Published in final edited form as:

IEEE Trans Biomed Eng. 2009 May ; 56(5): 1318. doi:10.1109/TBME.2009.2014243.

Automatically Generated, Anatomically Accurate Meshes for Cardiac Electrophysiology Problems

Anton J. Prassl¹, Ferdinand Kikinger², Helmut Ahammer³, Vicente Grau [Member, IEEE]⁴, Jürgen E. Schneider⁵, Ernst Hofer³, Edward J. Vigmond [Member, IEEE]⁶, Natalia A. Trayanova [Senior Member, IEEE]⁷, and Gernot Plank⁸

Anton J. Prassl: anton.prassl@medunigraz.at; Ferdinand Kikinger: ferdinand.kikinger@meshing.at; Helmut Ahammer: helmut.ahammer@medunigraz.at; Vicente Grau: vicente.grau@oerc.ox.ac.uk; Jürgen E. Schneider: sbhattac@well.ox.ac.uk; Ernst Hofer: ernst.hofer@meduni-graz.at; Edward J. Vigmond: vigmond@ucalgary.ca; Natalia A. Trayanova: ntrayanova@jhu.edu; Gernot Plank: gernot.plank@medunigraz.at

¹ Institute for Computational Medicine, Johns Hopkins University, Baltimore, MD 21218 USA. He is now with the Institute of Biophysics, Medical University of Graz, Graz A-8010, Austria

² CAE Software Solutions, Eggenburg 3730, Austria

³ Institute of Biophysics, Medical University of Graz, Graz 8010, Austria

⁴ Department of Engineering Science and Oxford e-Research Centre, University of Oxford, Oxford OX1 3PJ, U.K

⁵ Department of Cardiovascular Medicine, British Heart Foundation (BHF) Experimental Magnetic Resonance (MR) Unit, University of Oxford, Oxford OX3 7BN, U.K

⁶ Department of Electrical and Computer Engineering, University of Calgary, Calgary, AB T2N 1N4, Canada

⁷ Institute for Computational Medicine and the Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21218 USA

⁸ Institute for Computational Medicine, Johns Hopkins University, Baltimore, MD 21218 USA. He is now with the Institute of Biophysics, Medical University of Graz, Graz A-8010, Austria, and also with the University of Oxford, Oxford OX1 2JD, U.K

Abstract

Significant advancements in imaging technology and the dramatic increase in computer power over the last few years broke the ground for the construction of anatomically realistic models of the heart at an unprecedented level of detail. To effectively make use of high-resolution imaging datasets for modeling purposes, the imaged objects have to be discretized. This procedure is trivial for structured grids. However, to develop generally applicable heart models, unstructured grids are much preferable. In this study, a novel image-based unstructured mesh generation technique is proposed. It uses the dual mesh of an octree applied directly to segmented 3-D image stacks. The method produces conformal, boundary-fitted, and hexahedra-dominant meshes. The algorithm operates fully automatically with no requirements for interactivity and generates accurate volume-preserving representations of arbitrarily complex geometries with smooth surfaces. The method is very well suited for cardiac electrophysiological simulations. In the myocardium, the algorithm minimizes variations in element size, whereas in the surrounding medium, the element size is grown larger with the distance to the myocardial surfaces to reduce the computational burden. The numerical feasibility of the approach is demonstrated by discretizing and solving the monodomain and bidomain equations

on the generated grids for two preparations of high experimental relevance, a left ventricular wedge preparation, and a papillary muscle.

Index Terms

Bidomain model; individualized medicine; in silico heart model; mesh generation; multiscale modeling; octree method

I. Introduction

The heart is an organ of complex anatomy, where the structural complexity spans the spatial scales from the myocyte microstructure to the geometry and fiber architecture at the organ level. It is well documented that structural features of the heart play an important role in its electrophysiological behavior under normal and pathological conditions by defining conduction pathways [1]–[3] and altering conduction velocities [4] and wavefront curvature [5]. Structural heterogeneities in the heart cause local variations in the wavelength of a propagating action potential, which under certain disease conditions can present themselves as arrhythmogenic substrate, causing wavebreaks and reentrant activity [6].

Due to the complexity of structure and electrophysiology of the heart, scientific inquiry into the mechanisms of cardiac rhythm disorders has benefited tremendously from mathematical modeling and computer simulations. Current state-of-the-art electrophysiological models have already enabled the researchers to explore arrhythmogenesis on the organ level [7] and its relation to processes at the cellular and subcellular levels. However, all current models of the heart utilize simplified stylized representations of cardiac anatomy, where detail of the cardiac structure is often ignored. This shortcoming of the current state-of-the-art cardiac models stems mainly from technical difficulties associated with obtaining structurally accurate computational grids of the heart and the computational expenses associated with fine-grained representations. First, image acquisition and processing of high-resolution datasets of the entire heart is a challenging problem, and there are only very few imaging modalities that provide enough resolution to reconstruct cardiac structural detail [8]. Second, constructing an adequate mesh, on which the electrophysiological problem is to be solved, requires fine-grained discretizations to resolve geometrical detail, resulting in a large number of degrees of freedom and rendering the problem of significant computational expense. Finally, even if segmented high-resolution datasets are available and computational costs would be of no concern, the generation of high-resolution grids that faithfully represent cardiac anatomic detail at different spatial scales, from microscopic features (intercellular clefts, vascularization, fibrosis, Purkinje system) through macroscopic structural organization, such as fiber curvature and rotation as well as trabeculations and papillary muscles, to the gross anatomy of the heart, is an extremely challenging task.

Generating computational grids that represent suitable discretizations for the solution of a given cardiac electrophysiology problem is, in itself, associated with additional constraints and requirements.

1. To be generally applicable for both propagation and the delivery of external current to the heart, unstructured mesh generation techniques need to be employed. Jagged surfaces, as will inevitably occur when structured grids are employed, lead to spurious surface polarizations when defibrillation shocks are delivered to the heart.
2. In many cases, in addition to the heart, the surrounding volume conductor needs to be modeled. Modeling the passive surrounding medium requires significantly less spatial resolution as compared to the myocardial mass, where steep depolarization

wavefronts need to be resolved. Hence, it is desirable to continuously decrease the spatial resolution of a volume conductor grid with the distance to the cardiac surfaces. Spatially adaptive unstructured mesh generation techniques are ideally suited for this purpose, allowing to minimize the degrees of freedom required to represent the volume conductor while preserving the conformity of the mesh.

3. Within the myocardium, the variation in element size should be as small as possible. This is particularly important when coarser grids (element sizes $> 200 \mu\text{m}$) are employed. With coarser grids, the dependency of conduction velocity on grid granularity increases [9], [10], which, in turn, leads to a decrease in conduction velocity and even propagation block. Large variations in element size introduce spatial variations in local wavelength, which is a numerical artifact.

The vast majority of modeling efforts reported in the computational cardiac electrophysiology literature employed simple structured grids where mesh generation was not required and the desired spatial discretization was straightforwardly accomplished with standard discretization techniques such as the finite-difference method (FDM) [7], [11], [12], the finite-volume method (FVM) [13], and the finite-element method (FEM) [14], [15]. All studies where unstructured grids were employed relied on interactive mesh generation tools [16]–[19]. If one considers the most general case, i.e., bidomain simulations using unstructured grids where the tissue is immersed in a conductive fluid (bath, torso, blood-filled cavities), the number of reports in this category is even more limited [17]–[19]. The use of interactive mesh generation tools is, however, typically a time-consuming process with familiarization often taking weeks.

This paper presents a novel method for mesh generation that overcomes these shortcomings, allowing a fully automatic generation of 3-D FEM meshes. Segmented 3-D medical datasets serve directly as input for the mesh generator. The method relies on an octree-based meshing technique, and produces boundary-fitted, locally refined, and conformal FEM meshes consisting of either a variety of standard FEM elements, or alternatively, purely tetrahedral elements. Segmentation properties such as region tags, which are typically assigned with image processing techniques and thus are defined on a voxel grid, are preserved and seamlessly carried over onto the corresponding unstructured grid. This is also true for other types of data defined on a per voxel base, such as fiber orientations, which may be available through diffusion tensor MRI scans. Considering the particular importance of fiber orientations in cardiac modeling, a convenient and automatic way of transferring this information between the grids is a key issue.

