

Robust Graph-based Upscaling of Microscale Fibrotic Structures

Megan E. Farquhar¹, Kevin Burrage^{1,2}, Brodie A J Lawson^{1,3}

¹ ARC Centre for Mathematical and Statistical Frontiers, Queensland University of Technology, Brisbane, Australia

² University of Oxford, Oxfordshire, United Kingdom (Visiting Professor)

³ Queensland University of Technology Centre for Data Science, Brisbane, Australia

Abstract

Electrophysiological (and electromechanical) simulations are infamous for their computational runtimes, and remain a significant barrier to a full understanding of the electrical signalling process in the heart. These computational challenges are particularly impactful when it comes to understanding the important effects of small-scale heterogeneities, such as due to cardiac fibrosis, as refinement of cardiac meshes to even finer spatial resolutions incurs an insurmountable increase in simulation time. Therefore, methods that are able to appropriately upscale these micro-scale heterogeneities, that is, to appropriately represent their effects on conduction in a coarser mesh, are of critical importance to understanding how exactly, and under which circumstances, conditions such as cardiac fibrosis act to promote dangerous arrhythmia. In this work, we combine our recent ideas for upscaling using homogenisation via volume averaging and a unique graph-based approach to demonstrate a powerful tool for capturing some of the subtler and localisation-dependent effects of cardiac fibrosis with only minor impact on simulation runtimes.

1. Introduction

As with most tissue, cardiac tissue is far from homogeneous. Even ignoring for the moment the presence of fibroblasts and extracellular matrix, the cardiac myocytes that carry the signal responsible for the heartbeat can be viewed as disjoint islands, connected electronically by their gap junctions. As such, constructing spatially homogeneous continuum equations to describe the flow of charge, such as the ubiquitous monodomain and bidomain models, requires some kind of averaging over the conductivity properties of myocyte and gap junction [1]. This process, and the mathematical theory used to derive the appropriate averaging, is known as *homogenisation*, owing to the way a heterogeneous material is being replaced by a homogeneous one reflective of its bulk properties. Although

such homogenised models rely on an approximation (that can break down in some cases [2]), the undeniable success of the monodomain and bidomain models in enabling mechanistic simulation of arrhythmia demonstrates their broad potential.

Less well-established, although with much to be potentially gained, is the use of homogenisation for the capture of heterogeneities that are more complex in terms of shape, their effect on electrophysiological signalling, or both. It is easy enough to treat myocytes and their gap junctions as a periodically repeating pattern, that appears homogeneous both intuitively and mathematically as one “zooms out”, and manifests simply in an anisotropic conduction tensor [1]. On the other hand, the idea of capturing spatially heterogeneous structures that permit conduction only in one direction, driven by the phenomena of source/sink mismatch [3], simply by choosing upscaled conductivity tensors may seem ludicrous. Most works that have brought a homogenisation approach to upscaling in the context of obstacles to conduction associated with cardiac fibrosis have therefore limited themselves to considering spatially periodic patterns of obstruction, such that the mathematics underlying homogenisation apply nicely, and reasonable performance can be expected [4, 5]. However, recent work from the authors [6] suggests that by also tracking spatial variation in the *amount* of obstruction present, and selecting appropriate boundary conditions for homogenisation subproblems, some of the more complex and arrhythmia-relevant phenomena arising from fibrotic blockages may still be captured by upscaled models.

Additionally, we have demonstrated in work under submission that the travelling wave behaviour that defines electrical signalling in the heart permits a graph-based approach to homogenisation. By evaluating the time taken for excitation to travel through the heterogeneous, fine-scale structure to reach neighbouring nodes on the upscaled mesh, homogenised conduction tensors can be selected directly to exhibit the same timings, without any kind of implied assumptions of periodicity. Appealingly, in contrast to methods that take a graph-based

(e.g. Eikonal) approach to simulating the cardiac dynamics themselves, upscaled models produced via graph-based homogenisation are still monodomain (or bidomain) models that properly track depolarisation and repolarisation and hence are suitable for simulation of arrhythmia and not just the generation of activation maps.

