SpezialForschungsBereich F 32

Karl–Franzens Universität Graz
Technische Universität Graz
Medizinische Universität Graz

Anatomically accurate high resolution modeling of human whole heart electromechanics

C.M. Augustin A. Neic M. Liebmann
A.J. Prassl S.A. Niederer G. Haase
G. Plank

SFB-Report No. 2015-004 March 2015

A–8010 GRAZ, HEINRICHSTRASSE 36, AUSTRIA

Supported by the
Austrian Science Fund (FWF)

SFB sponsors:

- **Austrian Science Fund** (FWF)
- **University of Graz**
- **Graz University of Technology**
- **Medical University of Graz**
- **Government of Styria**
- **City of Graz**

Anatomically accurate high resolution modeling of human whole heart electromechanics

A strongly scalable algebraic multigrid solver method for nonlinear deformation

C.M. Augustin*, A. Neic, M. Liebmann, A.J. Prassl,
S.A. Niederer, G. Haase, G. Plank†

March 13, 2015

This research was supported by the grants F3210-N18 from the Austrian Science Fund (FWF), NIH 1RO1 HL 101196011 and the EU grant CardioProof 611232. We acknowledge PRACE for awarding us access to resource SuperMUC based in Germany at LRZ (grant CAMEL), and, partially, ARCHER based in the UK at EPCC (project e384).

Electromechanical (EM) models of the heart have been used successfully to study fundamental mechanisms underlying a heart beat in health and disease. However, in all modeling studies reported so far numerous simplifications were made in terms of representing biophysical details of cellular function and its heterogeneity, gross anatomy and tissue microstructure, as well as the bidirectional coupling between electrophysiology (EP) and tissue distension. One limiting factor is the employed spatial discretization methods which are not sufficiently flexible to accommodate complex geometries or resolve heterogeneities, but, even more importantly, the limited efficiency of the prevailing solver techniques which are not sufficiently scalable to deal with the incurring increase in degrees of freedom (DOF) when modeling cardiac electromechanics at high spatio-temporal resolution.

This study reports on the development of a novel methodology for solving the nonlinear equation of finite elasticity using human whole organ models of cardiac electromechanics, discretized at a high para-cellular resolution.

*Christoph M. Augustin, Institute of Biophysics, Medical University of Graz, Harrachgasse 21, 8010 Graz, Austria, christoph.augustin@medunigraz.at

†Gernot Plank, Institute of Biophysics, Medical University of Graz, Harrachgasse 21, 8010 Graz, Austria, gernot.plank@medunigraz.at

Three patient-specific, anatomically accurate, whole heart EM models were reconstructed from magnetic resonance (MR) scans at resolutions of 220 µm, 440 µm and 880 µm, yielding meshes of approximately 184.6, 24.4 and 3.7 million tetrahedral elements and 95.9, 13.2 and 2.1 million displacement DOF, respectively. The same mesh was used for discretizing the governing equations of both electrophysiology (EP) and nonlinear elasticity. A novel algebraic multigrid (AMG) preconditioner for an iterative conjugate gradient (CG) solver method was developed to deal with the resulting computational load. The AMG preconditioner was tailored to nonlinear mechanics under the primary design objective of achieving favorable strong scaling characteristics, as this is key for exploiting current high performance computing hardware.

Benchmark results using the 220 µm, 440 µm and 880 µm meshes demonstrate efficient scaling up to 1024, 4096 and 8192 compute cores which allowed the simulation of a single heart beat in 44.3, 87.8 and 235.3 minutes, respectively. The efficiency of the method allows fast simulation cycles without compromising anatomical or biophysical detail.

Keywords: Cardiac Electromechanics, Algebraic Multigrid, Parallel Computing, Whole Heart Model.

1 Introduction

Electromechanical function of the heart emerges from a complex cascade of processes, which interact across a broad range of spatial and temporal scales. Despite the large body of experimental and theoretical research accumulated, our current understanding of how alterations in subcellular function affect whole organ pumping performance or, conversely, how changes at the systemic level influence subcellular function, remains incomplete. Computational modeling bears high promise as a methodology for integrating experimental data into a mechanistic framework which enables the quantitative observation of complex cause-effect relationships across a broad range of spatio-temporal scales.

While the potential gains from modeling approaches are significant, the challenges to be addressed are daunting. Historically, the development of EP and mechanical models of cardiac function proceeded rather independently than in tandem. The vast majority of EP modeling studies ignored any effects due to mechanical deformation, and, vice versa, most mechanical modeling studies did not represent explicitly any feedback of deformation on EP. While such a simplified EM coupling has proven suitable for addressing a variety of questions [50, 56, 79], the assumption of a unidirectional coupling between EP and mechanical deformation is inaccurate as there is strong evidence that EP is modulated by tissue distension via mechano-electric feedback (MEF) mechanisms [55, 60, 61]. However, from a model development point of view, it is attractive to assume that MEF can be neglected, as this allows to develop EP and mechanical models independently which reduces the complexity of implementation. This effective split into

two separate sequentially executed solution steps is reflected in a notable divergence in the employed methodologies between EP and mechanics modeling communities.

State of the art organ scale EP finite element or finite volume models are discretized at a high spatio-temporal resolution using tetrahedral [109] or hybrid elements [97]. High resolutions are not only necessary to capture the fast transients and steep depolarization wavefronts, but also to resolve fine microscopic scale structural detail [90, 119] as well as functional heterogeneities [18, 54, 74]. In contrast, mechanical models build on the assumption that cardiac deformation is governed by smoother spatial and slower temporal scales, suggesting that the use of much coarser spatio-temporal discretizations may yield sufficient accuracy. Thus the use of higher order cubic Hermite elements became popular [24] as they allow tessellation of stylized ventricular anatomy using a very small number of elements and they avoid *volumetric locking*, however, their major drawback is the limited capacity to accommodate complex geometries. This issue is being addressed now with the use of tetrahedral $P_1 - P_0$, $P_2 - P_0$ or mixed formulation $P_2 - P_1$ elements, which have been introduced only very recently for modeling cardiac mechanics [32, 35, 41]. In general, finite element discretizations with tetrahedral and hexahedral elements of the same order showed similar results in terms of precision and efficiency [19].

Unsurprisingly, as a consequence the cost of model execution has been addressed quite differently as well. In EP modeling, due to a large number of DOF, a pressing need for strongly scalable iterative solvers has emerged [81, 94], whereas for the lower dimensional mechanical models this has not been the case. While high resolution models and corresponding scalable solver algorithms have been developed for vascular applications [3, 58], for cardiac mechanics direct solvers, executed on a very small number of compute cores [32], prevail.

The use of higher spatial resolutions and more flexible tetrahedral or hybrid finite element meshes is also driven by the needs of clinical modeling applications. Generation pipelines for individualized EM models rely increasingly on reconstruction from tomographic imaging [108] which provide an ever increasing level of anatomical detail [63]. Multimodal clinical imaging also facilitates more accurate tissue classification. Clinically important delineation of viable myocardium, fat deposits, cleft spaces and vascularization, as well as substrate abnormalities such as infarct scars [100] or the presence of fibrosis [72, 73] is becoming feasible. These are known to be important factors influencing cardiac deformation, however, the spatial resolutions and methodology currently used impede their elucidation in modeling studies.

Spatial resolution also becomes an important factor when considering bidirectionally coupled, anatomically accurate EM models. A bidirectional link between EP and mechanics demands information being passed back and forth between EP and mechanics solver components. With tomographically reconstructed models the bidirectional projection of data may constitute a significant technical challenge when markedly different spatio-temporal resolutions are used. Either one resorts to using simplified, stylized anatomical representations for both physics to ensure an exact overlap of both domains, a refined electrical mesh is derived from spatial refinement of finite elements forming the mechanical mesh [15], or one opts to use the same high resolution mesh for both physics and accepts the increase in computational load. The advantage of the latter approach

is that convergence studies with bidirectionally coupled EM models are feasible as discretization errors are governed by the spatio-temporal discretization of the EP problem. When attempting to build human whole organ EM models the use of high resolution meshes becomes particularly challenging due to the larger size of a human heart compared to other species used in many previous EM modeling studies [80], and the overall increase in DOF needed for discretizing both ventricles as well as the thin-walled atria [11].

This study reports on the development of a novel methodology for modeling of human whole organ cardiac electromechanics at a high para-cellular resolution. A clinical magnetic resonance (MR) scan was segmented to generate a patient-specific, anatomically accurate whole heart EM model. Tetrahedral finite element meshes were generated from the segmented image stack at resolutions of 220 µm, 440 µm and 880 µm, resulting in meshes with approximately 184.6, 24.4 and 3.7 million tetrahedral elements and 95.9, 13.2 and 2.1 million displacement DOF, respectively. The same mesh was used for discretizing the governing equations of both electrophysiology (EP) and nonlinear elasticity. A novel algebraic multigrid (AMG) preconditioner for an iterative CG solver method was developed to deal with the resulting computational load. The AMG preconditioner was tailored to nonlinear mechanics under the primary design objective of achieving favorable strong scaling characteristics, as this is key for exploiting current high performance computing hardware.

Strong scaling benchmarks using the 220 µm, 440 µm and 880 µm meshes with 95.5, 13.2 and 2.1 million displacement DOF demonstrated efficient scaling up to 8192, 4096 and 1024 compute cores which allowed the simulation of a single heart beat in 235.3, 87.8 and 44.3 minutes, respectively. The efficiency of the method allows fast simulation cycles without compromising in terms of anatomical or biophysical detail, thus enabling modeling studies which elucidate the role of fine scale anatomical features, substrate heterogeneities or mechanistic inquiries into bidirectional feedback loops governing cardiac adaptation processes at a fundamental level. Moreover, meaningful verification of EM models becomes feasible as spatial resolutions can be refined sufficiently to gauge convergence errors.

2 Governing Equations

2.1 Mechanical Model

Cardiac tissue is mechanically characterized as a hyperelastic, nearly incompressible, orthotropic material with a nonlinear stress-strain relationship. The deformation gradient $\mathbf{F}$ describes the deformation $\mathbf{U}$ of a body from the reference configuration $\Omega_0(\mathbf{X})$ to the current configuration $\Omega_t(\mathbf{x})$,

$$F_{ij} = \frac{\partial x_i}{\partial X_j}, \quad i,j = 1,2,3. \tag{1}$$

By convention, we denote $J = \det \mathbf{F} > 0$ and introduce the right Cauchy–Green tensor $\mathbf{C} = \mathbf{F}^\top \mathbf{F}$. The nearly incompressible behavior is modeled by a multiplicative decom-

position of the deformation gradient, see, e.g., [31], of the form

$$\mathbf{F} = J^{1/3}\overline{\mathbf{F}},\ \mathbf{C} = J^{2/3}\overline{\mathbf{C}},\ \text{with}\ \det\overline{\mathbf{F}} = \det\overline{\mathbf{C}} = 1, \tag{2}$$

where $J^{1/3}$ and $J^{2/3}$ are associated with volumetric deformations, while $\overline{\mathbf{F}}$ and $\overline{\mathbf{C}}$ are associated with isochoric processes. Mechanical deformation is governed by the stationary equilibrium equations given as

$$-\operatorname{Div}\left[\mathbf{F}\mathbf{S}(\mathbf{U},\mathbf{X})\right] = \mathbf{b}_0(\mathbf{X}) \quad \text{for } \mathbf{X} \in \Omega_0, \tag{3}$$

where $\mathbf{U}(\mathbf{X},t)$ is the unknown deformation, $\mathbf{S}(\mathbf{U},\mathbf{X})$ is the second Piola–Kirchhoff stress tensor, $\mathbf{b}_0(\mathbf{X})$ are the body forces and Div denotes the divergence operator in the reference configuration.

On the boundary $\partial\Omega_0 = \overline{\Gamma}_{0,\mathrm{D}} \cup \overline{\Gamma}_{0,\mathrm{N}}$ Dirichlet displacement boundary conditions are imposed

$$\mathbf{U}(\mathbf{X}) = \mathbf{U}_\mathrm{D}(\mathbf{X}) \quad \text{on} \quad \Gamma_{0,\mathrm{D}}, \tag{4}$$

with $\mathbf{U}_\mathrm{D}(\mathbf{X})$ being a prescribed displacement, and Neumann force boundary conditions

$$\mathbf{F}\mathbf{S}(\mathbf{U},\mathbf{X})\,\mathbf{N}_0(\mathbf{X}) = \mathbf{G}_0(\mathbf{X}) \quad \text{on} \quad \Gamma_{0,\mathrm{N}}, \tag{5}$$

where $\mathbf{N}_0(\mathbf{X})$ is the exterior normal vector and $\mathbf{G}_0(\mathbf{X})$ denotes a prescribed surface traction.

The deformation of cardiac tissue is governed by imposed external loads such as intracavitary and pericardial pressures, and active stresses intrinsically generated during contraction. The total stress $\mathbf{S}$ is additively decomposed according to

$$\mathbf{S} = \mathbf{S}_\mathrm{p} + \mathbf{S}_\mathrm{a}, \tag{6}$$

where $\mathbf{S}_\mathrm{p}$ and $\mathbf{S}_\mathrm{a}$ refer to the passive and active stresses, respectively. Passive stresses are modeled based on the constitutive equation

$$\mathbf{S}_p = 2\frac{\partial\Psi(\mathbf{C})}{\partial\mathbf{C}}, \tag{7}$$

where Ψ is an orthotropic, invariant-based strain-energy function [30, 47], which results in a frame-independent stress tensor. The strain-energy function Ψ is additively composed of three functions

$$\Psi(\mathbf{C}) = \Psi_\mathrm{vol}(J) + \overline{\Psi}_\mathrm{iso}(\overline{\mathbf{C}}) + \overline{\Psi}_\mathrm{aniso}(\overline{\mathbf{C}},\mathbf{f}_0,\mathbf{s}_0), \tag{8}$$

where $\Psi_\mathrm{vol}(J)$ is the volumetric contribution to the hyperelastic response while $\overline{\Psi}_\mathrm{iso}$ and $\overline{\Psi}_\mathrm{aniso}$ are two volume-preserving parts which relate to the isotropic and an anisotropic behavior, respectively. The prevailing myocyte orientation, referred to as the fiber axis, is denoted as $\mathbf{f}_0$ and the sheet axis, which is perpendicular to $\mathbf{f}_0$ and parallel to the collagen layers which envelop bundles of myocytes, is $\mathbf{s}_0$. Together with the sheet-normal axis $\mathbf{n}_0$,

defined to be orthogonal to the sheet and the fiber orientations this forms a right-handed orthonormal set of basis vectors [47]. The volumetric contribution is given as

$$\Psi_{\text{vol}}(J) = \frac{\kappa}{2} \ln(J)^2, \tag{9}$$

with $\kappa > 0$ being a penalty parameter for enforcing the nearly incompressible behavior of the tissue. The volume-preserving isotropic part is based upon an exponential model [25]

$$\Psi_{\text{iso}} = \frac{a}{2b} \exp\left[b(\overline{I}_1 - 3)\right], \quad \overline{I}_1 = \text{tr}(\overline{\mathbf{C}}), \tag{10}$$

where $a > 0$ is a stress-like and $b > 0$ a dimensionless material parameter, and the anisotropic function is described by $\overline{\Psi}_{\text{aniso}} = \overline{\Psi}_{\text{f}} + \overline{\Psi}_{\text{s}} + \overline{\Psi}_{\text{fs}}$ and

$$\overline{\Psi}_{\text{f}}(\overline{\mathbf{C}}, \mathbf{f}_0) = \frac{a_{\text{f}}}{2b_{\text{f}}} \left\{\exp[b_{\text{f}}(\overline{I}_{\text{f}} - 1)^2] - 1\right\}, \tag{11}$$

$$\overline{\Psi}_{\text{s}}(\overline{\mathbf{C}}, \mathbf{s}_0) = \frac{a_{\text{s}}}{2b_{\text{s}}} \left\{\exp[b_{\text{s}}(\overline{I}_{\text{s}} - 1)^2] - 1\right\}, \tag{12}$$

$$\overline{\Psi}_{\text{fs}}(\overline{\mathbf{C}}, \mathbf{f}_0, \mathbf{s}_0) = \frac{a_{\text{fs}}}{2b_{\text{fs}}} \left\{\exp[b_{\text{fs}}\overline{I}_{\text{fs}}^2] - 1\right\}, \tag{13}$$

with the invariants $\overline{I}_{\text{f}} := \mathbf{f}_0 \cdot \overline{\mathbf{C}}\mathbf{f}_0$, $\overline{I}_{\text{s}} := \mathbf{s}_0 \cdot \overline{\mathbf{C}}\mathbf{s}_0$ and $\overline{I}_{\text{fs}} := \mathbf{f}_0 \cdot \overline{\mathbf{C}}\mathbf{s}_0$. All parameters, the stress-like $a_{\text{f}}, a_{\text{s}}, a_{\text{fs}}$ as well as the dimensionless $b_{\text{f}}, b_{\text{s}}, b_{\text{fs}}$ are considered to be positive. Moreover, the anisotropic responses (11) and (12) only contribute for $\overline{I}_{\text{f}} > 1$ or $\overline{I}_{\text{s}} > 1$, respectively. Note that in the anisotropic model (11–13) sheet-normal properties and corresponding shear terms are omitted as they are either expressible in terms of other invariants or were shown to be of minor significance in shear tests, see [47, §5].

Stresses due to active contraction are assumed to act along the fiber orientation $\mathbf{f}_0$ in the reference configuration. Thus, the active Cauchy stress can be defined as

$$\mathbf{S}_{\text{a}} = S_{\text{a}} \left(\mathbf{f}_0 \cdot \mathbf{C}\mathbf{f}_0\right)^{-1} \mathbf{f}_0 \otimes \mathbf{f}_0, \tag{14}$$

where S_{a} is the scalar active stress induced in fiber direction $\mathbf{f}_0$, see Section 2.3.