To test the feasibility of the approach, high-resolution MRI scans are segmented and used as input for the mesh generation algorithm. Grids representing widely used experimental preparations such as papillary muscles [20], [21] and left ventricular (LV) wedge preparations [22] are generated at different average resolutions. Finally, the numerical feasibility of the approach is demonstrated by solving the cardiac monodomain and bidomain equations on these grids.

II. Methods

A. Image Processing

1) Image Segmentation—A high-resolution MRI scan [23] (isotropic resolution of $\sim 25 \mu\text{m}$) of the rabbit ventricles was acquired in a previous study [8] and served as a test case for the present study. Segmentation of the dataset was accomplished by the following steps. Due to the presence of a bias field and the overlap between associated intensity levels, a simple approach based on gray-level thresholding yielded unsatisfactory segmentation results. A more elaborate multistep segmentation approach had to be employed. We first estimated and removed the bias field from the datasets. Briefly, two tissue-free slices, one below the apex and one above the atria, were acquired, interpolated over the entire tissue volume, and

subtracted from the MRI datasets. Subsequently, the entire stack was segmented slice by slice, starting at the apex. In each slice, intensity thresholding, followed by a series of *ad hoc* morphological operations, was performed. The result of the segmentation of the previous slice was taken into account in the process to ensure a coherent 3-D result. Within the heart volume, differences in intensity levels in bias-corrected images were sufficient to allow correct segmentation between tissue, background, and extracellular cleft spaces or vessels.

2) Image Stack Preprocessing—A segmented 3-D image dataset, where each voxel's value is either "0" for background and "1" for tissue, served as input. A nonshrinkage Laplacian smoother [24] was used to eliminate the staircase and block boundary artifacts intrinsic to voxel-based data. Application of the filter directly on the original grid was shown to eliminate voxel boundary artifacts only partially, so the method was applied in a multiresolution framework. Filtering was carried out first at a coarse grid (original image resolution reduced by a factor of 2 on each axis), which provides a more complete elimination of voxel jaggedness. The results were interpolated and integrated into the full resolution grid, where the filtering results at full resolution were used to ensure that original volume and topology were preserved. This was done by not allowing original voxel values to change labels (i.e., to cross a predetermined threshold). This coarse-fine sequence was repeated iteratively, with results crossing between coarse and fine scales at each iteration. A trilinear function $\tilde{f}_v$ in space was finally defined by interpolating the values assigned to the voxel centers—this provides achieving subvoxel resolution in the subsequent steps. The values of $\tilde{f}_v$ at the element centers are not binary any more, but are in the range $0 \leq \tilde{f}_v \leq 1$, and the object's surface, implicitly defined as an isosurface of $\tilde{f}_v$, is smoothed.

B. Mesh Generation

1) Octree and Dual Mesh—The mesh generation algorithm employed an octree as the basic data structure. Initially, a coarse octree at a particular resolution, prespecified in multiples of the voxel size, is chosen to overlap the domain defined by the voxel stack. Each cell of an octree represented a cube in 3-D space. All voxels located along each edge of an octree box are scanned for intersections of the function $\tilde{f}_v$ with a chosen iso-value. Whenever an intersection is found, the octree is locally refined by subdividing a box into eight equally sized volumes. The refinement process in the tissue domain was stopped when the octree box size dropped below the voxel discretization. Further, in an additional sweep, octree boxes were refined to ensure that the difference in refinement level between adjacent octree boxes is one. The locally refined octree could have been used as a hanging node background mesh [see Fig. 1(a)]; however, the use of continuous FEM formulations required a conformal mesh, where two adjacent elements were interconnected by a common face with all nodes and edges shared between them, i.e., no hanging nodes or nonmatching edges were allowed. Hanging nodes are avoided here by using the dual mesh of the octree. The entities forming the dual mesh were related to the primal mesh (i.e., the octree cells) by the following relationships: $(\text{vertex})_p \rightarrow (\text{cell})_d$; $(\text{face})_p \rightarrow (\text{edge})_d$; $(\text{edge})_p \rightarrow (\text{face})_d$; $(\text{cell})_p \rightarrow (\text{vertex})_d$. Here, the subscripts p and d indicate *primal* and *dual*, respectively. The relationship between primal and dual meshes is illustrated in Fig. 1(b). The cells of the dual mesh are spanned by the centers of all boxes that share the corner of an octree box. For the sake of simplicity in visualization, the 2-D analog of an octree, a quadtree, was chosen in Fig. 1(b). Unfortunately, applying this simple approach in 3-D yields, besides the standard 3-D FEM elements such as hexahedra, prisms, tetrahedra, and pyramids [see Fig. 1(d)], additional nonstandard element types [Fig. 1(d), rightmost panel]. This difficulty is circumvented by modifying the dual mesh. Where coarse boxes meet fine boxes, a vertex is added [Fig. 1(c)]. In 3-D, the analogous operation is to add a single refinement layer of octree boxes, i.e., a coarse box is split into five subvolumes, for fine boxes that are connected to the hanging node, and one coarse box, filling half the volume of the coarse box. The advantage of the modified dual mesh is that it consists of standard element types only. An

additional property of this method is that no tetrahedra are generated in 3-D, i.e., the mesh consists of hexahedra, pyramids, and prisms only.

2) Tagging and Cutting—Every vertex of the modified dual mesh was tagged using the smoothed function $\tilde{f}_v$ defined over the octree box surrounding the vertex. When the smoothed image isosurface intersected a dual mesh cell, different tags were assigned to the vertices of this cell (Fig. 2). Isosurfaces can be computed on a per element basis for the dual mesh. Each vertex in the dual mesh is an element center in the primal mesh, i.e., the octree. Isosurfaces are simply found by checking whether a particular edge of a dual mesh element crosses the threshold defined by the isovalue.

The cutting procedure for the dual mesh cells involved the following steps. Element edges that intersected with the isosurface were subdivided into quarters. If the isosurface intersected with a cell edge within one of the two medial quarters of the edge, a new vertex was inserted. Otherwise, no additional vertex was added and the closest vertex was moved to the intersection with the isosurface. Thus, the object surface is projected immediately very close to the isosurface that prevents the formation of negative elements and yields very smooth surfaces without any explicit surface smoothing procedures.

Cells of the dual mesh intersected by the cutting procedure were filled up, in 3-D, with hexahedra, prisms, pyramids, and tetrahedra. This was achieved by solving a local mesh generation problem for the volume spanned by the general polyhedra that arose due to the initial cutting procedure. Initially, to each vertex in the dual mesh, a tag was assigned to indicate whether the value of $\tilde{f}_v$ at a particular vertex was above or below the prespecified isovalue. In a second iteration, for each vertex, the intersection of connected edges with the isosurface was determined. If an intersection was found with any of the connected edges within a distance of less than a quarter of the edge length from the vertex under consideration, the vertex was moved to the intersection. In these cases, another tag was assigned to this vertex to indicate that the isosurface runs through the vertex. After this procedure, all vertices were tagged to be either below, above, or exactly at the isovalue. In the case of a hexahedron, 3^8 cases have to be distinguished that can be reduced to 2^8 by taking advantages of symmetry. To speed up the mesh generation process, the topological information for all the 2^8 cases was precomputed and stored in tables such that the local mesh generation problems could be solved quickly by a simple table lookup. As a result, the object was discretized by a conformal volume mesh that approximated the real geometry fairly accurately, although some surface jitter was still present. Finally, the mesh surface was further smoothed by minimizing the functional $\int \|\kappa\|^2 ds$, where κ is the curvature of the object's surface.

3) Mesh Smoothing—A final smoothing of the entire volume mesh required the minimization of a nonlinear energy functional F for each cell locally

$$F(\mathbf{x}_t) = \frac{\left(\sum_{i=0}^5 l_i\right)^2}{v} \quad (1)$$

where l_i are the lengths of the element edges of the tetrahedra or of the virtual tetrahedra arising from subdivision of the other element types (hexahedra, pyramids, and prisms), $\mathbf{x}_t$ is a set of four vertices spanning the tetrahedron, and v is the element volume. The subdivision into tetrahedral elements relied on the same strategy as described in Section II-E. For elements where the volume turned out to be negative, the functional was stabilized to be able to repair these elements.

