)\nabla \phi_e + \sigma_i \nabla V_m) = I_{S_e}, \end{aligned} \tag{2}$$

where β is the cell surface to volume ratio, C_m is the membrane capacitance per unit area, σ_i is the intracellular conductivity tensor, σ_e is the extracellular conductivity tensor, I_{S_i} is an external stimulus applied to the intracellular space and I_{S_e} is an external stimulus applied to the extracellular space. I_{ion} is the ionic current, a function dependent on the cell model coupled to the bidomain model. η is a vector containing the state variables for the cell model.

2.2 The Graph Based Model

By considering the cardiac tissue as a connected graph, very fast approximations of activation sequence can be obtained. In this context, a graph consists of spatial nodes, connected by edges. Every edge in the graph is assigned a cost, based on

the time it takes the activation wavefront to traverse the corresponding path between two points in the tissue. Activation is initiated at one or several nodes, corresponding to the point or points where the tissue is initially stimulated. From there, the activation travels from node to node along the edges of the graph. At each node, an estimate of its activation time can be obtained by finding the accumulated cost of all edges traversed in order to reach it along a specific path. Typically, a very large number of paths can be taken between two nodes in the graph, so in order to obtain the best estimate of the activation time, the path with the lowest cost needs to be found.

Path Finding with the A* Algorithm. In this work, the A* algorithm has been used to find the lowest cost path though the network [6]. The A* works by keeping a priority queue of nodes to search. Each node n in the priority queue is assigned a score, based on the expected cost $f(n)$ of reaching the goal along a path passing through n, and the queue is ordered according to these scores. The score $f(n)$ is a sum of two other scores $g(n)$ and $h(n)$. $g(n)$ represents the accumulated cost of all the edges traversed in order to reach n, while $h(n)$ is a heuristic estimate of the cost that will be accumulated along the path from n to the goal.

Starting from a node s in the graph, a basic version of the A* algorithm can be described as follows:

1. Add s to the priority queue and calculate $f(s)$.
2. Select the node n on the queue whose value of f is smallest.
3. If n is the goal node, terminate the algorithm.
4. Otherwise, find all connected nodes n_i connected to n, calculate $f(n_i)$ for each and add them to the queue.
5. Goto step 2.

In the context of this work, two different heuristics have been employed. For the succession of slabs of decreasing resolution, a weighted Euclidean distance was used. For the more complex models, the heuristic estimate was based on another path finding procedure, performed on a down-sampled mesh, which itself used the heuristic $h(n) = 0$ for all nodes. The final estimates resulting from this procedure can thus be viewed as a refinement of initial coarser ones.

Graph Construction. As stated in the introduction, 9 different meshes were used in this work: 8 discretized slabs and 1 anatomically-based ventricular mesh. The nodes of the discretized slabs were regularly spaced in the shape of a Cartesian grid, while the nodes of the ventricular mesh had a less ordered distribution, conforming to the more complex geometry. For the regular meshes, the graphs presented here were constructed by successively processing cubic subgraphs of $2\times2\times2$ nodes, ensuring that all nodes in the subgraphs were 7-connected, so that an edge existed between any pair of nodes in the subgraph. For the anatomically-based ventricular geometry, the finite element mesh was used as the graph.

For each edge in the resulting graphs, costs of traversal was calculated according to

$$c_{i,i+1} = \sqrt{v_{i,i+1}^T M_{i,i+1} v_{i,i+1}}, \tag{3}$$

where $v_{i,i+1}$ is the vector from node n_i to node $n_i + 1$ and $M_{i,i+1}$ is a tensor describing conduction costs for that edge. For the 3D case, M can be written as

$$M = [m_l \lambda_l^2, m_t \lambda_t^2, m_n \lambda_n^2] \ . \tag{4}$$

Here, the set of vectors $\{m_l, m_t, m_n\}$ form an orthonormal system, describing the local fiber direction, while the scalars λ_l, λ_t, λ_n are the costs of traversal in the longitudinal, transversal and normal directions defined by this system.

In the results presented below, the conduction costs for the graphs were estimated using the full set of activation times resulting from the bidomain simulations. Starting from the bidomain simulation output, the activation time for each node was defined as the time when V_m changed sign from negative to positive. The three principal directions of conduction in the graph, corresponding to the m_l, m_t and m_n elements of M (i.e. the fibre direction), were the same as those used in the bidomain simulations. Conduction costs for these directions, corresponding to λ_l, λ_t and λ_n, were estimated using the Nelder-Mead simplex algorithm. As an objective function, the RMS error between the activation times estimated using the graph based method and those calculated from the bidomain simulations was used.

