

EFFICIENT MODELING OF CARDIAC TISSUE AND
RESULTING APPLICATIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MECHANICAL
ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Jonathan Joi-Mun Wong, Ph.D. Candidate

August 2012

Abstract

The heart is an essential heterogeneous organ that depends on strong coupling between electrical, chemical, and mechanical dynamics to properly function as a pump that supplies blood to the rest of the body. Cardiac arrhythmias are common disorders characterized by irregular beating of the heart that lead to serious clinical conditions. It is estimated that approximately 2.2 million adults in the United States are affected by atrial fibrillation, a prevalent arrhythmia. Unfortunately, a clinician often does not have enough information to diagnose a patient's heart condition to determine the optimal treatment procedure. This is an area that computational mechanics can address. While development of mechanical and electrophysiological models of cardiac tissue primarily started in the 1950s, fully-coupled models have only more recently been developed due to factors regarding computational cost, difficulty in quantifying material properties, and difficulty in integrating complex models in a cohesive and efficient manner. Therefore, in order for simulations tools to have impact in the clinical or experimental setting, these tools must be efficient, fast, robust, and accessible. The focus of this thesis is to develop methods of addressing the aforementioned issues and then illustrate how efficient electromechanical finite element models can be developed for the heart such that their use in the clinical and experimental setting can be realized in several examples.

In this thesis, a global-local variable splitting formulation borrowed from the field of plasticity, is used to address the issues of complex model integration, and maintain numerical stability at lower costs. Local internal variables are solved at the finite element integration points, while global variables are solved at the finite element degrees of freedom. To overcome the stability issues that arise from stiff non-linear strongly

coupled differential equations, a simple implicit first order time-stepping approach was taken. In the past, researchers have avoided using implicit schemes as they are perceived to be expensive; however this perception has led to prevalent use of ad-hoc schemes that ultimately reduce accuracy, increase computational cost, and result in numerical instabilities. Through careful examination of classical phenomenological models and detailed biophysical ionic models of the electrophysiology of the heart, almost all models can be reformulated into this global-local splitting framework. The numerical properties of cost-expensive ionic models are briefly analyzed within the context of this framework. Surprisingly the local bandwidth of the local tangent matrix is fairly low, and results in reduced computational cost. Use of implicit-time stepping in tandem with a simple iteration and error tolerance based adaptive time-stepping algorithm allows for reduction of computation time from hours to minutes. Comparison between an 18-variable strongly coupled ionic model and a simple two-variable phenomenological model within this framework, only shows a 3-fold increase in computational cost for the more accurate ionic model.

Flexibility and modularity of the framework are illustrated through the development of electrical, electro-chemical, electro-chemical-mechanical, and opto-electro-mechanical models of cardiac tissue. The heart is modeled efficiently using custom finite element ventricular cell models for physiological electrical simulations and large deformation excitation-contraction dry-pumping simulations of the heart. The models are developed in a modular way within the framework of the global-local split and built on top of one another in a monolithic implementation. The results accurately model the physiological condition of the heart, as well as certain arrhythmogenic conditions. The flexibility and multi-scale nature of the framework is also leveraged in developing novel optical-induced cardiac cell excitation models of new genetically engineered Channelrhodopsin-2 (ChR2) cardiac myocytes. These cells excite and contract when stimulated by a particular wavelength of blue light. An ionic model was developed for this particular bio-engineered stem cell, calibrated with experimental data from collaborators, and was able to predict electrical excitation behavior of the cells to a reasonable degree of accuracy. This ionic model was then combined with

self-pacing ionic pacemaker cell models and also with ventricular cell models into respective finite elements to simulate experiments and predict future therapies using ChR2 genetically modified cardiac tissue.

The thesis also addresses difficulties relating to identification and characterization of material parameter identification in inhomogeneous cardiac tissue. A metric for determining smoothness in electrical conduction in tissue cultures was validated with stochastic finite element models of microelectrode array cell conduction experiments. The results indicate that the metric is useful in characterizing different conduction patterns based on two metrics borrowed from wavelet theory: energy and contrast. In silico experiments corroborate the effects seen experimentally in tissue after geometrical differences between cells were taken into account. Difficulties in obtaining structural fiber data from clinical images were addressed by developing an algorithmic method for designating approximate physiologically accurate fiber distributions for the heart using only geometrical information obtained from MRI scans of the surfaces of the heart. Poisson interpolation is used and results in a smooth continuous rotating fiber description that matches experimentally obtained fiber directions from MRI scans. The algorithm is formulated carefully to also allow for smooth interpolation of experimentally obtained fiber measurements should they exist. The main benefits of this algorithm are its simplicity of implementation, physiologically accuracy, and generality in interpolating fiber distributions.

Lastly, the thesis demonstrates possible benefits of GPU computing in order to achieve near-real-time electrical simulations of arrhythmias in the heart. While GPU computing generally addresses high-performance computing needs, it is used within the context of this thesis to obtain the necessary performance increase within the limits of consumer-accessible hardware. The complete assembly and solver routines from the finite element code, FEAP from Berkeley, were ported to the GPU using CUDA. Even with a relatively minimally optimized proof-of-concept, the GPU-only finite element code achieves performance comparable to that of a twelve cores using only one GPU. The added benefit is realized when running real-time simulations, as sensor data can be integrated and post-processing of results can be performed by the CPU concurrently, while the GPU takes care of the scientific finite element simulation.

This effectively quadruples the performance of a consumer machine by fully utilizing the available hardware. To increase the overall efficiency of the method, current sparse matrix vector multiplication GPU algorithms are analyzed, and possible alternative algorithms are developed specifically with unstructured finite element meshes in mind.

Altogether, the different methods developed in this thesis have been shown to be effective in addressing issues related to efficiency, numerical stability, modularity, and flexibility in real computational applications of the heart. Special consideration was taken in designing the different methods to be compatible with one another, such that a majority of the methods could be integrated and the benefits of each method could be leveraged with each other to gain maximum efficiency. While these developed methods can still be improved, the thesis work as a whole serves to demonstrate and highlight future uses for computational models within experimental and clinical settings.

Preface

This Stanford Ph.D. dissertation topic was conceived more or less in 2006, nearly 6 years ago, by a naive and ambitious Mechanical engineering undergraduate student researcher in the Berkeley Orthopaedic Biomechanics Laboratory. After having worked in a productive biomechanics laboratory that did a mix of experimental and computational work, I decided I liked computational work more and that soft-tissue mechanics would pose a different set of theoretical and computational challenges that would be worthy of a Ph.D. dissertation topic. Roughly around that time, my mother started having episodes of paroxysmal atrial fibrillation, and being the caring, naive, ambitious son, I decided to make that the topic of my research statement paragraph' at the end of my Statement of Purpose when applying to universities for graduate school.

Stanford was a natural choice for me as it was the only school the had a hospital conveniently on campus, and also one that had multiple groups conducting research in the cardiovascular mechanics, computational mechanics and numerics. However, I was only admitted as a general M.S. student and did not receive any funding. Through a long arduous journey with several serendipitous events, I met my current adviser and gradually became fully funded after two long years. In a strange way, I thank God for my arduous funding experiences, as it tested my interested in research and helped validate my motivation in developing better computational models of the heart. I hope that the breadth and depth of this thesis reflects my motivation and joy in developing, applying, and testing the various techniques.

I had initially planned to write this dissertation as a primer on electromechanical heart modeling, but my zeal and interest in conducting ongoing research has prevented

me from doing so to a certain extent. Instead, I have decided to provide a general primer-like basis in the first four chapters to introduce the necessary information required to understanding the following journal publications that have been included. However, I have tried to include some unpublished information in relevant sections of my dissertation that may be helpful in understanding the development of a model or algorithm, or accessory information that may highlight different aspects of models not emphasized within the journal publications. Ultimately, I hope this dissertation will be entertaining and will spark someone's interest in this area, as the work that follows merely outlines future possible avenues of extending computational mechanics into the clinical and experimental setting.

Acknowledgments

I would like to thank my mother and father for their continual support and encouragement in pursuing academic interests. As a child who once failed math and then subsequently English in grade school, I must thank my parents for their excellent encouragement and tutoring. My parents put a lot of effort in teaching me to be comfortable making mistakes and to actively improve upon my weaknesses. I thank my mother for helping me build a sense of self-confidence and an impressive work ethic in approaching challenging problems, such that perhaps, now, I may tend to be overly ambitious. I thank my father for instilling scholarly traits in me at a young age, by pedagogically spending 3 hours to explain the mechanistic workings of a bicycle to a 3rd grade student. Details are important, but I have also learned the importance of adjusting presentations appropriately for the particular audience. As a child of a grade school teacher, and an optical laser engineer, I also must thank my parents for exposing me to a variety of subjects that extend beyond computation, mathematics, and science while growing up. These insights have been extremely helpful both in presenting work and in approaching the different interdisciplinary problems I have encountered during my Ph.D.

I'd like to thank my principal adviser, Ellen, for giving a self-funded M.S. student a chance at completing a Ph.D. at Stanford, or the extremely prompt communications, and patience in mentoring me in terms of research focus. Ellen has been an excellent adviser to me, and I wish I had more time to learn about the many areas Ellen has expertise in. She has always been very supportive of my work, even though initially she did not have expertise in electromechanical models of the heart. She has always been generous even though the lab is fiscally conservative. I must thank her for the

opportunities to visit Europe, Australia, and several other conferences for "academic" work. While, I did not complete my goal in visiting every single continent, I am extremely grateful to have been exposed to different research cultures.

I must also thank Serdar Göktepe for the interesting conversations and more than "supplemental" mechanics/numerical tutoring sessions. He has provided me with a solid foundation for my work and taught me balance enjoying and spending time in learning, while producing good results quickly. I have very much enjoyed working with you for a short three years during the beginning of my Ph.D.

I'd also like to thank Professor Eric Darve and Professor Pinsky. I am grateful for the opportunity of working with Eric on GPU algorithms. I have greatly enjoyed working with Eric and it has exposed me to a different mentoring style than the one I have experienced with Ellen. It was also very comforting to know that someone else at Stanford, appreciates designing and developing "frivolous" and geeky projects as much as I do. Likewise, I am thankful for the support Peter has given me over the years, and for his excellent classes on the Finite Element Method. Peter's brilliant sense of wit and British humor have definitely added a sense of flavor to the relative mundane environment students in computational mechanics trudge through.

I'd also like to thank my other friends and family members. I would like to especially acknowledge Michael Chen and Oscar Abilez for great collaborative experiences. I appreciate the opportunity to apply my computational models to exciting research at the forefront of science. I'd also like to acknowledge my current and former lab mates for their support and help in research. Special thanks to Alex Zoellner for the preparation of a multitude of meshes. I'd also like to thank Manuel, Adrian, and Mohan for the engaging conversations, we've had. Special mention must be given to students and postdocs from Professor Lew's lab for the engaging conversations on various topics ranging from the trivialities of lunch to the philosophical implications of continuum mechanics.

Lastly, I am very thankful for the two funding sources I have had through the Biomedical Computation Graduate Training Grant 5T32GM063495 and the SGF Sang Samuel Wang Stanford Graduate Fellowship. I have hopefully substantiated the funding awarded to me by the work that I have done.