4) Tetralyzation—Most codes for the solution of the mono-or bidomain equations in 3-D that employ FEM for spatial discretization rely on the use of tetrahedral elements [18], [25]. To ensure that the presented method is of general applicability, an additional postprocessing step was applied to split all non-tetrahedral elements into tetrahedra, and thus, arrive at a purely tetrahedral mesh. Although this can be achieved without adding any additional vertices, it is nontrivial. Therefore, a less elaborate approach was undertaken, although it inserted new vertices, and thus increased the computational burden as compared to the hybrid mesh, but was easier to code and maintain. The algorithm tried, nonetheless, to minimize the number of additional vertices. Briefly, all nontetrahedral elements, i.e., elements with at least one or several quadrilateral faces, were split into triangles along the shorter diagonal, forming tetrahedra. In the case of a pyramid, this was straightforward. With prisms, eight different cases had to be considered: in six of them, the prism was subdivided into three tetrahedra, while in two cases, an additional vertex had to be inserted, which split the prism into eight tetrahedra. Finally, with hexahedral elements, 64 cases had to be considered, where in 46 of them, the hexahedra were split into 5, 6, or 8 tetrahedra without inserting additional vertices. In 18 cases, an additional vertex was required. Since in the majority of cases a vertex insertion could be avoided, the increase in degrees of freedom was moderate.

C. Mesh Quality Metrics

To assess the quality of the meshes generated with the presented method, an artificial image stack of a torus shell immersed in a medium was generated at different resolutions (see Fig. 3). A torus can be considered as a more generic test case than, for instance, a sphere, since it contains all combinations of curvature directions and different curvature radii. Furthermore, choosing a torus shell allowed us to study the behavior of the algorithm when thin structures are to be discretized. To select the most general and challenging setup, the normal vector of the circle that generated the torus was chosen to be $\mathbf{n} = [1, 1, 1]^T$, thus avoiding possible parallel alignment with any of the coordinate axes. The radius of the generating circle r_1 was chosen to be 11.2, the radius of the revolving circle generating the outer hull of the torus r_2 was set to 5.3, and the radius generating the inner hull of the torus shell r_3 was 7.3. Using this geometrical description, image stacks were generated by discretization at a resolution H_v of 1, 0.5, 0.25, and 0.125, the volume within the bounding box circumscribing the object, and with the edges of the bounding box aligned with the coordinate axes. Each voxel was set either to 1, if the center of a voxel was within the volume of the torus shell, or 0 otherwise. Partial volume effects, as typically arise with imaging modalities such as MRI, were not simulated by this process since it causes surfaces to be even more jagged than in real world image stacks. The choice of radii and H_v avoided side effects, which could have potentially arisen if the object geometry descriptors were multiples of H_v . For each H_v , meshes were generated where the target mesh resolution H_m was chosen to be $\times 2$, $\times 1$, and $\times 0.5 H_v$. The mesh quality was assessed by computing the following commonly used metrics: 1) the skewness ε_e , as defined by $(V_{\text{opt}} - V_e)/V_{\text{opt}}$ of a tetrahedron where V_{opt} is the optimal volume of an equilateral tetrahedron of the same circumradius and V_e is the volume of a tetrahedral element with index e in the mesh. For nontetrahedral elements, skewness was determined by splitting elements into tetrahedra and the worst skewness value of all subtetrahedra was assigned to the element; 2) the accuracy of the volume V_δ , as defined by $|(V_g - V_{\text{mesh}})/V_g|$ where V_g is the volume of the geometry and V_{mesh} the volume of the mesh; 3) the accuracy of the surfaces S_δ , as defined by $|(S_g - S_{\text{mesh}})/S_g|$ where S_g is the surface of the geometry and S_{mesh} the surface of the mesh; and 4) the minimum and maximum angles $\alpha_{\min}$ and $\alpha_{\max}$ respectively, found in the mesh where the angles are measured between two connected edges of an element.

Besides $\hat{\varepsilon}$, $\alpha_{\min}$, and $\alpha_{\max}$, we computed the maximum aspect ratio $AR_{\max}$ for all biological test cases. The aspect ratio quantifies the degree that mesh elements are stretched, and is defined

per element by $l_{\max}/l_{\min}$, where $l_{\max}$ refers to the length of the longest edge and $l_{\min}$ to the shortest edge length.

D. Selection of Test Cases

Two widely used cardiac experimental preparations, the left anterior papillary muscle and the LV wedge, were chosen as test cases. The segmented image stack of the rabbit heart was imported into MATLAB (MathWorks Inc., Natick, MA). Simple geometric shapes were used to extract the regions of interest. In the case of the papillary muscle, a cone-like shape turned out to be well suited; for the wedge, a pie-sector was used. The chosen geometric shapes were intersected with the geometry of the image stack and a Boolean image stack was generated. The volumes inside and outside the geometric shapes were tagged with ones and zeros, respectively. Multiplication of the Boolean stack with the gray-level image stack allowed a straightforward delineation of the regions of interest [Fig. 4(A) and (B)]. Finally, the image stack was cropped, leaving a 12-pixel ($\approx 500 \mu\text{m}$) thick conductive layer on every boundary side to minimize the bath size.

E. Hybrid FEM Discretization and Numerical Techniques

The computational domain constructed as described before comprised hybrid macroelements, e.g., hexahedra, pyramids, and prisms. In all our simulations, the FEM method employed localized Ansatz functions [26]. To construct these Ansatz functions, the macroelements in the entire domain had to be subdivided into tetrahedra. As shown in [26], the potential within each subtetrahedron varies linearly and is continuous between tetrahedra. Thus, the potential and shape functions are linear and piecewise continuous within each element. The macroelements were subdivided based on the number of sides of each face. A virtual point was added at the center of each quadrilateral face to allow subtetrahedra to be constructed. This resulted in four subtetrahedra for a pyramid. At the center of these faces, the shape function associated with each node of the face was assigned a value of $1/4$. For hexahedra and prisms, the centroid of the element was also considered a virtual point, with subdivision resulting in 24 and 14 subtetrahedra, respectively. At the centroid of these two element types, each shape function was equal to $1/n$, where n was the number of vertices. The employed numerical technique has been described in detail previously [18].

F. Numerical Tests

To demonstrate the numerical feasibility, monodomain and bidomain simulations were performed in the $50\text{-}\mu\text{m}$ discretization model of the LV wedge preparation and the $35\text{-}\mu\text{m}$ discretization model of the papillary muscle. For the description of the active membrane behavior, a rabbit ventricular action potential model [27] was implemented. To mimic experimental setups as closely as possible, for the wedge preparation, a point electrode located at the top of the wedge, approximately in the middle of the epicardial edge of the preparation, delivered pacing pulses via transmembrane current injection (monodomain), while plate electrodes in the bath adjacent to the epi- and endocardial ventricular surfaces were chosen for the delivery of a $\approx 3\text{-V/cm}$ shock (bidomain). For the papillary muscle, two points at a distance of 3.75 mm between each other, located above the branchings of the papillary muscle, served as sites of shock application (bidomain), while a transmembrane stimulus was delivered via a point electrode close to the base of the papillary muscle (monodomain). The electrode configurations are illustrated in Fig. 5. Stimuli of 5 ms duration were administered via the pacing electrodes to elicit action potentials in both preparations, and propagation was monitored for another 15 ms . Model parameters of interest were degrees of freedom, distribution of edge lengths, and average execution times per millisecond of activity.

Simulations were executed using the CARP simulator [28], [29] running on the HPCx supercomputing facilities [30] with 128 CPUs. CARP was set up to employ a solver strategy

described previously [31] with the only modification that a Crank–Nicholson approach was used to update the parabolic partial differential equation (PDE). Briefly, the elliptic PDE was solved using an algebraic multigrid preconditioner with an iterative conjugate gradient solver. The parabolic problem was solved using a block–Jacobi preconditioner together with an incomplete Cholesky ICC (0) subblock preconditioner with a zero fill-in level to preserve the sparsity pattern. For both PDEs, an absolute error tolerance of $1e^{-6}$ was used, measured as the L_2 -norm of the unpreconditioned residual.

III. Results

A. Segmentation and Substack Selection

The regions of interest were delineated from the stack, as shown in Fig. 4. The segmentation preserved all important macroanatomical features, such as the geometry of the preparations, the endocardial structures including papillary muscles and trabeculae, and major parts of the vascularization. Finer microscopic details, such as the free-running Purkinje network, were preserved as well. The segmented image substacks were binarized and fed directly into the mesh generation algorithm without any further modifications.