2.2 Electrophysiology Model

The spread of electrical activation and repolarization is described by the bidomain equations. Cast in the elliptic-parabolic form, these are given in material coordinates by

$$-\text{Div}\left[J\mathbf{F}^{-1}\left(\boldsymbol{\sigma}_{\text{i}} + \boldsymbol{\sigma}_{\text{e}}\right)\mathbf{F}^{-\top}\nabla\,\phi_{\text{e}}\right] - \text{Div}\left[J\mathbf{F}^{-1}\boldsymbol{\sigma}_{\text{i}}\,\mathbf{F}^{-\top}\nabla V_{\text{m}}\right] = 0, \tag{15}$$

$$-\text{Div}\left[J\mathbf{F}^{-1}\boldsymbol{\sigma}_{\text{i}}\,\mathbf{F}^{-\top}\nabla V_{\text{m}}\right] - \text{Div}\left[J\mathbf{F}^{-1}\boldsymbol{\sigma}_{\text{i}}\,\mathbf{F}^{-\top}\nabla\phi_{\text{e}}\right] = -\beta I_{\text{m}} + \beta I_{\text{i}}, \tag{16}$$

$$\frac{\partial \boldsymbol{\eta}}{\partial t} = \boldsymbol{g}(V_{\text{m}}, \boldsymbol{\eta}, \lambda), \tag{17}$$

where $\boldsymbol{\sigma}_{\text{i}}$ and $\boldsymbol{\sigma}_{\text{e}}$ are the intracellular and extracellular conductivity tensors, respectively;

$$V_{\text{m}} = \phi_{\text{i}} - \phi_{\text{e}} \tag{18}$$

is the transmembrane voltage, with ϕ_i and ϕ_e being the intracellular and extracellular potentials, respectively; β is the bidomain surface to volume ratio;

$$I_\mathrm{m} = C_\mathrm{m}\frac{\partial V_\mathrm{m}}{\partial t} + I_\mathrm{ion}(V_\mathrm{m}, \boldsymbol{\eta}, \lambda) \tag{19}$$

is the transmembrane current density; C_m is the membrane capacitance per unit area; I_ion is the membrane ionic current density, which depends on V_m, a set of state variables $\boldsymbol{\eta}$, and the stretch ratio $\lambda = \mathbf{f}_0^\top \mathbf{C}\mathbf{f}_0$, defined as the ratio between current muscle length and reference length in the unloaded state measured along the fiber direction; I_i is an intracellular current density, specified on a per membrane area basis, which serves as a stimulus to initiate propagation; and ∇ is the gradient operator in the reference configuration.

A computationally less costly monodomain model can be derived from the bidomain model under the assumption that intracellular and interstitial conductivity tensors can be related by $\boldsymbol{\sigma}_\mathrm{i} = \alpha\boldsymbol{\sigma}_\mathrm{e}$. In this scenario Eq. (15) can be neglected and the intracellular conductivity tensor in Eq. (16) is replaced by the harmonic mean tensor

$$\boldsymbol{\sigma}_\mathrm{m} = \boldsymbol{\sigma}_\mathrm{i}\boldsymbol{\sigma}_\mathrm{e}(\boldsymbol{\sigma}_\mathrm{i} + \boldsymbol{\sigma}_\mathrm{e})^{-1}.$$

That is, the monodomain equation is given by

$$-\operatorname{Div}\left[J\mathbf{F}^{-1}\boldsymbol{\sigma}_\mathrm{m}\,\mathbf{F}^{-\top}\nabla V_\mathrm{m}\right] = -\beta I_\mathrm{m} + \beta I_\mathrm{tr}, \tag{20}$$

where I_tr is a transmembrane stimulus current.

At tissue boundaries, no-flux boundary conditions are imposed on ϕ_i and ϕ_e, i.e.,

$$\int_{\partial\Omega} \boldsymbol{\sigma}_\mathrm{i}\nabla\phi_\mathrm{i}\,\mathrm{d}\Gamma = \int_{\partial\Omega} \boldsymbol{\sigma}_\mathrm{e}\nabla\phi_\mathrm{e}\,\mathrm{d}\Gamma = 0 \tag{21}$$

holds. Note that for the numerical examples in Section 6 any effects of deformation upon the distribution of potentials is neglected. Under the conditions considered in here, feedback mechanisms of deformation upon current flow are of minor importance as the period in time where the magnitude of gradients $\|\boldsymbol{\sigma}_\mathrm{i}\nabla\phi_i\|$ and $\|\boldsymbol{\sigma}_\mathrm{e}\nabla\phi_e\|$ is large, is temporally disjoint from deformation [86]. Thus, in Eqs. (15–17), $\mathbf{F}$ is replaced by the identity matrix $\mathbf{I}$ and $J = \det(\mathbf{F})$ is replaced by 1.

2.3 Active Stress Model

Most physiological models of active stress generation are based on sets of ordinary differential equations of the form

$$\frac{\partial\boldsymbol{\zeta}}{\partial t} = \mathbf{h}(\boldsymbol{\zeta}, \boldsymbol{\eta}, \lambda, \dot{\lambda}) = \mathbf{h}(\boldsymbol{\zeta}, \boldsymbol{\eta}, \mathbf{U}, \dot{\mathbf{U}}) \tag{22}$$

$$S_\mathrm{a} = f(\boldsymbol{\zeta}, \boldsymbol{\eta}, \lambda, \dot{\lambda}) = f(\boldsymbol{\zeta}, \boldsymbol{\eta}, \mathbf{U}, \dot{\mathbf{U}}), \tag{23}$$

where the cellular mechanisms of active stress generation are represented by the state vector $\boldsymbol{\zeta}$, which depends on the electrophysiological state of the cell $\boldsymbol{\eta}$, stretch ratio λ and stretch rate $\dot{\lambda}$.

The dependency on λ and $\dot{\lambda}$ accounts for length and velocity dependence of active stress generation while the dependency on the electrophysiological state $\boldsymbol{\eta}$, in particular on the cytosolic calcium concentration $[Ca^{2+}]$, models excitation-contraction coupling.

Conversely, deformation exerts influence over cellular EP via mechano-electric feedback (MEF) mechanisms which are mediated through the length dependence of $\boldsymbol{\eta}$ through Eq. (17), mainly via λ dependent modulation of troponin C binding affinity for calcium, the stretch dependent modulation of ion channel conductivity through Eq. (19), and changes in intracellular and interstitial current flow through Eqs. (15–16). If any of these MEF mechanisms is under study, coupling between EP and deformation is bidirectional. In such scenarios, often referred to as *strong coupling*, the two physics cannot be solved sequentially; they must be solved together.

3 Finite Element Method for Cardiac Electromechanics

3.1 Temporal and Spatial Discretization

Temporal dynamics of cardiac EP and mechanical deformation are governed by very different time scales, suggesting that the use of different time steps for the two physics is advantageous to obtain a computationally efficient numerical scheme. Even when considering the EP problem alone, the time scales involved differ significantly between reaction and diffusion terms. Fast transients such as the upstroke of the action potential are governed by time constants in the µs range whereas slower processes occur at the order of tens or hundreds of µs up to ms. Mechanical deformation can be assumed to occur at similar time scales as slower EP processes such as the cytosolic calcium transients. As momentum is not considered here, there is no explicit time dependence in the stationary equilibrium equation (3). Efficient numerical schemes aim to exploit these differences in temporal resolution by using appropriate update intervals [89].

In this study, we conform to the following conventions. Time is discretized according to $t^k = k \cdot \Delta t$ where Δt is considered to be a global synchronization clock. Time steps used in the individual subproblems are multiples n of Δt with $n \in \mathbb{N}$. In our numerical scheme, three different time steps are used: Δt_{ode}, to discretize the systems of ODEs in Eq. (17) and Eq. (22) representing EP reaction terms and active stress generation; Δt_{par}, to discretize the parabolic diffusion terms given by the system of PDEs in Eq. (16); and Δt_{mech}, to update the current configuration given by the PDE in Eq. (3). Unless otherwise noted, throughout this paper we use $\Delta t_{\text{ode}} = \Delta t_{\text{par}} = \Delta t$ and $\Delta t_{\text{mech}} = n \cdot \Delta t$. Time is discretized by $t^k = k \cdot \Delta t$ for solving the bidomain equations and $t^\ell = \ell \cdot \Delta t_{\text{mech}}$ for referring to the time-varying states of deformation. The instants $t^k = k$ and t^ℓ are equal for $k = n \cdot \ell$.

A global time step of $\Delta t = 5\,\mu\text{s}$ is used for time discretization of the bidomain equations while a larger time step of $\Delta t_{\text{mech}} = 1\,\text{ms}$ is used to solve the mechanical model, owing to its slower dynamics. Thus, the system of ODEs (17) and (22) as well as elliptic PDE

(37) and parabolic PDE (41) are integrated in n time steps of size Δt from $t_{\mathrm{m}}^{\ell-1}$ to t_{m}^{ℓ} using λ and $\dot{\lambda}$ derived from $\mathbf{U}$ and $\dot{\mathbf{U}}$ at time $t^{\ell-1}$.

Piecewise continuous linear ansatz functions are employed for both the electric potentials in Eqs. (37–41) as well as the displacements in Eq. (31). To account for the near incompressibility condition (9) the idea of *static condensation* is applied [45, 101]. Using tetrahedral elements for domain tessellation, this leads to $P_1 - P_0$ elements which are commonly used for cardiac electromechanics applications [5, 98]. The same spatial discretization is used for both physics (see Fig. 4). Hence, the projection of coupling variables between electrical and mechanical grid is avoided.

3.2 Variational Formulation and Discretization of Nonlinear Elasticity Problems

In the numerical examples considered in this study no Neumann boundary conditions were applied (see Section 5.3), and body forces $\mathbf{b}_0$ were neglected. We introduce $\mathbf{w} = [\boldsymbol{\eta}; \boldsymbol{\zeta}]^\top$ which is the combined electrophysiological (17) and myofilament (22) state vector and $\mathbf{q} = [\mathbf{g}; \mathbf{h}]^\top$ which is the respective right-hand side term. Then, the boundary value problem (3–5) combined with the cellular EP model (17) and the myofilament model (22) is formally equivalent to the equations

$$\frac{\mathrm{d}\mathbf{w}}{\mathrm{d}t} = \mathbf{q}(\mathbf{w}, V_{\mathrm{m}}, \mathbf{U}), \tag{24}$$

$$\langle \mathcal{A}_0(\mathbf{U}), \mathbf{V} \rangle_{\Omega_0} = \mathbf{0}, \tag{25}$$

which is valid for all smooth enough vector fields $\mathbf{V}$ vanishing on the Dirichlet boundary $\Gamma_{0,\mathrm{D}}$ [20, Theorem 2.6-1]. The nonlinear operator $\mathcal{A}_0$ is induced by the stress tensor representation (6) and $\langle \cdot, \cdot \rangle_{\Omega_0}$ is the related duality pairing. The left hand side of the variational equation (25) has the physical interpretation of the rate of internal mechanical work and is given by

$$\langle \mathcal{A}_0(\mathbf{U}), \mathbf{V} \rangle_{\Omega_0} := \int_{\Omega_0} \mathbf{S}(\mathbf{w}(\mathbf{U}), \mathbf{U}) : \boldsymbol{\Sigma}(\mathbf{U}, \mathbf{V}) \, \mathrm{d}\mathbf{X}, \tag{26}$$

with the second Piola–Kirchhoff stress tensor

$$\mathbf{S}(\mathbf{w}(\mathbf{U}), \mathbf{U}) = \mathbf{S}_{\mathrm{p}}(\mathbf{U}) + \mathbf{S}_{\mathrm{a}}(\mathbf{w}(\mathbf{U}), \mathbf{U}) \tag{27}$$

and $\boldsymbol{\Sigma}(\mathbf{U}, \mathbf{V})$ is the directional derivative of the Green–Lagrange strain tensor, see [46]. Choosing a time discretization as described in Section 3.1 we compute the stress tensor (27) at time step ℓ using

$$\mathbf{S}(\mathbf{w}^*(\mathbf{U}^\ell), \mathbf{U}^\ell) = \mathbf{S}_{\mathrm{p}}(\mathbf{U}^\ell) + \mathbf{S}_{\mathrm{a}}(\mathbf{w}^*(\mathbf{U}^\ell), \mathbf{U}^\ell) \tag{28}$$

where $\mathbf{U}^\ell$ is the displacement at time step ℓ and $\mathbf{w}^*(\mathbf{U}^\ell)$ is an approximation of $\mathbf{w}^\ell$ which is computed using a forward Euler predictor step as

$$\mathbf{w}^*(\mathbf{U}^\ell) = \mathbf{w}^\alpha + \mathbf{q}(\mathbf{w}^\alpha, t^\alpha, \mathbf{U}^\ell)\Delta t$$

where t^α is the previous electrical time step, i.e.

$$t^\alpha = \ell \Delta t_{\text{mech}} - \Delta t = (\ell \cdot n - 1)\Delta t.$$

To apply the finite element method (FEM) we consider an admissible decomposition of the computational domain $\Omega \subset \mathbb{R}^3$ into M tetrahedral elements and introduce a conformal finite element space

$$X_h \subset H^1(\Omega_0), \;\; N = \dim X_h$$

of piecewise polynomial continuous basis functions φ_i. Using Newton's method and a Galerkin finite element discretization results in solving the nonlinear discretized system to find $\delta \mathbf{U}_h \in [X_h]^3$, $\delta \mathbf{U}_h = \mathbf{0}$ on $\Gamma_{0,\mathrm{D}}$ such that

$$\begin{aligned} \langle \delta \mathbf{U}_h, \mathcal{A}_0'(\mathbf{U}_{\nu,h}^\ell) \mathbf{V}_h \rangle_{\Omega_0} &= -\langle \mathcal{A}_0(\mathbf{U}_{\nu,h}^\ell), \mathbf{V}_h \rangle_{\Omega_0}, \\ \mathbf{U}_{\nu+1,h}^\ell &= \mathbf{U}_{\nu,h}^\ell + \delta \mathbf{U}_h \end{aligned} \tag{29}$$

holds for all $\mathbf{V}_h \in [X_h]^3$, $\mathbf{V}_h = \mathbf{0}$ on $\Gamma_{0,\mathrm{D}}$ at time step ℓ and ν is the index of the Newton iteration. The tangential term, omitting $(\cdot)_h$ to increase readability, is computed by

$$\langle \delta \mathbf{U}, A_0'(\mathbf{U}_\nu^\ell)\, \mathbf{V} \rangle_{\Omega_0} = \int_{\Omega_0} \nabla \delta \mathbf{U}\, \mathbf{S}(\mathbf{U}_\nu^\ell) : \nabla \mathbf{V} \, \mathrm{d}\mathbf{X} + \int_{\Omega_0} \mathbf{F}^\top \nabla \delta \mathbf{U} : \mathbb{C}(\mathbf{U}_\nu^\ell) : \nabla \mathbf{V} \, \mathrm{d}\mathbf{X}. \tag{30}$$

For more details see [2, 3, 46]. The standard finite element method now yields a linear system of equations which is equivalent to the discretized variational formulation (29). Finally, we have to solve

$$\mathbf{K}'(\underline{U}_\nu^\ell)\, \delta \underline{U} = -\underline{K}(\underline{U}_\nu^\ell), \quad \underline{\mathbf{U}}_{\nu+1}^\ell = \underline{U}_\nu^\ell + \delta \underline{U}, \tag{31}$$

with the solution vector $\underline{U}_\nu^\ell \in \mathbb{R}^{3N}$ at the ν–th Newton step and the increment $\delta \underline{U} \in \mathbb{R}^{3N}$. The tangent stiffness matrix $\mathbf{K}' \in \mathbb{R}^{3N \times 3N}$ is calculated according to

$$\mathbf{K}'(\underline{U}_\nu^\ell)[j,i] := \langle \varphi_i, \mathcal{A}_0'(\mathbf{U}_{\nu,h}^\ell)\, \varphi_j \rangle_{\Omega_0}, \tag{32}$$

and the terms of the right hand side are constructed by

$$\underline{K}(\underline{U}_\nu^\ell)[i] := \langle \mathcal{A}_0(\mathbf{U}_{\nu,h}^\ell), \varphi_i \rangle_{\Omega_0}. \tag{33}$$

It is possible to show that the operator $\mathcal{A}_0'(\mathbf{U}_{\nu,h}^\ell)$ is self-adjoint, $[H_0^1(\Omega_0, \Gamma_{0,\mathrm{D}})]^3$-elliptic and bounded [2, 84]. The linearized system (29) therefore admits a unique solution $\delta \mathbf{U}_h$ and the tangent stiffness matrix $\mathbf{K}'$ is symmetric and positive definite, allowing the use of the CG method as an iterative solver.

3.3 Variational Formulation and Discretization of the Bidomain Equation

The corresponding weak form to the elliptic-parabolic system (15–17) is to find sufficiently smooth $(V_\mathrm{m}, \phi_e, \boldsymbol{\eta})$ such that

$$\int_{\Omega_0} (\boldsymbol{\sigma}_\mathrm{i} + \boldsymbol{\sigma}_\mathrm{e}) \nabla \phi_\mathrm{e} \nabla q_1 \,\mathrm{d}\mathbf{X} + \int_{\Omega_0} \boldsymbol{\sigma}_\mathrm{i} \nabla V_\mathrm{m} \nabla q_1 \,\mathrm{d}\mathbf{X} = 0 \tag{34}$$

for all $q_1 \in H^1(\Omega_0)$ with $\int_{\Omega_0} q_1(\mathbf{X}) \,\mathrm{d}\mathbf{X} = 0$,

$$\int_{\Omega_0} \beta C_\mathrm{m} \frac{\partial V_\mathrm{m}}{\partial t} q_2 \,\mathrm{d}\mathbf{X} + \int_{\Omega_0} \boldsymbol{\sigma}_\mathrm{i} (\nabla V_\mathrm{m} + \nabla \phi_\mathrm{e}) \nabla q_2 \,\mathrm{d}\mathbf{X} + \int_{\Omega_0} \beta I_\mathrm{ion}(V_\mathrm{m}, \boldsymbol{\eta}) q_2 \,\mathrm{d}\mathbf{X} = 0 \tag{35}$$

for all $q_2 \in H^1(\Omega_0)$ and

$$\int_{\Omega_0} \frac{\partial \boldsymbol{\eta}}{\partial t} q_3 \,\mathrm{d}\mathbf{X} + \int_{\Omega_0} \boldsymbol{g}(V_\mathrm{m}, \boldsymbol{\eta}, \lambda) q_3 \,\mathrm{d}\mathbf{X} = \mathbf{0}, \tag{36}$$

for all $q_3 \in L^2(\Omega_0)$.