Contents

Abstract	v
Preface	ix
Acknowledgments	xi
1 Introduction	1
2 Background	4
2.1 Physiology	4
2.1.1 Structure and function of the heart: A multi-scale, multi-field overview	4
2.1.2 Electrophysiology of the cell	6
2.1.3 Electrophysiology of the heart	8
2.1.4 Mechanical Physiology - cellular and tissue level	10
2.2 Medical Conditions	11
2.3 Continuum Mechanics	13
2.3.1 Kinematics	14
2.4 Finite Elements and Numerics	17
2.4.1 Weak-form	18
2.4.2 Linearization	19
2.4.3 Lie Derivatives, \mathcal{L}	20
2.4.4 Spatial Discretization	21

3	Modeling overview	24
3.1	Electrophysiology	24
3.1.1	Phenomenological Models	26
3.1.2	Ionic Models	30
3.1.3	Conduction	33
3.2	Mechanics	34
3.2.1	Passive Mechanical Behavior Modeling	35
3.2.2	Active Mechanical Behavior Modeling	36
4	Monolithic Algorithmic Framework	38
4.1	General Algorithmic Framework	38
4.2	Phenomenological Electrophysiological formulation	40
4.3	Ionic Electrophysiological formulation	41
4.4	Coupled electromechanical formulation	45
4.5	Algorithmic Splitting Implications	45
4.5.1	Time Stepping Methods	46
5	Applications in modeling ventricular tissue	50
5.1	Introduction	51
5.2	Continuous problem of electrochemistry	55
5.2.1	Electrical problem - Partial differential equation	56
5.2.2	Chemical problem - System of ordinary differential equations	56
5.3	Discrete problem of electrochemistry	58
5.3.1	Electrical problem - Global discretization on the node point level	59
5.3.2	Chemical problem - Local discretization on the integration point level	60
5.4	Continuous model problem for human ventricular cardiomyocytes	62
5.4.1	Specification of sodium concentration, currents, and gating variables	64
5.4.2	Specification of potassium concentration, currents, and gating variables	67

5.4.3	Specification of calcium concentration, currents, and gating variables	70
5.4.4	Specification of sarcoplasmic reticulum calcium concentration, currents, and gating variables	73
5.5	Discrete model problem for human ventricular cardiomyocytes	74
5.6	Examples	75
5.6.1	Electrochemistry in a human ventricular cardiomyocyte	75
5.6.2	Electrochemistry in the human heart	78
5.7	Discussion	81
6	Applications for cell culture experiments	94
6.1	Introduction	95
6.2	Methods	97
6.2.1	In vitro conduction analysis using MEAs	97
6.2.2	In silico conduction analysis using FEM	99
6.2.3	Co-occurrence matrix	104
6.2.4	Co-occurrence energy and contrast	106
6.3	Results	108
6.3.1	In vitro conduction analysis using MEAs	108
6.3.2	In silico conduction analysis using FEM	110
6.3.3	Co-occurrence matrix	110
6.3.4	Co-occurrence energy and contrast	111
6.4	Discussion	112
6.5	Conclusion	115
7	Applications for optogenetically modified tissue cultures	123
8	Applications in virtual treatments using Chr2 cells	145
8.1	Motivation	146
8.2	Mathematical model of channelrhodopsin photocycle	150
8.3	Mathematical model of ionic currents	153
8.4	Mathematical model of ionic concentrations	155

8.5	Mathematical model of action potential propagation	156
8.6	Computational model of light-activated cells	157
8.6.1	Chemical problem - Local discretization on the integration point level	158
8.6.2	Electrical problem - Global discretization on the node point level	159
8.7	Computational model of a human heart	160
8.7.1	Atrioventricular node pacing through photostimulation	161
8.7.2	Apical pacing through photostimulation	162
8.7.3	Biventricular pacing through photostimulation	162
8.8	Discussion	163
8.8.1	Comparison with the literature	163
8.8.2	Conclusion	165
9	Designating fibers in the heart	175
9.1	Introduction	175
9.2	Methods	178
9.2.1	Poisson interpolation of scalar-valued feature θ	178
9.2.2	Local cardiac coordinate system $\{\mathbf{n}, \mathbf{c}, \mathbf{z}\}$	179
9.2.3	Nodal fiber and sheet orientations \mathbf{f} and \mathbf{s}	180
9.3	Results	182
9.3.1	Generic bi-ventricular heart model	182
9.3.2	Patient-specific human heart model	183
9.3.3	Sensitivity analysis for randomly distributed boundary conditions	184
9.3.4	Sensitivity analysis for uniformly distributed boundary conditions	185
9.4	Discussion	187
9.5	Conclusion	189
10	Electromechanically-coupled Applications	194
10.1	Introduction	195
10.2	Continuous problem of chemo-electro-mechanics	198
10.2.1	Kinematic equations	198
10.2.2	Balance equations	199

10.2.3	Constitutive equations	199
10.3	Discrete problem of chemo-electro-mechanics	201
10.3.1	Temporal discretization	201
10.3.2	Spatial discretization	202
10.3.3	Linearization	203
10.4	Model problem of chemo-electro-mechanics	203
10.4.1	Electrical flux	203
10.4.2	Electrical source	204
10.4.3	Momentum flux	207
10.5	Examples	209
10.5.1	Chemo-electro-mechanical coupling in a single cell	209
10.5.2	Chemo-electro-mechanical coupling in the human heart	211
10.6	Discussion	215
11	GPU applications	225
11.1	Introduction	225
11.2	Description of physical problem and finite element breakdown	227
11.2.1	Computational cost of the Finite Element method	229
11.3	GPU Architectural Organization	230
11.4	Description of current GPU methods	232
11.5	Development of novel GPU SPMV algorithms	236
11.5.1	WarpKernel1 (WPK1)	238
11.5.2	Warpkernel1 - reordered (WPK1r)	241
11.5.3	Warpkernel1 - reordered sorted (WPK1rs)	242
11.5.4	Warpkernel2 (WPK2)	243
11.5.5	Warpkernel2 variations (WPK2r) and (WPK2rs)	243
11.6	Benchmarking sparse-matrix vector multiplication methods	247
11.7	Performance Analysis	250
11.7.1	Effects of optimizations taken	251
11.7.2	Finite Element Comparison between GPU and multi-core PETSc258	
11.8	Discussion	259

11.8.1 Synthetic Benchmarks	260
11.8.2 Optimizations and Variations	262
11.8.3 Finite Element Simulations	262
11.8.4 Improvements	263
11.9 Conclusion	264
12 Concluding remarks and future work	274
Bibliography	278

List of Tables

4.1	General algorithmic treatment of global-local variable split for cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. Global variables, Φ , are solved at the nodal points of the mesh, while internal variables, r_i , are solved at the integration point level.	41
4.2	Algorithmic treatment of global-local variable split for the case of phenomenological electrophysiological finite element models. The transmembrane variable, ϕ_J , is solved at the nodal points of the mesh, while internal variable, r_i , are solved at the integration point level.	42
4.3	Algorithmic treatment of electrochemical coupling in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the chemical unknowns, the two sets of gating variables g_{gate}^I and g_{gate}^{II} and the ion concentrations c_{ion} are introduced locally on the integration point level.	47
4.4	Table of the dependence of each gating variable on the transmembrane potential, ϕ , and the ionic concentrations, c_{ion}	48
4.5	Dependency table for each ionic current on each ionic concentration species.	48
4.6	Dependency table of each ionic concentration residual, R^i , on ionic concentration species.	49

5.1	Algorithmic treatment of electrochemical coupling in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the chemical unknowns, the two sets of gating variables g_{gate}^I and g_{gate}^{II} and the ion concentrations c_{ion} are introduced locally on the integration point level.	85
5.2	Material parameters of the proposed human ventricular cardiomyocyte model [68, 116, 149, 188].	86
5.3	Algorithmic performance. Characteristic quadratic convergence of global Newton Raphson iteration, illustrated in terms of the representative residuals of the relative error during five different phases of the cardiac cycle.	86
6.1	Examples of characteristic activation plots from Figure 10.3 and their corresponding co-occurrence energy and contrast. The energy is highest for images with sharp fronts like the spiral wave and lowest for uniform images like the random noise. The contrast is highest for images with large spatial variations like the random noise and lowest for homogeneous images like the linear gradient.	107
9.1	Algorithm for creating smoothly varying fiber and sheet orientations \mathbf{f} and \mathbf{s} within a human heart, for a given mesh and prescribed fiber angles α on the surfaces of the heart.	181
9.2	Table of parameters for random and subdivision examples.	186
10.1	Chemo-electrical material parameters of human ventricular	208
10.2	Electro-mechanical material parameters of human cardiac tissue.	212

11.1	Algorithmic treatment of transmembrane voltage in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the phenomenological recovery variable, r , is introduced locally on the integration point level. Element assembly is highlighted in green, and the solution update is highlighted in red.	228
11.2	Matrix Benchmark Information. The number of nonzeros (nz), the number of rows (nrows), bytes, minimum row length (minrow), maximum row length (maxrow), and average number of non-zero values per row (nz/nrows) is reported for each benchmark matrix.	247
11.3	Benchmark best MGPU valuesPerThread parameter.	252
11.4	Benchmark best WPK1 blocksize parameter.	253
11.5	Benchmark WPK2 best threshold and blocksize parameters.	254
11.6	Benchmark data regarding the effects of sorting for WPK1. Padding difference designates the extra padding that results when rows are not sorted. The time difference percentage is defined as the WPK1 kernel time subtracted from the non-sorted kernel time divided by the non-sorted kernel time.	254
11.7	Comparison of reordering times and WPK1 and cusparse kernel SPMV times.	256
11.8	Comparison of reordering times and WPK2 and cusparse kernel SPMV times.	257

List of Figures

2.1	Diagram of the heart	5
2.2	Schematic diagram of an cardiac cell	7
2.3	Diagram of an conduction system in the heart	8
2.4	Continuum mapping relations	16
3.1	The generalized FitzHugh-Nagumo model with $\alpha = -0.5$, $a = 0$, $b = -0.6$, $c = 50$. The phase portrait depicts the trajectories for distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to a stable limiting cycle. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$ (left). Self-oscillatory time plot of the non-dimensional action potential ϕ and the recovery variable r (right).	28
3.2	The Aliev-Panfilov model with $\alpha = 0.05$, $\gamma = 0.002$, $b = 0.15$, $c = 8$, $\mu_1 = 0.2$, $\mu_2 = 0.3$. The phase portrait trajectories for distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to stable equilibrium point. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$ (left). Non-oscillatory time plot of the non-dimensional action potential ϕ and the recovery variable r is triggered by external stimulation $I=30$ from the steady state $\phi_0=r_0=0$ (right).	29

3.3 Phase portrait of classical FitzHugh-Nagumo model with $a=0.7$, $b=0.8$, $c=3$ (left). Trajectories for the distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to steady state. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$. Circuit diagram of corresponding tunnel-diode nerve model (right). 30

5.1 Electrochemistry in a human ventricular cardiomyocyte. Time dependent evolution of transmembrane potential ϕ . The characteristic action potential consists of five phases. Phase 0: The rapid upstroke is generated through an influx of positively charged sodium ions through fast sodium channels. Phase 1: Early, partial repolarization is initiated through the efflux of positively charged potassium ions through transient outward channels. Phase 2: During the plateau, the net influx of positively charged calcium ions through L-type sodium channels is balanced by the efflux of positively charged potassium ions through inward rectifier channels, rapid and slow delayed rectifier channels, and transient outward channels. Phase 3: Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. Phase 4: Throughout the interval between end of repolarization and the beginning of the next cycle the cell is at rest. 52

5.2 Ionic model of a human ventricular cardiomyocyte [12, 68, 116, 188]. In this model, the electrochemical state of the cardiomyocyte is characterized in terms of $n_{\text{ion}} = 4$ ion concentrations, the free intracellular sodium, potassium, and calcium concentrations and the free calcium concentration in the sarcoplasmic reticulum, $c_{\text{ion}} = [c_{\text{Na}}, c_{\text{K}}, c_{\text{Ca}}, c_{\text{Ca}}^{\text{sr}}]$. Ion concentrations are controlled through $n_{\text{crt}} = 15$ ionic currents, $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]$. Their channels are governed by $n_{\text{gate}} = 13$ gating variables $g_{\text{gate}} = [g_{\text{m}}, g_{\text{h}}, g_{\text{j}}, g_{\text{xK1}\infty}, g_{\text{xr1}}, g_{\text{xr2}}, g_{\text{xs}}, g_{\text{r}}, g_{\text{s}}]$ which are functions of the current membrane potential ϕ 54

5.3 Electrochemistry in a human ventricular cardiomyocyte. Voltage dependent evolution of time constants τ_m, τ_h, τ_j and steady state values $g_m^\infty, g_h^\infty, g_j^\infty$ for sodium activation and fast and slow sodium inactivation gates g_m, g_h, g_j . Calcium concentration dependent evolution of steady state value g_{fCa}^∞ for intracellular calcium dependent calcium inactivation gate g_{fCa} . Voltage dependency evolution of time constants τ_d, τ_f and steady state values g_d^∞, g_f^∞ for L-type calcium activation and inactivation gates g_d, g_f . Voltage dependency evolution of time constants τ_r, τ_s and steady state values g_r^∞, g_s^∞ for transient potassium outward activation and inactivation gates g_r, g_s . Voltage dependency evolution of time constants $\tau_{xs}, \tau_{xr1}, \tau_{xr2}$ and steady state values $g_{xs}^\infty, g_{xr1}^\infty, g_{xr2}^\infty$ for slow delayed rectifier gate and rapid delayed rectifier activation and inactivation gates g_{xs}, g_{xr1}, g_{xr2} 87

