Impacts of Boundary Conforming Meshes on Electrical Cardiac Simulation

Darrell J. Swenson[1,2], Joshua A. Levine[2], Jess D. Tate[1,2], Ross T. Whitaker[2,3], and Rob S. MacLeod[1,2]

[1] Department of Bioengineering, University of Utah, Salt Lake City, UT, USA
[2] Scientific Computing and Imaging Institute, Salt Lake City, UT, USA
[3] Department of Computer Science, University of Utah, Salt Lake City, UT, USA

Summary. Computational simulation has become an indispensable tool in the study of both basic mechanisms and pathophysiology of all forms of cardiac electrical activity. Such simulations depend heavily on geometric models that are either realistic or even patient specific. These models consist of a connected mesh of sometimes millions of polygonal elements that must capture the complex external shapes and internal boundaries among regions of the heart. The resulting meshes can be *non-conforming*, *i.e.,* they have element faces that fail to align with the tangents of the surfaces or boundaries and consequently the elements are a poor approximation of these smooth surfaces and boundaries. We hypothesize that such jagged, non-conforming meshes, which are often preferred, as they are easier to create, produce local artifactual concentrations of current that lead to simulation errors large enough to distort the resulting potential fields and generate misleading results. We tested this hypothesis on two types of numerical approximation used in bioelectric simulations: bidomain, and reaction-diffusion bidomain. Comparison with gold standard results for the monodomain and bidomain simulations showed that errors within a few elements (3-5) of the surface could be as large as 10-32%. The root mean squared error over the entire mesh was more modest, ranging from 1-6%. In the case of reaction diffusion simulations, by contrast, such meshing errors accounted for only an insignificant component of overall simulation uncertainty. These findings lead to the conclusion that while non-conforming meshes are certainly less costly to produce, their use can result in substantial local errors that depend highly on the specific problem of interest and the numerical approximation approach.

1 Introduction

Recent improvements in computing continue to fuel a rapid increase in the use of patient specific models to help predict and study disease [17, 30]; however, technical challenges impede achieving the full potential of this technology. One example of rapid progress is in the field of electrocardiology, in which bioelectrical activity can be modeled, both within the active myocardium and

through the surrounding passive volume conductor, using efficient implementations of the mesh-based computational strategies of finite and boundary elements. While the use of such models in clinical cardiology settings is growing, for example, in such domains as ventricular arrhythmias [32], implantable cardioverter defibrillator (ICD) placement [29, 13], and detecting atrial activation sequences [19, 7], these approaches have yet to achieve widespread clinical application.

Each stage of the typical simulation pipeline [17] presents modeling decisions that trade off fidelity to reality against computational cost. The focus of this study was to evaluate the appropriateness of boundary conforming meshes in the unique context of computational cardiac electrophysiology. While the conclusions are useful to the biomedical community, they also provide application driven criteria that the meshing community can use to improve meshing algorithms.

Specific to computational cardiology are the needs to capture the complex shape of the heart including many small features such as papillary muscle and endocardial surface irregularities or myocardial walls that can be as thin as 2 mm as well as large features like the thorax and its internal structures, *e.g.,* bone, lungs, muscle, fat, and other soft tissues. Representing these tissues in a computer model is challenging due to intricate or small external features, but also due to the complex internal interfaces formed as two or more tissues (generically, *materials*) share points, edges and surfaces. When only two materials are involved in the segmentation, the interface between them is by definition a manifold surface *i.e.,* at least locally it resembles the surface of a sphere in that it separates two regions. However, when three materials interface, the interfaces of each pair of materials come together in a "T"-shaped, non-manifold configuration. And if more than three materials come together in the same region, the complexity of their interfaces grows, creating more sophisticated non-manifold interfaces. These types of complex interfaces create special challenges for meshing algorithms, which typically assume only two materials. To bypass this difficulty, some meshing techniques simply ignore or approximate such complex internal material boundaries in a way that preserves simple topology. Such methods are "non-conforming" because the vertices of the mesh are not forced to align with the material interfaces. One consequence of non-conforming approaches to meshing is that each element is assigned tissue properties based only on the location of its centroid (or some other measure of its center of mass) relative to but not aligned with the original material boundary, which creates a more or less jagged representation of the actual interface.

Mesh refinement strategies, or adaptive meshing, is often used to increase mesh resolution in order to better approximate the boundary. This technique reduces the Hausdorff distance between surfaces but does not necessarily improve the jagged representations of smooth boundaries. Small elements may not deviate as far from the intended surface, but the tangents of the faces or edges do not align with or approximate the tangents of the intended surface.

One result is that the total surface area of the approximated boundary may be substantially larger than the intended boundary. There will also be sharp corners and edges introduced into the surface not intended by the boundary definition. Despite these limitations, mesh refinement remains the primary way in which non-conforming boundaries are handled. It is well understood in finite element models that increasing the mesh resolution is necessary for simulating fields with high spacial gradients. However, what is unclear is the extent to which jagged surfaces affect the simulation results particularly when coupled with mesh refinement.

In some instances the properties between two materials are blurred at the material transitions, primarily to improve the the convergence of the linear system. This has also been used when applying material properties to non-conforming meshes, though not necessarily improving the simulation accuracy. While blurred boundaries may be appropriate in some settings there are many instances in which they are not [33]. In particular many computational electrophysiology simulations require a sharp change in material properties to simulate the intended physics. Blurring the boundary is also problematic when setting boundary conditions where a region may be set to a constant voltage or current, such as an electrode. In many cardiac electrophysiology simulations, different ionic models are defined based on the tissue type, where the ionic model becomes a source or a sink during specific time instances of the simulation. The transition of ionic models should be based on physiological understanding of the physiology being modeled rather than a correction for a poorly represented boundary.