Here, we demonstrate and extend this graph-based approach by also incorporating the volume fraction of conductive tissue as per our work in [6], and improving both the generation of appropriate upscaled grids and the numerical approach to solving the monodomain model on these grids. As we demonstrate, this makes the graph-based approach both competitive with traditional homogenisation when it is appropriate, and applicable in many cases where the traditional approach is not.

2. Methods

2.1. Simulation of Cardiac Activity

We use the monodomain model for the spatiotemporal evolution of cardiac excitation,

$$C_m \frac{\partial V}{\partial t} = \nabla \cdot (\mathbf{D} \nabla V) - \mathbf{I}_{\text{ion}} - \mathbf{I}_{\text{stim}}, \quad (1)$$

where V is the transmembrane potential, \mathbf{D} is a combination of the conductivity tensor with some physical constants, and I_{ion} is a function that defines the current flow through the cell membrane in terms of a set of state variables \mathbf{s} that evolve according to their own set of ordinary differential equations. External stimulus, as provided to initiate excitation, is provided by the I_{stim} term. The monodomain model is a simplification (appropriate in many cases [7]) of the bidomain model, which can be derived from fundamental physical principles.

In this context, \mathbf{D} is a spatially-varying quantity, set to $\mathbf{D} = \mathbf{0}$ in regions of blockage (for example, non-conductive collagenous deposits in fibrosis). Here tissue conducts isotropically, $\mathbf{D} = D\mathbf{I}$ with $D = 2.5 \times 10^{-4}$ mS. For an ionic model we choose the reduced ten Tusscher model [8], which represents all of the major channels for ion transport in ventricular cells, but uses a set of simplifying assumptions to reduce computational cost.

The homogenised version of equation (1) that we use here is derived in [6], and takes the form

$$\phi C_m \frac{\partial V}{\partial t} = \nabla \cdot (\phi \mathbf{D}_{\text{eff}} \nabla V) + \phi (I_{\text{ion}} + I_{\text{stim}}), \quad (2)$$

where ϕ is the volume fraction of conductive tissue (a spatially-varying quantity) and \mathbf{D}_{eff} are the upscaled tensors obtained as described subsequently.

2.2. Graph-based Upscaling

Equations (1) and (2) generate travelling waves of excitation. This has motivated the use of graph-based simplified models that track only the moments at which the different points on the graph activate, by making use of the Eikonal equation [9]. Here, we use this approach as a way to define the effective tensors \mathbf{D}_{eff} that describe the upscaled model, selecting them so that the Eikonal activation times calculated on the finescale grid match those on the coarser upscaled grid. This way, upscaling that reflects the meaningful effect of finescale heterogeneities is achieved without implicit assumptions about structure.

With a mesh in place, a graph is already defined by treating mesh nodes as the nodes of the graph, and the edges of elements as graph edges. What must be calculated is the time ‘‘cost’’ associated with moving along each edge, connecting say nodes i and j , which can be approximated by the proportionality [9]

$$c_{ij} \propto \sqrt{\mathbf{v}_{ij}^T \mathbf{D}^{-1} \mathbf{v}_{ij}}, \quad (3)$$

with \mathbf{v}_{ij} the vector joining nodes i and j . We briefly note that where the diffusion tensor differs between the two elements on either side of an edge, some kind of modelling choice must be made (for example, selecting the smaller of the two resultant costs). Here however, we use a consistent tensor across all conductive elements and instead heterogeneity is introduced by the inclusion of wholly non-conductive elements. We treat any edge attached to at least one conductive element to be a valid connection.

The upscaled, irregular triangular mesh is constructed using one of MATLAB’s built-in meshing tools, first running `geometryFromMesh` to grab the problem geometry from the finescale mesh, then using `generateMesh` to create a new irregular mesh that tiles the geometry with a specified desired edge length to control the degree of upscaling. This produces a new set of nodepoints that do not necessarily coincide with those of the finescale grid, and in locations potentially buried in fibrosis and hence unreachable on the corresponding graph. To resolve this, we shift all of the upscaled nodes to the closest reachable node on the fine grid. Reachable nodes are those belonging to the largest connected component of the graph.