5.4 Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of sodium activation gate g_m , fast sodium inactivation gate g_h , slow sodium inactivation gate g_j , L-type calcium activation gate g_d , L-type calcium inactivation gate g_f , intracellular calcium dependent calcium inactivation gate g_{fCa} transient outward activation gate g_r , transient outward inactivation gate g_s , slow delayed rectifier gate g_{xs} , rapid delayed rectifier activation gate g_{xr1} , rapid delayed rectifier inactivation gate g_{xr2} , inward rectification factor g_{K1}^∞ , and calcium-dependent inactivation gate g_g 88

5.5 Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of the fast sodium current I_{Na} , the background sodium current I_{bNa} , the sodium potassium pump current I_{NaK} , and the sodium calcium exchanger current I_{NaCa} , the inward rectifier current I_{K1} , the rapid delayed rectifier current I_{Kr} , the slow delayed rectifier current I_{Ks} , the plateau potassium current I_{pK} , the transient outward current I_{t0} , the L-type calcium current I_{CaL} , the background calcium current I_{bCa} , the plateau calcium current I_{pCa} , the leakage current I_{leak} , the sarcoplasmic reticulum uptake current I_{up} , and the sarcoplasmic reticulum release current I_{rel} 89

5.6 Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of intracellular sodium concentration c_{Na} , potassium concentration c_K , calcium concentration c_{Ca} , and calcium concentration in the sarcomplastic reticulum c_{Ca}^{sr} . The sodium concentration increases rapidly from 11.60 mM to 11.61 mM within the first 5 ms to initiate the fast upstroke of the action potential which then, in turn, affects the voltage-gated calcium and potassium channels. Accordingly, the calcium concentration increases quickly to 1.0 μ M and then decreases gradually back to its resting value of 0.08 μ M. The potassium concentration decreases slowly to 138.29 mM until the beginning of the resting phase at after 0.28s and then gradually returns back to its initial value of 138.30 mV. In this last phase, the sodium concentration which had decreased to 11.585 mM increases gradually to its initial value of 11.60 mM. 90

5.7 Electrochemistry in the human heart. Spatio-temporal evolution of the membrane potential ϕ and the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} during the depolarization phase of the cardiac cycle. Depolarization is initiated through an increase in the intracellular sodium concentration c_{Na} which reflects itself in the rapid depolarization of the cell characterized through an increase in the membrane potential ϕ from -86 mV to +20 mV. This affects the voltage-gated potassium and calcium channels and initiates a decrease in the intracellular potassium concentration c_{K} and an increase in the intracellular calcium concentration c_{Ca} . After approximately 30 ms, both ventricles of the heart are fully depolarized. 91

5.8 Electrochemistry in the human heart. Spatio-temporal evolution of the membrane potential ϕ and the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} during the repolarization phase of the cardiac cycle. Repolarization is characterized through a smooth decrease in the membrane potential ϕ from its excited value of +20 mV back to its resting value of -86 mV. At the same time, the intracellular calcium concentration c_{Ca} decreases smoothly to its resting value. Both sodium c_{Na} and potassium c_{K} respond more slowly and reach minimum concentrations only after 280 ms before increasing gradually back to their initial values at the end of the cycle after 1000 ms. 92

5.9 Algorithmic performance. Number of iterations and time step size for non-adaptive and adaptive time stepping schemes. For both algorithms, we typically find convergence within five to six Newton Raphson iterations during the upstroke phase, and within three to four iterations during all other phases of the cardiac cycle. This results in a total run time of 3845.74 s for the non-adaptive scheme with a fixed time step size of $\Delta t = 0.125$ ms and 8000 time increments throughout the cardiac cycle of $t=1000$ ms, calculated on a single core of an i7-950 3.06 GHz desktop with 4GB of memory. The adaptive time stepping scheme automatically increases the time step size during the plateau phase, between $t=50$ ms and $t=275$ ms, and during the resting phase, after $t=350$ ms. Adaptive time stepping with a maximum time step size of $\Delta t^{\max} = 8.0$ ms reduces the number of increments to 492, and the overall run time to 395.46 s. 93

6.1 Microelectrode array analysis of in vitro activation patterns. **A.** The array of microelectrodes is wire-bonded to a circuit board carrier. A 35 mm Petri dish with an open center is adhered to the circuit board using bio-compatible epoxy. **B.** A 6×6 array of microelectrodes is fabricated on a glass substrate for electrical recording. The electrical propagation pattern can be represented through lateral isochrones maps or as through velocity vector maps. **C.** The lateral isochrones map is an interpolation of the time delay between electrodes, which is displayed by varying shades of color, overlaid on a representation of the 6×6 electrodes. The depolarization wave is initiated in the blue region at an activation time = 0 s, and propagates towards the red region, activation time = 0.04s. **D.** The velocity vector map is calculated by grouping the recording electrodes by three, and solving for the direction and magnitude of the time delay. 98

6.2	Algorithm to calculate the co-occurrence matrix based on counting the number of incidents that each local activation time occurs in an activation plot. Within a given activation plot, A , each entry is paired with its nearest neighbor, B , and noted, C . The values within each pair correspond to the coordinates of the co-occurrence matrix, and the number of occurrences is tallied, D . Steps B-D are repeated for each entry in the original matrix in A such that the final co-occurrence matrix is the sum of all these matrices.	104
6.3	Examples of characteristic activation plots (top) and their co-occurrence matrices (bottom). Activation plots consist of 300×300 pixels and exhibit grayscale intensities from 0 to 255. Smooth images are characterized through co-occurrence matrices that are clustered around the diagonal (left). The more the image is disordered, the farther away from the diagonal the co-occurrence matrix is populated (right). . . .	106
6.4	Optical micrographs of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically increase the cardiomyocyte-to-fibroblast ratio from 0:100 to 100:0 in increments of 5. As the cardiomyocyte-to-fibroblast ratio increases, here shown for 0:100, 60:40, 70:30, 80:20, 90:10, and 100:0, the co-culture system becomes increasingly dominated by circular-shaped cardiomyocytes (bottom right) and less dominated by spindle-shaped fibroblasts (top left). . . .	108
6.5	Fluorescence expression of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts, where HL-1 cardiomyocytes are transfected with green fluorescence protein (GFP) through a lentiviral vector. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Qualitatively, GFP expression decreases with decreasing cardiomyocyte-to-fibroblast ratio, here shown for 100:0, 90:10, 70:30, and 50:50 (top). Quantitatively, GFP expression decreases approximately linearly with decreasing cardiomyocyte-to-fibroblast ratios, shown in terms of averages and standard deviations for groups of $n=3$ (bottom).	116

6.6 Action potential profiles of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts, displayed for all 32 electrodes on the microelectrode array. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. The smoothness of the conduction pattern decreases with decreasing cardiomyocyte-to-fibroblast ratio, here shown for 100:0, 80:20, 70:30, and 50:50. The 100:0 plain cardiomyocyte culture displays a homogeneous conduction pattern (left). The 80:20 and 70:30 co-cultures demonstrate an intermediate behavior between smooth conduction and isolated spontaneous beating (middle). The 50:50 co-culture displays spontaneous beating, indicated through isolated signals in specific channels, with no sign of homogenous conduction (right). 117

6.7 Velocity vector maps (top) and percent active area (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. The active area fraction decreases with decreasing cardiomyocyte-to-fibroblast ratio, as the electrical activity becomes less synchronized. Pure 100:0 cultures of cardiomyocytes (left) support continuous, almost unidirectional, propagation throughout the homogeneous population. Heterogeneous 80:20 populations (right) exhibit a non-smooth conduction pattern impeded by non-conductive fibroblasts. The active area fraction (bottom) reveals that action potentials can only be recorded above a critical 50:50 ratio threshold. 118

6.8 Lateral isochrones map of in silico co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts, simulated with our finite element model. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5, here shown for 100:0, 90:10, 70:30, 50:50, 30:70, and 10:90. For the first set of simulations, 1:1 aspect ratio (top row), we model cardiomyocytes and fibroblasts to be of equal size. For the second set of simulations, 1:2 aspect ratio (bottom row), we model fibroblasts twice as long as cardiomyocytes. Continuous electrical conduction paths do not form below a critical cardiomyocyte-to-fibroblast threshold of 50:50. Similar to the lateral isochrones map in Figure 10.1C, early stage depolarization is shown in blue, while late depolarization is shown in red. 119

6.9 Isolated activation plots (top) and co-occurrence matrices (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Representative activation plots display the local activation times as a depolarization wave travels from the white regions to the dark regions. Pure 100:0 cardiomyocyte cultures (left) support gradual, smooth activation patterns, indicated through a single unique propagation front. More heterogeneous 90:10 (middle) and 80:20 (right) populations exhibit non-smooth activation patterns, resulting in a more diffuse signal propagation. With decreasing cardiomyocyte-to-fibroblast ratios, the co-occurrence matrices become less focused along their diagonals, and more populated in off-diagonal regions. 120

6.10	Co-occurrence energy (top) and co-occurrence contrast (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. We exclude results below the critical cardiomyocyte-to-fibroblast threshold of 80:20, for which less than half of the microelectrodes record an electrical signal. Energy and contrast values for each co-culture group are displayed in terms of averages and standard deviations for groups of n=3, normalized to the average value of the 100:0 group. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating the gradual disruption of the electrical conduction system.	121
6.11	Co-occurrence energy (top) and co-occurrence contrast (bottom) of in silico co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts, simulated with our finite element model. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Energy and contrast values for each co-culture group are displayed in terms of averages and standard deviations for groups of n=5, normalized to the average value of the 100:0 group. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating the gradual disruption of the electrical conduction system.	122
7.1	Channelrhodopsin-2 (ChR2) is a light-gated cation channel native to the green alga <i>Chlamydomonas reinhardtii</i> . It consists of seven transmembrane proteins and absorbs light through its interaction with retinal. Here, we induce channelrhodopsin coupled to yellow fluorescent protein (eYFP) into undifferentiated human embryonic stem cells via a lentiviral vector and differentiate these cells into cardiomyocytes. . .	124

- 7.2 Channelrhodopsin-2 (ChR2) is activated by photoisomerization of all-*trans* retinal to 13-*cis* retinal at wavelengths of 470 nm. After photoisomerization, the covalently bound retinal spontaneously relaxes to all-*trans* in the dark, providing closure of the ion channel and regeneration of the chromophore. 126
- 7.3 Three-state model for the channelrhodopsin photocycle. Upon photo absorption, molecules in the closed state g_c undergo a fast transition into the open state g_{ChR2} . After for some time, molecules spontaneously turn into the recovering state g_r where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After a recovery period, the molecules finally return to the closed state g_c , ready to undergo a new photocycle when subjected to light. 130
- 7.4 Ionic model of genetically engineered light sensitive cardiac cell. The electrochemical state of the cell is characterized in terms of $n_{\text{ion}} = 8$ ion concentrations, $c_{\text{ion}} = [c_{\text{Na}}^e, c_{\text{K}}^e, c_{\text{Ca}}^e, c_{\text{Na}}^i, c_{\text{K}}^i, c_{\text{Ca}}^i, c_{\text{Ca}}^{\text{up}}, c_{\text{Ca}}^{\text{rel}}]$, the extra-cellular and intracellular sodium, potassium, and calcium concentrations, and the sarcoplasmic reticulum calcium uptake and release. Ion concentrations are controlled through $n_{\text{crt}} = 12$ ionic currents, $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{f}}, I_{\text{NaCa}}, I_{\text{NaK}}, I_{\text{ChR2}}, I_{\text{K}}, I_{\text{bK}}, I_{\text{CaL}}, I_{\text{CaT}}, I_{\text{up}}, I_{\text{rel}}]$, where the baseline autorhythmic cell model [49] has been enhanced with the channelrhodopsin current I_{ChR2} , here shown in blue. The channels are governed by $n_{\text{gate}} = 10$ gating variables $g_{\text{gate}} = [g_{\text{m}}, g_{\text{h}}, g_{\text{ChR2}}, g_{\text{y}}, g_{\text{x}}, g_{\text{dL}}, g_{\text{fL}}, g_{\text{fCa}}, g_{\text{dT}}, g_{\text{rT}}]$ which may be functions of the current membrane potential ϕ 131