The guiding hypothesis for this study was that conforming boundaries produce significantly more accurate results at much lower resolutions compared to non-conforming boundaries, and that the errors resulting from non-conforming meshes are large enough to alter the interpretation of simulation results. In addition to verifying this hypothesis, we were able to show that simply refining the mesh resolution has only limited effect on removing artifacts since the non-conforming surfaces never truly become smooth even at high resolutions. This hypothesis was tested on two common numerical methods in cardiac bioelectric simulations: single time point bidomain and reaction-diffusion bidomain. Errors from non-conforming meshes were particularly evident in simulations of pathophysiology that were highly dependent on internal tissue boundaries, such as modeling the effects of discrete regions of damaged heart tissue. The results indicate that while creating conforming meshes is more challenging and time consuming, it is often necessary to capture local electrical behavior in the heart such as regional myocardial ischemia or focal ectopic activity. At the same time, the complexity and computational cost of conforming meshing may not be justified when simulating features that are more global in nature, such as the activation time of a multi-material model of the whole heart.

2 Related Work

The goal of all simulations of cardiac electrophysiology at the tissue and whole-heart scale is to incorporate relevant behavior from the smaller scales, *i.e.,* cellular, membrane, and molecular, into a tractable formulation that can capture the meaningful aspects of electrical activity observable at this scale. It is impossible to predict whole-heart behavior The two simulation approaches described below differ in how they carry out this approximation and which assumptions they include. Because of this diversity of approach, it is reasonable to assume that the impact of mesh structure may affect each method differently. In order to explore the role of mesh structure on at least a small sampling of applications, we also selected two timely simulation problems from cardiac electrophysiology.

2.1 Modeling Acute Myocardial Ischemia - Single Time Point Bidomain

The bidomain represents two continuous volume conductor domains that provide independent pathways for current to flow. The two domains are coupled using differential equations that capture the associated physics. In the heart, current can flow through the intracellular space or the extracellular space which are coupled by differential equations that model the behavior of the cell membrane. In the simplest form of the bidomain, a subset of potentials is known, *e.g.,* transmembrane potentials, and the goal becomes to compute the remaining unknowns, *e.g.,* extracellular potentials.

The additional complexity of the bidomain is required to capture electrical activity in the active domain of myocardial tissue, although with some simplification such as static assumptions. Elevations and depressions of what is known as the ST segment are commonly used markers in the electrocardiogram (ECG) that are indicators of myocardial perfusion deficits leading to what is known as ischemia or infarction. Such conditions arise relatively slowly in the heart and so represent relatively static electrophysiological conditions suitable for this simplified simulation approach. Motivation for simulation comes from the fact that while ST segment shifts are very commonly used clinical indicators, their underlying mechanisms are poorly understood, especially when the region of tissue involved does not span the full thickness of the heart [16]. By carrying out the simulation under static assumptions, it is possible to select a single time instance during the ST segment and supply approximate transmembrane potentials to evaluate the effects of variations in location and extent of the ischemic tissue [9, 12].

2.2 Modeling Cardiac Activation and Repolarization - Reaction Diffusion Bidomain

In the more complete form of the bidomain a cardiac myocyte is excited, and current begins to flow across the cell membrane and diffuses in the extracel-

lular and intracellular spaces, subsequently exciting neighboring cells. This chain reaction produces a wave that propagates through the myocardium. This process can be modeled as a reaction diffusion bidomain simulation that consists of two steps. First there is a reaction, in this case, current crossing the cell membrane modeled as a set of differential equations. In the subsequent diffusion step the ions move more or less freely through each continuous domain, which is approximated by a partial differential equation. The whole system is solved iteratively over time to simulate the entire excitation wave that occurs.

A full bidomain approach is required in simulation of the heart when there is a need to incorporate both spatial and temporal features, *i.e.,* when the electrical state of the heart depends on the time evolution of propagating waves. This generality of approach also makes the time evolving bidomain suitable for capturing a full range of cardiac behavior, with special focus on abnormal heart rhythms or "cardiac arrhythmias" [20, 11]. The bidomain approach includes details of myocardial structure, electrical conductivity in both intracellular and extracellular domain, and the behavior of the cellular membrane, but is also capable of high efficiency so that it has become the dominant approach for carrying out simulations across enormous ranges of scale and for subject specific settings [20].

2.3 Boundary Conforming Meshes

One feature of all these applications of simulation, which is also common to a wide range of disciplines, is the need to create subject specific geometric models in the form of nodes and polygons that support the application of numerical methods. Mesh generation is a challenging problem defined as the problem of creating in some automated way a tiling/tessellation (a *mesh*) of some simulation domain using simple geometric shapes such as tetrahedra, pyramids, or hexahedra. Popular approaches for mesh generation consist of tetrahedral meshes constructed with stenciling [14], optimization of vertex locations [1, 18], and iterative refinement [5, 22].

To represent geometric structures faithfully, one goal of mesh generation is to capture surfaces, both within the volume and on the outer surfaces, with great fidelity through careful choice of node density and location and element orientation [3, 6, 18]. The resulting tradeoff in mesh construction is between conforming to the original internal boundaries and preserving the quality of the resulting mesh elements, as well as the often considerable computational cost in the algorithms that identify and preserve their shape [25]. While somewhat ambiguously defined, the quality of mesh elements is determined by factors like the aspect ratios failing to meet some prescribed measurement.