With these upscaled nodes now coinciding with those of the finescale grid, we can use the fast marching method [9] to find the shortest-time paths on the graph connecting the nodes of the upscaled mesh. Briefly, the fast marching method finds the shortest paths on graphs, but in a manner that respects the observation that excitation can also travel *across* elements, as opposed to just running along their edges [9]. With the costs associated with travelling from node to node on the upscaled mesh found, we then

calculate upscaled tensors

$$\mathbf{D}_{\text{eff}}^{-1} = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$$

by expanding equation (3) for each edge of an upscaled element, and solving the resultant three linear equations (one per triangle edge) for the three unknown tensor elements a , b and c . Where this system cannot be inverted, an element is treated as non-conductive. Resultant tensors that are not positive definite are forced to be so by zeroing any negative eigenvalues, as described in [6].

As the microscopic arrangement of obstacles varies from element to element of the upscaled mesh, the effective tensors \mathbf{D}_{eff} vary from element to element. Volume fractions ϕ are also element-specific, estimated by Monte Carlo integration using uniform samples over each triangular element Ω 's rectangular bounding box,

$$\phi = \frac{\int_{\Omega} \chi_{\text{conductive}} d\Omega}{\int_{\Omega} 1 d\Omega} \approx \frac{\sum \chi_{\text{conductive}} \chi_{\text{in element}}}{\sum \chi_{\text{in element}}}.$$

Here $\chi_{\text{condition}}$ is the indicator function, taking value one when its condition is satisfied, and zero when it is not.

2.3. Numerical Approach

Equation (2) is solved by using the control volume finite element scheme described in [6], but adapted for irregular triangular meshes. This scheme is vertex-centred, enabling (2) to be discretised without any kind of averaging of volume fractions or diffusion tensors. Timestepping uses a second-order generalisation of the Rush–Larsen method [10], with $\Delta t = 0.05$ ms.

3. Results

3.1. Activation Maps in Fibrotic Tissue

We use a regular quadrilateral mesh for the finescale grid and to define the arrangement of fibrotic obstacles, which are represented by marking elements as non-conductive at random with some fixed probability $\rho = 0.45x/L$. This creates a patterning of diffuse/replacement fibrosis that slows conduction, with the extent of fibrosis increasing the further into the tissue away from the stimulus location on the left. To minimise numerical effects associated with upscaling and thus better test specifically our homogenisation, we choose a relatively small square domain ($L = 5$ mm) that provides a finescale grid with high enough resolution to represent fibrotic obstructions [6].

Some fractionation of the wavefront is apparent, and similar fractionation appears in the activation maps produced by the upscaled models (Figure 1). Both approaches




Figure 1. Under conditions of diffuse fibrosis (black pixels), to which traditional (volume averaged) homogenisation is suited, graph-based upscaling offers comparable performance. Brightness of colour indicates lateness of activation, with teal regions those that were not activated during the fixed simulation time. Stimulus is provided in a small region surrounding the centre of the left boundary.




Figure 2. Relative error values in activation times demonstrate the comparable performance of traditional volume averaged homogenisation and graph-based upscaling, on a problem where the former is appropriate.

perform about equally as well as each other in this scenario, and can offer significant reductions in the number of nodes without incurring too much inaccuracy (Figure 2). This scenario is well suited to traditional homogenisation, as the positioning of obstacles is approximately homogeneous when considered on a larger spatial scale, outside of the increasing frequency of obstruction along the horizontal direction. Even in such a scenario however, graph-based upscaling still remains competitive.

3.2. A Pernicious Example

The advantages of graph-based upscaling emerge when considering patterns of fibrosis that create more complex structures. Here we consider an example constructed to deliberately challenge traditional homogenisation, somewhat representative of intricate paths of activation taken through even small regions of fibrotic tissue [3]. Again for numerical reasons, we use a small $2.5 \text{ mm} \times 2.5 \text{ mm}$ patch of tissue. The problem is visualised in Figure 3.

The linear boundary conditions chosen for traditional homogenisation here (following their strengths observed in [6]) allow excitation to pass through the non-conductive spiral arms, resulting in a greatly underestimated transit time for excitation passing through the structure. Tests



Figure 3. Graph-based homogenisation creates upscaled problems that respect the effects of intricate patterns of fibrotic obstruction (in black). Colours are as in Figure 1, with the stimulus located similarly. Traditional (volume averaged) homogenisation produces an upscaled model that faintly represents the spiral structure in its activation map, but obstacles are leaky and results wholly incorrect.

with other boundary conditions for traditional homogenisation also proved incapable of recapturing the complex pattern of activation (results not shown). On the other hand, our graph-based approach completely turns off elements that would otherwise cause excitation to flow through the spiral walls. As such, the approach can be viewed as a way to intelligently re-mesh domains in a way that ensures their fundamental conductive property is preserved. The method is not simply a re-meshing, however, as the effective diffusivity tensors selected also represent the impacts of the small obstacles scattered throughout the domain.