7.5 Undifferentiated human embryonic stem cells (hESC) stably transduced with a ChR2-eYFP lentiviral vector (hESC^{ChR2+}) remain pluripotent and can differentiate into cardiomyocytes (hESC^{ChR2+}-CM). **a.** PCR shows that hESC^{ChR2+} express the pluripotent Oct-4 gene (169 bp, lane 4) and Nanog gene (154 bp, lane 5), solid blue box. In addition, amplification within the ChR2 gene (174 bp, lane 6), across the ChR2-eYFP gene (197 bp, lane 7), and within the eYFP gene (187bp, lane 8), confirms stable transduction of the ChR2-eYFP lentivirus in undifferentiated hESC^{ChR2+}, solid yellow box. A ladder (100 bp, lane 1) confirms the predicted sizes of PCR products. Non-template control (lane 2) and GAPDH (152 bp, lane 3) serve as negative and positive controls, respectively. **b.** Pluripotent hESC^{ChR2+} stain positive for alkaline phosphatase (red). **c.** Fluorescence microscopy shows hESC^{ChR2+} positive eYFP signal (green). **d.** hESC^{ChR2+}-CM have positive TnI signals (red), consistent with a CM phenotype. DAPI staining (blue) demonstrates the position of nuclei. **e.** Transmission electron microscopy shows sarcomeres with associated z-lines (z) and mitochondria (m) in hESC^{ChR2+}-CM. **f.** Light microscopy shows three hESC^{ChR2+}-CM colonies (dashed white circles) on a multielectrode array. 141

7.6 Experimental and computational sensitivity of hESC^{ChR2}-CM with respect to light intensity. With the light turned on, the photocurrent I_{ChR2} increases rapidly, peaks, and decays towards a characteristic plateau value. With the light turned off, the photocurrent I_{ChR2} drops rapidly and decays to zero. Light intensity is varied from 12.5% to 25%, 50%, and 100% (top). Whole cell voltage clamp reveals an increased photocurrent I_{ChR2} as the light intensity increases. The computational hESC^{ChR2}-CM model captures the light sensitivity and displays increased photocurrents I_{ChR2} with increased light intensity (bottom). . 142

7.7 Experimental and computational sensitivity of of hESC^{ChR2}-CM with respect to stimulation frequency. Light stimulation (blue) evokes field potentials (black) which translate into mechanical contractions (red). Light stimulation at 100% intensity is performed at 0.5 Hz (top), 1.0 Hz (middle), and 1.5 Hz (bottom). Evoked signals during light stimulation (center) are markedly different from pre- and post-stimulation signals at all frequencies (left and right). The computational hESC^{ChR2}-CM model (green) captures the electrical signal at all frequencies, both during light stimulation (center), and pre and post stimulation (left and right). 143

7.8 Virtual activation sequences of light-paced hearts. Atrioventricular (AV) node photostimulation (top) is initiated through hESC^{ChR2}-CM, virtually injected into the basal region of the septum, while all other regions are modeled is standard ventricular CM. A depolarization wave forms at the AV node, travels down the septum, and activates the left and right ventricles. Bi-ventricular (bi-VP) photostimulation (bottom) is initiated through hESC^{ChR2}-CM, virtually injected into the lateral wall of the left and right ventricles, while all other regions are modeled is standard ventricular CM. A depolarization wave forms at the lateral left and right ventricular walls, travels along the ventricles, and activates the apex and the septum. The color code indicates the magnitude of the transmembrane potential ϕ varying from -90mV (blue) to +20mV (red). 144

8.1	Channelrhodopsin-2 is a light-gated cation channel native to the green alga <i>Chlamydomonas reinhardtii</i> . It consists of seven transmembrane proteins and absorbs blue light through its interaction with retinal. Photoisomerization of retinal opens the channel to sodium ions, which have a higher concentration outside than inside the cell. To make our cells responsive to light, and allow sodium concentrations to equilibrate, we induce channelrhodopsin into a conventional cardiac muscle cell model.	148
8.2	Channelrhodopsin-2 is activated by photoisomerization of all- <i>trans</i> retinal to 13- <i>cis</i> retinal at wavelengths of 470 nm. After photoisomerization, the covalently bound retinal spontaneously relaxes to all- <i>trans</i> in the dark, providing closure of the ion channel and regeneration of the chromophore.	148
8.3	Multiscale model for the photoelectrochemistry of living systems. Optical stimulation opens the cation channel channelrhodopsin g_{ChR2} . This initiates a photocurrent I_{ChR2} increasing the chemical concentration of sodium ions c_{Na} inside the cell. Concentration changes evoke changes in the electrical potential ϕ , which propagates across the tissue system in the form of smooth excitation waves.	149
8.4	Three-state model for the channelrhodopsin photocycle. Upon photoabsorption, molecules in the closed state g_{closed} undergo a fast transition into the open state g_{ChR2} . The molecules spontaneously turn into the refractory state g_{refrac} where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After the refractory period, the molecules return to the closed state g_{closed} , ready to undergo a new photocycle when subjected to light [133, 138].	151

8.5 Genetically engineered light sensitive cardiac cell. The electrophysiology of the cell is characterized in terms of $n_{\text{ion}} = 4$ ion concentrations, the intracellular sodium, potassium, and calcium concentrations and the calcium concentration in the sarcoplasmic reticulum. Ion concentrations are controlled through $n_{\text{ert}} = 16$ ionic currents, where we have enhanced the conventional cell model [188, 197] with the channelrhodopsin photocurrent I_{ChR2} , here shown in blue [1]. The channels are governed by $n_{\text{gate}} = 14$ gating variables, where we have added the channelrhodopsin gating variable g_{ChR2} to characterize the cell's response to photostimulation. 152

8.6 Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of sodium activation gate g_m , fast sodium inactivation gate g_h , slow sodium inactivation gate g_j , L-type calcium activation gate g_d , L-type calcium inactivation gate g_f , intracellular calcium dependent calcium inactivation gate g_{fCa} transient outward activation gate g_r , transient outward inactivation gate g_s , slow delayed rectifier gate g_{xs} , rapid delayed rectifier activation gate g_{xr1} , rapid delayed rectifier inactivation gate g_{xr2} , inward rectification factor g_{K1}^∞ , calcium-dependent inactivation gate g_g , and channelrhodopsin activation gate g_{ChR2} . The gating dynamics for the electrically stimulated cell have been delayed by 34 ms for the purposes of comparison against the optically stimulated cell. 167

8.7 Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of fast sodium current I_{Na} , background sodium current I_{bNa} , sodium potassium pump current I_{NaK} , sodium calcium exchanger current I_{NaCa} , inward rectifier current I_{K1} , rapid delayed rectifier current I_{Kr} , slow delayed rectifier current I_{Ks} , plateau potassium current I_{pK} , transient outward current I_{t0} , L-type calcium current I_{CaL} , background calcium current I_{bCa} , plateau calcium current I_{pCa} , leakage current I_{leak} , sarcoplasmic reticulum uptake current I_{up} , sarcoplasmic reticulum release current I_{rel} , and channelrhodopsin current I_{ChR2} . The current dynamics for the electrically stimulated cell have been delayed by 34 ms for the purposes of comparison against the optically stimulated cell. 168

8.8 Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of intracellular sodium c_{Na} , potassium c_K , calcium c_{Ca} concentrations, and calcium concentration c_{Ca}^{sr} in the sarcomplastic reticulum. The chemical concentration dynamics for the electrically stimulated cell have been delayed by 34 ms for the purposes of comparison against the optically stimulated cell. 169

8.9 Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of transmembrane potential ϕ . The characteristic action potential consists of five phases. Phase 0: The rapid upstroke is generated through an influx of positively charged sodium ions. Phase 1: Early, partial repolarization is initiated through the efflux of positively charged potassium ions. Phase 2: During the plateau, the net influx of positively charged calcium ions is balanced by the efflux of positively charged potassium ions. Phase 3: Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. Phase 4: Throughout the interval between end of repolarization and the beginning of the next cycle the cell is at rest. The transmembrane potential for the electrically stimulated cell has been delayed by 34 ms for the purposes of comparison against the optically stimulated cell. This delay agrees nicely with the time delay of activation of 19.7 ± 3.4 ms reported in the literature [?]. 170

8.10 Virtual injection of genetically engineered light sensitive cardiac cells into a human heart. Magnetic resonance imaging generates a sequence of two-dimensional images at different depths (top, left). We segment cardiac muscle tissue semi-manually using standard image processing techniques (bottom, left). Thresholding and binary masking convert the raw grayscale images to monochrome images with sharply defined boundaries (top, right). From these slices, we create a preliminary triangular surface mesh and converted it into the final tetrahedral volume mesh consisting of 3,129 nodes and 11,347 tetrahedral elements (bottom, right). Last, we virtually inject photosensitive cells into different regions of the heart and stimulated with light (middle). 171

8.11 Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for atrioventricular node paced heart. Photosensitive cells are virtually injected into the basal region of the septum, while all other regions are modeled as conventional cardiac muscle cells. A depolarization wave forms at the atrioventricular node, travels down the septum, and activates the left and right ventricles. 172

8.12 Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for apically paced heart. Photosensitive cells are virtually injected into the apex, while all other regions are modeled as conventional cardiac muscle cells. A depolarization wave forms at the apex, travels up, and activates the septum and both ventricles simultaneously. 173

8.13 Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for bi-ventricularly paced heart. Photosensitive cells are virtually injected into the lateral walls of the left and right ventricles, while all other regions are modeled as conventional cardiac muscle cells. Two depolarization waves form in the lateral left and right ventricular walls to travel along the ventricles and activate the apex and the septum. 174

9.1 (Left) Poisson interpolation results of fiber orientation angle with respect to circumferential fiber. Epicardial and right ventricular septal surfaces are assigned as -70° , while the endocardial surfaces are assigned as 80° as boundary conditions. (Right) Resulting interpolated fiber directions throughout the heart. 183

9.2	(Left) Fiber directions at various slices in the heart are shown. (Right) The normal directions at the corresponding slices are shown. The colors at each point of the cross-sectional slices correspond to the fiber angle orientation with blue representing -70° and red representing $+80^\circ$	184
9.3	(Left) Poisson interpolation results of fiber orientation angle with respect to circumferential fiber. Epicardial and right ventricular septal surfaces are assigned as -70° , while the endocardial surfaces are assigned as 80° as boundary conditions. (Right) Resulting interpolated fiber directions throughout the heart.	185
9.4	(Left) Fiber directions at various slices in the heart are shown. (Right) The normal directions at the corresponding slices are shown.	186
9.5	Representative images/fiber of angle fiber error given random subsets of boundary conditions for the fiber interpolation algorithm. Fibers highlighted in blue have no error (0°) whereas red fibers are closer to perpendicular (90°). Increasing sizes of random subsets are chosen and are shown starting from 10% (top left), 20% (top center), 30% (top right), 40% (left), 50% (center), 60% (right), 70% (bottom left), 80% (bottom center), 90% (bottom right).	191
9.6	Average angle error when 10% to 100% of the surface nodes used as boundary conditions to generate fiber orientation distribution.	192
9.7	(left) Fiber direction color coded by angle error. (right) Fiber angle error over the heart without fibers. Small hexagonal patterns are formed, because the nodes at the coarsest mesh are used as the boundary condition set, and therefore the new nodes on the surfaces are not part of this boundary condition set.	192
9.8	Histogram of fiber angle error in first subdivision mesh compared to second subdivision mesh. The histogram is organized into bins of 5° error.	193

10.1	Multiscale model of the human heart. At the molecular level, gating variables g_{gate} and ion concentrations c_{ion} characterize the bio-chemical response. At the cellular level, ionic currents I_{crt} and the transmembrane potential ϕ characterize the chemo-electrical response. At the organ level, the propagation of the electrical potential ϕ and the deformation φ characterize the electro-mechanical response.	197
10.2	Human ventricular cardiomyocyte model with 15 ionic currents resulting from ten transmembrane channels, one exchanger, and one pump. Three additional currents characterize ionic changes inside the sarcoplasmic reticulum, shown in grey. Sodium currents are indicated in red, potassium currents in orange, and calcium currents in green. . .	205
10.3	Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of the transmembrane potential ϕ and of the intracellular sodium potassium and calcium concentrations c_{Ca} , c_{K} , and c_{Ca} . The influx of positively charged sodium ions generates a rapid upstroke in the transmembrane potential. At peak, the efflux of positively charged potassium ions initiates an early, partial repolarization. During the plateau, the influx of positively charged calcium ions balances the efflux of positively charged potassium ions. Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. The cell is at rest throughout the interval between the end of repolarization and the beginning of the next cycle.	218
10.4	Mechanical contraction in a human ventricular cardiomyocyte. Temporal evolution of the active force F^{act} throughout an excitation cycle. The rapid increase in the intracellular calcium concentration c_{Ca} initiates a rapid increase in the active force. After reaching its peak value, the force gradually returns to zero.	219