In biological models, meshes are often made from segmented images which are discrete voxel representations of continuous and usually smooth geometry of organs and parts of the body captured at the resolution of the imaging modality. The simplest of all reconstructions would be to assume each voxel

represented a hexahedral, resulting in a mesh that matched the original image but with surfaces that had a ragged or stair-stepped profile. However, in most biological tissues the true boundaries are not stair-stepped but are smooth. Consequently, creating smooth surfaces is a goal of most image based meshing algorithms such as marching cubes, Laplacian smoothing, curvature limiting smoothing, or spline based surface fitting.

3 Methods

In both bioelectric simulations non-conforming and conforming meshes created with a range of resolutions. The gold standard for the simulations was a conforming mesh created at the highest resolution that was tractable and error metrics quantified both global and local variation. A further test was based on changes not only on the quantitative results, but also the interpretation of those results in the context of the simulation goals such as the change in predicted shock value in the defibrillation models, change in the epicardial elevations and depressions of the ischemia models, and change in activation patterns and times for the reaction diffusion models.

3.1 Mesh Generation

As described in the introduction our goal was not to evaluate the quality of fit between the smoothed surface and the original anatomy, but to determine the extent of the error introduced into some sample simulations by representing an anatomically smooth surface as a non-boundary fitting, or non-conforming volume mesh. Creating a ground truth is challenging because of a lack of both analytical representations of realistic anatomy and analytical solutions to simulations of interest. Instead, we first built a very high resolution (0.18 mm average element edge length), smooth representation of each geometry and generated all meshes from this.

Both the conforming and non-conforming meshes were then based upon the smoothed reference geometry. There are many different meshing techniques that are considered non-conforming; however, in this study we focused on unstructured tetrahedral meshes that conformed to the exterior surface but not the internal boundaries. The non-conforming meshes were created by filling the volume defined by the exterior conforming surface with Delaunay tetrahedra at prescribed resolutions, but ignoring the material interfaces. After the mesh was generated, the elements were grouped by material type based on the location of each element relative to the reference surfaces. In contrast, the conforming meshes were generated in a way that preserved the internal boundaries of the reference geometry.

Here we leverage a variational meshing algorithm, BioMesh3D [21], to create conforming external and internal material interfaces. BioMesh3D creates meshes by first distributing particles on the surfaces to which the mesh must

conform [18], using variation to drive an energy minimization process to find ideal placements for the particles. These particles are then treated as mesh vertices, and the remaining mesh elements are constructed using a three-dimensional Delaunay triangulation implemented in the open source software TetGen [23]. For this study, we created meshes that both conformed to, and ignored various shapes of internal boundaries in a series of simulations from cardiac electrophysiology and examined carefully the resulting errors. Our goal was to establish the extent and the nature of the improvements that result from creating meshes that conform faithfully to internal tissue boundaries, at least in this domain of bioelectric field simulation.

To formally describe the conforming meshes, we define a volumetric mesh $M = (V, E, F, C)$ to be a collection of vertices V, edges E, faces F, and volumetric cells C and let Σ be a surface embedded in three dimensions. In this context, M is called *boundary conforming* if it explicitly represents Σ with the following conditions:

1. A subset of mesh vertices, V_Σ lie on Σ.
2. A set of mesh faces (in our context, triangular faces of tetrahedra) approximate Σ. In particular, these triangles will have their vertices as part of V_Σ and the plane of each of these triangles will be an approximation of the tangent space of Σ at each vertex of the triangle.

While there are stricter definitions for conforming meshes, for example those that require a homeomorphism between Σ and the triangular faces which approximate it [6], applying such requirements would be unrealistic since the true Σ is not known, but rather only approximated from the image data. Thus, we think of conforming meshing as both requiring a geometric approximation in terms of distances to the surface boundaries as well as requiring the tangent space of the surfaces to be well approximated by the linear elements. Within this framework of conforming meshing, the specific geometric model and the numerical approach was different for each of the three examples, as described below.

3.2 Static Simulation of Acute Myocardial Ischemia

Geometric Model

The geometric model used for the simulation of electrocardiographic fields that arise during acute myocardial ischemia consisted of the heart and blood that were derived from an MRI of the same canine heart as used in the associated experimental study. The goal of these studies was to determine from measurement of intramyocardial potentials the distribution of ischemic regions in the heart and then simulate from them the heart surface potentials and eventually the ECG [2, 27]. In addition to representing the blood and the myocardium, these models contained one or more ischemic regions modeled as seen in Figure 1. The shape and location of these regions were derived

from the three-dimensional distributions of extracellular potentials captured during ischemic phases of the experiment. Simulations based on cardiac tissue also require a description of local muscle fiber direction, which was acquired using diffusion tensor MRI and applied to each mesh as anisotropic conductivity tensors assigned to each element. The magnitude of these tensors was assigned based on microdomain simulations of ischemia by Stinstra *et al.*[26].

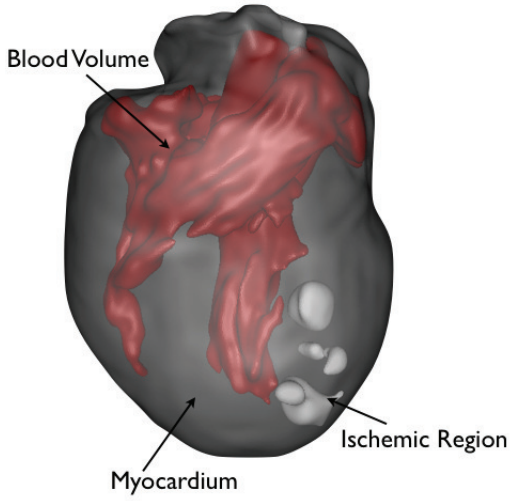

Fig. 1. Model used to simulate myocardial ischemia. The model includes three regions: normal myocardium, blood, and ischemic myocardium and the figure shows the high quality visualization possible from meshes that conform to the material boundaries.