Recently, van Dam et al. [7] used a graph based method to estimate activation times on a ventricular surface. Although superficially similar, the method differs from ours in several important ways. While the previous method only considers nodes on the cardiac surface, our method also takes into account nodes that are situated within the cardiac walls. Additionally, while the previous method uses a fully connected graph, our method uses a sparser graph with a comparatively small number of connections for each node, allowing it to be applied to larger graphs. Furthermore, our method treats fibre directions in a more explicit way, using information about the the principal conduction directions in each node when determining the conduction costs for the graph. Finally, our method relies on a heuristic path finding algorithm, leading to very low computational costs for estimating the activation time in individual nodes.

3 Results

Activation times obtained using the graph-based method were compared to corresponding bidomain simulations in each of the three different settings: The succession of seven slabs of decreasing resolution, the slab incorporating fiber rotation and conduction heterogeneity, and the full ventricular mesh. The slab succession provides means to compare how the two methods perform for different mesh resolutions, while the two more complex settings elucidate how well the graph based method is able to mimic a given bidomain activation pattern.

All bidomain simulations presented were done using the Chaste software package [8]. The relative errors were calculated as $E_i = (\hat{AT}_i - AT_i)/AT_{\text{tot}}$, where E is relative error, i is the node index, $\hat{AT}$ is the approximate activation time, AT is the correct activation time, and AT_{tot} is the time it takes for all the nodes in the considered mesh to become activated. For the tissue slab with varying mesh resolution, the activation times from the finest mesh was considered correct, while for the remaining simulations, the activation times calculated from the bidomain simulations were considered correct, giving the errors a subtly different interpretation.

3.1 Method Comparison for Different Mesh Resolutions

The first set of bidomain simulations were performed on a cuboid geometry, measuring $1.80 \times 0.73 \times 0.29$ cm. The cuboid was discretized at seven different levels, producing seven tetrahedral meshes with homogeneous element size. The average node distance for each of the 7 meshes was 684 μm, 622 μm, 442 μm, 306 μm, 179 μm, 132 μm and 75 μm respectively.

For the bidomain simulations, activation of cardiac tissue was initiated by point stimulation at one of the corners of the mesh. After stimulation, activation time at each mesh node was calculated. These activation times were also estimated using the graph based approach, resulting in another set of activation times. For the bidomain simulation at a mesh resolution of 75 μm mesh, AT_{tot} was 34 ms.

For comparison between the different discretizations and methods, a grid of 250 points was chosen, with the points evenly distributed throughout the volume. The activation time values in these points, obtained from the bidomain simulations and the graph based approximations, were found using the built in linear interpolation functionality in MATLAB. Results are shown in Fig. 1.

From the histograms and graph in Fig. 1 it can be seen that the bidomain simulations have a lower relative error than the graph based method up to a discretization of 300 μm, while for coarser discretizations the graph based method has a lower relative error. The results also show that while the average relative error of the bidomain equations greatly increases for discretizations above 400 μm, the corresponding error for the graph based method remains almost constant.

3.2 Method Comparison in a Slab of Heterogeneous, Anisotropic Tissue

Next, propagation in a tissue slab of 4×4×1 cm (for the x,y and z directions respectively), with a spatial discretization of 300 μm was simulated. The conductivity was anisotropic, with fiber directions according to the Streeter model [9]. For the current geometry this implies that fibres were oriented in planes of homogeneous fibre orientation, perpendicular to the z-axis, with the fibre direction gradually rotating 120° between the planes $z = 0$ and $z = 1$. Additionally, a region of low conductivity was incorporated at the centre of the tissue, centred at the point $(x, y, z) = (2, 2, 0)$ and defined by two ellipsoids. The conductivity

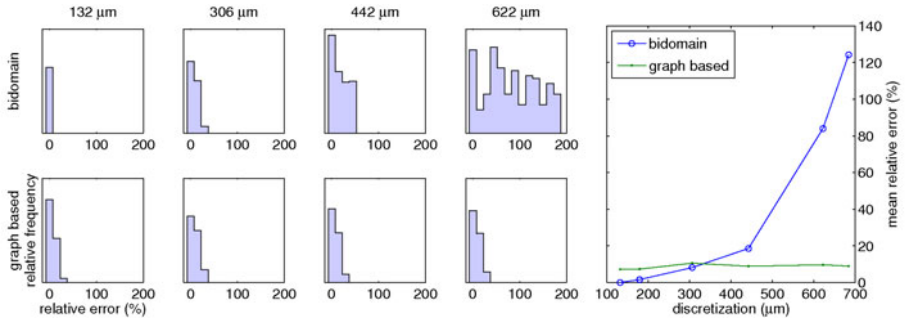

Fig. 1. Left panel shows histograms of relative errors in activation time for bidomain (upper row) and the graph based method (lower row). The columns correspond to different mesh discretizations, with the node spacing increasing from left to right. Right panel shows the mean relative error in activation time of bidomain (blue line, points marked with circles) and graph based approximations (green line, points marked with dots).