10.5	Human heart model. The finite element discretization is generated from magnetic resonance images, left. The mesh consists of 46,896 linear tetrahedral elements, 13,831 nodes, and 55,324 degrees of freedom. The fiber orientation is created from feature interpolation using a Lagrangian interpolant, right. Fiber directions vary gradually from -70° in the epicardium, the outer wall shown in blue, to $+80^\circ$ in the endocardium, the inner wall shown in red.	220
10.6	Chemo-electro-mechanical coupling in the human heart. Spatio-temporal evolution of the fiber contraction λ_{ff} , the transmembrane potential ϕ , the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} , and the calcium concentration in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$ during the rapid depolarization phase of the cardiac cycle. .	221
10.7	Chemo-electro-mechanical coupling in the human heart. Spatio-temporal evolution of the fiber contraction λ_{ff} , the transmembrane potential ϕ , the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} , and the calcium concentration in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$ during the gradual repolarization phase of the cardiac cycle. Changes in the individual ion concentrations initiate a slow decrease in the transmembrane potential ϕ from $+20$ mV to -86 mV. A decrease in the intracellular calcium concentration c_{Ca} initiates mechanical relaxation with λ_{ff} returning gradually to 0%. During the filling phase, the apex moves away from the base and the heart undergoes a counterclockwise rotation back to its original position.	222
10.8	Mechanical contraction in a human heart. Temporal evolution of apical lift δ and of ventricular torsion ϑ . Shortly after the onset of excitation, the apex lifts rapidly towards the base moving upward by approximately 14mm, left. Simultaneously, the heart twists rapidly about its long axis rotating clockwise by approximately 8° and 17° , with the amount of torsion increasing from the fixed base to the free apex, right. Both apical lift and ventricular torsion then decrease gradually as the heart returns to its original position.	223

10.9	Algorithmic performance. Time step size and number of iterations for adaptive time stepping scheme. The algorithm typically converges within four Newton Raphson iterations. For more iterations, the adaptive algorithm automatically decreases the time step size, for example, during the rapid upstroke phase before $t = 0.05$ s and during the repolarization phase between $t = 0.25$ s and $t = 0.32$ s. For less iterations, the adaptive algorithm automatically increases the time step size, for example during the plateau phase, between $t = 0.05$ s and $t = 0.25$ s and during the resting phase after $t = 0.32$ s. The total number of time increments is 1,288, and the overall run time is 51.97 hours, calculated on a single core of an i7-950 3.06 GHz desktop with 12GB of memory.	224
11.1	Percentage of computation time dedicated to finite element assembly (red) and the CG solver (blue) respectively on the CPU (left) and GPU (right).	230
11.2	The proportion of time spent on different finite element assembly tasks within our GPU implementation for different refinements is shown (top left). Proportions of SPMV with relation to other tasks for the CG solver is shown (top right). Total proportional time spent on SPMV is shown (bottom).	231
11.3	A histogram of the non-zeros per row in a small 3129 row stiffness matrix for a patient-specific simulation of the heart.	237
11.4	Sort rows of the matrix from longest to shortest.	238
11.5	Arrange rows into groups of warp size and then pad accordingly. In this figure, we use a warp size = 4 only for purposes of illustration.	239
11.6	Reorder the first warp in Figure 11.5 in a coalesced column-major ordering.	239
11.7	A threshold value (purple dashed line) is given and the original padding in WPK1 is padded and grouped into new warps.	244

11.8	Coloring represents the data for a particular thread. In this instance, 4 threads are used to process the representative row. Data is arranged in column-major order, and then a parallel reduction step occurs resulting with the final sum in the first entry of the output row.	245
11.9	Benchmark results over matrices for 1 iteration.	249
11.10	Benchmark results over matrices for 50 iteration.	250
11.11	Benchmark results over matrices for 1200 iteration.	251
11.12	Effective bandwidth [GB/s] comparison over 1, 50, and 1200 iterations for cusp-csr, cusp-hyb, and cusparse implementations.	265
11.13	Effective bandwidth comparison over 1, 50, and 1200 iterations for mgpu, WPK1, and WPK2 implementations.	266
11.14	Comparison of coalesced "column-major" data vs. non-coalesced data ordering kernel times.	267
11.15	Comparison of WPK1 vs WPK1r kernel times.	267
11.16	Comparison of WPK2 vs WPK2r kernel times.	268
11.17	Comparison of kernel times to investigate the cost of the final remapping to the original reordering vs WPK1r.	268
11.18	Comparison of WPK1r vs WPK1rs kernel times.	269
11.19	Comparison of WPK2r vs WPK2rs kernel times.	269
11.20	Comparison of WPK2 vs WPK2 without shared memory buffering register.	270
11.21	Running time in [s] of the a single cardiac cycle (1 second) for different finite element tested configurations: GPU CG solver, GPU finite element implementation, and multi-core PETSc implementation. . . .	270
11.22	The resulting factor of increase in speed is shown for the GPU CG Solver and GPU finite element method with a single core PETSc block-jacobi CG solver as reference. PETSc implementation speedups are also shown.	271
11.23	TDP normalized resulting factor of increase in speed is shown for the GPU CG Solver and GPU finite element method with a single core PETSc block-jacobi CG solver as reference.	272

11.24	Speed up factor for GPU reordered WPK1 and WPK2 CG and finite element implementations.	273
-------	------------------------------------------------------------------------------------------------	-----

Chapter 1

Introduction

While modern science has revealed much about the workings of the heart, many matters of the heart are complicated as different pathologies are manifested in a variety of ways due to the intrinsic inter-dependent nature of cardiac tissue. Simulation tools can be leveraged to help elucidate underlying causes to complex pathologies. Much work has been done to develop accurate models of the heart; however, in order for simulations tools to have a meaningful impact in clinical and experimental settings, these tools must be efficient, fast, robust, and accessible. The focus of this thesis is to develop various methods of addressing these issues, and to highlight examples and applications of these methods.

Many researchers have been tackling issues regarding the modelling of the mechanics and electrophysiology of tissue [88], as well as other related aspects of interests in cardiac tissue including: metabolic demands [41] and perfusion[35], and coupling between the different multi-field properties[71] . The field of biophysics has provided detailed models of the biological underpinnings of cardiac electromechanics. Mechanicians on the other hand have utilized continuum mechanics as a tool for describing the macroscopic behavior of tissues. While each aspect of the heart - electrical, mechanical, and chemical - provide individual challenges in modeling and elucidating its behavior in both healthy and diseased physiological states, the strongly-coupled behavior necessary for healthy heart function provides additional challenges.

Over the years, researchers studying the heart have employed a divide-and-conquer

approach; different physical phenomena are studied separately, and then combined to determine the behavior of the whole. While scientifically and philosophically sound, this strategy ignores the inherent difficulty in model integration that is robust and computational efficient. Lack of numerical emphasis has resulted in ad-hoc schemes where collaborating researchers attempt to 'glue' together their respective contributions through use of different complex staggered numerical schemes. Ultimately, many of these schemes have large trade-offs between performance, accessibility of implementation, and robustness when newer related problems are presented.

Therefore, the aim of this dissertation is to develop a cohesive approach for modeling both the individual and coupled aspects of the heart while concurrently addressing the issues of flexibility, accessibility, efficiency, and robustness. Several approaches are taken in developing an efficient, yet modular finite element framework for cardiac modelling. The first is to develop a reformulation of current cell models that is modular and efficient. The second is to carefully optimize the implementation of models mathematically and numerically. The final approach is to utilize multi-core and Graphics Processing Unit (GPU) technology to speed up calculations while retaining scientific accuracy. By taking advantage of the benefits from different approaches, it is hoped that the aforementioned issues will be sufficiently addressed such that cardiac simulation tools can impact clinical and experimental settings.

The dissertation is organized into the following sections. First a brief background on topics related to the thesis will be given with relation to the physiological aspects of the heart, description and brief overview of differences between healthy and diseased conditions of the heart. A brief explanation of continuum mechanics with regards to this work, and a section on finite elements and related numerical techniques will also be provided. Secondly a historic overview of past and current models and their use in modeling different aspects of cardiac tissue will be given with a brief comparison of computational cost and utility of the discussed models. Models discussed can be organized into two groups: electrophysiological and mechanical models. Their benefits and short-comings will be discussed. Thirdly, a monolithic coupling framework is described and proposed, of which the following portions of this dissertation is based upon. Applications of this monolithic coupling framework will be given in regards

to phenomenological and biophysical, and electrical and electro-mechanically coupled problems. In addition to heart simulations, simulations of microelectroarray experiments are used to demonstrate the suitability of co-occurrence matrices for characterizing the smoothness in conduction in cardiac tissue. The modeling framework is also used to develop novel cell models of recently bioengineering-ChR2 optically paceable cells. The models are then used to investigate future uses of ChR2 as a biological alternative to electric pacemakers.

The remaining sections relate to different improvements to current modeling techniques. A novel fiber interpolating algorithm is proposed and analyzed. The algorithm allows modelers to build approximate fiber distributions for meshes when fiber data is not directly measured, but can also be used to interpolate existing fiber data measurements on corresponding meshes. Lastly a chapter is dedicated to improving the computational performance of the algorithmic framework by utilizing massively parallel graphics processors (GPUs) using CUDA. Lastly, this thesis ends with concluding remarks regarding further applications or improvements to this body of work. An integrative proof-of-concept using consumer available hardware is described that utilizes the different aspects studied in the dissertation and is framed in an environment experimentalists and clinicians can intuitively understand.

Chapter 2

Background

The purpose of this chapter is to introduce fundamental concepts and techniques used in developing and analyzing electromechanical models of the heart in the context of this dissertation. While a general overview will be given for the physiology of the heart, epidemiological statistics on heart diseases, continuum kinematics, and finite elements, it may not be sufficient to fully understand Chapters 5 to ?? as electromechanical modeling of the heart is a broad interdisciplinary area spanning biology, mathematics, computation, and mechanics.

2.1 Physiology

2.1.1 Structure and function of the heart: A multi-scale, multi-field overview

The primary function of the heart is to pump blood throughout the body to supply the proper nutrients and oxygen to different biological tissues. The human heart is composed of four pumping chambers: the left and right atria at the top, and the left and right ventricles at the bottom (Figure 2.1). Arteries are vessels that carry blood away from the heart, while veins are those that carry blood back to the heart. Atria move blood received from veins into their respective ventricles, where blood is pushed throughout the pulmonary or systemic systems via arteries. The heart also contains 4

valves: tricuspid, pulmonary, mitral, and aortic valves. These valves control the flow of blood within the four chambers, and help produce pressure gradients necessary for moving blood through the lungs and the body back into the atria.

The flow of blood starts at the Superior and Inferior Vena Cava where deoxygenated blood is received from the body into the right atrium. Next the tricuspid valve releases the blood into the right ventricle, where it is then subsequently pumped through the pulmonary system (lungs). Blood is oxygenated while traveling through the lungs and flows into the left atrium through four pulmonary veins. Blood is then released into the left ventricle by the mitral valve. The left ventricle is substantially more muscular than the right ventricle because the resistance due to pressure in the systemic system is far greater than pumping resistance throughout the lungs. Oxygenated blood is then pumped through the aortic valve to the rest of the body through the Aorta. These are illustrated in Figure 2.1.

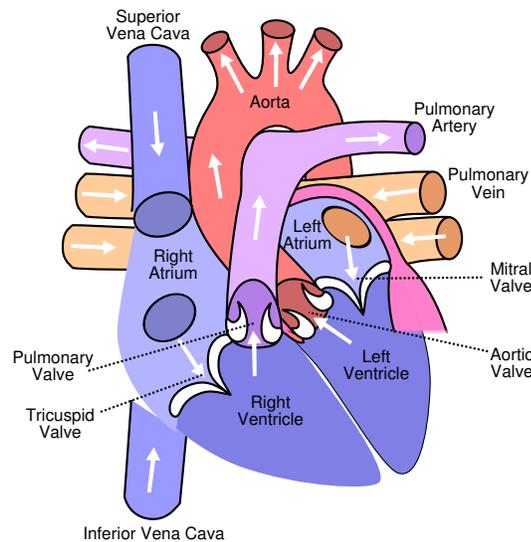


Figure 2.1: Diagram of the heart¹

The coordinated pumping behavior of the atria and ventricles is dependent on the synchronization of electrical, chemical, and mechanical dynamics within cardiac

¹Patrick J. Lynch, medical illustrator; C. Carl Jaffe, MD, cardiologist

tissue, and also dependent upon the structural electrical and mechanical properties of the heart. Cardiac tissue is composed primarily of two types of cells: myocytes, and fibroblasts. Cardiac myocytes are cells that are electrically active and mechanically contractile. Cardiac fibroblasts, on the other hand, are only electrically conductive and are not contractile, although they contribute to the passive mechanical behavior within the tissue. Myocytes are the predominant cell type in the healthy heart; therefore this thesis will mainly focus on electromechanical modeling issues regarding cardiac myocytes. Myocytes are connected to neighboring myocytes through cell-to-cell junctions, which allow for transport of materials, and most importantly conduction of electrical charges. Lastly, certain myocytes have a preferred contractile orientation that coincides with their preferred anisotropic electrical conduction direction [3, 63]. This anisotropy helps organize contraction and relaxation of cardiac muscle in each ventricle.