For more details including regularity assumptions and the proof of existence and uniqueness of a solution of the weak system (34–36), we refer to [62].

To reduce the complexity of the problem, the elliptic (34) and parabolic portions (35) of the bidomain equations are decoupled and solved sequentially [115]. First, the elliptic system (34) is solved, then an operator splitting approach is employed to subsequently solve the parabolic system [93, 103, 105]. This allows us to treat reaction and diffusion terms separately. The reaction term represented by the ODE system (36) is solved using an accelerated Rush–Larsen technique, see [89], and the diffusion problem is solved using a Crank–Nicolson scheme. A Galerkin finite element approach and a time discretization as described in Section 3.1 with $t^k = k\Delta t$ yields the following linear system of equations to obtain $(\underline{\phi}_\mathrm{e}^{k+1}, \underline{V}_\mathrm{m}^{k+1}, \underline{\eta}^{k+1}) \in \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^N$ such that

$$\mathbf{K}_{\mathrm{i+e}} \underline{\phi}_\mathrm{e}^{k+1} = -\mathbf{K}_\mathrm{i} \underline{V}_\mathrm{m}^k, \tag{37}$$

$$\underline{\eta}_\mathrm{f}^{k+1} = \underline{\eta}_\mathrm{f}^k \, e^{-\frac{\Delta t}{\tau}} + \underline{\eta}_{\mathrm{f}\infty} \left(1 - e^{-\frac{\Delta t}{\tau}}\right) \tag{38}$$

$$\underline{\eta}_\mathrm{s}^{k+1} = \underline{\eta}_\mathrm{s}^k + \underline{g}(V_\mathrm{m}^{k^*}, \underline{\eta}_\mathrm{s}^k, \lambda^{\ell-1}) \Delta t \tag{39}$$

$$\underline{V}_\mathrm{m}^{k+\frac{1}{2}} = \underline{V}_\mathrm{m}^k - \frac{\Delta t}{C_\mathrm{m}} I_\mathrm{ion}(\underline{V}_\mathrm{m}^k, \underline{\eta}^{k+1}, \underline{\lambda}^{\ell-1}), \tag{40}$$

$$(\widetilde{\mathbf{M}}_\mathrm{i} + \frac{1}{2}\mathbf{K}_\mathrm{i}) \underline{V}_\mathrm{m}^{k+1} = -\mathbf{K}_\mathrm{i} \left(\frac{1}{2} \underline{V}_\mathrm{m}^{k+\frac{1}{2}} + \underline{\phi}_\mathrm{e}^k\right) + \widetilde{\mathbf{M}}_\mathrm{i} \underline{V}_\mathrm{m}^{k+\frac{1}{2}}, \tag{41}$$

where

$$\mathbf{K}_{\mathrm{i+e}}[j,i] := \int_{\Omega_0} (\boldsymbol{\sigma}_\mathrm{i} + \boldsymbol{\sigma}_\mathrm{e}) \nabla \varphi_i \nabla \varphi_j \,\mathrm{d}\mathbf{X}, \tag{42}$$

$$\mathbf{K}_\mathrm{i}[j,i] := \int_{\Omega_0} \boldsymbol{\sigma}_\mathrm{i} \nabla \varphi_i \nabla \varphi_j \,\mathrm{d}\mathbf{X}, \tag{43}$$

$$\widetilde{\mathbf{M}}_\mathrm{i}[j,i] := \frac{\beta C_\mathrm{m}}{\Delta t} \int_{\Omega_0} \varphi_i \varphi_j \,\mathrm{d}\mathbf{X}, \tag{44}$$

$\underline{\eta} = \{\underline{\eta}_\mathrm{f}, \underline{\eta}_\mathrm{s}\}$ with $\underline{\eta}_\mathrm{f}$ being fast acting gating variables solved for by evaluating an analytical solution, τ and $\underline{\eta}_{\mathrm{f}\infty}$ are functions of the rate coefficients which govern channel gating kinetics, and $\underline{\eta}_\mathrm{s}$ are slower acting states, integrated with a forward Euler method [89]. From this construction and $\boldsymbol{\sigma}_\mathrm{i} = \boldsymbol{\sigma}_\mathrm{i}^\top$, $\boldsymbol{\sigma}_\mathrm{e} = \boldsymbol{\sigma}_\mathrm{e}^\top$ we conclude that the matrices (42–44) are symmetric and positive definite. The CG method is therefore applicable to solve the linear systems of equations (37) and (41).

4 Linear Solvers

All linear solvers used in this study were implemented in the publicly available numerical package Parallel Toolbox (pt) [42, 69, 70]. The cardiac bidomain equations and the nonlinear deformation equation were solved in parallel using the Cardiac Arrhythmia Research Package (CARP) [114] as an outer framework which was interfaced with the pt library. Technical aspects have been described elsewhere [76]. CARP makes use of the MPI based library PETSc [7] as the basic infrastructure for handling parallel matrices and vectors. Parallel partitioning of unstructured grids relies upon the graph based domain decomposition ParMetis [51, 52].

Details on the methods employed for solving the bidomain equations in parallel have been reported previously [76, 88, 115]. Briefly, the elliptic portion of the bidomain equations is solved using the algebraic multigrid preconditioner ptAMG [69] with an iterative CG solver method (ptAMG-PCG). The parabolic portion of the bidomain equation is solved using an ω-Jacobi preconditioner, also in combination with the iterative CG solver (ωJ-PCG), using an ω of $2/3$ and two iterations per CG iteration [76]. The ptAMG-PCG method used for solving the scalar elliptic PDE turned out to be of very limited efficiency for solving the nonlinear deformation problem, thus motivating the development of an AMG preconditioner that is adapted for nonlinear biomechanical applications. An extensive description of the underlying concepts is presented in the following.

4.1 Algebraic Multigrid Design

For a detailed introduction to the AMG method we refer to [104, 111]. The initial design of the ptAMG preconditioner was geared towards solving linear systems arising from the discretization of scalar PDEs such as the elliptic PDE embedded in the bidomain equations. To be suitable for systems of equations where each node in a finite element mesh is associated with more than one DOF, the AMG setup had to be altered.

We assume the equation system $\mathbf{K}\underline{u} = \underline{b}$ with $\mathbf{K} \in \mathbb{R}^{N\times N}$ and $\underline{u}, \underline{b} \in \mathbb{R}^N$. The total number of unknowns N can be factorized into $N = k \cdot d$, with k representing the number of finite element grid points and d the number of physical unknowns per grid point. Classical multigrid coarsening methods would compute the coarse grid selection based on the discrete operator $\mathbf{K}$, taking into account the matrix graph alone [9], or additionally also the magnitude of matrix entries [99]. However, in our application, this would lead to different coarse grids for the physical unknowns of one grid point which can translate into bad convergence behavior [38].

This problem is circumvented by computing an auxiliary matrix $\tilde{\mathbf{K}} \in \mathbb{R}^{k \times k}$, defined as

$$\tilde{\mathbf{K}}[i,j] := \|K_{ij}\|_F \tag{45}$$

where $K_{ij} \in \mathbb{R}^{d \times d}$ denotes the subblock of $\mathbf{K}$ which couples the DOF associated with a finite element grid point i with the DOF of the grid point j, and $\|\cdot\|_F$ is the Frobenius matrix norm. A coarsening scheme inspired by the Ruge-Stüben scheme [99] is applied to the auxiliary matrix using

$$\tilde{\mathbf{K}}[i,j] > \epsilon\ \tilde{\mathbf{K}}[i,i] \tag{46}$$

as strong connection criterion for the edge (i,j). The set of FE grid point indices $T := \{0,\dots,k-1\}$ is split into the disjoint subsets of coarse grid $C \subset T$ and fine grid indices $F \subset T$. The prolongation operator $\mathbf{P}$ given by

$$\mathbf{P} := \begin{pmatrix} \mathbf{I}_{\mathrm{CC}} \\ \mathbf{P}_{\mathrm{FC}} \end{pmatrix} \tag{47}$$

interpolates from C onto T. The coarse grid $C \subset T$ can be mapped directly via the identity $\mathbf{I}_{\mathrm{CC}}$ while the fine grid $F \subset T$ needs to be interpolated from the coarse grid via $\mathbf{P}_{\mathrm{FC}}$. The interpolation weight n_i is computed as the cardinality of the set of all coarse grid points interpolating any $i \in F$:

$$n_i := |\{j \in C : \tilde{\mathbf{K}}[i,j] > \epsilon\ \tilde{\mathbf{K}}[i,i]\}| \tag{48}$$

Then, the coarse to fine interpolation $\mathbf{P}_{\mathrm{FC}}$ is defined as

$$(\mathbf{P}_{\mathrm{FC}})_{ij} := \begin{cases} (1/n_i)\, I_{d\times d} & \text{if } \tilde{\mathbf{K}}[i,j] > \epsilon\ \tilde{\mathbf{K}}[i,i] \\ 0 & \text{otherwise} \end{cases} \tag{49}$$

with $i \in F$ and $j \in C$. Finally, the coarse grid operator $\mathbf{K}_{\mathrm{c}}$ is defined as the triple matrix product

$$\mathbf{K}_{\mathrm{c}} := \mathbf{P}^T \mathbf{K} \mathbf{P}. \tag{50}$$

By default, the multigrid coarsening in ptAMG is continued until no strong connections between the remaining unknowns are found. The inverse of the resulting coarse grid system can therefore be approximated by the inverse $d \times d$ block diagonal. Alternatively, the coarsening process can be configured to finish earlier at a finer level. The resulting coarse grid system needs to be inverted by a direct solver. While using a direct coarse-grid solver can greatly reduce the number of solver iterations, the overall setup and solving time was consistently worse in our numerical benchmarks. One pre- and post-smoothing iteration with a Jacobi-smoother is used in a multigrid V-cycle as a preconditioner step. The inverse $d \times d$ block diagonal of the system matrix is used as the approximate inverse in the Jacobi iteration.

4.2 Strongly Scalable Parallelization

Devising iterative solvers and suitable preconditioners with favorable strong scaling properties, i.e., where execution time reduces when the number of compute cores used increases, is of pivotal importance when attempting to model cardiac electromechanics at high spatial resolutions. Departing from an initial implementation of the AMG-PCG method described above, strong scaling experiments were performed, in which an entire single heart beat was simulated with a varying number of processors. Parallel performance and scalability was gauged based on the following performance metrics:

1. *Real time lag factor* ξ_{r}: Execution time divided by the time span simulated.

2. *Speedup* S: The reduction in execution time achieved when using N_{p} parallel processes relative to a reference run using N_{pr} processes

$$S = \frac{T_{\mathrm{pr}}}{T_{\mathrm{p}}}. \tag{51}$$

3. *Parallel efficiency* E: The algorithm's ability to reduce execution time as N_{p} increases, defined as the ratio of the fractional reduction in execution time to the fractional increase in N_{p}

$$E = \frac{T_{\mathrm{p}}}{T_{\mathrm{pr}}} \cdot \frac{N_{\mathrm{pr}}}{N_{\mathrm{p}}}, \tag{52}$$

   where T_{p} refers to the execution time when using N_{p} processes, and T_{pr} refers to the same quantity measured in a reference simulation using N_{pr} processes. An E value of 1 is referred to as linear scaling, because computation time is reduced by the same factor as N_{p} is increased. E values below 1 are known as sub-linear scaling whereas those above 1 are known as super-linear scaling.

Whenever parallel efficiency degraded, parallel profiling data was generated and analyzed using the Scalasca toolset [33] (Score-P, CUBE4, Scalasca) to identify the scaling bottlenecks. Based on these analyses, the following parallelization techniques were implemented:

1. Data exchange on the interfaces between individual subdomains was reworked to better load balance the cost of accumulation of interface data, as shown previously for solving a scalar elliptic PDE [71].

2. Parallel data layout of the multigrid hierarchy was modified to increase the effective network bandwidth of the data exchange across subdomain interfaces. Parallel profiling has shown that a combination of small average message sizes and dense communication patterns significantly decrease the effective network bandwidth on the intermediate multigrid levels. A strategy to counter this phenomenon is to redistribute the parallel data of the intermediate multigrid levels onto fewer active processes [6], thus reducing communication density and increasing average message size at the same time. The reasoning of this approach is, that the increase in effective network bandwidth outweighs the loss of computational units.

The approach implemented in ptAMG is to systematically reduce the number of active processes on each multigrid level. A constant group size $g > 1$ is introduced. On any multigrid level, the range of active processes is split into groups of size g. One process of each group, declared to be the master, will be active on the next coarser multigrid level, while the other processes are declared idle. The idle processes redistribute their current coarse grid system to their group's master, thus reducing the number of active processes on any multigrid level to $1/g$ of the previous level.

The main advantage of this approach is that there is no need for tuning the data redistribution process to individual use cases (e.g. problem size, number of computational units). The optimal choice for g depends only on the rate by which the unknowns are reduced during coarsening. By default g is set to 2.

3. The overall AMG algorithm was split into steps which depend on the matrix sparsity pattern (and thus on the FE graph) and those which depend on the actual matrix entries. When solving Eq. (29), the system matrix is updated in each iteration of the Newton solver, but the sparsity structure remains unchanged as our method is not h-adaptive. Therefore only those setup steps which depend on the matrix values have to be executed.

5 Four Chamber Electromechanical Model

5.1 Generation of Whole Heart Anatomy Model

A processing pipeline consisting of multiple stages was implemented, incorporating segmentation, tissue classification, geometry smoothing and resampling, FE mesh generation and tag transfer, and the assignment of orthotropic eigenaxes. The development of the model generation pipeline aimed at achieving maximum flexibility, allowing the generation of complete four chamber models of the heart including the connected major vessels. Such models are preferable from a modeling point of view as they allow the application of more realistic mechanical boundary conditions [32].

5.1.1 Image Acquisition

The anatomical model was derived from an end diastolic 3D balanced steady state free precession (SSFP) cardiac magnetic resonance imaging acquisition in a sagittal orientation with whole-heart coverage and an isotropic resolution of 1.3 mm. The image was processed using automatic model based segmentation of the left and right ventricle chambers, left and right atrial chambers and the left ventricle myocardium wall [87].

5.1.2 Segmentation and Tissue Classification

As the thin walled atrial walls and the right ventricular wall could not be delineated directly from the image stacks, these walls were introduced by dilating the atrial blood pool

by 2 voxels and a 4 voxel right ventricular wall was created by dilating the corresponding chamber [21]. Classification tags were added to further discriminate between different anatomical regions and tissue types relevant to the subsequent model parametrization. As there is no region-specific signature in the acquired imaging data which would allow for an automatic discrimination, a semi-automatic approach was implemented using various tools such as CAIPI (Fraunhofer, MEVIS) and itksnap (University of Pennsylvania, USA).

5.1.3 Anatomy Smoothing

Depending on the resolution of a given medical image stack, the geometric appearance of the heart's surfaces in a generated mesh can be jagged, which is ideal neither for visual representation nor for computational accuracy. These staircasing artefacts were reduced by smoothing and upsampling the segmented image stack using a variational technique [44].

Briefly, a marching cubes algorithm, as implemented in the VTK library [116], was used to generate a triangulated surface representation of the heart's anatomy. In an iterative variational optimization procedure, the generated surfaces were smoothed. The method avoided volumetric shrinking and limited the maximum nodal displacements to ± 0.5 voxels. The smoothed surface mesh was resampled at an isotropic resolution of 110 µm using the 2D Cairo library [16] and stacked together to generate the smoothed 3D volume. Note that only the anatomical mask was smoothed, while previously assigned classification tags were preserved.

5.1.4 Finite Element Mesh Generation

Segmentation and tagged image stacks were fed into the image based unstructured mesh generation software Tarantula (CAE Software Solutions, Eggenburg, Austria) which builds fully unstructured, boundary fitted, locally refined, hex-dominant hybrid tessellations consisting of hexahedra, prisms, pyramids and tetrahedra [92].

Hybrid meshes were converted into purely tetrahedral element meshes. Classification tags were carried over from the tagged image stacks to the finite element mesh to be used later for the assignment of electrical and mechanical tissue properties during finite element matrix assembly. For the sake of numerical testing, three finite element meshes of the same whole heart geometry were generated at different spatial resolutions in multiples of the image stack resolution of 110 µm. A fine model (FM) was discretized at an average $\bar{h}$ of ≈ 220 µm, a medium model (MM) at $\bar{h} \approx 440$ µm, and a coarser model (CM) at $\bar{h} \approx 880$ µm.

While the CM with a resolution $\bar{h} \approx 880$ µm is, to our knowledge, the highest resolution mesh used for modeling cardiac mechanics, for solving the bidomain equations with acceptable accuracy its resolution is far too coarse. With physiological EP models as used here, predictions of conduction velocity would carry unacceptably large errors. At slower conduction velocities, as they govern conduction in a transmural direction, even an artificial conduction block may occur, thus rendering such a model unsuitable for

bidirectionally coupled EM modeling. Nonetheless, this coarse model is deliberately included in this study to report performance and scalability limits since such smaller models are suitable for being used in weakly coupled electromechanical applications where EP and mechanics can be treated in isolation. Further, results obtained with the CM model are informative for simulating smaller animal models with appropriately higher spatial resolution.

Cardiac tissues are orthotropic in both electrical [17] as well as mechanical properties [47]. Since measuring the eigenaxes of the orthotropic tensor *in vivo* is not yet clinical routine and generates sparse and often noisy data sets, a previously developed Laplace–Dirichlet rule based method (LDRB) is employed [8] to assign the principal eigenaxes along fiber direction $\mathbf{f}_0$, sheet direction $\mathbf{s}_0$ and sheet-normal direction $\mathbf{n}_0$ to each finite element in the ventricles. In the atria no fibers were assigned. The atria were modeled as an unexcitable, non-contracting, isotropic hyperelastic material, that had the principal role of providing more physiological boundary conditions in this model.