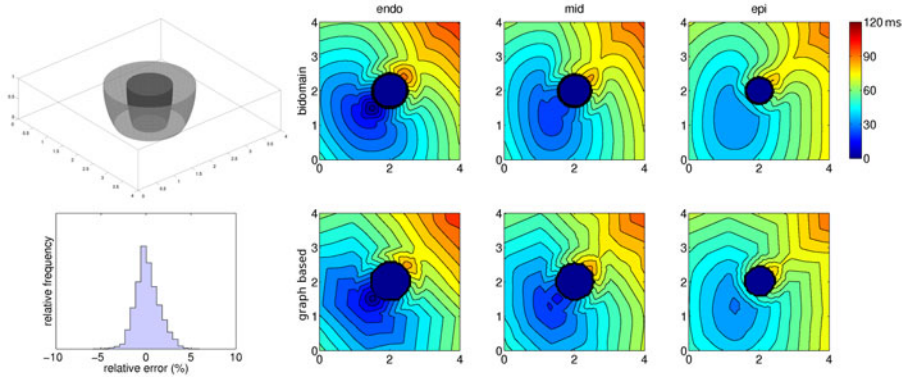

Fig. 2. The upper left panel illustrates the tissue slab with its central low conductivity region. The lower left panel shows a histogram of the relative errors of the graph based activation time approximations. The six right panels show activation isochrones of the six slab simulations. The upper row corresponds to bidomain results, while the lower row shows graph based approximations. The columns from left to right correspond to endocardium, mid-myocardium and epicardium.

inside the larger ellipsoid was 10 times lower than outside. Additionally, in the region overlapping the smaller ellipsoid, the tissue was modelled using passive diffusion only. Conductivity in the slow region was isotropic. A bidomain simulation of the activation was performed, with the initially stimulated area situated just outside the low conductivity region. Again, activation times were computed for each node, and corresponding activation times were also estimated using the graph based method. The total activation time for the bidomain simulation was

120 ms. The geometric model, along with a histogram of relative errors and activation isochrones for 3 cross sections of the slab, are shown in Fig. 2.

From the histogram in Fig. 2 it can be seen that the errors of the graph based activation time approximations keep within 5% of the total activation time. Considering the activation isochrones on the right side of Fig. 2, it can be observed that although polygonal in shape, the activation isochrones from the graph based method (lower row) show strong similarities in shape to the isochrones from bidomain (upper row).

3.3 Method Comparison in a Rabbit Ventricular Mesh

Finally, a realistic mesh of the rabbit heart ventricles, derived from MRI [10], was simulated. The mesh resolution was 420 μm, giving a total of 82619 nodes. The model was activated at several endocardial points in the lower half of the ventricles, emulating purkinje activation. Subsequently, the activation times for each node as obtained from the bidomain simulations were compared to the corresponding approximations obtained using the graph based method. The total activation time for the model was 71 ms. Results are shown in Fig. 3.

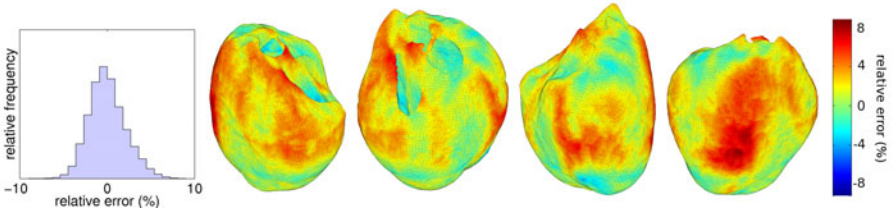

Fig. 3. Left panel shows a histogram of relative errors in activation time for graph based approximations. Right panel shows four rotated views of the distribution of relative errors in activation time on the epicardium of the rabbit ventricular mesh. The geometry is rotated 90 degrees between consecutive frames.