B. Mesh Generation

The extracted image substacks were sufficiently complex to serve as challenging test cases for the mesh generation algorithm. Finite-element discretizations of the LV wedge and left anterior papillary muscle were generated by specifying the desired mean spatial resolutions of the mesh as input. For each model, three different resolutions were tested, namely 100, 50, and 35 μm for the LV wedge (meshes w100, w50, and w35), and 75, 35, and 25 μm for the papillary muscle (meshes pm75, pm35, and pm25). The execution of the mesh generation algorithm lasted between 4.18 and 64.90 min for the smallest and the largest meshes, respectively. Mesh generation time depended linearly on the number of generated elements with 2.29 min needed per one million elements (Table I).

The number of vertices required to resolve the geometry at a particular resolution varied between ~ 1 million for the coarsest and ~ 24 million for the finest mesh. The average edge lengths of the generated elements and their variations are summarized in Table II. As expected, due to the octree-based data structure underlying the mesh generation algorithm, the predominating finite-element types were hexahedra (37.0%–52.3%), followed by tetrahedra (20.4%–26.4%), pyramids (19.4%–25.3%), and prisms (7.9%–11.2%). Splitting the hybrid elements to arrive at an entirely tetrahedral mesh increased the number of vertices somewhat by 3% in the best case and 4% in the worst case. Fig. 6 shows the resultant meshes generated at the medium resolution for both preparations.

C. Mesh Adaptivity

As outlined in Section II, the algorithm attempts to minimize variations in element sizes within the myocardial volume, but decreases the resolution quickly with the distance from the myocardial surfaces outside the myocardium. This effect is clearly visible in Fig. 6.

Relative reduction of the number of vertices in the bath depended, to a certain degree, on the mesh resolution. Although the bath volume was small, significant gains in vertex reductions could be achieved at both the finest resolution (reduction of 81.5% for the wedge and 91.2% for the papillary muscle) and the coarsest resolution (67.5% and 79.6%, respectively).

The distribution of vertices between tissue and bath in the two example cases is summarized in Table II for the meshes of finest resolution, w35 and pm25. Note that the vertices located at the interface between tissue and bath are counted in both domains, which explains the total

number of vertices being slightly more than 100%. Clearly, more than 80% of the vertices and 66.5%–70.5% of the finite elements are in the tissue. Edge lengths in the tissue vary in the range of 8.0–64.5 μm for the w35 mesh and in the range of 3.6–46.0 μm for the pm25 mesh. Elements larger than twice the desired element size were not observed. In the bath, the variation in edge length increases significantly due to the adaptive nature of the mesh.

D. Mesh Quality Metrics

As outlined in Section II, artificial image stacks of a torus shell were generated at different resolutions H_v , and fed into the mesh generator. Meshes were generated at three different target mesh resolutions H_m , and several standard mesh quality metrics were computed. A summary of results is found in Tables III and IV.

E. Numerical Tests

Monodomain and bidomain simulations were performed in the 50- μm discretization model of the LV wedge preparation and the 35- μm discretization model of the papillary muscle to demonstrate the numerical feasibility of the approach. We applied the same numerical techniques as in previous studies of ours, where purely tetrahedral meshes, generated with commercial packages, were used. The same time stepping schemes and the same time steps could be used here without any adjustments of our solver code. Convergence behavior and execution speed were essentially unaffected, yielding values very similar to previous results [18]. Differences between hybrid and tetrahedral meshes were also found to be insignificant in terms of both electrophysiological results and solver performance.

The parabolic portion of the bidomain and the monodomain equations could be integrated using a standard time step of 20 μs . Execution times are summarized in Table V. Distributions of transmembrane voltage V_m for both test cases and for monodomain and bidomain simulations are shown in Fig. 5. For bidomain simulations, distributions of extracellular potentials Φ_e are shown as well.

IV. Discussion

In this study, we present a novel implementation of an octree-based mesh generation technique that is applied directly to segmented image stacks. The mesh generation procedure produces conformal, boundary-fitted, and hexahedra-dominant meshes. The algorithm operates fully automatically, with no requirements for interactivity, and generates accurate volume-preserving representations of arbitrarily complex geometries with smooth surfaces. Different mesh generation strategies are employed for the myocardium and the surrounding medium. In the myocardium, the algorithm minimizes variations in element size, whereas in the surrounding medium, the element size is grown larger with the distance to the myocardial surfaces to reduce the computational burden in simulating activity in the extracellular domain. Finally, the numerical feasibility of the approach is demonstrated by discretizing and solving the monodomain and bidomain equations on the generated grids for two preparations of high experimental relevance.

A. Comparison With Other Mesh Generation Techniques

While automatic unstructured mesh generation is a relatively new field, due to its importance and practical applicability, it has been growing rapidly and a variety of approaches have been developed over the last few years. Although a thorough discussion of the strengths and weaknesses of the most important techniques is beyond the scope of this study, a short summary of the different techniques and how they compare to the presented approach seems appropriate here. More thorough surveys on mesh generation techniques are found elsewhere [32], [33]. In general, depending on the application at hand, the quality of the mesh generation algorithm

is assessed on the basis of: 1) the “quality” of the generated elements (where “quality” typically refers to a set of geometric properties, such as skewness); 2) the required resources in terms of the number of elements and vertices and the memory footprint of the data structures; 3) the smoothness and fidelity of the surface representation; and 4) the time required to execute the algorithm.

The most frequently used techniques can be classified as block-structured, advancing front, or grid-based methods. With block-structured techniques, meshes of excellent quality can be generated; mesh parameters can be controlled very well, and all types of elements can be used. However, block-structured techniques are not amenable to automatization, and difficulties in subdividing an object into appropriate blocks increase rapidly with the increase in complexity of the object’s geometry. For instance, in [16], a model of the human atria was developed using this technique, where the geometrical complexity of the object rendered the application of the block-structured approach a daunting endeavor in terms of the time and effort required. Block-structured methods appear rather poorly suited when highly realistic representations of biological structures are desired, and in particular, when one is interested in high-throughput meshing with mesh generation times on the order of seconds and not weeks.

Advancing front techniques [34]–[36] discretize volumes starting from the surfaces and work toward the inside of the volume. The main advantage is the control over the meshing process, and if desired, coarse meshes can be generated. On the other hand, advancing front techniques are very sensitive to errors in the input data. As a result of such data errors, it may not be possible to define a unique surface normal. As an example, for an object consisting of two cones that share a common tip, no unique surface normal can be defined at the tip. Such situations inevitably arise in segmented image data stacks, where adjacent object voxels are interconnected by a single shared vertex. Also, thin and crossing surfaces, as are encountered frequently in biomedical image datasets, are known to be problematic [34].

Further, the computational burden with advancing wavefront methods is significant, particularly in avoiding intersections of elements when advancing fronts of the mesh are to collide. Finally, as a starting point of the method, triangulated surfaces of sufficiently good quality are required, which may be nontrivial to obtain for geometries as complex as in this study.

The technique presented in this study can be classified as a grid-based mesh generation method that uses the octree as the main data structure. Octree-based mesh generation techniques were pioneered by Yerry and Shepard [37], who demonstrated their suitability for automatic mesh generation of typical geometry-based problems in engineering-type applications [38]. Essentially, in a first approximation, the volume of an object is filled up with a regular grid, which is then adapted at the boundaries to arrive at a smooth match to the actual geometry. The background grid can be chosen freely. For instance, with biological image stacks, the segmented voxel grid itself or any interpolation of it at a desired resolution can serve directly as the initial mesh. The approach presented here uses the dual mesh of the octree as the background grid and locally refines along the surfaces. Refinement is achieved by splitting of elements that intersect with the true geometry of the object. The employed splitting approach makes the method very tolerant to errors in input data. The mesh within the interior of an object is typically very good since the mesh topology is close to being regular. The disadvantage of the method is the complexities arising from the filling of the gap between the surface of the regular background mesh and the surfaces of the true geometry.