5.2 Modeling Electrophysiology and Active Stress Generation

Cellular dynamics were described by the recent Grandi–Pasqualini–Bers (GPB) model of the human ventricular myocyte [37], which was coupled to the Land–Niederer (LN) active stress model [66]. To allow for strong coupling, i.e., to account for length effects on the cytosolic calcium transient, the equation governing the binding of calcium to the low affinity regulatory sites on troponin in the myofilament model (see Eq. (114) in the Appendix of [37])

$$\frac{\mathrm{d}[\mathrm{TnC_L}]}{\mathrm{d}t} = k_{\mathrm{on_{TnC_L}}}[\mathrm{Ca_i^{2+}}]\left(\hat{B}_{\mathrm{TnC_L}} - [\mathrm{TnC_L}]\right) - k_{\mathrm{off_{TnC_L}}}[\mathrm{TnC_L}] \tag{53}$$

is replaced by the analogous Eq. (1) of the LN active stress model [66], given as

$$\frac{\mathrm{d}\,\mathrm{TRPN}}{\mathrm{d}t} = k_{\mathrm{TRPN}}\left(\frac{[\mathrm{Ca_i^{2+}}]}{[\mathrm{Ca_i^{2+}}]_{\mathrm{T50}}(\lambda)}\right)^{n_{\mathrm{TRPN}}}(1-\mathrm{TRPN}) - k_{\mathrm{TRPN}}\,\mathrm{TRPN}. \tag{54}$$

In the combined GPB+LN model (53) is replaced by (54), scaled by the total buffer concentration $\hat{B}_{\mathrm{TnC_L}}$ assumed to be $70\,\mu\mathrm{mol\,L^{-1}}$ in the GPB model. This scaling is necessary to correctly relate fractional occupancy, used by the LN myofilament model, with the concentration of troponin bound calcium, used by the GPB model. Parameters of the GPB model were left unaltered whereas parameters of the LN model were adapted to give human myocardium tension transients when coupled to the GPB calcium transient [107] (see Tab. 1). This study sought to find an optimal set of parameters for combining cytosolic calcium transients, as predicted by the GPB model, with active stress transients, as predicted by the LN model.

In single cell experiments, both the unaltered GPB model and the combined GPB+LN model were paced at a cycle length of 500 ms for a duration of 60 seconds. In the combined GPB+LN model the stretch ratio was kept constant at $\lambda = 1.0$ over the entire protocol. In both models, all state variable transients were plotted over the entire

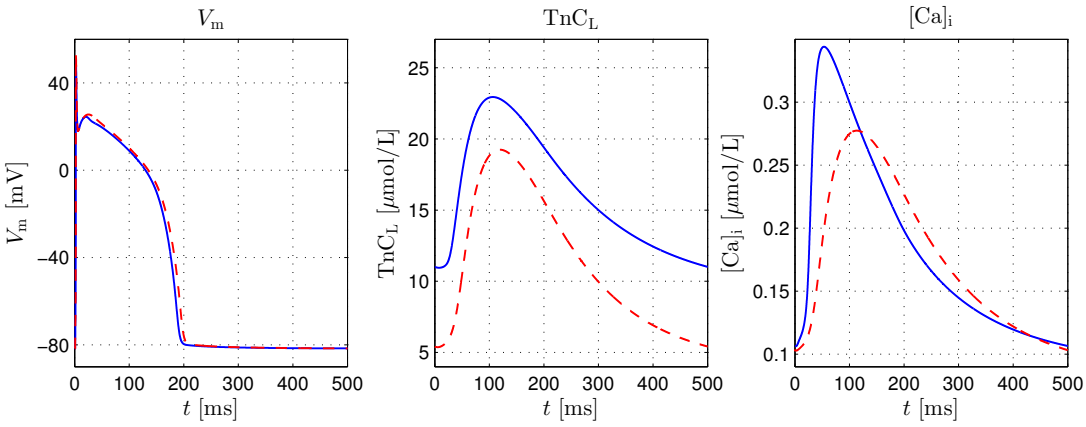

Figure 1: Comparison of V_m, TnC$_L$ and [Ca_i^{2+}] single cell traces between the GPB model (blue) and the GPB+LN model at steady state with $\lambda = 1$ (red).

pacing experiment to ensure that the system had settled close to a stable limit cycle. The effect of the altered troponin buffering equation in the GPB+LN model relative to the GPB model is illustrated in Fig. 1. The effect of strong coupling in the GPB+LN model under fixed stretch conditions in the range $\lambda \in [1.0, 1.2]$, with the same pacing protocol, is illustrated in Fig. 2. The state vector at the end of the pacing protocol is stored to be used later for initializing the organ model. The same GPB+LN model is used throughout the ventricles, known electrophysiological heterogeneities remained unaccounted for.

5.3 Mechanical Boundary Conditions

Homogeneous Dirichlet boundary conditions were applied at the termini of the meshed pulmonary veins and superior and inferior vena cava (Fig. 3). In addition, an elastic apical boundary condition was applied by attaching a small fragment of soft material to the apex of the heart, and applying homogeneous Dirichlet boundary conditions to the opposite side of this material.

The heart was modeled without attachment to the vascular system, with left and right ventricles remaining unpressurized over the entire cardiac cycle. This setup therefore modeled a Langendorff perfused preparation of a human heart [75].

5.4 Material Parameters

The bidomain equations were parametrized to produce conduction velocities of $0.6\,\mathrm{m\,s^{-1}}$, $0.4\,\mathrm{m\,s^{-1}}$ and $0.2\,\mathrm{m\,s^{-1}}$ along the orthotropic eigenaxes $\mathbf{f}_0$, $\mathbf{s}_0$ and $\mathbf{n}_0$ of the tissue. While simulating action potential propagation in the FM ($h = 220\,\mu\mathrm{m}$) was possible with convergence errors that were well within the bounds of measurement uncertainties (2.8 %, 4.5 % and 19.4 % for v_f, v_s and v_n with the numerical settings used in this study), significant errors incurred at coarser mesh resolutions, particularly at slower conduc-

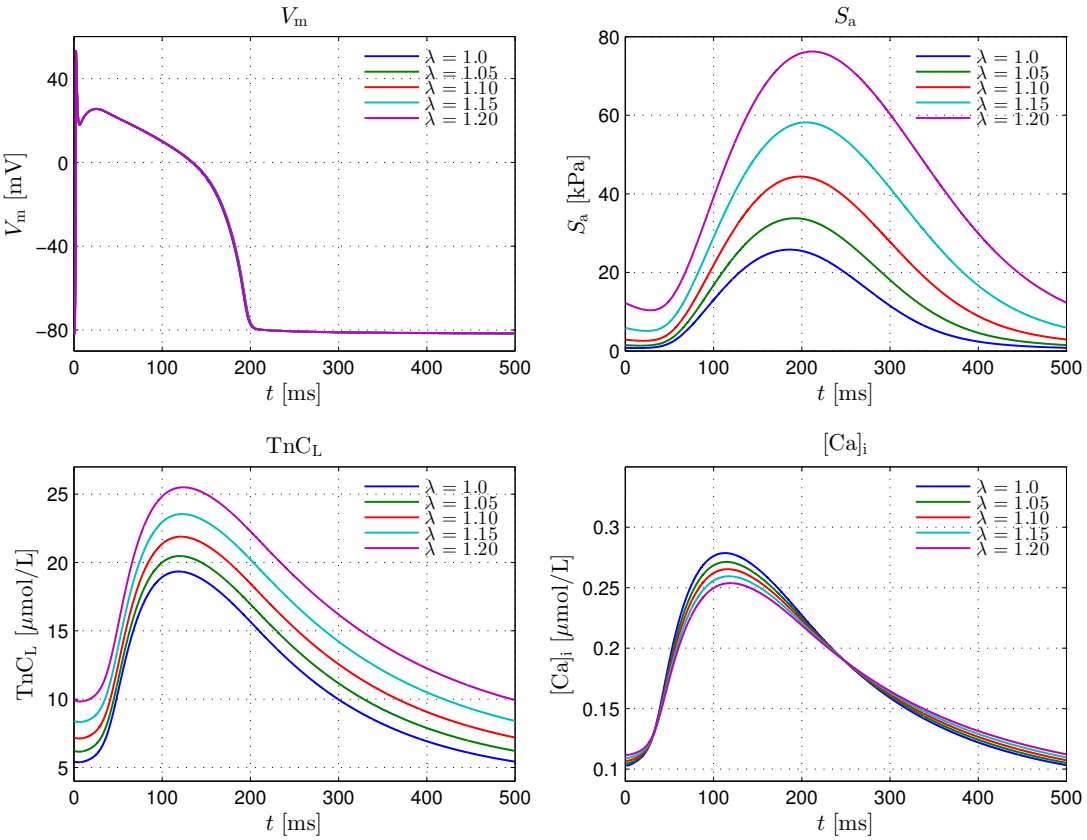

Figure 2: Steady state traces of V_m, isometric tension, TnC_L and $[Ca_i^{2+}]$ in a single cell pacing experiment under varying stretch ratios.

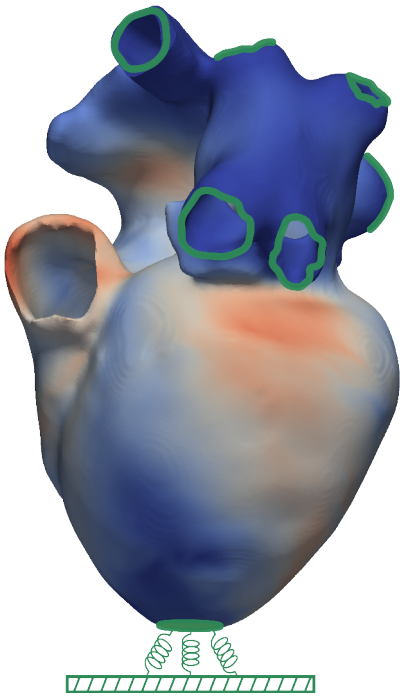

Figure 3: Mechanical boundary conditions: Homogeneous Dirichlet boundary conditions were enforced at the termini of the meshed pulmonary veins and venae cavae, and at the bottom of the soft material attached to the apex.

tion velocities. Without appropriate adjustments of conductivities, artificial conduction block would occur in the CM for propagation along the $\mathbf{n}_0$ axis. Therefore, to minimize differences in mechanical response induced by deviations in activation pattern intracellular conductivities were varied as a function of h to arrive at comparable conduction velocities for FM, MM and CM [23]. Using these tuned conductivity values (Tab. 1) led to the same conduction velocities for each mesh resolution under planar wavefront propagation conditions. It is worth noting that the necessary variations were well within the uncertainty of conductivity measurements reported in the literature [22, 95, 96] for both the FM and MM model, but not the CM model as an $\bar{h} = 880\,\mu\text{m}$ is insufficient when using a biophysically detailed model of cellular dynamics.

5.5 Benchmark Setup

The combined electromechanical four chamber model was used to benchmark performance and scalability of the AMG solver. At time $t = 0$ the endocardium of both ventricles was stimulated with a transmembrane current of $200\,\mu\text{A}\,\text{cm}^{-2}$ for 2 ms, initiating the transmural activation of the ventricles. Electromechanical activity was simulated for 500 ms to cover all phases of an entire heart beat. The same protocol was applied to all three models (CM, MM and FM), with the only alteration being the use of different conductivities (Tab. 1) to arrive at sufficiently similar electrical activation patterns. Simulations were performed with varying numbers of processes (N_p) for each model, with the range of N_p scaled by overall problem size (32–2048, 128–4096, and 512–8192 processes for CM, MM and FM, respectively). For all runs, time spent on solving the individual components of the elliptic PDE (bidomain equations), parabolic PDE and ODE systems was recorded. For the mechanics solver, time spent on matrix assembly and solver time was measured. Furthermore, the number of iterations was stored for both linear and nonlinear solves. The nonlinear mechanical problem was solved using Newton's method with an absolute ℓ^2 norm error reduction of the residual (33) of $\varepsilon = 10^{-6}$. For all linear subproblems, we used the preconditioned CG method with an absolute error reduction of $\varepsilon = 10^{-8}$.

6 Results

6.1 Model Generation and Baseline Simulations

Finite element meshes were generated from the segmented image stack at three different average mesh resolutions $\bar{h}$ of 220 µm, 440 µm and 880 µm. The resulting FM, MM and CM geometry models are visualized in Fig. 4 and metrics for these meshes are summarized in Tab. 2.

Prior to numerical benchmarking, baseline simulations were carried out for each model using the same settings except for the bidomain conductivities, which were modified as a function of $\bar{h}$ (Tab. 1). Maps of local activation time (T_act) were computed from V_m to validate that the electrical activation sequence of the ventricles was consistent between models. As shown in Fig. 5, the total activation time of the ventricles was around 45 ms.

Table 1: Summary of electrical and mechanical material parameters. Bidomain conductivities were varied as a function of mean resolution $\bar{h}$ to keep conduction velocities at prescribed values.

Used for	Parameter		
*P*assive stress	$\kappa = 1000\,\mathrm{kPa}$,	$a_\mathrm{f} = 8.535\,\mathrm{kPa}$,	$b_\mathrm{f} = 5.972\,[-]$,
	$a = 1.333\,\mathrm{kPa}$,	$a_\mathrm{s} = 2.564\,\mathrm{kPa}$,	$b_\mathrm{s} = 10.446\,[-]$,
	$b = 9.242\,[-]$,	$a_\mathrm{fs} = 0.417\,\mathrm{kPa}$,	$b_\mathrm{fs} = 11.602\,[-]$,
*A*ctive stress	$T_\mathrm{ref} = 117.1\,\mathrm{mN\,mm^{-2}}$,	$[\mathrm{Ca}^{2+}]_\mathrm{T50} = 0.52\,\mu\mathrm{mol\,L^{-1}}$,	$\mathrm{TRPN}_{50} = 0.37\,[-]$,
	$n_\mathrm{TRPN} = 1.54\,[-]$,	$k_\mathrm{TRPN} = 0.14\,\mathrm{ms^{-1}}$,	$n_\mathrm{xb} = 3.38\,[-]$,
	$k_\mathrm{xb} = 4.9 \times 10^{-3}\,\mathrm{ms^{-1}}$,		
*B*idomain, 220 µm	$\beta = 1400\,\mathrm{cm^{-1}}$,	$\sigma_\mathrm{il} = 0.234\,15\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{el} = 0.22\,\mathrm{S\,m^{-1}}$,
	$C_\mathrm{m} = 1.0\,\mu\mathrm{F\,cm^{-2}}$,	$\sigma_\mathrm{it} = 0.086\,03\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{et} = 0.13\,\mathrm{S\,m^{-1}}$,
		$\sigma_\mathrm{in} = 0.019\,33\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{en} = 0.13\,\mathrm{S\,m^{-1}}$,
*B*idomain, 440 µm	$\beta = 1400\,\mathrm{cm^{-1}}$,	$\sigma_\mathrm{il} = 0.294\,31\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{el} = 0.22\,\mathrm{S\,m^{-1}}$,
	$C_\mathrm{m} = 1.0\,\mu\mathrm{F\,cm^{-2}}$,	$\sigma_\mathrm{it} = 0.140\,47\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{et} = 0.13\,\mathrm{S\,m^{-1}}$,
		$\sigma_\mathrm{in} = 0.035\,21\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{en} = 0.13\,\mathrm{S\,m^{-1}}$,
*B*idomain, 880 µm	$\beta = 1400\,\mathrm{cm^{-1}}$,	$\sigma_\mathrm{il} = 1.083\,96\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{el} = 0.22\,\mathrm{S\,m^{-1}}$,
	$C_\mathrm{m} = 1.0\,\mu\mathrm{F\,cm^{-2}}$,	$\sigma_\mathrm{it} = 0.760\,53\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{et} = 0.13\,\mathrm{S\,m^{-1}}$,
		$\sigma_\mathrm{in} = 0.090\,40\,\mathrm{S\,m^{-1}}$,	$\sigma_\mathrm{en} = 0.13\,\mathrm{S\,m^{-1}}$.

Figure 4: The CM (left), MM (middle) and FM (right) models of a four chamber heart geometry, discretized with average resolutions of 880 µm, 440 µm and 220 µm respectively. Insets illustrate level of detail in geometry resolution and smoothness of organ surfaces.

Table 2: Finite element meshes of human four chamber heart geometry.

	FM	MM	CM
Resolution	220 µm	440 µm	880 µm
Elements	184 554 709	24 593 765	3 698 356
Nodes	31 971 596	4 413 283	709 878

These activation times which are shorter than anticipated for a physiological activation sequence in a human, are due to the simple endocardial stimulation profile which covered the entire endocardia of both ventricles. This is different from a physiological case where the Purkinje system stimulates the endocardia only over apical and mid segments, but not over the basal segments which are not covered by the Purkinje system. Visual comparison of activation isochrones for later activation times around 35–45 ms suggests that the choice of $\bar{h}$-dependent conductivities led to activation sequences which were indeed very similar between the models (Fig. 5). As expected for the case of closely matching activation sequences, close agreement was also found between models for the electrophysiological quantities V_m, ϕ_e and $[\mathrm{Ca}]_\mathrm{i}$. Similarly, the mechanical quantities $\|\mathbf{u}\|$, $\mathbf{S}_\mathrm{a}$, $\mathbf{S}_\mathrm{p}$ and λ, computed from (25) at a temporal resolution of 1 ms were also compared visually between models. Differences in global mechanical performance of the left ventricle were gauged by tracking the time course of the left ventricular cavity volume $V_\mathrm{LV}(t)$ over an entire heartbeat. A representative result is shown for the FM model in Fig. 6.