The histogram in Fig. 3 shows a fairly narrow distribution of errors, keeping well within 10% of the total activation time. The spatial distribution of errors show a largely irregular pattern, with no strong correlation between the distance from the area of initial stimulation and the magnitude of the errors. The rightmost view show a central accumulation of relatively high positive errors, which might be due to the greater wall thickness of this area, accentuating effects of wavefront curvature in the bidomain propagation.

4 Conclusions

We have presented a fast, graph based method for cardiac activation time estimation under steady state conditions. The comparison between the graph based method and the succession of increasingly coarse slabs show that while the average relative error for the bidomain greatly increases at discretizations above

300 μm, the corresponding error for the graph based method stays almost constant for the entire range of resolutions investigated. Additionally, results from both the slab with anisotropic, heterogeneous conduction and from the ventricular mesh show that the graph based algorithm is able to approximate activation times to well within 10% of the total activation time. This suggest that the method provides a viable alternative for situations where need for fast performance justifies limited loss of accuracy.

In the simulations presented here, the wall time for approximating activation times of isolated nodes in the most complex meshes were of the order of 10 ms on a standard desktop computer, making them 10^4 - 10^5 times faster than the bidomain simulations for that particular task. The authors are aware of the existence of alternative fast methods for activation time approximations, such as the fast marching method [11]. However, the primary aim of this work has been to investigate the accuracy of the proposed method. For this purpose the bidomain equations, representing the state of the art in cardiac electrophysiological modelling, provide a good basis of comparison. A full comparison of all alternative methods in terms of computational performance is beyond the scope of this manuscript. It should however be noted that the computational complexity of Dijkstra-like methods, such as the fast marching method (FMM), is $O(NlogN)$ for N nodes, while for the A* algorithm, this complexity provides a very rare worst case scenario. Additionally, each node expansion is likely to be much more computationally expensive for FMM than for the A* algorithm. While these considerations do not quantify the difference in performance, they strongly suggest that the proposed method is less computationally expensive than the bulk of competing fast methods.

Acknowledgements

This work was supported by the European Commission as part of the euHeart project FP7-ICT-2007-224495 and the preDiCT grant (DG-INFSO-224381). BR holds a UK Medical Research Council Career Development Award.

References

1. Keener, J., Sneyd, J.: Mathematical Physiology, 2nd edn. Springer, New York (1998)
2. Plonsey, R., Barr, R.C.: Mathematical modeling of electrical activity of the heart. Journal of electrocardiology 20(3), 219–226 (1987)
3. Ye, P., Entcheva, E., Grosu, R., Smolka, S.A.: Efficient modeling of excitable cells using hybrid automata. In: Proc. of CMSB, pp. 216–227 (2005)
4. Tomlinson, K.A., Hunter, P.J., Pullan, A.J.: A finite element method for an eikonal equation model of myocardial excitation wavefront propagation. SIAM J. Appl. Math. 63(1), 324–350 (2002)
5. Pashaei, A., Sebastian, R., Zimmerman, V., Bijnens, B.H., Frangi, A.F.: A meshless approach for fast estimation of electrical activation time in the ventricular wall. In: Computers in Cardiology 2009, pp. 209–212 (2010)

6. Hart, P., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics 4(2), 100–107 (1968)
7. van Dam, P.M., Oostendorp, T.F., Oosterom, A.: Application of the fastest route algorithm in the interactive simulation of the effect of local ischemia on the ECG. Medical & Biological Engineering & Computing 47(1), 11–20 (2008)
8. Pitt-Francis, J., Pathmanathan, P., Bernabeu, M.O., Bordas, R., Cooper, J., Fletcher, A.G., Mirams, G.R., Murray, P., Osborne, J.M., Walter, A., et al.: Chaste: a test-driven approach to software development for biological modelling. Computer Physics Communications 180(12), 2452–2471 (2009)
9. Streeter, D.D., Spotnitz, H.M., Patel, D.P., Ross, J., Sonnenblick, E.H.: Fiber orientation in the canine left ventricle during diastole and systole. Circulation Research 24(3), 339–347 (1969)
10. Corrias, A., Jie, X., Romero, L., Bishop, M.J., Bernabeu, M., Pueyo, E., Rodriguez, B.: Arrhythmic risk biomarkers for the assessment of drug cardiotoxicity: from experiments to computer simulations. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 368(192), 3001–3025 (1921)
11. Sethian, J.A., Vladimirsky, A.: Fast methods for the eikonal and related Hamilton Jacobi equations on unstructured meshes. Proceedings of the National Academy of Sciences of the United States of America 97(11), 5699–5703 (2000)