The following sections provide an overview of the different material properties observed in cardiac tissue, which are functionally and clinically relevant for diagnosis of heart conditions.

2.1.2 Electrophysiology of the cell

The heart is an electrically active material, where the electrical material properties serve as synchronizing mechanisms for coordinating the pumping motion of the heart. The heart is composed of an assortment of electrically active cell types; the location and electrical properties of these cells form an electrical circuit that has important timing characteristics for proper heart function. The electrical properties have multi-scale influence at the cellular, and macroscale levels.

At the cellular level, the temporal electrical waveform of a cell (action potential) is dependent on the ionic species that move in and out of the cell. The main ions of interest within cardiac cells are potassium, c_K , sodium c_{NA} , and calcium, c_{Ca} , concentrations within the cell. These ions are used both as signal regulators for other cellular processes and as also for modulating the action potential of a cell directly. Ion channels are pores in the cellular membrane that control the flow of different

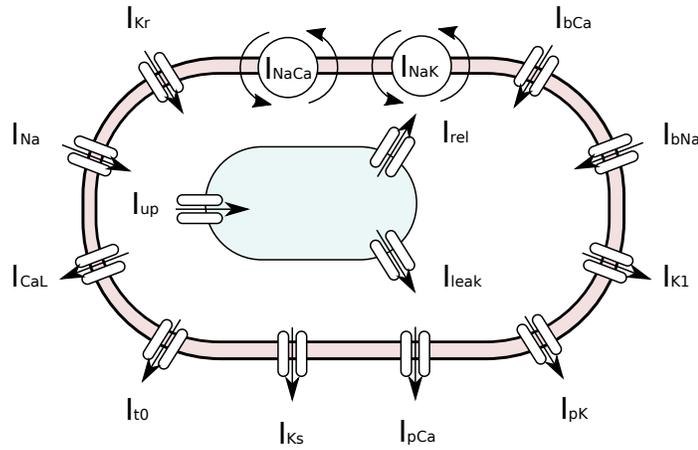


Figure 2.2: Schematic diagram of an cardiac cell

ions across the cell membrane (Figure 2.2). Ion channels can be passive, or dynamically change between physical configurations depending on the material properties the ionic gate is dependent upon. Many of the ion channels of interest in cardiac cells are voltage-dependent; the potential across the cellular membrane determines how closed or open a particular gate is in allowing for ion transport. However, other types of ion channels are also important for cardiac myocyte function and interesting examples will be given later in Chapters 7, 8. Ionic pumps are another mechanism for ion transport and are used to exchange a particular ion species for another, usually at the cost of adenosine triphosphate (ATP), the subcellular energetic currency within cells.

Intracellular potassium, sodium, and calcium concentrations are found within the sarcolemma, the cellular membrane of a cardiomyocyte. Sacromeres, which form the basis of the cellular contractile mechanisms exist within the sarcolemma and are modulated by levels of intracellular calcium. During relaxation, and pre-contractile states, intracellular calcium is accumulated within the sarcoplasmic reticulum (SR) which holds a large amount of intracellular calcium (blue compartment in Figure 2.2). The dynamic calcium behavior of the sarcoplasmic reticulum provides fine-tune control over muscle contraction, as the SR is spatially located close to sarcomeres and can quickly adjust the amount of calcium available within the sarcolemma. Lastly, the

gating dynamics of each of the transmembrane ion channels is directly determined by their internal gating states. While each particular type of ion channel has certain characteristics, their combined behavior poses restrictions on the rate at which these cells can generate an electrical pulse or become excited electrically. Their collective behavior with regards to difference pacing frequencies at different time points in each cell's action potential describes the recovery or restitution behavior of an excitable cell. Certain cardiac cells take longer than others to recover their excitability properties (to reset the ion channel gates) and this contributes to a global phenomenon called overdriving, where cells with faster intrinsic pacing rates override the self-exciting ability of slower self-pacing cells.

2.1.3 Electrophysiology of the heart

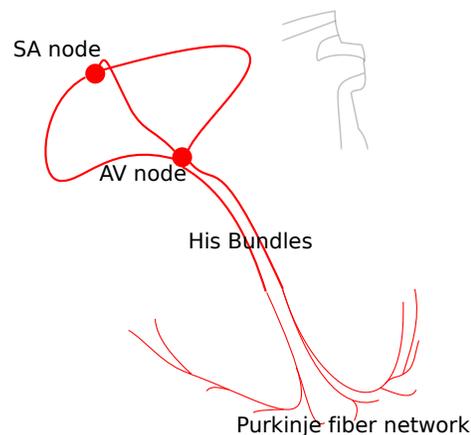


Figure 2.3: Diagram of an conduction system in the heart

The contraction of the heart is controlled by a network of electrically active cell types each with unique electrical properties important for their particular purpose (Figure 2.3). A heart beat is initiated at the sino-atrial (SA) node which is a group of cells located in the right atrium. These cells are the primary pacemakers of the

heart and are self-exciting. In healthy humans, the SA node is affected by the central nervous system so that the SA node paces at roughly 70 beats per minute (b.p.m.). Electrical potential conduction travels from the SA node throughout the atria to the atrioventricular node (AV) causing atrial tissue to excite and contract. The AV node is located between the atria and ventricles near the midseptum and has relatively slow conduction properties that effectively slow down the electrical pulse such that the propagating electrical wavefront causes atrial cells to contract. The AV node has a self-pacing frequency of about 60 b.p.m., but this phenomenon is generally overridden by the SA node, which has a faster pacing rate. From the AV node, the electrical conduction activates ventricular tissue and travels through the left and right His bundles. Electrical charge then subsequently passes through Purkinje cells which, together, form a high speed electrical distribution network throughout the rest of the ventricles such that contraction of the higher pressure and higher volume ventricles happens quickly and forcefully. His bundles and Purkinje cells have an intrinsic pacing rate of 30-40 b.p.m. The His-Purkinje system is distributed proportionally relative to the muscularity of the left and right ventricles [2]. Lastly, the electrical wavefront is conducted through ventricular tissue causing simultaneous contraction of the ventricles. Atrial and ventricular cells are not automatic phenotypes, and rely on the automatic electrical network to properly synchronize atrial and ventricular electromechanical pumping.

The majority of electrically active cells in the heart have preferred orientation directions. This provides structural anisotropic electrical and mechanical properties and is responsible for macroscopic torsion. Atrial and ventricular cells have anisotropic transverse conduction due to differences between axial and transverse electrical resistance [172]. This results in roughly a 2:1 or 3:1 ratio of axial vs. transverse conduction velocity. The His-Purkinje system acts as a point-to-point system where interactions within ventricular tissue happen only at specialized Purkinje fiber ventricular tissue insertion points.

Lastly, changes in individual cardiac myocytes show that the duration of each cell's action potential is strongly correlated to its activation time [58]. This implies a rough general relation between spatial and temporal properties due to the conductive

nature of cardiac tissue; cells activated later will be farther from the Purkinje fibers than those that are closer to the Purkinje fibers. This correlates with the relative difference in density of purkinje fiber insertion points between the endocardium and the epicardium, where the endocardium is excited first and subsequently at the epicardium [2]. Together the electrical network of heart cells coordinate macroscopic temporal and spatial pumping in the heart.

2.1.4 Mechanical Physiology - cellular and tissue level

The physiology of mechanical aspects of contractile cells is also important in understanding the mechanical function of the heart, as well as the coupling between electrical, chemical, and mechanical properties of individual cells, and of the tissue as a whole.

On the subcellular level, groups of sacromeres are the contractile apparatus of cardiac myocytes. Sacromeres are composed of different subcellular components, of which actin, myosin, and Z-discs are the most important. Z-discs are anchored to the sacrolemma. Actin filaments are smaller components that are attached to Z-discs and are situated in an orthogonal manner to the Z-disc plane. Myosin filaments are attached to the M-line which is mid-way between two Z-discs. Myosin heads attach and crawl along the actin attachments to contract the cell. The process by which myosin heads crawl along the actin attachment points is an electro-chemo-mechanically coupled phenomenon. At rest, tropomyosin molecules shield myosin head attachments, such that myosin cannot attach to actin. During cell depolarization, intracellular calcium is released into the sacroplasm and causes conformational changes to tropomyosin, allowing the myosin heads to crawl along the actin rails. As levels of intracellular calcium decrease, fewer myosin heads attach, and allow for relaxation of the cell. Thus, cellular contractile mechanics are dependent on an electrical - chemical - mechanical signal relay.

At the tissue level, the general preferred direction of cells correspond to the alignment of the sacromeres. The physical preferred ventricular cell orientations, which form structures similar to fiber arrangements in skeletal muscle, organize mechanical

contraction of cardiac tissue. These fibers serve both anisotropic electrical and mechanical functions. Macroscopically, ventricular cardiomyocytes have a special known spatial orientation[173]. Cells on the epicardium are -70° degrees with respect to a right-hand circumferential direction. The fiber angle varies transmurally till it reaches an 80° degree orientation at the endocardium. The right ventricle has a similar distribution, except the endocardium has two separate and distinct orientations. Along the septal wall, diving the two ventricles, it has a -70° angle orientation, while on the free wall, it has an 80° degree angle orientation. The organized torsional preference observed in the left ventricle most likely is related to the necessity in generating high systemic pressures in the left ventricle, as opposed to the lower pressure gradients needed in pumping blood through the pulmonary system.

2.2 Medical Conditions

The unique structure and organization in the healthy heart leads to reliability in its function as a pump. For example, the over-driving behavior of the Sinoatrial node provides proper pacing of about 70 b.p.m. at rest in healthy humans. Should the Sinoatrial nodes miss a beat or modulate in frequency due to sympathetic or parasympathetic modulation by the nervous system, the AV node and subsequent downstream His-Purkinje network serve to keep the heart pumping properly. Likewise, the heart is affected by the systemic and diastolic pumping resistances of the cardiovascular system which it is coupled to. The closed loop of cardiovascular and cardiac remodeling have made diagnosis of origins of disease more difficult for clinicians. Likewise, the coupled electrical, chemical, and mechanical properties of cardiac tissue have provided difficulties in understanding the underlying mechanisms behind electromechanical disorders. However, the relatively robust nature of the healthy heart has also made physiological and epidemiological studies of health and disease meaningful.

Physiologically, proper heart function has several characteristics, almost all of which have some mechanical significance. For example, for people in good health, a heart rate of 70 b.p.m. at rest provide an indication for the blood flow necessary to

provides cells and tissues throughout the body with enough nutrients and oxygen. Normal blood pressure in humans is about 115 mmHg (systolic) over 75 mmHg (diastolic) mean arterial pressure. Cardiac muscle perfusion relies on the diastolic pressure gradient to provide blood flow throughout the heart muscle. The 40 mmHg pressure difference between systolic and diastolic pressures indicate the amount of resistance the systemic arterial system in the body is providing. The difference is also important in calculating the work the heart is generating during pressure-volume loops, which are used to help clinicians determine the efficiency of a particular patient's heart.

Epidemiology, on the other hand, provides mainly statistical information over the health of a particular population, and may not give much insight into the underlying physical mechanisms behind the disease. However, epidemiology highlights the importance of studying particular diseases and can be used to motivate more accurate computational models to analyze complex electromechanically-coupled physiological phenomena. For example, a study from 2006 shows that heart disease is a leading cause of death resulting in roughly 600,000 deaths in the United States [81]. The study also indicates that heart disease accounted for 26% of deaths that year and occurs equally among men and women. The estimated cost of heart disease is about 315.4 billion per year [113].

While these statistics include an assortment of heart diseases, several are particularly interesting to computational modelers and represent the type of highly coupled behaviors that make pinpointing and analyzing the underlying mechanisms difficult for clinicians. Arrhythmias are diseases characterized by irregular beating of the heart. This can be due to changes in electrical pacing, electrical conduction properties, and even electromechanical changes due to other pathologies. Atrial fibrillation (AF) is the most prevalent arrhythmia and affecting over 2.66 million people in the US and the incidence also increases with age. AF is the primary cause for roughly 11,438 deaths and 461,000 hospital discharges in 2010. While ventricular fibrillation is only the cause of 1014 deaths, it is believed to be responsible for a large portion of the 310,000 sudden cardiac deaths per year [113]. Thus, arrhythmias are a primary example of coupled diseases that can be difficult to pin down to a specific cause, and may benefit from added computational analysis.