6.2 Solver Configuration

Solver parameters were chosen based on the experience gained during the baseline simulations. The parameter ϵ, which is used during the AMG setup to decide which nodes are strongly connected, was set to 0 for both the electrics and mechanics solvers. With $\epsilon = 0$, the AMG setup is based on the matrix graph alone. While this choice had a negative impact on the linear solver convergence, it significantly improved parallel load balancing and, consequently, overall performance. Compared to $\epsilon = 0$, choosing $\epsilon > 0$ yields a larger parallel variability in size of the linear systems arising at the various levels of the multigrid hierarchy, particularly for uneven parallel distributions of the Dirichlet boundaries. These variations in linear system size led to poorly balanced computational load and thus increased MPI waiting time. The number of pre- and post-smoothing sweeps was set to 1. A larger number of smoothing sweeps would improve convergence but overall solving time would be longer as the reduction in iterations is insufficient to compensate for the increased computational cost of the additional smoothing sweeps. The most important factor for improving strong scalability was the grouping of parallel data at the intermediate multigrid levels onto fewer active computing units, as described in Section 4.2. The reduction in communication density and the increase in average message size turned out to be beneficial for $N_\mathrm{p} > 320$. This effect was even greater for $N_\mathrm{p} > 640$, above which scalability stalled with a group size of $g = 1$, meaning

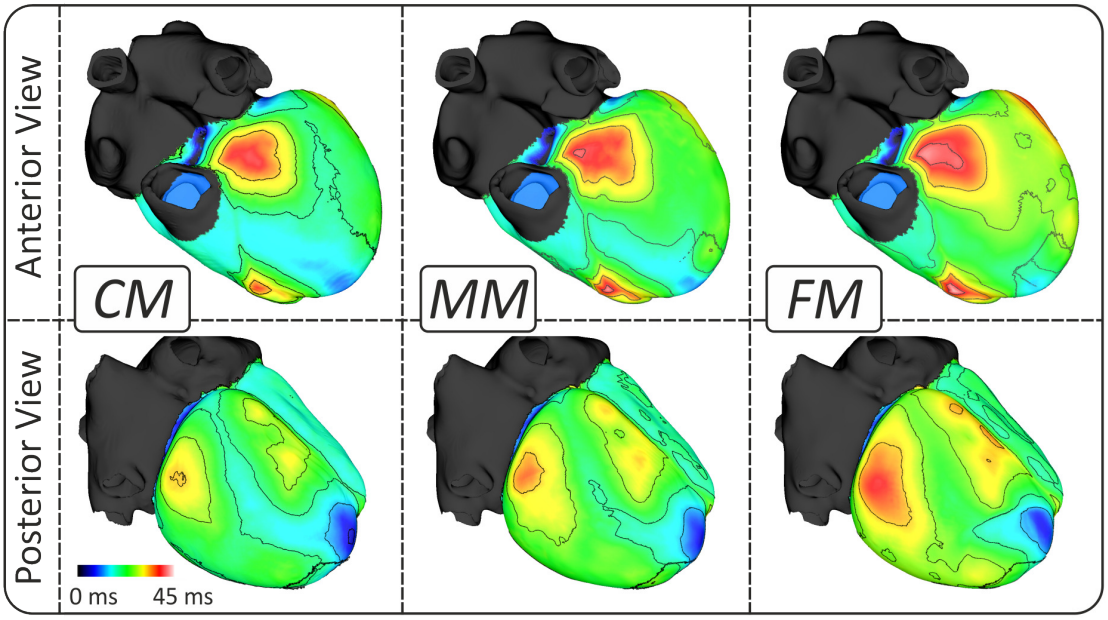

Figure 5: Electrical activation patterns for the CM, MM and FM models, shown from anterior and posterior views.

data redistribution was turned off (Fig. 7). In terms of overall solving time, a group size of $g = 2$ proved to be optimal; a more aggressive reduction of the computing units $g = 3, 4$ was not beneficial.

The performance of the linear solver in terms of convergence and cost of individual solver steps was monitored by recording iteration count, N_{it}, and time per iteration, T_{it}. Over the range of processes N_{p} tested, only a very minor dependency of N_{it} was observed. Thus only one representative graph for $N_{\text{p}} = 512$ is shown to illustrate the course of N_{it} over an entire cardiac cycle (Fig. 8). N_{it} increased with the overall size of the system to be solved, which shows that ptAMG does not fully succeed in preserving convergence independently of the spatial resolution. Nevertheless, the increase in iteration counts is moderate compared to the increase in spatial resolution and thus linear system size from CM to FM (also compare Table 2). The number of nonlinear Newton iterations $N_{\text{N}_{\text{it}}}$ was monitored as well. Depending on the phase of deformation $N_{\text{N}_{\text{it}}}$ varied in the range between 3 to 7 Newton iterations, largely independent of the problem size (Fig. 8).

6.3 Strong Scaling Benchmarks

The final baseline setup, which simulated EP and mechanics of an entire heartbeat over a basic cycle length of 500 ms, was used for performing a strong scaling benchmark. The goal was to determine the strong scaling performance of the major contributing factors (finite element matrix assembly time and solver time) as well as the overall scaling. The initial number of used processes N_{p} was chosen sufficiently large to be

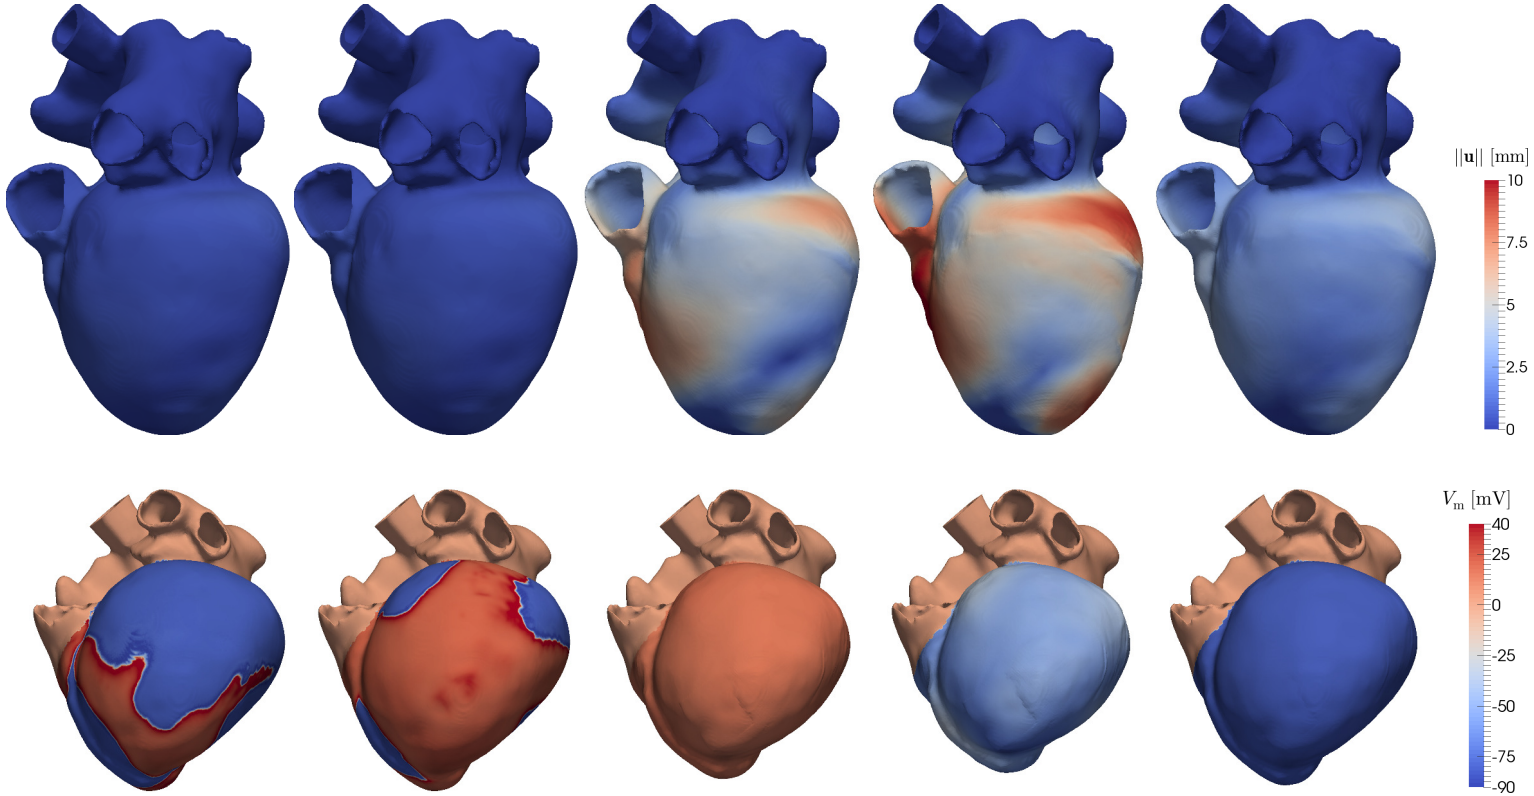

Figure 6: Electromechanical simulation of a human heartbeat. Top row shows displacement $\|\mathbf{u}\|$ and bottom row transmembrane voltage V_{m} at instants 0 ms, 40 ms, 80 ms, 150 ms and 350 ms after the stimulus delivered to the endocardia.

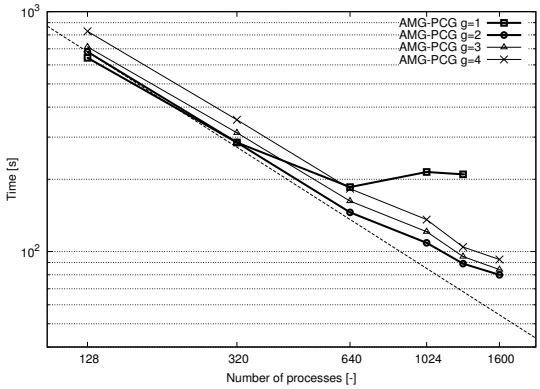

Figure 7: Comparison between different multigrid redistribution strategies used in the AMG-PCG solver. The time needed to solve the linear systems in each newton step of the MM model during a 20 ms simulation was measured. The variable g denotes the group size of the AMG redistribution as described in Section 4.2. With $g = 1$ redistribution is turned off, while starting with $g = 2$, multigrid redistribution is activated.

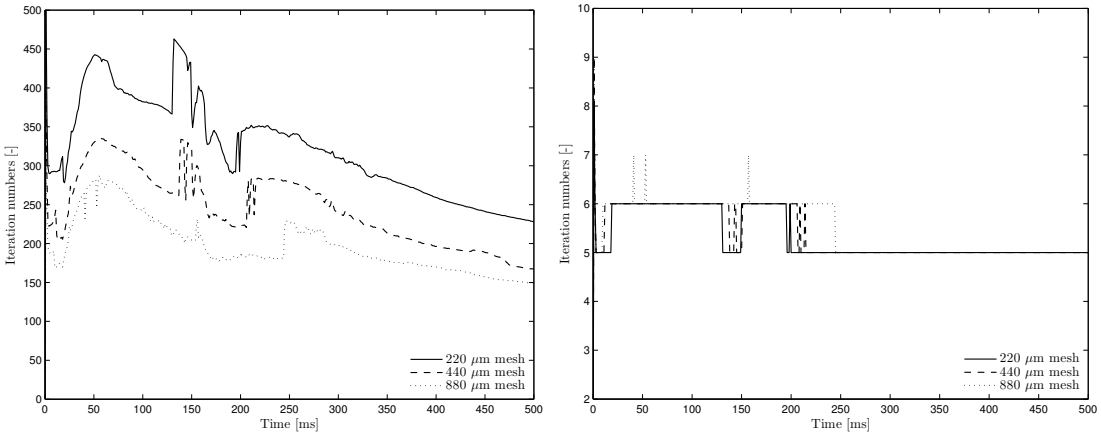

Figure 8: Comparison of the average number of linear solver iterations (left panel) per linear solver step and the number of nonlinear Newton iterations (right panel) over a full heart beat for simulation runs using $N_\mathrm{p} = 512$.

able to fit the electromechanical model into memory. The required N_p was 32, 128 and 512 for CM, MM and FM respectively. Due to its highly parallel nature, the matrix assembly showed linear scaling over almost all N_p probed. With the chosen ptAMG-PCG solver configuration for the deformation problem, strong scaling with high efficiency was achieved over a wide range of processor counts (Fig. 9, left panel). For the CM model a parallel efficiency above 69 % was achieved as far as 512 processes. From 1024 processes parallel efficiency started to degrade, which can be explained by low computational load on each of the very small local partitions ($\approx$ 2500 DOF per process). The MM model showed an efficiency above 64 % up to 2048 processes. Parallel efficiency saturated at $N_\mathrm{p} = 4096$ ($\approx$ 4000 DOF per partition). The FM model showed strong scaling for all

N_p probed. At $N_p = 8192$ (with $\approx 13\,000$ DOF per partition) parallel efficiency was still at 58 %. Further increasing $N_p > 8192$ could not be tested, due to an unresolved fault of the MPI library on the available HPC system. Measured timings and parallel efficiency values are listed in Table 3. Overall, for electromechanical simulations using the monodomain equations to represent EP, computing costs are clearly dominated by solving the deformation equation which is between 3× and 7× as expensive as the matrix assembly costs. In comparison, the costs of solving the parabolic PDE and the system of ODEs which make up the monodomain model are negligible.

When considering a bidomain model to represent EP in an electromechanical simulation, the distribution of costs is markedly different. As shown in the right panel of Fig. 9, computational costs of solving the elliptic portion of the bidomain equation alone is comparable to solving an electromechanical problem with a monodomain EP model. In addition, solving of parabolic and elliptic PDEs can be seen to be far less scalable as solving the deformation problem on the same grid. This is due to the smaller number of degrees of freedom per node which leads to a linear system to be solved that is smaller by a factor of 3× and a computational load that is smaller by a factor of 9×. Therefore, the critical surface to volume ratio of a local partition which translates to a communication to compute load ratio, is hit earlier, at $\approx N_p/3$. This confirmed by Fig. 9, where scaling efficiency can be seen to drop markedly at $N_p = 512$, 2048 and 4096 for CM, MM and FM respectively.

6.4 Impact of Length Dependence upon Calcium

The effect of length dependent calcium buffering of troponin C on cytosolic calcium was investigated by comparing two sets of simulations using the same settings except for the reference value $[Ca^{2+}]_{T50}(\lambda)$ in Eq. (54). In the strong coupling scenario $[Ca^{2+}]_{T50}(\lambda)$ was used as described in [66], i.e.,

$$[Ca^{2+}]_{T50}(\lambda) = [Ca^{2+}]^{ref}_{T50}(1 + \beta_1(\lambda - 1)) \tag{55}$$

whereas in the weak coupling case $[Ca^{2+}]_{T50}(\lambda)$ was kept constant at the value corresponding to $\lambda = 1.0$, i.e., $[Ca]_{T50}(\lambda) = [Ca^{2+}]^{ref}_{T50}$. Fig. 10 illustrates the effect of length dependent calcium binding affinity on calcium transients, stretch ratio λ and active stress S_a.

7 Discussion

This study describes novel methodology for modeling human whole heart electromechanics at a high spatial resolution. Such high resolutions are beneficial for several reasons: (i) cardiac anatomy can be faithfully represented with high geometric fidelity; (ii) achieving fully, or at least nearly, converged simulation results becomes feasible in strong coupling scenarios where mechanical deformation influences the electrophysiological state of cells via MEF mechanisms; and (iii) both fine scale structural heterogeneities such as vessels, fat deposits, scarred tissue and fibrosis as well as functional heterogeneities

Table 3: Strong scaling benchmark results for a monodomain electromechanical simulation of a heartbeat with CM, MM and FM model: Table columns are number of cores N_p, and timings for matrix assembly (MA), solving (Solve), monodomain (Mono), total time and parallel efficiency E.

	N_p	MA	Solve	Mono	Total	E
CM	32	5746 s	25 088 s	12 852 s	43 687 s	1.00
	64	2928 s	13 054 s	6381 s	22 363 s	0.98
	28	1488 s	6946 s	3290 s	11 724 s	0.93
	256	771 s	4006 s	2009 s	6787 s	0.80
	512	404 s	2372 s	1175 s	3951 s	0.69
	1024	213 s	1633 s	812 s	2658 s	0.51
MM	128	9503 s	50 122 s	20 042 s	79 667 s	1.00
	256	4835 s	25 752 s	10 683 s	41 270 s	0.97
	512	2484 s	13 545 s	5541 s	21 572 s	0.92
	1024	1293 s	8163 s	3030 s	12 486 s	0.80
	2048	716 s	5212 s	1889 s	7817 s	0.64
	4096	461 s	3563 s	1246 s	5270 s	0.47
FM	512	17 572 s	79 413 s	34 583 s	131 568 s	1.00
	1024	8896 s	40 842 s	17 760 s	67 498 s	0.97
	2048	4596 s	23 388 s	9952 s	37 936 s	0.87
	4096	2506 s	14 086 s	5078 s	21 670 s	0.76
	8192	1743 s	9363 s	3014 s	14 120 s	0.58

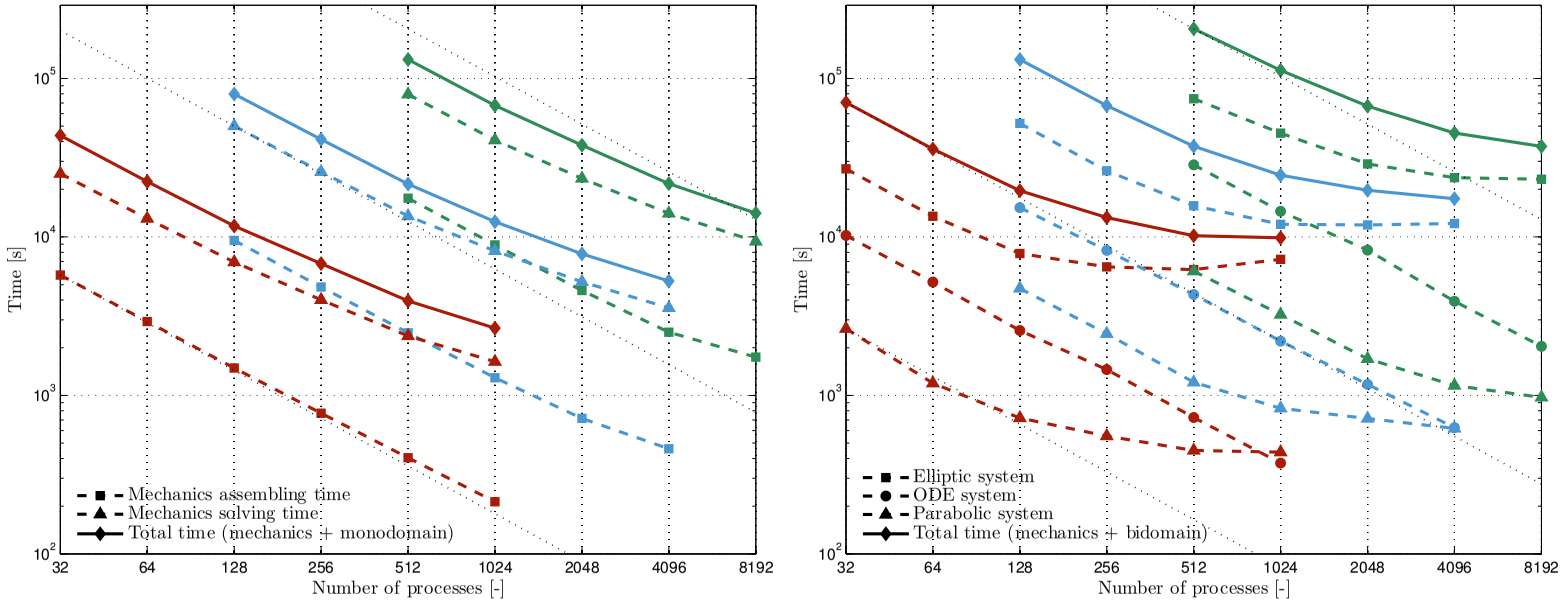

Figure 9: Strong scaling profile of an electromechanical model using a monodomain (left) or bidomain (right) EP model. Red, blue and green traces refer to CM, MM and FM simulation runs respectively, over a time frame of 500 ms. On the left hand side the assembly time for mechanics (32–33) and the solve time (i.e. preconditioner setup and solving) for all linearized systems of equations (31) is given. In addition, we show the total solving time of a monodomain electromechanical simulation. Depicted in the right plot is the computational time for all elliptic systems (37), the computational time for all ODE systems (38–39) and the computational time for all parabolic systems (41). In these three cases the computational time includes assembling, preconditioner setup and solving times. Solid lines correspond to the total solving time of a bidomain electromechanical simulation.