We do observe that if graph-based upscaling is pushed too far, no conductive elements within the spiral-shaped channel are found and the whole structure is rendered non-conductive in the upscaled model. Possibly, further refining the method that selects which finescale nodes are brought into the upscaled model could avoid this issue. That said, the upscaled model pictured in Figure 3 already represents a ~ 65 fold reduction in nodecount and thus very considerable speed up of monodomain model simulation.

4. Conclusions

In this work we have presented our latest efforts towards a method for representing the effects of very finely-structured obstructions due to cardiac fibrosis in upscaled cardiac electrophysiology simulations. These techniques are critically important for mechanistically understanding the complex interplay of different impacts on signalling in cardiac fibrosis [3] that manifest in the disease’s deleterious effect on heart function. By combining the concept of graph-based calculation of effective tensors to avoid questionable assumptions made by traditional homogenisation, and the volume-averaged monodomain formulation that can better capture source/sink unbalance, we have created and demonstrated a robust type of upscaling for cardiac electrophysiology. We hope to continue this work by applying the method in the context of three-dimensional

simulations on realistic cardiac anatomies, and/or using histology-inspired arrangements of obstruction.

Acknowledgments

The authors would like to acknowledge the Australian Research Council (ARC), who supported this work through the ARC Centre of Excellence for Mathematical and Statistical Frontiers (grant number CE140100049).

References

- [1] Keener J, Sneyd J. *Mathematical Physiology I: Cellular Physiology*. Springer-Verlag New York, 2009.
- [2] Costa CM, Silva PAA, dos Santos RW. Mind the gap: A semicontinuum model for discrete electrical propagation in cardiac tissue. *IEEE Trans Biomed Eng* 2016;63:765–774.
- [3] Nguyen TP, Qu Z, Weiss JN. Cardiac fibrosis and arrhythmogenesis: the road to repair is paved with perils. *J Mol Cell Cardiol* 2014;70:83–91.
- [4] Davidović A, Coudière Y, Bourgault Y. Image-based modeling of the heterogeneity of propagation of the cardiac action potential. example of rat heart high resolution mri. In Pop M, Wright GA (eds.), *Functional Imaging and Modelling of the Heart*. Cham: Springer International Publishing. ISBN 978-3-319-59448-4, 2017; 260–270.
- [5] Gokhale Tanmay A, Asfour H, Verma S, Bursac N, Henriquez CS. Microheterogeneity-induced conduction slowing and wavefront collisions govern macroscopic conduction behaviour: A computational and experimental study. *PLoS Comp Biol* 2018;14:e1006276.
- [6] Lawson BAJ, dos Santos RW, Turner IW, Bueno-Orovio A, Burrage P, Burrage K. Homogenisation for the monodomain model in the presence of microscopic fibrotic structures. *arXiv* 2020;2012:05527.
- [7] Potse M, Dubé B, Richer J, Vinet A, Gulrajani RM. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Trans Biomed Eng* 2006;53:2425–2435.
- [8] Ten Tusscher KHWJ, Panfilov AV. Cell model for efficient simulation of wave propagation in human ventricular tissue under normal and pathological conditions. *Phys Med Biol* 2006;51:6141–6156.
- [9] Wallman M, Smith NP, Rodriguez B. A comparative study of graph-based, eikonal and monodomain simulations for the estimation of cardiac activation times. *IEEE Trans Biomed Eng* 2012;59(6):1739–1748.
- [10] Perego M, Veneziani A. An efficient generalization of the Rush-Larsen method for solving electro-physiology membrane equations. *Electron Trans Numer Anal* 2009;35:234–256.

Address for correspondence:

Megan Farquhar
 2 George Street, BRISBANE QLD 4000, Australia
 me.farquhar@qut.edu.au