Coronary heart disease (CHD), on the other hand, has a prevalence of 17.6 million where those above the age of 20 with a 3.6% prevalence of myocardial infarction. In 2006, Coronary heart disease was responsible for 1 in every 6 deaths in the US[113]. In CHD, the coronary arteries that supply blood to cardiac tissue become clogged or start forming thromboembolisms, such that part of the heart dies due to lack of oxygen and nutrient supply. Since the electrical, chemical, and mechanical properties change with the onset of CHD, this is another area that cardiac modelers can perhaps provide additional scientific insight.

Lastly, there are several imaging and measurement modalities clinicians can use to obtain spatial and temporal varying information about the heart. Electrocardiograms (ECG) use multiple leads that sense the changes in electromagnetic fields along lines formed by the lead electrodes during the heart cycle. From ECGs, clinicians can infer how the heart is depolarizing and repolarizing, and whether the heart is functioning properly in terms of excitation and rhythm regularity. The convention taken in most ECGs is that the upward direction corresponds with the electrical flux towards the measuring lead (away from ground). With enough leads, one can determine whether the approximated projected conduction pattern in the heart is organized or irregular. ECGs can therefore be used to spot abnormalities in the frequency of the characteristic heart recording that may indicate arrhythmic behavior, or alternatively the unhealthy repolarizing behavior of cells.

2.3 Continuum Mechanics

As continuum mechanics is a broad area with a relatively lengthy and steep learning curve, this section primarily serves to introduce notation and definitions used in the modelling sections later on. This section will briefly cover kinematics and balance laws of heart mechanics. Constitutive relations will be discussed later for each individual simulation problem type. Tensors and tensor calculus will be used throughout the rest of this dissertation, along with simplified Einstein index notation when convenient.

2.3.1 Kinematics

The kinematics of a material describe how points in a body or space move both individually and in relation to other points. The area of kinematics overlaps with areas in topology, differential geometry, and differential calculus. For the purposes of this dissertation, we are mainly interested in studying the motion of particles in a material body (cardiac tissue) and we are interested in the deformation of solid materials; thus we will start from a Lagrangian or referential setting. This is also a common setting for finite element simulations of solid materials. Also, we will make a continuum assumption that the behavior of biological materials varies in a relatively smooth manner, at least from a macroscopic tissue level where we use continuum mechanics. Lastly, for the sake of convenience we will stick to using a Cartesian basis as Cartesian coordinate systems are commonly used when imaging biological objects.

Suppose we know the location of particles in a material body in an initial reference configuration. The placement mapping, φ , is a function which maps the position of a point in the initial configuration to the new position of the point in another configuration. The deformation at a material point from the initial to secondary deformed (spatial) configuration is described by the deformation gradient, \mathbf{F} . This tensor maps tangent lines in the initial configuration to those in the secondary configuration. It is defined as follows:

$$\mathbf{F} = \frac{\partial \varphi}{\partial \mathbf{X}} = \partial_{\mathbf{X}} \varphi \quad (2.1)$$

From basic geometry, the tangent to a curve is simply the spatial derivative along that line. Therefore the deformation gradient is simply a higher dimensional analog and describes how the tangent vectors of a line of points have changed or “deformed” with respect to the tangent vectors in the initial configuration. For rigid body translation, no deformation occurs, and thus \mathbf{F} is unaffected. The Jacobian, J is defined as

$$J = \det(\mathbf{F}) \quad (2.2)$$

The Jacobian describes the infinitesimal change in volume at a material point between the initial configuration and the subsequent configuration.

The right Cauchy-Green deformation tensor in the reference configuration, \mathbf{C} , is a symmetric tensor that is convenient in some situations when objectivity and symmetry of constitutive relations are important. Likewise there is a spatial analog called the Finger tensor, \mathbf{b} , in the spatial configuration. The Green-Lagrange strain tensor, \mathbf{E} , and the Euler-Almansi strain tensor, \mathbf{e} , are measures of strain in their respective reference and spatial settings.

$$\mathbf{C} = \mathbf{F}^T \mathbf{F} \quad (2.3)$$

$$\mathbf{b} = \mathbf{F} \mathbf{F}^T \quad (2.4)$$

$$\mathbf{E} = \frac{1}{2}[\mathbf{C} - \mathbf{I}] \quad (2.5)$$

$$\mathbf{e} = \frac{1}{2}[\mathbf{i} - \mathbf{b}^{-1}] \quad (2.6)$$

The spatial velocity gradient tensor, \mathbf{l} , is defined as

$$\mathbf{l} = \partial_x v = \partial_X v \partial_x X = \dot{\mathbf{F}} \mathbf{F}^{-1} \quad (2.7)$$

The spatial rate of deformation tensor, \mathbf{d} , is defined as the symmetric part of the spatial velocity gradient, \mathbf{l} :

$$\mathbf{d} = \frac{1}{2}[\mathbf{l}^T + \mathbf{l}] = \frac{1}{2}[\dot{\mathbf{F}} \mathbf{F}^{-1} + \mathbf{F}^{-T} \dot{\mathbf{F}}^T] = \mathbf{F}^{-T} \dot{\mathbf{E}} \mathbf{F}^{-1} \quad (2.8)$$

These spatial tensors have many uses and in relation to the work in this dissertation, are mainly used for calculating rates and for total derivative quantities from linearizing non-linear quantities in the context of finite elements. \mathbf{d} is also the push-forward of the rate of strain, $\chi^*(\dot{\mathbf{E}})$.

Push-forward & Pull-back operations

Push-forward and pull-back operations allow one to represent quantities in the material or spatial setting in terms of the other. It is important to note that different tensors transform differently. A tensor can be either covariant, contravariant, or mixed depending on how the tensor transforms. To facilitate these operations, maps

have been developed [118] that provide a visual mnemonic for converting different material and spatial first and second order tensorial quantities.

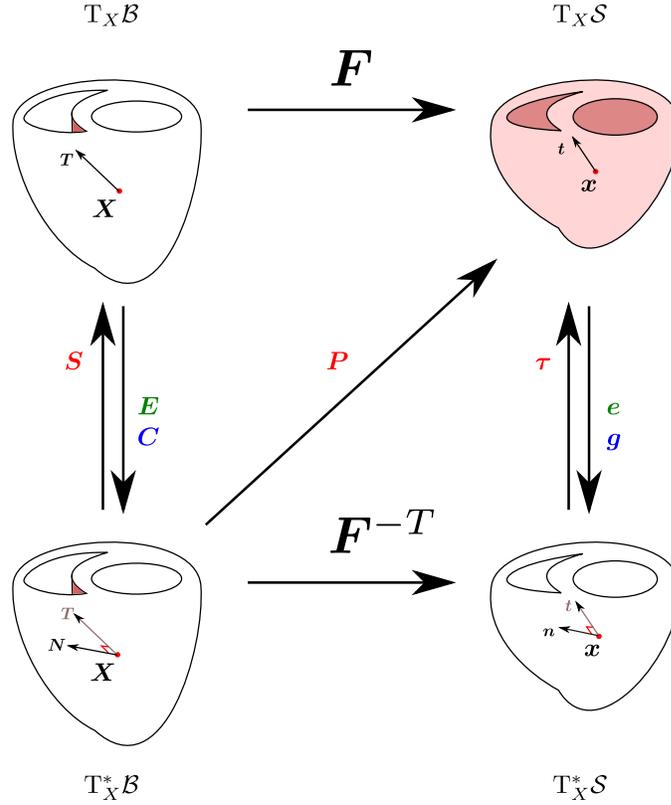


Figure 2.4: Continuum mapping relations

One can view the push-forward and pull-back operations as a series of mappings, and therefore, directly write down equivalent relations by use of the mapping directions in Figure 2.4. To move opposite to an arrow is to simply take the inverse mapping, $(\bullet)^{-1}$. In this diagram, there are four spaces: $T_X \mathcal{B}, T_X \mathcal{S}, T_X^* \mathcal{B}, T_X^* \mathcal{S}$. They represent the tangent spaces, $T_X(\bullet)$, between the material and spatial setting, and also the respective cotangential mappings, $T_X^*(\bullet)$. For example, the Second Piola Stress is defined as $S : T_X^* \mathcal{B} \rightarrow T_X \mathcal{B}$. Another equivalent mapping can be written in terms of the Kirkoff stress, τ , by taking another route from $T_X^* \mathcal{B}$ to $T_X \mathcal{B}$.

$$S = \chi_*(\tau) = F^{-1} \tau F^{-T} \quad (2.9)$$

Balance Laws

There are many two balance laws which are relevant to simulation of quasi-static mechanical and dynamic electrophysiological problems. The first is the balance of linear momentum:

$$\text{Div}(\mathbf{S}) + \mathbf{F}^\varphi = \mathbf{0} \quad (2.10)$$

where, \mathbf{F}^φ is a body force term, and \mathbf{S} is the Second-Piola stress tensor.

The other balance law is the ionic transport balance across the cell-membrane:

$$\dot{\phi} = \text{Div}(\mathbf{Q}) + F^\phi \quad (2.11)$$

where, F^ϕ , is the electrical source term, and \mathbf{Q} is the phenomenological material electrical flux, \mathbf{D}_o is the phenomenological material conduction tensor, and ϕ is the electrical transmembrane voltage:

$$\mathbf{Q} = \mathbf{D}_o \cdot \partial_{\mathbf{X}} \phi \quad (2.12)$$

These maps are particularly useful when dealing with first and second order tensor transformations.

2.4 Finite Elements and Numerics

While there are several different ways to solve boundary value problems with the given governing equations (2.10,2.11), the finite element method is used in this dissertation for several reasons. The finite element method is an established method used for solving boundary value problems in the field of solid mechanics [85]. Many of its properties have been well studied and many solutions have been carefully investigated in terms of modelling stiff problems such as incompressibility, and also issues related to geometric and material nonlinearity. In addition, the finite element method can naturally handle well-formed unstructured meshes because integration is performed per element and not over a stencil of nodes, as in finite difference methods. Because of these and other reasons, the standard non-linear finite element method is used

for both the mechanics and the electrical boundary value problems over arbitrary unstructured mesh domains (the heart).

2.4.1 Weak-form

As per the usual nonlinear finite element procedure, we will pose the strong form of the electrical (2.11) and mechanical (2.10) balance laws in the Galerkin-residual projection weak-form. The Galerkin residual of the aforementioned electrical and mechanical strong-forms are:

$$\dot{\phi} - \text{Div}(\mathbf{Q}) - \mathbf{F}^\phi = 0 \text{ in } \mathcal{B} \quad (2.13)$$

$$\text{Div}(\mathbf{S}) + \mathbf{F}^\varphi = \mathbf{0} \text{ in } \mathcal{B} \quad (2.14)$$

where $\mathcal{B} \subset \mathbb{R}^3$ is the region of material points in the reference configuration that corresponds to $\mathcal{S} \subset \mathbb{R}^3$ in the current spatial configuration. The essential and natural boundary conditions are

$$\phi = \bar{\phi} \text{ on } \partial\mathcal{S}_\phi \text{ and } \mathbf{q} = \bar{\mathbf{q}} \text{ on } \partial\mathcal{S}_q \quad (2.15)$$

$$\varphi = \bar{\varphi} \text{ on } \partial\mathcal{S}_\varphi \text{ and } \mathbf{t} = \bar{\mathbf{t}} \text{ on } \partial\mathcal{S}_t \quad (2.16)$$

where $\partial\mathcal{S}$ is the boundary surface domain of the current configuration composed of the union of the disjointed essential and natural domains.

The weak-forms are then to find $\phi \in H^1(\mathcal{B})$ and $\varphi_i \in H^1(\mathcal{B})$ such that

$$G^\phi(\partial\phi, \phi, \dots) = \int_{\mathcal{B}} \partial\phi \left(\dot{\phi} - \text{Div}(\mathbf{Q}) - \mathbf{F}^\phi \right) dV = 0 \quad (2.17)$$

$$G^\varphi(\partial\varphi, \varphi, \dots) = \int_{\mathcal{B}} \partial\varphi (\text{Div}(\mathbf{S}) + \mathbf{F}^\varphi) dV = \mathbf{0} \quad (2.18)$$

where $\partial\phi \in H^1(\mathcal{B})$ and $\partial\varphi_i \in H^1(\mathcal{B})$ are the square-integrable weighted functions that satisfy $\partial\phi = 0$ on $\partial\mathcal{S}_\phi$ and $\partial\varphi = \mathbf{0}$ on $\partial\mathcal{S}_\varphi$.