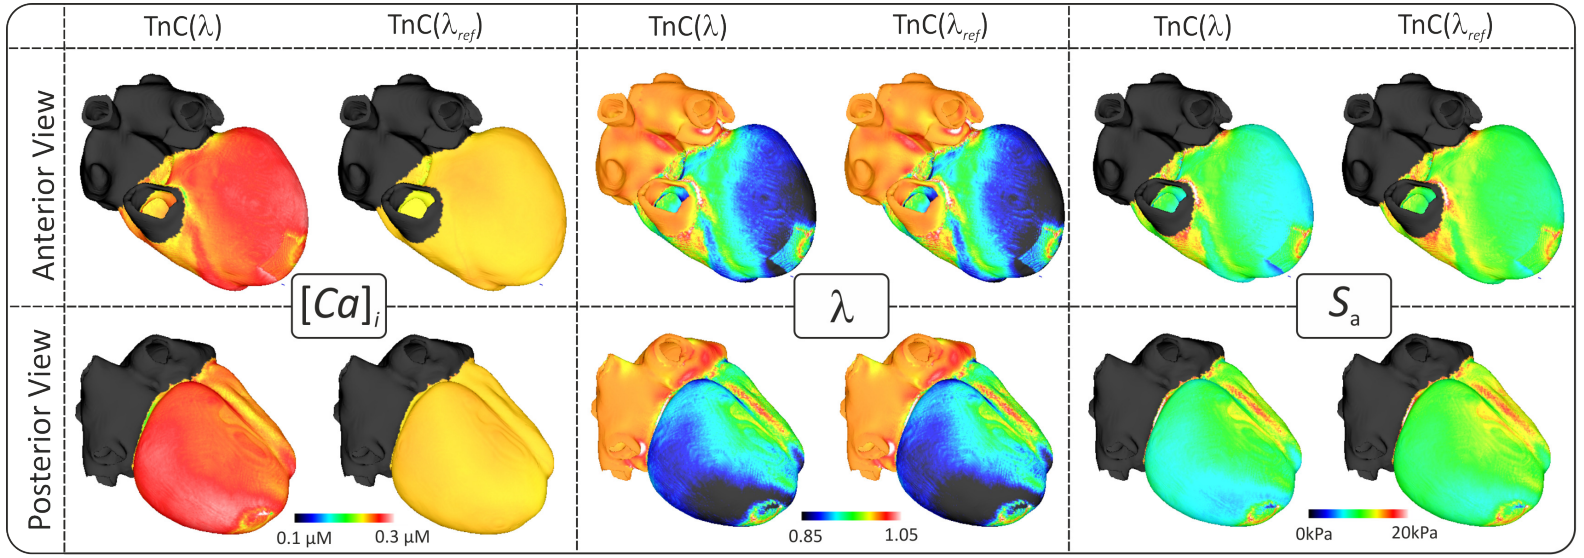

Figure 10: Effects of length dependence calcium binding affinity of troponin C (TnC). Shown are the distribution of intracellular calcium $[Ca^{2+}]_i$ and stretch ratio λ at $t = 130\,\mathrm{ms}$, and active stress S_a at $t = 190\,\mathrm{ms}$ in the strong coupling (TnC(λ)) and weak coupling (TnC(λ_{ref})) scenario.

due to spatial variation in micro-perfusion, metabolic state or protein expression can be explicitly included in a model. The computational burden imposed by such models is significantly larger when compared to lower dimensional models as they are standard in cardiac mechanics applications [57].

This issue of computational cost has been addressed by focusing on a massively parallel iterative solving method for large deformation problems, allowing us to exploit the full potential of modern HPC hardware. In particular, we developed a novel AMG preconditioner for an iterative CG solver method for application to nonlinear mechanics. The primary design objective was to achieve favorable strong scaling characteristics up to a sufficiently large core count such that execution times were reduced and simulations were computationally tractable. Depending on the chosen resolution, benchmark results demonstrated efficient scaling up to 1024, 4096 and 8192 compute cores which allowed the simulation of a single heart beat in 44.3, 87.8 and 235.3 minutes with the CM, MM and FM models respectively. The achieved efficiency makes sufficiently short simulation cycles feasible without compromise of anatomical or biophysical detail.

7.1 Computational costs

The computational burden imposed by high resolution electromechanical models demands the most efficient numerical approaches in order to keep simulations tractable. In the cardiac EP modeling community, two approaches are currently investigated, the use of spatio-temporal adaptivity [10, 26], polynomial adaptivity [1], or, as in this study, the use of strongly scalable solvers which reduce execution times by engaging a larger number of computational units, be it traditional CPUs [81, 94], or acceleration devices such as GPUs [76], or a combination of both parallelization and adaptivity [27]. Solving deformation problems in cardiac or vascular modeling applications relies more on reducing cost by resorting to lower dimensional EP models with slow upstroke velocities combined with simplified active stress models [5, 85] and weak unidirectional coupling to allow for coarse spatial discretization of mechanics [5, 41, 79]. In terms of solver methods used, most studies relied upon direct solvers which tend to become unfeasible for high resolution 3D problems due to linear system size and sparsity pattern [43]. Exceptions exist where strongly scalable solvers were employed, so far only in vascular modeling studies [3, 58] or for a small number of cores [41]. Another approach to reduce costs would be the use of grids of different resolution, a fine mesh for discretizing the bidomain equations to capture the steep wave fronts, and a coarser mesh for discretizing the deformation equation, as mechanical solutions are supposed to be much smoother in space.

Departing from our previously developed bidomain solver infrastructure for which strong scaling properties have been reported in detail [76, 81, 88], we aimed to achieve similar scalability for solving nonlinear deformation problems. As shown in Fig. 9, left panel, scaling efficiencies of 51 %, 47 % and 58 % (Tab. 3) could be achieved when using 1024, 4096 and 8192 cores for executing the CM, MM and FM model, respectively. Scalability of the two major cost factors, matrix assembly and solving of the linear systems, are shown in Fig. 9. As expected, matrix assembly scales almost linearly as

this is an embarrassingly parallel problem. Solver scalability levels off when hitting a critical surface-to-volume ratio between communication, $\propto$ to the surface of the local partition, and compute cost, $\propto$ to the volume of the local partition. When compared to the scalability of solving the monodomain equation, the much better scalability of the deformation solver becomes apparent. Noticeable degradation in scalability starts at lower N_p of 256, 1024 and 4096 cores for CM, MM and FM model, respectively. This is due to the less favorable compute load to communication ratio of a scalar PDE with one degree of freedom per node as compared to the three degrees of freedom per node in the deformation equation. Moreover, a significant portion of the mechanics solver work is dedicated to matrix assembly which is computationally quite expensive, but highly scalable as there is no communication cost involved in this step.

While for most electromechanical problems the combination of monodomain and deformation suffices, there are applications of practical relevance where a bidomain formulation may be more appropriate, such as investigating the influence of deformation on the T-wave of the ECG [53]. For this sake scaling data for bidomain simulations were added as well as a reference. As shown in Fig. 9, right panel, scalability of the elliptic and parabolic solve is comparable but solving the elliptic PDE is roughly one order of magnitude more expensive than solving the parabolic PDE. Solving the biophysically detailed GPB+LN model is more costly than the parabolic solve, but since state variables $\boldsymbol{\eta}$ do not diffuse, the ODE solver step is embarrassingly parallel, which is reflected in a linear scaling graph.

A note of caution is warranted when comparing execution times for EP and mechanics, as a much finer time step of $\Delta t = 5\,\mu\mathrm{s}$ was used for EP than for mechanics with $\Delta t = 1\,\mathrm{ms}$, that is, the EP problem was solved 200$\times$ more frequently than the deformation problem. Therefore, a single solver step of the deformation problem is 470–560$\times$ and 75–160$\times$ more expensive than a single solver step of the monodomain and bidomain equation, respectively.

Preliminary data suggest that further significant improvements in scalability are achievable by exploiting shared memory parallelism in a hybrid MPI-OpenMP approach. The main scaling benefits are a result of the shorter latencies of shared memory data access and the more efficient use of interconnect hardware in terms of latency and bandwidth as less messages are passed, each message being of larger size. Another potentially effective way of enhancing scalability is to employ latency hiding CG methods [34].

While both approaches are promising, a detailed exploration of their suitability for cardiac electromechanics applications is beyond the scope of this study.

7.2 Representing Cardiac Anatomy

Cardiac anatomy plays a role of pivotal importance in electrophysiological and mechanical function. Computer models therefore aim to represent anatomy as faithfully as possible, however due to limitations in image acquisition, image segmentation and registration techniques and the difficulty in tessellating anatomically complex 3D structures into discrete finite element models of sufficient mesh quality, most reported studies resorted to using simplified, idealized FE models. In earlier modeling studies, the left

ventricle was modeled in isolation as an idealized geometry of cylindrical or ellipsoidal shape [14, 77]. Later studies used more elaborate finite element models which incorporated anatomically realistic biventricular geometries based on histological sections [39, 102, 110, 113]. Due to the cost of generating detailed anatomical models, only one representative heart per species was available in these studies. While such an "one heart fits all" approach proved to be of great utility for studying generic mechanisms, this approach is limited in its capability to make predictions for individual patients, due to the wide anatomical inter-subject variability of the heart. Therefore, driven by major advancements in magnetic resonance (MR) and diffusion tensor (DT) MR imaging, a new trend emerged over the past few years towards models derived from tomographic medical images. Traditionally, the most widely used approach is based on using cubic Hermite finite elements to approximate the biventricular geometry [40, 82, 112]. Techniques were developed to register cubic Hermite template models with segmented medical image stacks [64]. More recently, the cubic Hermite method has also been applied to generating models of the thin walled atria, but significant processing efforts and manual steering is necessary [36]. The main advantage of this cubic Hermite element based method is its capability to approximate cardiac geometry with a limited number of finite elements, keeping compute and memory costs tractable even when using modest computing hardware. However, with regard to creating biophysically detailed and anatomically accurate models coarse mesh approaches pose three challenges. Firstly, they are limited in their ability to capture fine anatomical structures, including papillary muscle, thin walled atria and trabeculation. Secondly, They are limited in their ability to capture regional heterogeneities, including ion channel gradients, peri infarct zones, local fiber dispersion and vasculature. Finally, most approaches rely on the adaptation of template meshes which may not confirm to the breadth of patient anatomies. Moreover, finding an optimal tradeoff between geometric fitting accuracy, mesh quality and stability for coarse cubic hermite based models under large deformations can be challenging [65]. These issues can all be addressed through the use of fine resolution unstructured tetrahedral or hybrid meshes [41, 97], generated directly from image segmentations. The ability to generate models automatically from segmentations improves the robustness of the model generation process. This will be beneficial for future clinical applications and will pave the way towards making models available to non-expert users.

Advanced image acquisition and processing [13, 90] as well as unstructured mesh generation pipelines [92] have been developed for such purposes. These methods permit the automatic construction of anatomical FE models which represent individual geometries at an unprecedented level of geometric fidelity. So far, such detailed models have been used in electrophysiological modeling studies [12], but not yet in any mechanical modeling studies. Gurev and co-workers reported a method which combined an anatomically detailed biventricular tetrahedral FE geometry model for simulating electrophysiology, with a cubic Hermite FE model for mechanics [40]. The method allows wrapping a cubic Hermite template model around the detailed tetrahedral FE model. While efforts were made to minimize differences between the meshes, due to the significantly lower level of detailed of the mechanical model, achieving a perfect overlap was impossible. This increases the complexity, of the model generation and simulation process, particularly in

strongly coupled electro-mechanical simulations where projection operators would have to be constructed to transfer data between two FE spaces of imperfect overlap.

An alternative method based on CAD reconstruction of hearts, which uses a four chamber representation plus the attached large vessels similar to that used in this study, has been reported recently [5]. As details on the model generation process have not been published yet, it remains to be seen whether a CAD based approach is viable for generating tomographically reconstructed FE models. This study has similarities with ours, in particular that $P_1 - P_0$ tetrahedral elements were used, facilitating an accurate approximation of the gross anatomy and the use of the same FE mesh for both electrophysiology and mechanics. The major difference of this study from ours is that the low resolution used precludes the use of physiologically realistic models of cellular dynamics such as the GPB+LN model used in our study.

The four chamber heart model used in our study represents all details which can be delineated in a standard clinical image stack. The mesh generation methodology draws from previously developed tools [92], which permits, in theory, to represent any discernible structure of the heart down to the microscopic level [90]. All FE meshes in this study were generated fully automatically, without any manual intervention. Any anatomical deviations of the model are within the uncertainty of the segmentation. Generating the CM, MM and FM meshes on a desktop computer took 51, 81 and 205 minutes respectively.

7.3 Bidirectional Coupling

From a physiological point of view, strong or bidirectional coupling is important as this mediates acute and regulatory MEF, i.e., changes in electrophysiological state of a cell as a function of tissue deformation, measured as stretch ratio λ. From a modeling point of view, strong coupling adds technical complexity as electrophysiology and mechanics cannot be treated independently, necessitating the solution of a coupled mono- or bidomain and deformation equation system. Three major mechanisms mediate MEF, by which deformation alters the electrophysiological state of the tissue: (i) deformation alters geometry and thus electrical gradients $\nabla\phi_\mathrm{i}$ and $\nabla\phi_\mathrm{e}$ respectively; (ii) currents through stretch activated channels $I_\mathrm{SAC}(V_\mathrm{m}, \lambda)$ directly alter transmembrane voltage V_m which indirectly influences upon cellular signaling; (iii) stretch alters binding affinity of calcium to troponin C which influences cytosolic calcium transients and diastolic calcium levels. As calcium is involved in many important signaling pathways in a cell this effect may play important roles in both acute and long term regulation of contractility. While our framework would easily accommodate all three MEF pathways, for the sake of testing the scalability of the developed solver methods we restricted ourselves to the feedback through calcium binding affinity of troponin C. In the single heart beat scenario, effects on electrical gradients are very small as electrical gradients are large only prior to contraction, but are very small during contraction when cells are at the plateau of the action potential or slowly repolarize [117]. While the existence of stretch activated channels in cardiac muscle has been postulated [49], in models they are largely treated phenomenologically as a "black box", as the molecular mechanisms which relate stretch

to channel conductance remain poorly understood. Various mathematical structures for I_{SAC} currents have been published, but parametrized for other species than humans [106] or without parametrization at all [59].

As illustrated in Fig. 10, distribution of stretch ratio λ and active stress S_{a} during a heartbeat is quite heterogeneous. This is particularly apparent around the apex where the ventricles were elastically anchored, as well as along sites of transition between actively contracting myocardium and passive structures such as the superior vena cava or the interface between ventricles and atria which were modeled as mechanically passive material in this study. Differences between strong coupling case $\mathrm{TnC}(\lambda)$ and weak coupling case $\mathrm{TnC}(\lambda_{\mathrm{ref}})$ in terms of λ and S_{a} were rather minor (Fig. 10). Overall, stresses were slightly higher in the weak coupling case. This was due to the length dependent calcium binding affinity being higher compared to the $\mathrm{TnC}(\lambda)$ case where under the given loading conditions and in absence of any cavitary pressures the stretch λ was smaller than 1 during contraction, entailing a comparably lower calcium-bound [TnC] and thus less force generation. As length dependence of active stress in the LN model is modeled as $S_{\mathrm{a}} \propto h(\lambda) \times \mathrm{XB}([\mathrm{TnC}](\lambda))$ where $h(\lambda)$ and XB refer to filament overlap and fraction of force generating crossbridges, in the weak coupling case length dependent force generation is mediated only by filament overlap $h(\lambda)$.

The observed heterogeneity in $\lambda(\boldsymbol{x})$ translated into a marked heterogeneity in cytosolic calcium $[\mathrm{Ca}_{\mathrm{i}}^{2+}](\boldsymbol{x})$ only in the strong coupling case $\mathrm{TnC}(\lambda)$, whereas in the weak coupling case a $\mathrm{TnC}(\lambda_{\mathrm{ref}})$ heterogeneity at the same instant in time was marginal (compare left and right in $[\mathrm{Ca}_{\mathrm{i}}^{2+}]$ panel of Fig. 10). It is important to note that this effect on calcium heterogeneity can only be captured in a strongly coupled model. In most weak coupling scenarios, EP models such as the GPB model are coupled with stress models such as the LN model in a way where calcium buffered [TnC] is tracked by both EP and stress model using different equations. That is, the GPB model tracks TnC concentration based on Eq. (53) and the LN model based on Eq. (54). This is not only inconsistent from a physical point of view, such approaches may also miss a potentially important feedback mechanism. While this study is far from being an exhaustive evaluation the effects of strong coupling, as this was beyond our scope, these results suggest that such effects may be non-negligible, depending on the particular question under study.