Simulation

To simulate the electric potential on the surface of the heart from regions of acute ischemia within the ventricles, we used a simplified, static version of the bidomain formulation described previously [9, 12, 28], which can be written as

$$\nabla \cdot (\sigma_i + \sigma_e)\nabla \Phi_e = -\nabla \cdot \sigma_i \nabla \Phi_m, \tag{1}$$

where Φ_e and Φ_m are the extracellular and transmembrane potentials, and σ_e and σ_i are the extracellular and intracellular conductivities, respectively. The ischemic region was modeled as having reduced action potential amplitude, expressed as a 30 mV potential difference between ischemic and healthy tissue during the plateau phase of the action potential (ST segment of the ECG). The blood was modeled by setting the intracellular conductivity to zero and the

extracellular conductivity to 4 times that of the myocardium [26]. Neumann boundary conditions were applied to the outer surface of the heart, which simulated the open chest conditions of the experiments.

3.3 Propagation Modeling

Geometric Model

The simulation of propagation made use of the same geometry of the heart, blood, and ischemic regions as the ischemia model, but made use of a reaction diffusion simulation to predict the spread of cardiac electrical activity. Such a simulation requires a much more refined spatial resolution so that a full scale model of the heart would include tens of millions of nodes, beyond the scope of most computational resources. To reduce the size of the models to produce tractable simulations, we scaled the size of the geometry by a factor of 0.4, reducing the number of nodes by an order of magnitude.

The number of elements and nodes required to achieve a particular mesh resolution depends on the total surface area and volume of the geometry. For example, if the edge length of a hexahedral is cut in half, the result is eight new hexahedral and thus a cubic relationship. Tetrahedra are similar, in that there also exists a cubic relationship but they may not be precisely cubic if the tetrahedral are calculated by re-meshing rather than subdividing. Because hearts naturally vary in size and shape, and in our case due to a scaling factor, the total number of elements and nodes make very poor descriptors of mesh resolution. We have chosen instead to report the results in terms of average edge length which has a cubic relationship with the number of elements.

Simulation

To create realistic simulations of the spread of excitation in the heart, the full version of the time evolving bidomain [8] is the most commonly used approximation approach and has been implemented in widely available software (*e.g.,* CHASTE [4] and CARP [31]). This formulation also represents a reaction-diffusion system, expressed as the following system of coupled equations

$$\nabla \cdot (\sigma_i \nabla \Phi_i) = \beta I_m \tag{2}$$

$$\nabla \cdot (\sigma_e \nabla \Phi_e) = -\beta I_m \tag{3}$$

where Φ_e, Φ_m, σ_e, and σ_i are defined as above and I_m is the ionic current flowing through the membrane. The transmembrane current was solved using the Faber Rudy model [15]. To predict changes in the spread of excitation in response to ischemia, we modified the ionic concentrations and the ionic currents according to Jie *et al.* [11].

3.4 Evaluation of Error

The error was calculated by comparing each simulation to that generated using the high resolution reference conforming mesh. The global error was determined by first interpolating onto a sampling grid with a spacing of 0.3 mm using standard Barycentric coordinates to linearly interpolate the data. Then the root mean squared error was calculated over the entire sampling grid in comparison to the reference simulation. The maximum error was also reported for each simulation, representing a more local error measurement. These errors were recorded through a range of mesh resolutions as defined by edge length. In addition to the RMS and the maximum error, at each material boundary the mean and maximum error were reported as a function of distance. In the case of the reaction diffusion simulations, the RMS error was calculated at each time step and then averaged for the whole time series.

The errors in the simulations were also evaluated as to their possible influence on the interpretation of the results. These parameters included change in the epicardial elevations and depressions of the ischemia models and change in activation patterns and times for the reaction diffusion models.

4 Results

4.1 Ischemia ST segment - Non-time series Bidomain

Electrical potentials for the conforming meshes were smoothly distributed throughout the myocardium and on the epicardial surface Figure 2. The non-conforming meshes had regions that looked more discontinues and resulted in small patches of elevations and depressions not seen in the results from conforming meshes. These patches created errors even on the epicardial surface Figure 2 that were as large as 2 mV which is 30% of the largest elevations seen on the surface. Errors of this magnitude are large enough to change whether an elevation or depression is considered ischemic, or just natural variation. The location of the errors is collocated with the region being analyzed. Further away from the non-conforming surfaces bounding the ischemic zone, both meshes produced smooth distributions of voltage, however, the voltage maps were altered even at a distance when there were significant current concentrations near the ischemic zone boundaries, a situation to be expected as the so called “leakage” currents from ischemia arise at these boundaries.