B. Tetrahedral Meshes

Studies that have discretized the monodomain or bidomain equations on unstructured grids using FEM have overwhelmingly employed tetrahedral elements [18], [25]. In all cases,

interactive mesh generation tools were used to discretize the geometries. Compared to purely tetrahedral meshes, the advantages of hybrid elements are that a particular geometry can be represented with a significantly smaller number of elements, thus reducing memory requirements for storing element lists. For instance, the ratio in element numbers between hexahedral meshes and tetrahedral meshes is, depending on the approach, between 1/5 and 1/12, while the number of vertices is roughly the same. On the other hand, vertices in hybrid meshes are linked to more neighboring vertices, which tend to increase not only accuracy and but also memory requirements. For the meshes under study, the stiffness matrix associated with hybrid meshes required, on average, between 3.5 to 4.6 more entries per matrix row than their purely tetrahedral analog. Overall, in the implementation used in this study, there were memory savings with using hybrid meshes since shorter element lists outweighed the additional costs of slightly denser matrices.

One of the main advantages of using tetrahedral mesh generation techniques is that mesh smoothing is easier to achieve by simple procedures, such as edge swapping and vertex shifting [36]. Smoothing of hybrid meshes is a more involved task, but could potentially improve the mesh quality. Ansatz shape functions described in [26] for the FEM, as utilized in this study, have not been used for the cardiac bidomain equations yet. Hybrid elements, besides apparent savings in the number of elements, offer further numerical advantages over purely tetrahedral meshes [26]. The presented technique is very well suited for the generation of tetrahedral meshes as well. The overhead introduced, in terms of additional vertices, is very moderate (depending on the resolution, it is between 3% and 4%).

C. Accuracy Considerations and Mesh Quality Metrics

Unlike many engineering applications, the accuracy of the solution is not the most important factor when solving the bidomain equations. Errors of up to 10% are widely considered as acceptable. This is justified by considering that none of the tens to hundreds of parameters involved in the bidomain equations and ionic kinetics models are known with a better accuracy than 10%. Reported values on key parameters such as tissue conductivities and the ratio between the conductivities in both the intracellular and the extracellular/interstitial domain show a striking variance [39]. Further, these parameters are likely to show interspecies and individual variations. Even within the same heart, conductivities vary regionally within the ventricles and the atria due to differences in expression of gap junctions, in microscopic fiber arrangement, packing density of myocytes, the amount of connective tissue, etc. Under pathological conditions or age-related remodeling of the cardiac substrate, these parameters may dramatically change, i.e., for any computational study, some of these parameters have to be tuned to arrive at a good match with experimental observations for a particular problem. With this background, one has to define the requirements on mesh generation techniques in the following way: it is desirable that mesh resolution does not vary dramatically within the volume of the myocardium since variations in resolution will introduce variations in conduction velocity [9]. This is particularly true for coarser meshes, where mesh resolution approaches the spatial extent of the depolarization wavefront. For very fine meshes, as those generated in the present study, these effects are negligible, with modeling of pathologies, which lead to very slow decremental conduction, notwithstanding. Also, depending on the numerical technique applied to solve the parabolic PDE, the maximum time step of the entire system may be limited by a few small elements. This may render the intrinsically expensive cardiac modeling studies computationally intractable. Finally, it is important that grid-based mesh generation techniques preserve topological features of an object within the scale of interest of a study, and that a mesh could be straightforwardly adjusted to the scale of interest. The presented method was custom-tailored with these design goals in mind. The quality of the mesh was mainly assessed based on comparing the numerical expense of computations between hybrid meshes and purely tetrahedral meshes, which were first used by Vigmond *et al.* [40] to solve the cardiac bidomain

equations. In addition, several standard metrics commonly used to assess mesh quality were computed as well. As expected, with increasing resolution of the image stack H_v , the volume and surface of an object are represented with increasing *accuracy*. With coarser meshes, increasing the mesh resolution H_m may improve V_δ ; however, with finer meshes, this is not the case. Both H_v and H_m have to increase to ensure geometric convergence of the mesh. If one excludes the extreme case, where a mesh was generated from the coarsest image stack $H_v = 1$ with a mesh resolution of $H_m = 2 \times H_v$, all mesh quality metrics were within the desired range. The maximum skewness $\hat{\epsilon}$ of the elements varied in the range between 0.80 and 0.87, and the minimum and maximum angles were in the range of 5.89° – 26.07° and 136.28° – 153.44° , respectively.

The same quality metrics were computed for the biological test cases as well. The maximum aspect ratios were in the expected range. For instance, Shepard and Georges reported maximum aspect ratios between 2.66 and 49.3 for small meshes representing fairly simple geometric objects [38]. The maximum aspect ratios for the large meshes of substantially more complex geometry are noticeably smaller suggesting that the presented approach is competitive in this regard.

D. Choice of Resolution

The presented mesh generation technique allows for the desired resolution of the mesh to be chosen freely and independently of the input image stack. Ideally, computations should be carried out on a fine grid, but computational expenses increase exponentially with increasing resolution, i.e., in a real world context, a tradeoff has to be made between the resolution of the grid and the computational costs. Flexibility in choosing the resolution is a key factor since the manageable computational burden for a given study depends on several factors, such as the efficiency of the simulator, the size of the parameter space to be explored, and the availability and performance of high-performance computing platforms. The presented algorithm allows quick adjustments of the mesh size to a particular problem. With mesh generation times on the order of a few minutes and no user interaction required, meshes can be generated during the setup phase of a simulation, where a maximum number of degrees of freedom is supplied as an input parameter.

Clearly, severely constraining the maximum degrees of freedom in representing a geometry leads to underrepresentation of smaller structures. The mesh generator technique does not generate any elements to represent structures that are below a certain prespecified limit. This is evident in Fig. 7, where the intramyocardial structures, such as vascularization and interstitial heterogeneities, are exposed in two meshes of different resolution.

E. Adaptive Discretization of the Surrounding Medium

One of the key advantages of bidomain models is their capability to explicitly account for current flow in the medium surrounding the heart. Thus, problems such as extramyocardial injection of currents (defibrillation shocks) or loading effects of the bath on the myocardial excitation sequence can be studied. Since no steep gradients of the propagating wavefront need to be spatially resolved in the surrounding medium, the spatial resolution can be decreased quickly with the distance to the myocardial surfaces to reduce the overall computational load. With structured grids, it is not possible to achieve this using conformal FEM approaches, unless one resorts to discontinuous Galerkin methods, which are yet to be applied to the cardiac bidomain problem. For instance, in a recent study by Potse *et al.* [12], an insulated human heart was discretized on a structured grid of $200\ \mu\text{m}$ resolution using FDM. To resolve the geometry, 25 million vertices were required. Adding the heart's cavities, an epicardial fluid layer of only 1 mm thickness and a volume representing the atria added another 29 million vertices. As stated by the authors, it will be impossible to model the heart immersed in a torso using structured

grids. For such a modeling endeavor, the use of adaptive mesh generation is the key in keeping simulations tractable. Gains due to adaptive coarsening with the distance to the heart increase with the size of the surrounding volume conductor, allowing the discretization of a torso at a substantially lower cost.

The use of larger elements can cause numerical instabilities in the myocardial volume, but not in the surrounding volume conductor. The equation governing the bath potential is an elliptic PDE, and thus, unlike the parabolic PDE governing the activation spread within the myocardium, it imposes no stability constraints on the mesh granularity. Due to the low-pass filter effect of the bath, the potential distribution in the bath becomes increasingly smoother with the distance to the myocardial surface, and lower spatial sampling by coarser meshes suffices to faithfully represent the solution without any loss of information.

F. Segmentation

The quality of the segmentation step is an important factor determining the quality of the generated mesh. The presented algorithm generates meshes that represent faithfully the *segmented* geometry of a set of biological objects. Wherever errors occur in the segmentation step, they will be reflected in the mesh and cannot be corrected by the mesh generation approach. For instance, isolated voxels or groups of voxels that are not connected to the myocardium will be incorporated into the myocardial mesh. If required, a geometry cleanup of nonconnected myocardial tissue could be straightforwardly achieved. For instance, running a single simulation allows to identify nonconnected tissue fragments. Backprojecting the locations of isolated elements onto the image stack allows to correct such segmentation problems.

V. Conclusion

Although macroanatomically realistic meshes of the ventricles [17], [41] and the atria [16] have been developed over the last decade, the represented geometries were simplified and stylized. Moreover, the applied mesh generation techniques were interactive and required substantial resources in terms of time and effort.

In this study, a novel fully automatic mesh generation method was proposed that uses the dual mesh of an octree and which can be applied directly to segmented image data stacks. The suggested method generates accurate representations of arbitrarily complex geometries without any interactivity and is custom-tailored to account for the specifics of cardiac modeling. The number of parameters required to tune the algorithm is small, decreasing the required training effort to minimum. The generated high-resolution meshes of the cardiac anatomy, as shown in this study, are significantly more detailed than any other previously published models, but the mesh generation effort was a fraction of any other approach presented thus far.