Furthermore, it is assumed that deformation is governed by slower time scales and that solutions are much smoother in space compared to electrophysiology which is characterized by fast transients and steep wave fronts. This notion, if correct, suggests that coarser spatio-temporal discretization might be suitable, thereby reducing computational cost when solving the deformation equations. The exact choice of time and space step for mechanics and the resulting errors are likely to depend on the scenario to be modeled and the level of numerical accuracy one aims to achieve. While exhaustive convergence studies are missing for the complex models as used in this study, numerical work with simplified 1D strand models suggest that spatio-temporal discretization in a strongly coupled scenario also requires fairly fine time and space steps [117]. It therefore remains to be seen whether or not the use of a coarser mesh for mechanics pays off in terms of reduced execution time and overall modeling cost. Additional costs would be incurred due to the need of generating two meshes of different spatial resolution which, ideally,

perfectly overlap in space. In practice, this is difficult to achieve when modeling complex structures such as an anatomically realistic four chamber representation of the human heart. Due to the complex geometry, the surfaces of both meshes will be different, requiring extrapolation to transfer coupling variables between the two grids. In a parallel context, a coarser mesh for mechanics will be partitioned differently from the mesh used for EP, causing an increase in the communication cost of the projection operation. Using the same mesh for both EP and mechanics, as implemented in this study or as used by others [5], may be less efficient in weak coupling scenarios where a single static offline projection of data such as activation time or calcium transients can be used as an input for mechanics, but in strong coupling scenarios this may not be the case. Thorough numerical convergence testing studies are necessary in future work to address the issue of spatio-temporal discretization for strongly coupled electromechanical simulations. Such convergence studies will lead to a robust quantification of errors, allowing us to find a balanced tradeoff between computational cost and accuracy.

7.4 Convergence Testing

In an ideal numerical scheme for solving PDEs, the solution depends only on the chosen parameters, but not any other factors such as the spatial discretization step h. While it is feasible to construct such numerical schemes, in practice this is rarely the case. From an applied point of view, the relevant aspect is not to drive discretization errors to zero. Rather, it is more important to approximate the discretization error and contrast it with the typically large uncertainties in parameters and model assumptions. By considering all these factors, we are able to find a balanced tradeoff between computational cost, model parametrization and accuracy.

In the EP modeling communities this issue has been addressed, starting with very early pioneering studies [91], by spatio-temporal convergence testing to choose h and dt appropriately. An extension of this approach brought the introduction of N-version benchmarking to verify whether N different and independent implementations yield comparable results [78]. For solving the monodomain equations it has been demonstrated that all $N = 19$ independent codes used by the EP modeling community yielded the same result, if h and dt were chosen to be sufficiently small. For electromechanical models, convergence testing or N-version benchmarking has not yet been established. In a recent study an N-version benchmark was carried out for solving a cardiac deformation problem, but mechanical deformation was driven by a time triggered active stress models, i.e. a rather simplified weakly coupled model was considered. Moreover, with $N = 2$ the number of independent implementations was quite small. In a strong coupling context, as the modeling work in this study, testing for convergence of electromechanical simulations is more challenging as any errors in the driving EP solution will feed into the solution of the deformation problem. With biophysically detailed models of cellular dynamics such as the GPB model, very fine spatial resolutions are key to keep discretization errors within acceptable bounds. To minimize the impact of errors of the electrical activation pattern on mechanical deformation, conductivities were fitted to travel at prescribed velocities for each of the CM, MM and FM models. As shown in Fig. 5,

this approach indeed results in very similar activation sequences, though deviations are visible in the CM model, indicating that this approach may not be suitable for even coarser discretizations as done by others [5, 41]. While we did not perform a thorough convergence study, as the complex geometry used did not lend itself well to such an investigation, the effect of mesh resolution on global mechanical performance was evaluated by comparing the computed left ventricular cavity volume over time between models. As shown in Fig. 11, volumes at end-diastole are not discernibly different, but noticeable discrepancies arise during systole, which do not change monotonically with $\bar{h}$.

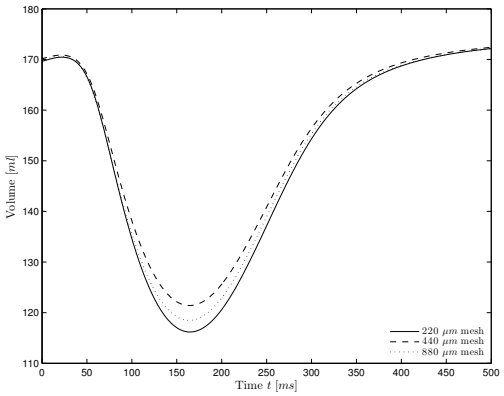

Figure 11: Volume changes of the left ventricle, $V_{\mathrm{LV}}(t)$, in milliliters.

In weak coupling cases using activation time triggered mechanics [41] or when simplified phenomenological models of FitzHugh–Nagumo type are used [5], electromechanical convergence is likely to be less of a concern. However, such models are not suitable for studying biophysical mechanisms of excitation-contraction coupling and MEF at the organ scale. The mesh resolutions used in such studies are unsuitable for use with physiologically detailed models of cellular dynamics such as the GPB. At such coarse resolutions artificial conduction block can occur, impeding action potential propagation. This is not the case with simple phenomenological models as the upstroke of the action potential is orders of magnitude slower. As shown in [5], the upstroke of the used EP model lasts around 50 ms whereas the upstroke duration of a biophysically realistic model last around 1 ms. Such slow upstroke velocities translate into less steep wave fronts, which can be resolved on coarser meshes, reducing computational costs significantly.

In the context of convergence testing of strongly coupled electromechanical models the presented numerical approach is of pivotal importance. The demonstrated excellent scaling properties enable the use of ultra-high spatial resolutions which are necessary for generating high fidelity reference simulation results. The complex models used in this study, while well suited for scaling benchmarks, are less suitable for rigorous convergence testing. Simpler geometries, as used in standard EP benchmarks [78], lend themselves more easily to such tests since spatio-temporal resolution can be varied over domains of exact overlap in space.

7.5 Limitations

The focus of this study was to design and implement a strongly scalable iterative solver which is sufficiently efficient to enable execution of high resolution strongly coupled whole heart EM models. Other issues which have not been addressed and may limit the applicability of the proposed method in different modeling scenarios include:

1. A penalty formulation and static condensation is used to ensure a nearly incompressible material behavior. Depending on the penalty parameter κ this may lead to very ill-conditioned problems. However, with the choice of $\kappa = 1000\,\mathrm{kPa}$ (Table 1) and the design of the AMG method as described in Section 4, we were able to solve the linearized systems within a reasonable number of iterations (Table 8). Nevertheless, as the iteration numbers of the linear solver increase with κ, a mixed approach [4, 68] may be preferable.

2. Simple linear tetrahedral elements were employed for both the electrophysiological and the mechanical models. While this is a valid and accepted approach for the bidomain equations, the resulting $P_1 - P_0$ elements for solid mechanics are known to be prone to volumetric and shear locking. A mixed formulation with quadratic shape functions for displacement and a linear hydrostatic pressure field, i.e. $P_2 - P_1$ tetrahedral elements, would help to avoid these problems [28, 32, 83]. Such an approach is well suited for large deformation problems and shows the same performance as a mixed formulation using $Q_2 - P_1$ hexahedral elements [19].

   However, as we wanted to stick to our principle of utilizing the same grid for both physics, the usage of higher order elements for mechanics would lead to a highly unbalanced relation between the size of the mechanical and electrical systems. Hence, a reasonable employment of $P_2 - P_1$ elements requires different spatial resolutions for mechanics and electrics and an appropriate projection of coupling data between the electrical and the mechanical grid. The construction of such projections and therefore the mixed approach are beyond the scope of this specific paper and left to future work.

3. A Langendorff perfused setup was modeled where the ventricles were not pressurized. This is different from a working heart model where isovolumetric phases are enforced when valves are closed and pressure-volume relations are governed by the cardiovascular system or the atria when valves are open.

4. *In vivo*, the heart is enveloped by the pericardium, which is elastically attached to surrounding tissue via the mediastinum. This important aspect of cardiac function has not been accounted for [32].

Addressing some of these limitations yields block-structured systems that require special treatment. Such systems can be solved iteratively using specially designed AMG methods [48, 67, 118] or properly chosen block preconditioners, generally based on a Schur complement approach [28, 29, 41]. While the solver approach presented here is not directly applicable to such problems, it still constitutes an important building block

for solving block systems. However, conceiving AMG solvers for block system or building Schur complement preconditioners is beyond the scope of this study.

Acknowledgements

This project has received funding from the Austrian Science Fund (FWF) within the SFB Mathematical Optimization and Applications in Biomedical Sciences and from from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 611232. We acknowledge PRACE for awarding us access to resource SuperMUC based in Germany at LRZ (grant CAMEL), and, partially, ARCHER based in the UK at EPCC (project e384).

References

[1] CJ Arthurs, MJ Bishop, and D Kay. "Efficient simulation of cardiac electrical propagation using high order finite elements". In: *J Comput Phys* 231.10 (May 2012), pp. 3946–3962.

[2] CM Augustin. "Classical and All-floating FETI Methods with Applications to Biomechanical Models". PhD thesis. Graz University of Technology, 2012.

[3] CM Augustin, GA Holzapfel, and O Steinbach. "Classical and All-floating FETI Methods for the Simulation of Arterial Tissues". In: *Int J Numer Methods Eng* 99.4 (2014), pp. 290–312.

[4] F Auricchio, LB da Veiga, C Lovadina, A Reali, RL Taylor, and P Wriggers. "Approximation of incompressible large deformation elastic problems: some unresolved issues". In: *Comput Mech* 52.5 (2013), pp. 1153–1167.

[5] B Baillargeon, N Rebelo, DD Fox, RL Taylor, and E Kuhl. "The Living Heart Project: A robust and integrative simulator for human heart function". In: *Eur J Mech A Solids* 48 (Nov. 2014), pp. 38–47.

[6] AH Baker, RD Falgout, H Gahvari, T Gamblin, W Gropp, KE Jordan, TV Kolev, M Schulz, and UM Yang. "Preparing Algebraic Multigrid for Exascale". In: *Lawrence Livermore National Laboratory, Tech. Rep. LLNL-TR-533076* (2012).

[7] S Balay, K Buschelman, V Eijkhout, WD Gropp, D Kaushik, MG Knepley, LC McInnes, BF Smith, and H Zhang. *PETSc Users Manual*. Tech. rep. ANL-95/11 - Revision 3.0.0. Argonne Nat. Lab., 2008.

[8] J Bayer, R Blake, G Plank, and N Trayanova. "A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models". In: *Ann Biomed Eng* 40.10 (2012), pp. 2243–2254.

[9] R Beck. *Graph-based algebraic multigrid for Lagrange-type finite elements on simplicial meshes*. ZIB, 1999.

[10] Y Belhamadia, A Fortin, and Y Bourgault. "Towards accurate numerical method for monodomain models using a realistic heart geometry". In: *Math Biosci* 220.2 (Aug. 2009), pp. 89–101.

[11] M Bishop, R Rajani, G Plank, N Gaddum, G Carr-White, M Wright, M O'Neill, and SA Niederer. "Three dimensional atrial wall thickness maps to inform catheter ablation procedures for atrial fibrillation". In: *Euro Pace* (2015).

[12] MJ Bishop and G Plank. "The role of fine-scale anatomical structure in the dynamics of reentry in computational models of the rabbit ventricles." In: *J Physiol* 590.Pt 18 (Sept. 2012), pp. 4515–4535.

[13] MJ Bishop, G Plank, RA Burton, JE Schneider, DJ Gavaghan, V Grau, and P Kohl. "Development of an anatomically detailed MRI-derived rabbit ventricular model and assessment of its impact on simulations of electrophysiological function". In: *Am J Physiol Heart Circ Physiol* 298.2 (2010), H699–H718.

[14] P Bovendeerd, T Arts, J Huyghe, D Van Campen, and R Reneman. "Dependence of local left ventricular wall mechanics on myocardial fiber orientation: a model study". In: *J Biomech* 25.10 (1992), pp. 1129–1140.

[15] M Buist, G Sands, P Hunter, and A Pullan. "A deformable finite element derived finite difference method for cardiac activation problems". In: *Ann Biomed Eng* 31.5 (2003), pp. 577–88.

[16] *Cairo website*. http://cairographics.org. 2015.

[17] BJ Caldwell, ML Trew, GB Sands, DA Hooks, IJ LeGrice, and BH Smaill. "Three distinct directions of intramural activation reveal nonuniform side-to-side electrical coupling of ventricular myocytes". In: *Circ Arrhythm Electrophysiol* 2.4 (2009), pp. 433–440.

[18] FO Campos, AJ Prassl, G Seemann, R Weber dos Santos, G Plank, and E Hofer. "Influence of ischemic core muscle fibers on surface depolarization potentials in superfused cardiac tissue preparations: a simulation study." In: *Med Biol Eng Comput* 50.5 (May 2012), pp. 461–472.

[19] É Chamberland, A Fortin, and M Fortin. "Comparison of the performance of some finite element discretizations for large deformation elasticity problems". In: *Comput Struct* 88.11-12 (2010), pp. 664–673.

[20] PG Ciarlet. *Mathematical elasticity. Vol. I.* Vol. 20. Studies in Mathematics and its Applications. Amsterdam: North-Holland, 1988.

[21] CIBC. Seg3D: Volumetric Image Segmentation and Visualization. Scientific Computing and Imaging Institute (SCI), Download from: http://www.seg3d.org. 2015.

[22] L Clerc. "Directional Differences of Impulse Spread in Trabecular Muscle from Mammalian Heart". In: *J Physiol* 255 (1976), pp. 335–346.

[23] CM Costa, E Hoetzl, BM Rocha, AJ Prassl, and G Plank. "Automatic Parameterization Strategy for Cardiac Electrophysiology Simulations". In: *Comput Cardiol (2010)* 40 (Oct. 2013), pp. 373–376.

[24] KD Costa, PJ Hunter, JS Wayne, LK Waldman, JM Guccione, and AD McCulloch. "A three-dimensional finite element method for large elastic deformations of ventricular myocardium: II–Prolate spheroidal coordinates". In: *J Biomech Eng* 118.4 (1996), pp. 464–472.

[25] H Demiray. "A note on the elasticity of soft biological tissues". In: *J Biomech* 5 (1972), pp. 309–311.

[26] P Deuflhard, B Erdmann, R Roitzsch, and G Lines. "Adaptive Finite Element Simulation of Ventricular Dynamics". In: *Comput Vis Sci* 12 (2009), pp. 201–205.

[27] T Dickopf, D Krause, R Krause, and M Potse. "Analysis of a lightweight parallel adaptive scheme for the solution of the monodomain equation". In: *SIAM J Sci Comput* 36.2 (2013), pp. C163–C189.

[28] A El maliki, M Fortin, N Tardieu, and A Fortin. “Iterative solvers for 3D linear and nonlinear elasticity problems: Displacement and mixed formulations”. In: *Int J Numer Methods Eng* 83.13 (2010), pp. 1780–1802.

[29] H Elman, VE Howle, J Shadid, R Shuttleworth, and R Tuminaro. “Block preconditioners based on approximate commutators”. In: *SIAM J Sci Comput* 27.5 (2006), pp. 1651–1668.

[30] TSE Eriksson, AJ Prassl, G Plank, and GA Holzapfel. “Influence of myocardial fiber/sheet orientations on left ventricular mechanical contraction”. In: *Math Mech Solids* 18.6 (2013), pp. 592–606.

[31] PJ Flory. “Thermodynamic relations for high elastic materials”. In: *Trans Faraday Soc* 57 (1961), pp. 829–838.

[32] T Fritz, C Wieners, G Seemann, H Steen, and O Dössel. “Simulation of the contraction of the ventricles in a human heart model including atria and pericardium : Finite element analysis of a frictionless contact problem”. In: *Biomech Model Mechanobiol* 13.3 (2014), pp. 627–641.

[33] M Geimer, F Wolf, BJ Wylie, E Ábrahám, D Becker, and B Mohr. “The Scalasca performance toolset architecture”. In: *Concurr Comput* 22.6 (2010), pp. 702–719.

[34] P Ghysels and W Vanroose. “Hiding global synchronization latency in the preconditioned Conjugate Gradient algorithm”. In: *Parallel Computing* 40.7 (2014), pp. 224–238.

[35] S Göktepe and E Kuhl. “Electromechanics of the heart: a unified approach to the strongly coupled excitation-contraction coupling”. In: *Comput Mech* 45 (2010), pp. 227–243.

[36] MJ Gonzales, G Sturgeon, A Krishnamurthy, J Hake, R Jonas, P Stark, WJ Rappel, SM Narayan, Y Zhang, WP Segars, and AD McCulloch. “A three-dimensional finite element model of human atrial anatomy: new methods for cubic Hermite meshes with extraordinary vertices”. In: *Med Image Anal* 17.5 (July 2013), pp. 525–537.

[37] E Grandi, FS Pasqualini, and DM Bers. “A novel computational model of the human ventricular action potential and Ca transient”. In: *J Mol Cell Cardiol* 48.1 (2010), pp. 112–121.

[38] M Griebel, D Oeltz, and MA Schweitzer. “An algebraic multigrid method for linear elasticity”. In: *SIAM J Sci Comput* 25.2 (2003), pp. 385–407.

[39] JM Guccione, KD Costa, and AD McCulloch. “Finite element stress analysis of left ventricular mechanics in the beating dog heart”. In: *J Biomech* 28.10 (1995), pp. 1167–1177.

[40] V Gurev, T Lee, J Constantino, H Arevalo, and NA Trayanova. “Models of cardiac electromechanics based on individual hearts imaging data”. In: *Biomech Model Mechanobiol* 10.3 (2011), pp. 295–306.

[41] V Gurev, P Pathmanathan, JL Fattebert, HF Wen, J Magerlein, RA Gray, DF Richards, and JJ Rice. “A high-resolution computational model of the deforming human heart”. In: *Biomech Model Mechanobiol* (2015), pp. 1–21.

[42] G Haase, M Liebmann, CC Douglas, and G Plank. “A Parallel Algebraic Multigrid Solver on Graphics Processing Units.” In: *HPCA (China), Revised Selected Papers*. Ed. by W Zhang, Z Chen, CC Douglas, and W Tong. Vol. 5938. Lecture Notes in Computer Science. Springer, 2009, pp. 38–47.

[43] KL Ho and L Greengard. “A fast direct solver for structured linear systems by recursive skeletonization”. In: *SIAM Journal on Scientific Computing* 34.5 (2012), A2507–A2532.