Figure 3 shows both root mean squared (RMS) errors and the maximum errors as functions of edge length. The potentials in this simulation ranged from 30 mV to -10 mV making the maximum errors of 10 mV very significant. Regions with the largest errors were located along to border between the ischemic and healthy tissue. Due to the proximity of the ischemic region to the epicardial surface (as low as 3-4 mm), many of the local errors appeared on the epicardial surface directly above the ischemic region. The errors resulted in small elevations of 4 mV and depressions of 5 mV in the coarser

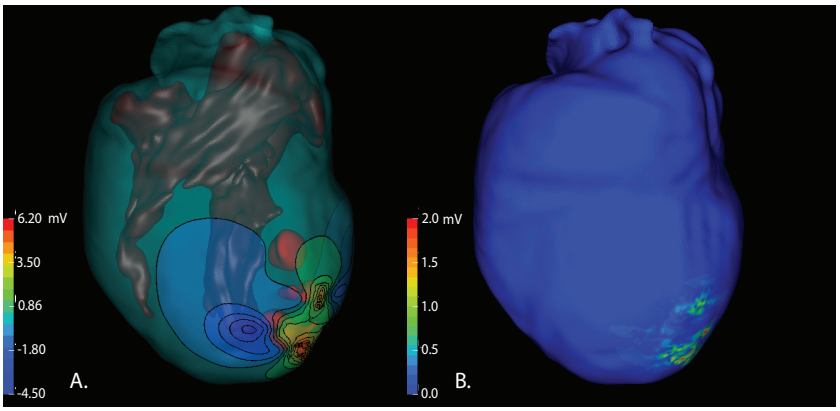

Fig. 2. Solved bidomain simulation of epicardial potentials during acute ischemia with the location of maximal errors. A: Is a voltage map on the epicardial surface due to injury currents of an ischemic region within the myocardium simulated with a conforming mesh while B: shows the location an magnitude of errors that arise when a non-conforming mesh is used.

non-conforming meshes. These errors improved with mesh resolution, but persisted until very high resolution meshes of 0.8 mm edge length or smaller. The maximum errors were substantial for both types of meshes at very low resolution. However, the error in the conforming meshes decreased very rapidly with improved resolution while there were only slight improvements in the non-conforming meshes.

Other non-conforming boundaries were much smaller sources of error than those surrounding the ischemic regions. The errors at the blood boundaries caused voltage shifts of up to 5% over the gold standard. The maximum errors at these boundaries showed slight improvements with increased mesh resolution, reducing the error to just over 3%. Current densities at blood boundaries were much smaller than those over the ischemic zones, consistent with the observation that the largest errors arise near regions of concentrated sources.

Globally, the conforming meshes produced an RMS error of at least 0.2 mV better than non-conforming meshes over the entire range of mesh resolutions which did not improve as the meshes became more refined. A second way to consider these results is that a non-conforming mesh with an average edge length of 0.65 mm would give the same accuracy as a conforming mesh with a larger edge length of 1 mm. For such models, this improvement would represent the difference between 1.7 million and 0.49 million nodes.

4.2 Reaction Diffusion Bidomain

Figure 4 shows the RMS error of potentials during activation as a function of model resolution for the spread of activation of a heart beat under conditions

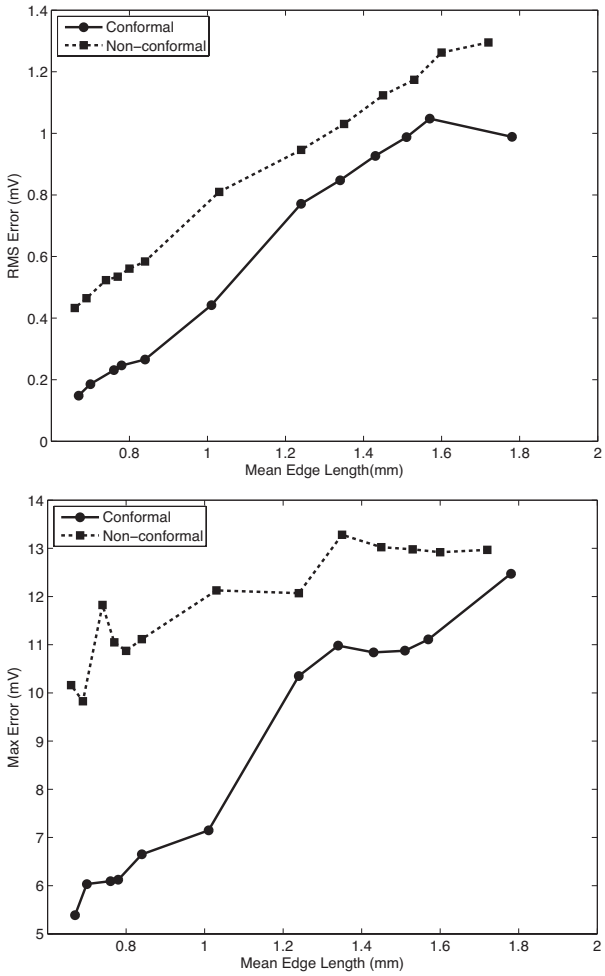

Fig. 3. RMS and max error in conforming and non-conforming ischemia models. The graph on the left indicates the RMS error for both a set of conforming and non-conforming meshes at varied resolutions. The graph on right indicates the maximum error seen in the same set of simulations.

of localized ischemia and suggests that there is no clear difference in performance between conforming and non-conforming meshes for this problem. The maximal errors indicated more substantial sensitivity to mesh choice than RMS as they showed values as large as 18 mV, which all occurred along the activation front, whereas errors along the material boundaries were an order of magnitude smaller, much like the errors found in the static bidomain simulations found in Figure 3, suggesting once again that regions of high source strength respond most strongly to choice of mesh type.