The quality of the elements generated within the myocardium volume, the adaptivity in the surrounding medium, and the short execution times render the presented method ideally suited for high-throughput cardiac modeling endeavors, such as simulations of subject- or patient-specific electrophysiology problems.

Acknowledgments

The authors thank the teams of Dr. P. Kohl and Dr. D. Gavaghan at the University of Oxford for access to data from their 3-D Heart Project, BBSRC Grant #E003443.

This work was supported in part by the Austrian Science Fund (FWF) under Grants R21-N04 and SFB F3210-N18 and in part by the Engineering and Physical Sciences Research Council (GR/S72023/01) and International Business Machines Corporation (IBM) under the Integrative Biology Project. The work of G. Plank was supported by the

European Commission under Grant MC-OIF 040190. The work of V. Grau was supported in part by the Research Councils U.K. The work of E. J. Vigmond was supported in part by the Natural Sciences and Engineering Research Council of Canada. The work of N. A. Trayanova was supported in part by the National Institutes of Health (NIH) under Award HL063195, Award HL082729, and Award HL067322.

References

1. Sanchez-Quintana D, Anderson RH, Cabrera JA, Climent V, Martin R, Farre J, Ho SY. The terminal crest: Morphological features relevant to electrophysiology. *Heart* 2002;88(4):406–411. [PubMed: 12231604]
2. Cabrera JA, Ho SY, Climent V, Sanchez-Quintana D. The architecture of the left lateral atrial wall: A particular anatomic region with implications for ablation of atrial fibrillation. *Eur Heart J* 2008;29(3): 356–362. [PubMed: 18245120]
3. Sakamoto S, Nitta T, Ishii Y, Miyagi Y, Ohmori H, Shimizu K. Interatrial electrical connections: The precise location and preferential conduction. *J Cardiovasc Electrophysiol* 2005;16(10):1077–1086. [PubMed: 16191118]
4. Kim YH, Xie F, Yashima M, Wu TJ, Valderrabano M, Lee MH, Ohara T, Voroshilovsky O, Doshi RN, Fishbein MC, Qu Z, Garfinkel A, Weiss JN, Karagueuzian HS, Chen PS. Role of papillary muscle in the generation and maintenance of reentry during ventricular tachycardia and fibrillation in isolated swine right ventricle. *Circulation* 1999;100(13):1450–1459. [PubMed: 10500048]
5. Cabo C, Pertsov AM, Davidenko JM, Jalife J. Electrical turbulence as a result of the critical curvature for propagation in cardiac tissue. *Chaos* 1998;8(1):116–126. [PubMed: 12779715]
6. Weiss JN, Qu Z, Chen PS, Lin SF, Karagueuzian HS, Hayashi H, Garfinkel A, Karma A. The dynamics of cardiac fibrillation. *Circulation* 2005;112(8):1232–1240. [PubMed: 16116073]
7. Ten Tusscher KH, Hren R, Panfilov AV. Organization of ventricular fibrillation in the human heart. *Circ Res* 2007;100(12):e87–e101. [PubMed: 17540975]
8. Burton RA, Plank G, Schneider JE, Grau V, Ahammer H, Keeling SL, Lee J, Smith NP, Gavaghan D, Trayanova N, Kohl P. Three-dimensional models of individual cardiac histology: Tools and challenges. *Ann N Y Acad Sci* 2006;1080:301–319. [PubMed: 17132791]
9. Pollard AE, Burgess MJ, Spitzer KW. Computer simulations of three-dimensional propagation in ventricular myocardium. Effects of intramural fiber rotation and inhomogeneous conductivity on epicardial activation. *Circ Res* 1993;72(4):744–756. [PubMed: 8443866]
10. Cherry EM, Greenside HS, Henriquez CS. A space-time adaptive method for simulating complex cardiac dynamics. *Phys Rev Lett* 2000;84(6):1343–1346. [PubMed: 11017514]
11. Cherry, EM.; Greenside, HS.; Henriquez, CS. Efficient simulation of three-dimensional anisotropic cardiac tissue using an adaptive mesh refinement method; *Chaos* [Online]. 20023. p. 853-865. Available: <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?cmd=prlinks&dbfrom=pubmed&retmode=ref&id=12946177>, 2009
12. Potse M, Dube B, Vinet A, Cardinal R. A comparison of monodomain and bidomain propagation models for the human heart. *Proc IEEE Eng Med Biol Soc Conf* 2006;1:3895–3898.
13. Trew M, Le Grice I, Smaill B, Pullan A. A finite volume method for modeling discontinuous electrical activation in cardiac tissue. *Ann Biomed Eng* 2005;33(5):590–602. [PubMed: 15981860]
14. Rogers JM, McCulloch AD. A collocation–Galerkin finite element model of cardiac action potential propagation. *IEEE Trans Biomed Eng* Aug;1994 41(8):743–757. [PubMed: 7927397]
15. Plank G, Leon LJ, Kimber S, Vigmond EJ. Defibrillation depends on conductivity fluctuations and the degree of disorganization in reentry patterns. *J Cardiovasc Electrophysiol* 2005;16(2):205–216. [PubMed: 15720461]
16. Harrild D, Henriquez C. A computer model of normal conduction in the human atria. *Circ Res* 2000;87(7):E25–E36. [PubMed: 11009627]
17. Trayanova NA, Eason JC, Aguel F. Computer simulations of cardiac defibrillation: A look inside the heart. *Comput Vis Sci* 2002;4:259–270.
18. Plank G, Liebmann M, dos Santos RW, Vigmond EJ, Haase G. Algebraic multigrid preconditioner for the cardiac bidomain model. *IEEE Trans Biomed Eng* Apr;2007 54(4):585–596. [PubMed: 17405366]

19. Vigmond EJ, Clements C. Construction of a computer model to investigate sawtooth effects in the Purkinje system. *IEEE Trans Biomed Eng* Mar;2007 54(3):389–399. [PubMed: 17355050]
20. Hofer E, Urban G, Spach MS, Schafferhofer I, Mohr G, Platzer D. Measuring activation patterns of the heart at a microscopic size scale with thin-film sensors. *Amer J Physiol* 1994;266(5):H2136–H2145. [PubMed: 8203613]
21. Fleischhauer J, Lehmann L, Kleber AG. Electrical resistances of interstitial and microvascular space as determinants of the extracellular electrical field and velocity of propagation in ventricular myocardium. *Circulation* 1995;92(3):587–594. [PubMed: 7634473]
22. Fast VG, Sharifov OF, Cheek ER, Newton JC, Ideker RE. Intramural virtual electrodes during defibrillation shocks in left ventricular wall assessed by optical mapping of membrane potential. *Circulation* 2002;106(8):1007–1014. [PubMed: 12186808]
23. Schneider JE, Böse J, Bamforth SD, Gruber AD, Broadbent C, Clarke K, Neubauer S, Lengeling A, Bhattacharya S. Identification of cardiac malformations in mice lacking Ptdsr using a novel high-throughput magnetic resonance imaging technique. *BMC Dev Biol* 2004;4(16)
24. Vollmer J, Mencl R, Müller H. Improved Laplacian smoothing of noisy surface meshes. *Comput Graphics Forum* 1999;18(3):131–138.
25. Rodriguez B, Li L, Eason JC, Efimov IR, Trayanova NA. Differences between left and right ventricular chamber geometry affect cardiac vulnerability to electric shocks. *Circ Res* 2005;97(2):168–175. [PubMed: 15976315]
26. Kickinger, F. Johannes-Kepler-Universität Linz Series. Linz, Austria: Trauner Verlag; 1997. Tools for the Numerical Simulation of 3D Magnetic Field Problems: Construction, Analysis and C++ Implementation.
27. Mahajan A, Shiferaw Y, Sato D, Baher A, Olcese R, Xie LH, Yang MJ, Chen PS, Restrepo JG, Karma A, Grafinkel A, Qu Z, Weiss JN. A rabbit ventricular action potential model replicating cardiac dynamics at rapid heart rates. *Biophys J* 2008;94:392–410. [PubMed: 18160660]
28. Vigmond E, Hughes M, Plank G, Leon L. Computational tools for modeling electrical activity in cardiac tissue. *J Electrocardiol* 2003;36:69–74. [PubMed: 14716595]
29. Vigmond, E.; Plank, PG. Available: <http://carp.medunigraz.at/>
30. Plank, G.; Clayton, R.; Boyd, D.; Vigmond, EJ. Integrative biology: Modelling heart attacks with supercomputers; HPCx Capability Comput J. 2006. p. 4-8.[Online]Available: <http://www.hpcx.ac.uk/>
31. Vigmond EJ, Santos RWD, Prassl AJ, Deo M, Plank G. Solvers for the cardiac bidomain equations. *Prog Biophys Mol Biol* Aug;2007 96(1–3):3–18. [PubMed: 17900668]
32. Owen, S. A survey of unstructured mesh generation technology. *Proc. 7th Int. Meshing Roundtable; Dearborn, MI.* 1998. p. 26-28.
33. Teng SH. Unstructured mesh generation: theory, practice, perspectives. *Int J Comput Geom Appl* 2000;10(3):277–266.
34. Lohner R. Progress in grid generation via the advancing front technique. *Eng Comput* 1996;12:186–210.
35. Lo S. 3D triangulation by advancing wavefront approach. *Comput Struct* 1991;39:501–511.
36. Schoeberl J. NETGEN—An advancing front 2D/3D-mesh generator based on abstract rules. *Comp Vis Sci* 1997;1:41–52.
37. Yerry MA, Shephard MS. Three-dimensional mesh generation by modified Octree technique. *Int J Num Methods Eng* 1984;20:1965–1990.
38. Shephard MS, Georges MK. Three-dimensional mesh generation by finite Octree technique. *Int J Num Methods Eng* 1991;32:709–749.
39. Roth BJ. Electrical conductivity values used with the bidomain model of cardiac tissue. *IEEE Trans Biomed Eng* Apr;1997 44(4):326–328. [PubMed: 9125816]
40. Vigmond E, Aguel F, Trayanova N. Computational techniques for solving the bidomain equations in three dimensions. *IEEE Trans Biomed Eng* Nov;2002 49(11):1260–1269. [PubMed: 12450356]
41. Vetter F, McCulloch A. Three-dimensional analysis of regional cardiac function: A model of rabbit ventricular anatomy. *Prog Biophys Mol Biol* 1998;69(23):157–183. [PubMed: 9785937]