[44] M Holler and K Bredies. “A variational method for smooth upsampling of low resolution discrete volume data”. In: *under review* (2014).

[45] GA Holzapfel. “Structural and numerical models for the (visco)elastic response of arterial walls with residual stresses”. In: *Biomechanics of Soft Tissue in Cardiovascular Systems*. Ed. by GA Holzapfel and RW Ogden. Wien, New York: Springer, 2003.

[46] G Holzapfel. *Nonlinear Solid Mechanics. A Continuum Approach for Engineering*. Chichester: John Wiley & Sons Ltd, 2000.

[47] G Holzapfel and R Ogden. “Constitutive modelling of passive myocardium: a structurally based framework for material characterization”. In: *Philos Trans R Soc Lond Ser A Math Phys Eng Sci* 367 (2009), pp. 3445–3475.

[48] J Hron and S Turek. “A Monolithic FEM/Multigrid Solver for an ALE Formulation of Fluid-Structure Interaction with Applications in Biomechanics”. In: *Fluid-Structure Interaction*. Ed. by HJ Bungartz and M Schäfer. Vol. 53. Lecture Notes in Computational Science and Engineering. Springer Berlin Heidelberg, 2006, pp. 146–170.

[49] H Hu and F Sachs. “Stretch-activated ion channels in the heart.” In: *J Mol Cell Cardiol* 29.6 (June 1997), pp. 1511–1523.

[50] X Jie, V Gurev, and N Trayanova. “Mechanisms of mechanically induced spontaneous arrhythmias in acute regional ischemia”. In: *Circ Res* 106.1 (2010), pp. 185–192.

[51] G Karypis and V Kumar. “A fast and high quality multilevel scheme for partitioning irregular graphs”. In: *SIAM J Sci Comput* 20.1 (1998), 359–392 (electronic).

[52] G Karypis, K Schloegel, and V Kumar. “Parmetis: Parallel graph partitioning and sparse matrix ordering library”. In: *Version 1.0, Dept. of Computer Science, University of Minnesota* (1997).

[53] DUJ Keller, O Jarrousse, T Fritz, S Ley, O Dossel, and G Seemann. “Impact of physiological ventricular deformation on the morphology of the T-wave: a hybrid, static-dynamic approach”. In: *IEEE Trans Biomed Eng* 58.7 (July 2011), pp. 2109–2119.

[54] DUJ Keller, DL Weiss, O Dossel, and G Seemann. “Influence of I(Ks) heterogeneities on the genesis of the T-wave: a computational evaluation”. In: *IEEE Trans Biomed Eng* 59.2 (Feb. 2012), pp. 311–322.

[55] JC Kentish and A Wrzosek. “Changes in force and cytosolic Ca2+ concentration after length changes in isolated rat ventricular trabeculae”. In: *J Physiol* 506.2 (1998), pp. 431–444.

[56] RCP Kerckhoffs, J Lumens, K Vernooy, JH Omens, LJ Mulligan, T Delhaas, T Arts, AD McCulloch, and FW Prinzen. “Cardiac resynchronization: insight from experimental and computational models”. In: *Prog Biophys Mol Biol* 97.2-3 (2008), pp. 543–561.

[57] RCP Kerckhoffs, ML Neal, Q Gu, JB Bassingthwaighte, JH Omens, and AD McCulloch. “Coupling of a 3D finite element model of cardiac ventricular mechanics to lumped systems models of the systemic and pulmonic circulation”. In: *Ann Biomed Eng* 35.1 (Jan. 2007), pp. 1–18.

[58] A Klawonn and O Rheinbach. “Highly scalable parallel domain decomposition methods with an application to biomechanics”. In: *ZAMM Z Angew Math Mech* 90.1 (2010), pp. 5–32.

[59] PSF Kohl. “Mechanoelectric feedback in cardiac cells”. In: *Phil. Trans. R. Soc: Lond. A* 359 (2001), pp. 1173–1185.

[60] P Kohl, P Hunter, and D Noble. "Stretch-induced changes in heart rate and rhythm: clinical observations, experiments and mathematical models". In: *Prog Biophys Mol Biol* 71.1 (1999), pp. 91–138.

[61] P Kohl, C Bollensdorff, and A Garny. "Effects of mechanosensitive ion channels on ventricular electrophysiology: experimental and theoretical models". In: *Exp Physiol* 91.2 (2006), pp. 307–321.

[62] K Kunisch and M Wagner. "Optimal control of the bidomain system (II): Uniqueness and regularity theorems for weak solutions". In: *Ann Mat Pura Appl* 192.6 (2013), pp. 951–986.

[63] P Lamata, R Casero, V Carapella, SA Niederer, MJ Bishop, JE Schneider, P Kohl, and V Grau. "Images as drivers of progress in cardiac computational modelling." In: *Prog Biophys Mol Biol* 115.2-3 (Aug. 2014), pp. 198–212.

[64] P Lamata, SA Niederer, DA Nordsletten, DC Barber, I Roy, DR Hose, and N Smith. "An accurate, fast and robust method to generate patient-specific cubic Hermite meshes". In: *Med Image Anal* 15.6 (Dec. 2011), pp. 801–13.

[65] P Lamata, I Roy, B Blazevic, A Crozier, S Land, SA Niederer, DR Hose, and NP Smith. "Quality metrics for high order meshes: analysis of the mechanical simulation of the heart beat." In: *IEEE Trans Med Imaging* 32.1 (Jan. 2013), pp. 130–138.

[66] S Land, SA Niederer, JM Aronsen, EKS Espe, L Zhang, WE Louch, I Sjaastad, OM Sejersted, and NP Smith. "An analysis of deformation-dependent electromechanical coupling in the mouse heart". In: *J Physiol* 590.Pt 18 (2012), pp. 4553–4569.

[67] U Langer and H Yang. "Partitioned solution algorithms for fluid–structure interaction problems with hyperelastic models". In: *Journal of Computational and Applied Mathematics* 276 (2015), pp. 47–61.

[68] P Le Tallec. "Numerical methods for nonlinear three-dimensional elasticity". In: *Handb Numer Anal* 3 (1994), pp. 465–622.

[69] M Liebmann. "Efficient PDE solvers on modern hardware with applications in medical and technical sciences". PhD thesis. University of Graz, 2009.

[70] M Liebmann. *Parallel Toolbox web page*. http://paralleltoolbox.sourceforge.net. 2015.

[71] M Liebmann, G Haase, and A Neic. "A Balanced Accumulation Scheme for Parallel PDE Solvers". In: *Comput Vis Sci* 16.1 (2014), pp. 33–40.

[72] DR Messroghli, S Nordmeyer, T Dietrich, O Dirsch, E Kaschina, K Savvatis, D O h-Ici, C Klein, F Berger, and T Kuehne. "Assessment of diffuse myocardial fibrosis in rats using small-animal Look-Locker inversion recovery T1 mapping". In: *Circ Cardiovasc Imaging* 4.6 (2011), pp. 636–640.

[73] DR Messroghli, A Radjenovic, S Kozerke, DM Higgins, MU Sivananthan, and JP Ridgway. "Modified Look-Locker inversion recovery (MOLLI) for high-resolution T1 mapping of the heart". In: *Magn Reson Med* 52.1 (2004), pp. 141–146.

[74] MA Muñoz, J Kaur, and EJ Vigmond. "Onset of atrial arrhythmias elicited by autonomic modulation of rabbit sinoatrial node activity: a modeling study." In: *Am J Physiol Heart Circ Physiol* 301.5 (Nov. 2011), H1974–H1983.

[75] K Nanthakumar, J Jalife, S Massé, E Downar, M Pop, J Asta, H Ross, V Rao, S Mironov, E Sevaptsidis, J Rogers, G Wright, and R Dhopeshwarkar. "Optical mapping of Langendorff-perfused human hearts: establishing a model for the study of ventricular fibrillation in humans". In: *Am J Physiol Heart Circ Physiol* 293.1 (2007), H875–H880.

[76] A Neic, M Liebmann, E Hoetzl, L Mitchell, EJ Vigmond, G Haase, and G Plank. "Accelerating cardiac bidomain simulations using graphics processing units". In: *IEEE Trans Biomed Eng* 59.8 (2012), pp. 2281–2290.

[77] D Nickerson, N Smith, and P Hunter. "New developments in a strongly coupled cardiac electromechanical model". In: *Europace* 7 Suppl 2 (2005), pp. 118–127.

[78] SA Niederer, E Kerfoot, AP Benson, MO Bernabeu, O Bernus, C Bradley, EM Cherry, R Clayton, FH Fenton, A Garny, E Heidenreich, S Land, M Maleckar, P Pathmanathan, G Plank, JF Rodríguez, I Roy, FB Sachse, G Seemann, O Skavhaug, and NP Smith. "Verification of cardiac tissue electrophysiology simulators using an N-version benchmark". In: *Philos Trans A Math Phys Eng Sci* 369.1954 (Nov. 2011), pp. 4331–4351.

[79] SA Niederer, G Plank, P Chinchapatnam, M Ginks, P Lamata, KS Rhode, CA Rinaldi, R Razavi, and NP Smith. "Length-dependent tension in the failing heart and the efficacy of cardiac resynchronization therapy." In: *Cardiovasc Res* 89.2 (2011), pp. 336–343.

[80] SA Niederer and NP Smith. "The role of the Frank-Starling law in the transduction of cellular work to whole organ pump function: a computational modeling analysis." In: *PLoS Comput Biol* 5.4 (Apr. 2009), e1000371.

[81] S Niederer, L Mitchell, N Smith, and G Plank. "Simulating human cardiac electrophysiology on clinical time-scales". In: *Front Physiol* 2 (2011).

[82] PM Nielsen, IJ Le Grice, BH Smaill, and PJ Hunter. "Mathematical model of geometry and fibrous structure of the heart." In: *American J Physiol* 260.4 Pt 2 (Apr. 1991), H1365–78.

[83] F Nobile, A Quarteroni, and R Ruiz-Baier. "An active strain electromechanical model for cardiac tissue". In: *Int J Numer Method Biomed Eng* 28.1 (2012), pp. 52–71.

[84] RW Ogden. *Non-linear elastic deformations*. Courier Dover Publications, 1997.

[85] AV Panfilov, RH Keldermann, and MP Nash. "Drift and breakup of spiral waves in reaction-diffusion-mechanics systems". In: *Proc Natl Acad Sci U S A* 104.19 (May 2007), pp. 7922–7926.

[86] P Pathmanathan and JP Whiteley. "A numerical method for cardiac mechanoelectric simulations". In: *Ann Biomed Eng* 37.5 (2009), pp. 860–873.

[87] J Peters, O Ecabert, C Meyer, H Schramm, R Kneser, A Groth, and J Weese. "Automatic whole heart segmentation in static magnetic resonance image volumes". In: *Med Image Comput Comput Assist Interv* 10.2 (2007), pp. 402–410.

[88] G Plank, M Liebmann, R Weber dos Santos, EJ Vigmond, and G Haase. "Algebraic multigrid preconditioner for the cardiac bidomain model". In: *IEEE Trans Biomed Eng* 54.4 (2007), pp. 585–96.

[89] G Plank, L Zhou, JL Greenstein, S Cortassa, RL Winslow, B O'Rourke, and NA Trayanova. "From mitochondrial ion channels to arrhythmias in the heart: computational techniques to bridge the spatio-temporal scales". In: *Philos Trans R Soc Lond Ser A Math Phys Eng Sci* 366.1879 (2008), pp. 3381–3409.

[90] G Plank, RA Burton, P Hales, M Bishop, T Mansoori, MO Bernabeu, A Garny, AJ Prassl, C Bollensdorff, F Mason, et al. "Generation of histo-anatomically representative models of the individual heart: tools and application". In: *Philos Trans R Soc Lond Ser A Math Phys Eng Sci* 367.1896 (2009), pp. 2257–2292.

[91] AE Pollard, MJ Burgess, and KW Spitzer. "Computer simulations of three-dimensional propagation in ventricular myocardium. Effects of intramural fiber rotation and inhomogeneous conductivity on epicardial activation." In: *Circ Res* 72.4 (Apr. 1993), pp. 744–756.

[92] AJ Prassl, F Kickinger, H Ahammer, V Grau, JE Schneider, E Hofer, EJ Vigmond, NA Trayanova, and G Plank. "Automatically generated, anatomically accurate meshes for cardiac electrophysiology problems". In: *IEEE Trans Biomed Eng* 56.5 (2009), pp. 1318–1330.

[93] Z Qu and A Garfinkel. "An advanced algorithm for solving partial differential equation in cardiac conduction." In: *IEEE Trans Biomed Eng* 46.9 (1999), pp. 1166–8.

[94] DF Richards, JN Glosli, EW Draeger, AA Mirin, B Chan, JL Fattebert, WD Krauss, T Oppelstrup, CJ Butler, JA Gunnels, V Gurev, C Kim, J Magerlein, M Reumann, HF Wen, and JJ Rice. "Towards real-time simulation of cardiac electrophysiology in a human heart at high resolution". In: *Comput Methods Biomech Biomed Engin* 16.7 (2013), pp. 802–805.

[95] DE Roberts, LT Hersh, and AM Scher. "Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog". In: *Circ Res* 44.5 (1979), pp. 701–12.

[96] DE Roberts and AM Scher. "Effect of tissue anisotropy on extracellular potential fields in canine myocardium in situ". In: *Circ Res* 50.3 (1982), pp. 342–51.

[97] BM Rocha, F Kickinger, AJ Prassl, G Haase, EJ Vigmond, RW dos Santos, S Zaglmayr, and G Plank. "A macro finite-element formulation for cardiac electrophysiology simulations using hybrid unstructured grids". In: *IEEE Trans Biomed Eng* 58.4 (2011), pp. 1055–1065.

[98] S Rossi, T Lassila, R Ruiz-Baier, A Sequeira, and A Quarteroni. "Thermodynamically consistent orthotropic activation model capturing ventricular systolic wall thickening in cardiac electromechanics". In: *Eur J Mech A Solids* (2013).

[99] J Ruge and K Stüben. "Algebraic multigrid". In: *Multigrid methods* 3 (1987), pp. 73–130.

[100] A Schmidt, CF Azevedo, A Cheng, SN Gupta, DA Bluemke, TK Foo, G Gerstenblith, RG Weiss, E Marbán, GF Tomaselli, JAC Lima, and KC Wu. "Infarct tissue heterogeneity by magnetic resonance imaging identifies enhanced cardiac arrhythmia susceptibility in patients with left ventricular dysfunction". In: *Circulation* 115.15 (2007), pp. 2006–2014.

[101] J Simo, R Taylor, and K Pister. "Variational and projection methods for the volume constraint in finite deformation elasto-plasticity". In: *Comput Methods Appl Mech Eng* 51.1 (1985), pp. 177–208.

[102] C Stevens, E Remme, I LeGrice, and P Hunter. "Ventricular mechanics in diastole: material parameter sensitivity". In: *J Biomech* 36.5 (2003), pp. 737–748.

[103] G Strang. "On the construction and comparision of difference scheme". In: *SIAM J Numer Anal* 5 (1968), pp. 506–517.

[104] K Stüben. "A review of algebraic multigrid". In: *J Comput Appl Math* 128.1 (2001), pp. 281–309.

[105] J Sundnes, GT Lines, and A Tveito. "An operator splitting method for solving the bidomain equations coupled to a volume conductor model for the torso". In: *Math Biosci* 194.2 (2005), pp. 233–48.

[106] P Tavi, C Han, and M Weckström. "Mechanisms of stretch-induced changes in [Ca2+]i in rat atrial myocytes: role of increased troponin C affinity and stretch-activated ion channels." In: *Circ Res* 83.11 (Nov. 1998), pp. 1165–1177.

[107] K Tøndel, S Land, SA Niederer, and NP Smith. "Quantifying inter-species differences in contractile function through biophysical modelling". In: *J Physiol* (Dec. 2014).

[108] NA Trayanova, J Constantino, and V Gurev. "Electromechanical models of the ventricles". In: *Am J Physiol Heart Circ Physiol* 301.2 (2011), H279–H286.

[109] M Trew, I Le Grice, B Smaill, and A Pullan. "A finite volume method for modeling discontinuous electrical activation in cardiac tissue". In: *Ann Biomed Eng* 33.5 (May 2005), pp. 590–602.

[110] T Usyk, I LeGrice, and A McCulloch. "Computational model of three-dimensional cardiac electromechanics". In: *Comput Visual Sci* 4 (2002), pp. 249–257.

[111] P Vassilevski. *Multilevel Block Factorization Preconditioners: Matrix-based Analysis and Algorithms for Solving Finite Element Equations*. Springer, 2008.

[112] FJ Vetter and AD McCulloch. "Three-dimensional analysis of regional cardiac function: a model of rabbit ventricular anatomy". In: *Prog Biophys Mol Biol* 69.2-3 (Jan. 1998), pp. 157–83.

[113] F Vetter and A McCulloch. "Three-dimensional stress and strain in passive rabbit left ventricle: a model study". In: *Ann Biomed Eng* 28.7 (2000), pp. 781–92.

[114] E Vigmond, M Hughes, G Plank, and L Leon. "Computational tools for modeling electrical activity in cardiac tissue". In: *J Electrocardiol* 36 (2003), pp. 69–74.

[115] E Vigmond, R Weber dos Santos, AJ Prassl, M Deo, and G Plank. "Solvers for the cardiac bidomain equations". In: *Prog Biophys Mol Biol* 96.1 (2008), pp. 3–18.

[116] *VTK - The Visualization Toolkit*. http://vtk.org. 2015.

[117] JP Whiteley, MJ Bishop, and DJ Gavaghan. "Soft tissue modelling of cardiac fibres for use in coupled mechano-electric simulations." In: *Bull Math Biol* 69.7 (Oct. 2007), pp. 2199–2225.

[118] H Wobker and S Turek. "Numerical studies of Vanka-type smoothers in computational solid mechanics". In: *Adv Appl Math Mech* 1.1 (2009), pp. 29–55.

[119] J Zhao, TD Butters, H Zhang, IJ LeGrice, GB Sands, and BH Smaill. "Image-based model of atrial anatomy and electrical activation: a computational platform for investigating atrial arrhythmia". In: *IEEE Trans Med Imaging* 32.1 (Jan. 2013), pp. 18–27.