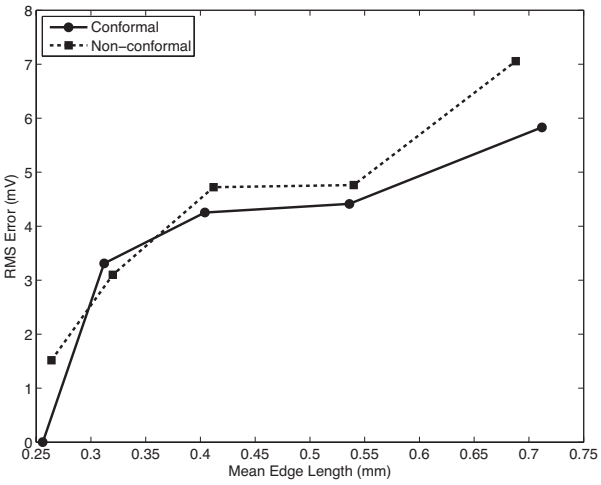

Fig. 4. RMS error from reaction diffusion bidomain simulation of the spread of activation. The RMS error over the entire activation of the ventricles comparing conforming with non-conforming meshes as a function of mesh resolution (edge length).

Figure 5 shows the first 200 ms of the RMS curves of the simulated electrograms for two different mesh resolutions. The RMS curves produced by the lower resolution meshes had decreased slopes and increased widths of the activation wave form, which is indicative of conduction slowing. This trend held through all mesh resolutions, confirming that the propagation velocity does change across mesh resolutions, but also that there is very little difference between the conforming and non-conforming meshes when they had similar resolutions. The activation times computed for both mesh types were very similar, within 3% demonstrating that there was no clear advantage to one or the other. These findings suggests that the selection of mesh type plays no substantial role in simulation accuracy.

5 Discussion

The hypothesis that non-conforming meshes could lead to artifacts that significantly affect the results of simulations was supported in the static bidomain studies, but less convincing for the spread of activation in the reaction-diffusion simulation using the full bidomain. In the first case, overall errors were higher when using non-conforming meshes. Moreover, the errors were largest in regions with large local sources, which were the tissues that were of direct interest to users of the simulations. For the static simulation of ischemia, the errors related to meshing strategy were large, 30%, on the outer

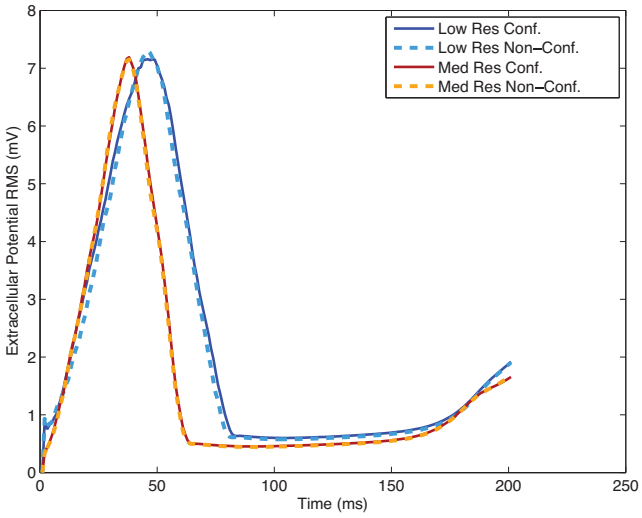

Fig. 5. Extracellular RMS electrograms from a reaction diffusion bidomain simulation for each time step. This figure shows for two different mesh resolutions, the RMS error in voltage as a function of time for beats computed using both conforming and non-conforming meshes.

(epicardial) surface centered over a non-transmural ischemic region. This variation changed the magnitude of several of the elevations such that they would not meet the inclusion criteria that distinguishes between natural variation in potentials and elevations arising from ischemia.

In contrast to the reaction diffusion bidomain simulation did not show a clear preference for either mesh type, at least in part due to the strong dependency of the simulations on parameters common to both conforming and non-conforming meshes, primarily the spatial resolution in the region of the activation wave front. The transition from resting to fully stimulated cells at the wave front extends over only approximately 1 mm so that simulations of propagation require substantially sub-millimeter mesh resolution. At this fine scale, the differences between conforming meshes driven by fixed anatomical surfaces and non-conforming meshes based on imaging orientation disappear and so do not play a role in simulation error. A further critical structural factor of cardiac tissue is the anisotropic nature of current flow along the long axis of heart cells and the fibers they form. While it is possible to imagine a conforming meshing strategy based on fiber orientation, our conforming meshes sought instead to respect larger scale anatomical boundaries and so would be unlikely to perform any better than a non-conforming mesh with regard to anisotropy.

From these findings we make the following observations which can be used to determine the appropriate structure of the underlying meshes:

1. Sources and sinks should be meshed with boundary conforming meshes. This may be obvious in the traditional case of electrodes. However, in bidomain simulations regions regions with differing ionc models act as ether a source or sink at particular instances of the cardiac cycle. For this reason each tissue represented by a distinct ionc model should be modeled conformally.
2. Calculation of activation times were heavily dependent on element size, much more so that boundary conforming mesh elements. In this simulation, uniform element sizes greatly improves simulation results even as it adversely affects the fidelity of the mesh to the material interfaces.
3. Conformal meshes should be used when there is a close proximity of the material interfaces to the areas of the models most relevant for subsequent analysis and interpretation of the results. For example, in the simulation of the bioelectric effects of myocardial ischemia, the epicardial surface was a place where measurements and analysis of electric potentials occurs in experiments and occasionally even clinical practice [24]. A lack of fidelity nearby surfaces, as in non-conforming meshes, could be expected to lead to highly relevant localized errors, which our results were able to substantiate. In contrast, in the case of simulating the spread of activation in a bidomain model of the ventricles, there was little relationship between anatomical structure and the wave front and hence little benefit to a conforming meshing approach based on gross anatomical elements.