Biographies

Anton J. Prassl received the M.Sc. and Ph.D. degrees in electrical engineering from the Institute of Biomedical Engineering, Graz University of Technology, Graz, Austria, in 2003 and 2008, respectively.

He is currently a Postdoctoral Fellow at the Institutes of Physiology and Biophysics, Medical University of Graz, Graz. During 2006–2007, he was a Research Scholar at the Johns Hopkins University. His current research interests include computational modeling of the electrical and mechanical cardiac activity and the underlying finite-element models.

Dr. Prassl is a member of the Austrian Society for Biomedical Engineering.

Ferdinand Kicking received the M.Sc. degree in industrial mathematics from the Johannes Kepler University of Linz, Linz, Austria, in 1996.

During 1998–1999, he was an Assistant Professor at the Johannes Kepler University of Linz. During 1999–2004, he was a Senior Software Engineer at AVL List GmbH, Austria. During 2005–2009, he was a Lecturer of scientific computing at St. Pölten University of Applied Sciences, Austria. In 2004, he founded the company CAE Software Solutions, Eggenburg,

Austria. His current research interests include algorithms in computational fluid dynamics with a special focus on mesh generation.

Helmut Ahammer received the M.Sc. and Ph.D. degree in experimental physics from the University of Graz, Graz, Austria, in 1990 and 1996, respectively.

He is currently an Associate Professor of medical physics and biophysics at the Institute of Biophysics, Department of Physiological Medicine, Medical University of Graz, Graz, where he is involved in the field of image processing and quantitative image analysis. He was responsible for the organization, administration, and the radioprotection of a multiuser

radionuclide laboratory at the University of Graz. His current research interests include fractals, nonlinear methods, and quantitative methods in order to analyze digital images of biological objects.

Vicente Grau (M'08) received the M.S. and Ph.D. degrees in telecommunication engineering from the Universidad Politecnica de Valencia, Valencia, Spain, in 1994 and 2000, respectively.

From 2002 to 2003, he was a Research Fellow at the Surgical Planning Laboratory, Brigham and Women's Hospital, Harvard University. From 2003 to 2004, he was with the optic nerve head (ONH) Biomechanics Laboratory, Louisiana State University (LSU) Health Sciences

Center. Since 2004, he has been with the Department of Engineering Science, University of Oxford, Oxford, U.K., where he is currently a Research Councils U.K. (RCUK) Academic Fellowship jointly with the Oxford e-Research Centre. His current research interests include the combination of biomedical image analysis with computational models, with application on cardiovascular and pulmonary medicine.

Dr. Grau is a member of the Medical Image Computing and Computer Assisted Intervention (MICCAI) Society and the British Machine Vision Association.

Jürgen E. Schneider received the Diploma in physics and the Ph.D. degree in physics from the University of Würzburg, Würzburg, Germany, in 1996 and 2000, respectively.

In 2001, he joined Oxford University, Oxford, U.K., as a Postdoctoral Research Fellow, where he is currently a British Heart Foundation (BHF) Basic Science Lecturer in the Department of Cardiovascular Medicine.

Ernst Hofer received the M.Sc. and Ph.D. degrees in electrical engineering from the Technical University of Graz, Graz, Austria, in 1977 and 1985, respectively.

From 1985 to 1990, he was an Assistant Professor at the Karl-Franzens-University Graz, Graz. Since 1990, he has been an Associate Professor of medical physics and biophysics in the Department of Biophysics, Medical University of Graz, Graz. His current research interests include development of scientific instruments and measurement systems for cardiac electrophysiology, specifically for measurement and analysis of microscopic excitation spread in cardiac tissue.

Dr. Hofer has been the President of the Austrian Society for Biomedical Engineering since 2005.

Edward J. Vigmond (S'96–M'97) received the B.A.Sc. degree in electrical and computer engineering from the University of Toronto, Toronto, ON, Canada, in 1988, and the M.A.Sc. and Ph.D. degrees from the Institute of Biomedical Engineering, University of Toronto, in 1991 and 1996, respectively.

He is currently an Associate Professor in the Department of Electrical and Computer Engineering, University of Calgary, Calgary, AB, Canada, where he is modeling biological electrical activity. He was a Postdoctoral Fellow at the University of Montreal (1997–1999) and Tulane University (1999–2001). His current research interests include numerical computation of electrical fields, biomedical signal processing, and modeling of nonlinear biosystems.

Natalia A. Trayanova (A'90–SM'91) received the Ph.D. degree from the Bulgarian Academy of Sciences, Sofia, Bulgaria.

She is currently a Professor in the Department of Biomedical Engineering and the Institute for Computational Medicine, Johns Hopkins University, Baltimore, MD. She was a Postdoctoral Fellow in the Department of Biomedical Engineering, Duke University, Durham, NC. During 1995–2006, she was a faculty member in the Department of Biomedical Engineering, Tulane University. Her current research interests include understanding the normal and pathological electromechanical behavior of the heart, with emphasis on the mechanisms for cardiac arrhythmogenesis and defibrillation, the role of cardiac tissue structure, and the mechano-

electric feedback in the heart. She has authored or coauthored extensively. She is a member of the Editorial Board of the journal *Heart Rhythm*.

Prof. Trayanova is the recipient of numerous awards, including the Excellence in Research and Scholarship Award (2005) and the Outstanding Researcher Award (2002) from Tulane University, and the Fulbright Distinguished Research Award (2002), to name just a few. She has also received awards for excellence in teaching. She was an Associate Editor of the IEEE Transactions of Biomedical Engineering during 1997–2005, and is an Area Editor of the IEEE Reviews in Biomedical Engineering. She is currently a member of the National Institutes of Health (NIH) Electrical Signaling, Ion Transport, and Arrhythmias Study Section. She was the Vice Chair (in 2007) and is currently the Chair (2009) of the Gordon Research Conference on Cardiac Arrhythmia Mechanisms. She has presented papers at a large number of international meetings

Gernot Plank received the M.Sc. and Ph.D. degrees in electrical engineering from the Institute of Biomedical Engineering, Technical University of Graz, Graz, Austria, in 1996 and 2000, respectively.

He is currently an Assistant Professor at the Institute of Biophysics, Medical University of Graz, Graz. He is also an Academic Fellow with the Computational Biology Group, University of Oxford, Oxford, U.K. He was a Postdoctoral Fellow with the Technical University of Valencia, Spain (2000–2002) and the University of Calgary, Calgary, AB, Canada (2003). He was a Marie Curie Fellow with the Johns Hopkins University (2006–2008). His current research

interests include computational modeling of cardiac bioelectric activity, microscopic mapping of the cardiac electric field, and defibrillation.

Fig. 1.

Octree-based domain discretization. For the sake of simplicity, the basic principles are outlined by using a quadtree, the 2-D analog of a 3-D octree. (a) *Quadtree structure*: the tissue (gray quads, material tag "1") is discretized at a higher resolution, whereas a coarser resolution is used to discretize the bath (white quads, material tag "0"). The solid red line represents the surface of the object. The quadtree is nonconformal, with hanging nodes arising wherever larger cells meet smaller cells. (b) *Dual mesh generation*: the dual mesh is derived from the primal mesh, i.e., the octree cells, by the following transformation rules. The entities vertex, edge, face, and cell in the primal mesh are translated to cell, face, edge, and vertex in the dual mesh, respectively. (c) *Modified dual mesh*: in 3-D, the basic dual mesh approach may also yield nonstandard element types. The insertion of nodes at the transition between larger and smaller boxes eliminates this undesirable property, resulting in the generation of "standard" elements only. (d) *Element types*: the four "standard" elements generated by the algorithm and an example of a "nonstandard" element that could arise in 3-D without modifications of the dual mesh are shown.

Fig. 2.

Marking: the image isosurface approximating the surface of the object (red line) subdivides cells of the dual mesh along the object's boundary. In cells that are not intersected, all vertices pertinent to the cell are marked with the tag assigned to the enclosing cell of the octree.

Cutting: to better approximate the object's surface, cells of the dual mesh that have intersecting edges with the isosurface need to be split. Edges with intersections are subdivided into quarters. If the intersection with a cell edge is located in the second or third quarter of its length, a new vertex is inserted, otherwise the closest vertex is chosen as the new boundary node. The surface of the mesh after splitting (blue line) approximates the object's surface better than the octree (border between light and dark gray areas).

Fig. 3.

Test case for determining mesh quality metrics. (Left panel) Torus geometry is defined by the radius of the generating circle $r_1 = 11.2$, the radius of the revolving circle generating the outer hull $r_2 = 5.3$, and the radius generating the inner hull of the torus shell $r_3 = 7.3$. The generating circle was chosen to lie in a plane defined by the normal vector $\mathbf{n} = [1, 1, 1]^T$. (Right panel) Example of a generated mesh for $H_v = 1$ and $H_m = 1/2$.

Fig. 4.

LV wedge and an anterior papillary muscle were selected by defining masks of simple geometric shapes (shown in transparent red): (A) a pie-sector shape for the wedge and (B) a cone-like shape for the papillary muscle. The segmented gray-level image substacks are shown in the rightmost panels (a and b).

Fig. 5.

Numerical tests using meshes generated with the presented approach for bidomain simulations (eliciting action potential propagation via field stimulation with electrodes located in the surrounding bath) and monodomain simulations (transmembrane stimulation) for both preparations. Results for the medium mesh resolutions are shown. Locations of the stimulus (I_{stim}) and grounding (GND) electrodes are indicated (see text for further detail).

Fig. 6.

Leftmost panels show meshes of (a) the LV wedge and (b) the left anterior papillary muscle at the medium resolutions of 50 and 35 μm , respectively. The gray rectangles outline the inset, presenting a magnified view of the portion of the mesh in the middle panels. The rightmost panels show adaptive mesh generation in the surrounding bath.

Fig. 7. Effects of reducing the target resolution from 100 μm in (a) down to 50 μm in (b) in the wedge preparation. Both the endocardial surface and a cut through the myocardium showcase the difference in structural detail of the mesh.

TABLE I
General Statistics on Mesh Generation: Generated Element Types Are Hexahedra (Hx), Prisms (Pr), Pyramids (Py) and Tetrahedra (Tt)

Model	w100	w50	w35	pm75	pm35	pm25
Number of Vertices ($\times 1e6$)	1.31	8.94	23.98	0.96	7.13	18.07
Number of Elements ($\times 1e6$)	2.22	14.37	36.56	1.69	11.03	27.09
Hx						
Pr						
Py						
Tt						
Mesh Generation Time	<i>min</i>					
Mesh Generation Time Per Element	<i>min/1e6</i>					

The models were generated using two AMD Opteron 2222s (for a total of four cores) running at 3 GHz that are equipped with 32 GB of memory.

TABLE II
Comparison of the Element, Vertex, and Edge Length (EL) Distribution in Bath and Tissue for the Two Selected Models

		w35		pm25	
		Tissue	Bath	Tissue	Bath
Vertices	%	85.6	27.6	81.4	29.7
Hexahedra	%	44.4	6.2	44.3	8.0
Prisms	%	3.7	4.9	2.6	5.3
Pyramids	%	10.8	9.2	9.2	10.2
Tetrahedra	%	11.6	9.3	10.4	10.0
Overall	%	70.5	29.5	66.5	33.5
Min EL	μm	8.0	9.3	3.6	3.6
Mean EL	μm	33.4	35.5	24.0	28.6
Max EL	μm	64.5	1590.0	46.0	2459.1

TABLE III

Mesh Quality Metrics on the Basis of a Torus Shell

H_v	H_m	V_δ	S_δ	$\hat{\epsilon}$	$\alpha_{\min} [^\circ]$	$\alpha_{\max} [^\circ]$	#Cells
1	2	0.173965	0.017025	0.99	5.89	158.83	18586
	1	0.024894	0.034053	0.87	19.92	153.44	84449
	1/2	0.004054	0.013356	0.87	20.41	153.44	376736
1/2	1	0.012797	0.049057	0.85	23.25	136.28	83519
	1/2	0.020560	0.014814	0.85	19.77	140.33	383027
	1/4	0.016683	0.004409	0.85	19.53	140.93	1698038
1/4	1/2	0.000495	0.012335	0.86	26.07	141.50	372961
	1/4	0.001212	0.003189	0.86	20.38	138.38	1731658
	1/8	0.001244	0.000100	0.86	19.57	139.33	8107507
1/8	1/4	0.000503	0.003244	0.80	21.03	140.01	1681390
	1/8	0.000518	0.000715	0.81	20.25	138.99	8148058
	1/16	0.000555	0.000816	0.81	19.04	140.33	33690721

For three given voxel resolutions H_v , mesh resolutions H_m were varied between twice and half of H_v . The volumetric accuracy V_δ , the surface accuracy S_δ , the maximum skewness $\hat{\epsilon}$, the minimum angle $\alpha_{\min}$, and the maximum angle $\alpha_{\max}$ were compared.

TABLE IV

Mesh Quality Metrics for All Electrophysiological Models of the Study

Model	AR_{max}	$\hat{\varepsilon}$	$\alpha_{min}[^{\circ}]$	$\alpha_{max}[^{\circ}]$
w100	6.95	0.93	14.10	158.67
w50	32.11	0.99	2.75	168.73
w35	8.12	0.93	12.85	163.07
pm75	22.61	0.98	9.29	168.59
pm35	7.23	0.96	15.37	160.88
pm25	15.08	0.98	13.68	160.05

The maximum aspect ratio AR_{max} , the maximum skewness $\hat{\varepsilon}$, the minimum angle α_{min} , and the maximum angle α_{max} , were computed for the different resolutions of the papillary muscle and the wedge preparation.

Solution Methods MD (monodomain) or BD (bidomain), Total Time for Elliptic Solution T_{ELL} , Total Time for Parabolic Solution T_{PAR} , Total Time for ODEs T_{ODE} , and Total Execution Time T_t for Simulating Action Potential Propagation

TABLE V

Model		w50		pm35	
Solution		MD	BD	MD	BD
T_{ell}	s/ms	0.0	856.2	0.0	694.6
T_{par}	s/ms	63.2	70.7	57.8	57.4
T_{ODE}	s/ms	3.3	3.1	2.5	2.5
T_t	s/ms	115.4	1096.9	111.2	890.0