Algorithmic and computational costs are key drivers in any discussion of meshing (or modeling) strategies and they must be part of the interpretation of our results. We found no case of a conforming mesh of similar resolution producing *worse* results than a non-conforming mesh and thus it would be tempting to propose conforming meshing as a general strategy. However, in many cases, non-conforming meshes can take seconds to generate compared to hours for comparable conforming meshes of the same size. Not conforming to the boundaries allows for computationally efficient meshing strategies such as regular grids to be implemented. In the case of the ischemia model, an edge length of 0.65 mm in the non-conforming mesh was equivalent to a conforming mesh with an edge length of 1 mm which would reduce the number of nodes in the model from 1.7 million to 0.49 million. Computational speedup depends on specifications of the computer used, however, we observed improvements that were at least proportional to the reduction in size of the model and often much greater, particularly for the larger meshes. The smaller meshes also decreased the time and memory needed to process and visualize the computed solutions.

An important finding of this study was that non-conforming meshes performed just as well as the conforming meshes in the reaction-diffusion bidomain simulations. Reaction-diffusion simulations, and the bidomain approach

specifically, already represent a simplification, typically a mathematical homogenization, that seek to achieve the efficiency necessary to carry out simulations of large structures, for example, the whole heart rather than just a small, presumably representative block of tissue. They are intrinsically a compromise driven by computational resources so that any improvement in memory usage or calculations is critically important. Non-conforming meshes are always faster to create and to refine, and often show advantages in numerical approximations. Of specific current interest are their potential advantages in the application of parallel algorithms, either CPU or GPU (graphical processing unit) based [10]. It is somewhat reassuring that non-conforming meshes performed very well in our comparisons with conforming meshes of similar size in this setting.

One observation that can tie our findings to the more traditional evaluation based on mesh quality metrics is that element shape tends to improve as the mesh resolution becomes finer because the elements better approximate regions of high curvature. To reduce the impact of errors associated with element quality the scaled Jacobian, was monitored during the mesh creation so that all poorly shaped elements could be improved. In an effort to reduce the differences in element quality between conforming and non-conforming meshes we used the same Delaunay based meshing algorithm to create both and produced similar mesh element qualities at each resolution. Additionally, local mesh refinement schemes adversely affect element quality as the elements have to transition from small elements to large ones. For this reason we do not apply mesh refinement locally, only globally. Local refinement would help reduce the number of elements needed in the models for both the conforming and non-conforming meshes, but would adversely affect element quality.

In conclusion, these experiments suggest a complex and application dependent role of mesh structure on simulation accuracy. This study has demonstrated settings in which conforming meshes outperformed non-conforming meshes of similar size but at least one major application domain in which there was no difference in simulation accuracy. Simulation scientists are faced with many decisions in creating functional pipelines and our findings suggest that careful selection of mesh generation approaches, perhaps preceded by numerical experiments, may be necessary in order to optimize these decisions for the available computational resources and desired simulation scope and accuracy. It appears unlikely that selecting the closest software to hand, or even the algorithms that have proven utility in one setting, will yield the best compromise in a new application domain.

References

1. Pierre Alliez, David Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun. Variational tetrahedral meshing. *ACM Transactions on Graphics*, 24(3):617–625, 2005.

2. K.K. Aras, S. Shome, D.J. Swenson, J.G. Stinstra, and R.S. MacLeod. Electrocardiographic response of the heart to myocardial ischemia. In *Computers in Cardiology 2009*, pages 105–108, 2009.
3. Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005.
4. Rafel Bordas, Bruno Carpentieri, Giorgio Fotia, Fabio Maggio, Ross Nobes, Joe Pitt-Francis, and James Southern. Simulation of cardiac electrophysiology on next-generation high-performance computers. *Philos Transact A Math Phys Eng Sci*, 367(1895):1951–69, May 2009.
5. Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Tathagata Ray. Quality meshing for polyhedra with small angles. *Int. J. Comput. Geometry Appl.*, 15(4):421–461, 2005.
6. Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Tathagata Ray. Sampling and meshing a surface with guaranteed topology and geometry. *SIAM J. Comput.*, 37(4):1199–1227, 2007.
7. P.S. Cuculich, Y. Wang, B.D. Lindsay, M.N. Faddis, R.B. Schuessler, J.r. Damiano RJ, L. Li, and Y. Rudy. Noninvasive characterization of epicardial activation in humans with diverse atrial fibrillation patterns. *Circ.*, 122(14):1364–1372, Oct 2010.
8. C S Henriquez. Simulating the electrical behavior of cardiac tissue using the bidomain model. *Crit Rev Biomed Eng*, 21(1):1–77, 1993.
9. B. Hopenfeld, J.G. Stinstra, and R.S. MacLeod. The effect of conductivity on ST segment epicardial potentials arising from subendocardial ischemia. *Annal. Biomed. Eng.*, 33(6):751–763, 2005.
10. W.-K. Jeong and R.T. Whitaker. A fast iterative method for eikonal equations. *SIAM J. Scientific Computing*, 30(5):2512–2534, 2009.
11. Xiao Jie, Viatcheslav Gurev, and Natalia Trayanova. Mechanisms of mechanically induced spontaneous arrhythmias in acute regional ischemia. *Circ Res*, 106(1):185–92, Jan 2010.
12. PR Johnston and D Kilpatrick. The effect of conductivity values on ST segment shift in subendocardial ischaemia. *IEEE Trans Biomed. Eng.*, 50:150–158, 2003.
13. M. Jolley, J. Stinstra, J. Tate, S. Pieper, R. Macleod, L. Chu, P. Wang, and J.K. Triedman. Finite element modeling of subcutaneous implantable defibrillator electrodes in an adult torso. *Heart Rhythm. J.*, 7(5):692–698, May 2010.
14. François Labelle and Jonathan Richard Shewchuk. Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. *ACM Trans. Graph.*, 26:1–10, July 2007.
15. C.M. Lloyd, P.J. Hunter, and P.F. Nielsen. The cellml model repository. *Bioinformatics*, 24(18):2122–2123, September 2008.
16. R.S. MacLeod, S. Shome, J.G. Stinstra, B.B. Punske, and B. Hopenfeld. Mechanisms of ischemia-induced ST-segment changes. *J. Electrocardiol.*, 38:8–13, 2005.
17. R.S. Macleod, J.G. Stinstra, S. Lew, R.T. Whitaker, D.J. Swenson, M.J. Cole, J. Kruger, D.H. Brooks, and C.R. Johnson. Subject-specific, multiscale simulation of electrophysiology: a software pipeline for image-based models and application examples. *Phil. Trans. Royal Soc.*, 367(1896):2293–2310, Jun 2009.
18. Miriah D. Meyer, Ross T. Whitaker, Robert M. Kirby, Christian Ledergerber, and Hanspeter Pfister. Particle-based sampling and meshing of surfaces in multimaterial volumes. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1539–1546, 2008.

19. B. Pfeifer, G. Fischer, F. Hanser, M. Seger, C. Hintermuller, R. Modre-Osprian, T. Trieb, and B. Tilg. Atrial and ventricular myocardium extraction using model-based techniques. *Methods Inf Med*, 45(1):19–26, 2006.
20. G. Plank, L. Zhou, J.L. Greenstein, S. Cortassa, R.L. Winslow, B. ORourke, and N.A. Trayanova. From mitochondrial ion channels to arrhythmias in the heart: computational techniques to bridge the spatio-temporal scales. *Phil. Trans. Royal Soc.*, 366:3381–3409, 2008.
21. Scientific Computing and Imaging Institute (SCI). BioMesh3D: Quality Mesh Generator for Biomedical Applications. http://www.biomesh3d.org.
22. Jonathan Richard Shewchuk. Tetrahedral mesh generation by Delaunay refinement. In *Proceedings of the 14th Symposium on Computational Geometry*, pages 86–95, 1998.
23. Hang Si. TetGen: A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator. http://tetgen.berlios.de/.
24. E. Sosa, M. Scanavacca, A. D'Avila, F. Oliviera, and J.A.F Ramires. Nonsurgical transthoracic epicardial ablation to treat recurrent of ventricular tachycardia. *J. Am. Coll. Cardiol.*, 35(1):1442–1449, 2000.
25. Matthew L. Staten, Robert Kerr, Steven J. Owen, and Ted D. Blacker. Unconstrained paving and plastering: Progress update. In *Proceedings of the 15th International Meshing Roundatable*, pages 469–486, 2006.
26. J G Stinstra, S Shome, B Hopenfeld, and R S MacLeod. Modelling passive cardiac conductivity during ischaemia. *Med Biol Eng Comput*, 43(6):776–782, 2005.
27. D. Swenson, J. Stinstra, B. Burton, K. Aras, L. Healy, , and R.S. MacLeod. Evaluating the effects of border zone approximations with subject specific ischemia models. In Olaf Doessel and Wolfgang C. Schlegel, editors, *World Congress on Med. Phys. and Biomed. Eng.*, volume 25/IV, pages 1680–1683, Heidelberg, 2009. Springer.
28. D.J. Swenson, J.G. Stinstra, B.M. Burton, K. Aras, and R.S. MacLeod. The effect of finite element mesh quality on electrical bidomain simulations. In A. Murray, editor, *Computers in Cardiology 2010*, page (in press), 2010.
29. J.K. Triedman, M. Jolley, J. Stinstra, D.H. Brooks, and R.S. MacLeod. Predictive modeling of defibrillation using hexahedral and tetrahedral finite element models: recent advances. *J. Electrocardiol.*, 41(6):483–486, 2008.
30. F. Vadakkumpadan, H. Arevalo, A.J. Prassl, J. Chen, F. Kickinger, P. Kohl, G. Plank, and N.A. Trayanova. Image-based models of cardiac structure in health and disease. *WIREs Syst Biol Med*, 2(4):489–506, 2010. defibrillation, model building, meshing.
31. E. Vigmond, F. Vadakkumpadan, V. Gurev, H. Arevalo, M. Deo, G. Plank, and N.A. Trayanova. Towards predictive modelling of the electrophysiology of the heart. *J Exp. Physiol.*, 94(5):563–577, 2009.
32. Y. Wang, P.S. Cuculich, J. Zhang, K.A. Desouza, R. Vijayakumar, J. Chen, M.N. Faddis, B.D. Lindsay, T.W. Smith, and Y. Rudy. Noninvasive electroanatomic mapping of human ventricular arrhythmias with electrocardiographic imaging. *Sci Transl Med*, 3(98):98ra84, Aug 2011.
33. J. Xu and L.T. Zikatanov. On multigrid methods for generalized finite element methods. *Meshfree Methods for Partial Differential Equations*, pages 401–418, 2002.