2.4.2 Linearization

Since F^ϕ in (2.22) and \mathbf{S} in (2.18) are non-linear, the Newton-Raphson method is used in conjunction with the finite element method to iteratively solve for the solution that satisfies:

$$\mathbf{0} = G^{(\bullet)}(\dots) + \Delta G^{(\bullet)}(\dots) \quad (2.19)$$

where $G^{(\bullet)}$ is the the galerkin residual and $\Delta G^{(\bullet)}(\dots)$ is the increment which contains the linearized tangent and update terms.

$$\Delta G^\phi = D_{\Delta\phi}G^\phi + D_{\Delta\varphi}G^\phi \quad (2.20)$$

$$\Delta \mathbf{G}^\varphi = D_{\Delta\phi}\mathbf{G}^\varphi + D_{\Delta\varphi}\mathbf{G}^\varphi \quad (2.21)$$

In order for consistent algorithmic linearzation we first discretize in time. We then linearize and perform spatial discretization. For simplicity, we use a first-order implicit finite difference time-stepping scheme, which has the added benefit of oscillation dampening, and G^ϕ becomes

$$G^\phi(\partial\phi, \phi, \dots) = \int_B \partial\phi \left[\frac{\phi - \phi_n}{\Delta t} - \text{Div}(\mathbf{Q}) - F^\phi \right] dV = 0 \quad (2.22)$$

For clarification, ϕ refers to the current iteration solution, whereas ϕ_n refers to the previous timestep solution. Before linearization we will also divide G^ϕ and G^φ into internal and external parts and apply the Divergence Theorem.

$$\begin{aligned} G_{\text{int}}^\phi &= \int_B \delta\phi \frac{\phi - \phi_n}{\Delta t} dV + \int_B \text{Grad}(\delta\phi) \cdot \mathbf{Q} dV \\ G_{\text{ext}}^\phi &= \int_{\partial B} \delta\phi \mathbf{T}^\phi dA + \int_B \delta\phi F^\phi dV \\ G_{\text{int}}^\varphi &= \int_B \mathbf{S} : \delta\mathbf{E} dV = \int_B \mathbf{S} : \text{sym}(\mathbf{F}^T \text{Grad}(\delta\varphi)) dV \\ &= \int_B \boldsymbol{\tau} : \delta\mathbf{e} dV = \int_B \boldsymbol{\tau} : \text{grad}(\delta\varphi) dV \end{aligned} \quad (2.23)$$

We linearize the individual terms to obtain:

$$D_{\Delta\phi}G_{\text{int}}^{\phi} = \int_B \frac{\delta\phi\Delta\phi}{\Delta t} dV + \int_B \text{Grad}(\delta\phi) \cdot \mathbf{D}_0 \text{Grad}(\Delta\phi) dV \quad (2.24)$$

$$D_{\Delta\phi}G_{\text{ext}}^{\phi} = \int_B \delta\phi[\partial_{\phi}F_e^{\phi} + \partial_{\phi}F_m^{\phi}]\Delta\phi dV \quad (2.25)$$

$$D_{\Delta\phi}G_{\text{int}}^{\varphi} = \int_B \text{sym}(\mathbf{F}^T \text{Grad}(\delta\varphi)) : \partial_{\phi}\mathbf{S}^{\text{act}} \Delta\phi dV \quad (2.26)$$

$$D_{\Delta\varphi}G_{\text{int}}^{\varphi} = \int_B \text{sym}(\mathbf{F}^T \text{Grad}(\Delta\varphi)) : \mathbb{C} : \text{sym}(\mathbf{F}^T \text{Grad}(\delta\varphi)) dV \quad (2.27)$$

$$+ \int_B \mathbf{S} : \text{sym}(\text{Grad}(\Delta\varphi) \text{Grad}(\delta\varphi)) dV \quad (2.28)$$

where $\mathbb{C} = 2\partial_{\mathbf{C}}\mathbf{S}$ is the tangent moduli of \mathbf{S} with respect to \mathbf{E} . The first term is the non-linear material tangent (2.26), and the second term is the geometric non-linearity part (2.28) of the tangent.

2.4.3 Lie Derivatives, \mathcal{L}

Lie Derivatives are sometimes used to describe terms in the consistent algorithmic linearization of the Galerkin residuals of the governing equations in the spatial setting. They are used to describe the derivatives of spatial quantities along a particular incremental direction. However, one can also simply interpret Lie derivatives, as the spatial push-forward of the material pull-back of the spatial quantity. In other words,

$$\mathcal{L}(\bullet) = \chi^*(D_{\Delta}\chi_*(\bullet)) \quad (2.29)$$

For example, when one derives the tangent for a governing equation, such as $D_{\Delta\varphi} G_{\text{int}}^{\varphi}$ the following are equivalent:

$$\begin{aligned}
D_{\Delta\varphi} G_{\text{int}}^{\varphi} &= \int_B D_{\Delta\varphi} \mathbf{S} : \text{sym}(\mathbf{F}^T \text{Grad}(\delta\varphi)) dV \\
&+ \int_B \mathbf{S} : \text{sym}(\text{Grad}(\Delta\varphi)^T \text{Grad}(\delta\varphi)) dV \\
&= \int_B \mathcal{L}_{\Delta\varphi} \boldsymbol{\tau} : \text{grad}(\delta\varphi) dV \\
&+ \int_B \boldsymbol{\tau} : \text{grad}(\Delta\varphi)^T \text{grad}(\delta\varphi) dV
\end{aligned} \tag{2.30}$$

While it may be more tractable for many to formulate their equations in the material or mixed settings, for simulations it is particularly meaningful to use directly measurable spatial quantities. While one can push-forward referential quantities, the computational cost of the push-forward operations is not insignificant as they are composed of a series of matrix vector multiplications. Therefore, it is generally most beneficial to directly code the spatial version of the governing equations, while one uses the most tractable equivalent version for derivations. Lie derivatives should not be intimidating! They can generally be interpreted in the context of (2.30). Of course, the same notational chain rules that apply to directional derivatives also apply to Lie derivatives.

2.4.4 Spatial Discretization

Lastly, we perform spatial discretization on the incremental terms (2.24-2.28) to complete the consistent algorithmic linearization of the Galerkin residual tangents. As the problem is non-linear, we will use reduced order integration and simply use a centroid integration Gaussian integration scheme with linear Lagrangian shape functions. Thus we use the following spatial elemental discretization for the incremental

update and the respective weighting functions.

$$\delta\phi^h|_{\mathcal{B}^e} = \sum_{i=1}^{ne\phi} N_{\phi}^i \delta\phi_i^e \quad , \quad \delta\varphi^h|_{\mathcal{B}^e} = \sum_{i=1}^{ne\varphi} N_{\varphi}^i \mathbf{x}_i^e \quad (2.31)$$

$$\Delta\phi^h|_{\mathcal{B}^e} = \sum_{i=1}^{ne\phi} N_{\phi}^i \delta\Delta\phi_i^e \quad , \quad \Delta\varphi^h|_{\mathcal{B}^e} = \sum_{i=1}^{ne\varphi} N_{\varphi}^i \Delta\mathbf{x}_i^e \quad (2.32)$$

\mathcal{B}_e is the domain of the element, while $\delta\phi_i^e$ and $\delta\varphi_i^e$ are the discretized test function values at the nodes and $\Delta\phi_i^e$ and \mathbf{x}_i^e are the updates to the current iteration solution at the nodes. N_{ϕ}^i and N_{φ}^i are the shape functions and for this dissertation have been chosen to be the same. $ne\phi$ and $ne\varphi$ are the number of nodes for a given element respective of the potential, ϕ , and displacement, φ degrees of freedom. Likewise, the gradients are also discretized.

$$\text{grad}(\delta\phi^h)|_{\mathcal{B}^e} = \sum_{i=1}^{ne\phi} N_{\phi,x}^i \delta\phi_i^e = \sum_{i=1}^{ne\phi} \delta\phi_i^e \otimes N_{\phi,x}^i \quad (2.33)$$

$$\text{grad}(\delta\varphi^h)|_{\mathcal{B}^e} = \sum_{i=1}^{ne\varphi} N_{\phi,x}^i \mathbf{x}_i^e = \sum_{i=1}^{ne\varphi} \mathbf{x}_i^e \otimes N_{\varphi,x}^i \quad (2.34)$$

$$\text{grad}(\Delta\phi^h)|_{\mathcal{B}^e} = \sum_{i=1}^{ne\phi} \Delta\phi_i^e \otimes N_{\phi,x}^i \quad (2.35)$$

$$\text{grad}(\Delta\varphi^h)|_{\mathcal{B}^e} = \sum_{i=1}^{ne\varphi} \Delta\mathbf{x}_i^e \otimes N_{\varphi,x}^i \quad (2.36)$$

After pushing-forward the material quantities, we are left with the following elemental residuals. They are written in indicial notation for the convenience of implementation.

$$\mathbf{R}_i^{\varphi} = \mathbf{A}_{e=1}^{nel} \left[\int_{B_e^h} \text{grad}(N_{\varphi}^i) \cdot \boldsymbol{\tau} dV - \int_{\partial S_e^e} N_{\varphi}^i \mathbf{t} da \right] \quad (2.37)$$

$$\begin{aligned} \mathbf{R}_i^{\phi} = \mathbf{A}_{e=1}^{nel} & \left[\int_{B_e^h} N_{\phi}^i \left[\frac{\phi - \phi_n}{\Delta t} \right] + \text{grad}(N_{\phi}^i) \cdot \mathbf{q} dV \right. \\ & \left. - \int_{\partial S_e^e} N_{\phi}^i \mathbf{t}^{\phi} da - \int_{B_e^h} N_{\phi}^i F^{\phi} dV \right] \quad (2.38) \end{aligned}$$

Lastly, terms in the elemental residual, R_i^e and elemental tangent, K_{ij}^e can be computed and organized in the following manner into the global residual and corresponding tangent matrix.

$$-R_i^\varphi + K_{ij}^{\varphi\varphi} \Delta\varphi_j + K_{ij}^{\varphi\phi} \Delta\phi_j = 0 \quad (2.39)$$

$$-R_i^\phi + K_{ij}^{\phi\varphi} \Delta\varphi_j + K_{ij}^{\phi\phi} \Delta\phi_j = 0 \quad (2.40)$$

As seen in (2.39, 2.40), the tangent may or may not have coupled terms. Coupling between electrical, chemical, mechanical and optical components is handled primarily through the formulation of F^ϕ , the spatial conduction tensor \mathbf{D} , and the stress tensor \mathbf{S} , which encompasses passive, active, and mechano-electric feedback behavior. The following chapter will further discuss the details regarding the constitutive relations for the electrical, mechanical, and coupled biophysical phenomena.

Chapter 3

Modeling overview

This chapter serves as an introduction to the various constitutive models used in this dissertation work. It is by no means comprehensive, and will not go into too much detail about the actual algorithmic implementation of these constitutive continuum models, as that will be covered in the following sections where the finite element algorithmic framework is introduced. Electrophysiological and mechanical constitutive models will be briefly covered; however other constitutive models, such as those regarding modelling of Channelrhodopsin-2 and other models will be covered in their respective sections as they do not generally pertain to cardiac modelling.

3.1 Electrophysiology

There are several classes of electrophysiological models that try to model different electrical cell behaviors in the heart. There are phenomenological models, which have parameters that may not be directly related to experimentally measurable quantities. Biophysical models, conversely, are related to directly measurable phenomena. Biophysical models are based on different cell experiments - voltage clamping and patch clamping - from which the cell membrane is viewed as a small capacitor, and changes in transmembrane ionic flow produce changes in electrical voltage across the cell membrane. Cell models that only model the transmembrane potential are called monodomain models, while bi-domain models also model the diffusion of ions in the

extracellular space (3.1). They are given by the following equations [?, 44]

$$\begin{aligned}\dot{\phi} &= \operatorname{div}(\mathbf{D}_i \cdot \partial_{\mathbf{x}}\phi) + \operatorname{div}(\mathbf{D}_e \cdot \partial_{\mathbf{x}}\phi_e) + f^\phi(\phi, \mathbf{r}) \\ 0 &= -\operatorname{div}(\mathbf{D}_i \cdot \partial_{\mathbf{x}}\phi) - \operatorname{div}(\mathbf{D} \cdot \partial_{\mathbf{x}}\phi_e)\end{aligned}\tag{3.1}$$

There are two equations. The first is again the balance of electrical flux across the cell membrane which is similar to the mono-domain equation, while the second equation is the balance of flux between the extracellular and intracellular domain. There are now two different conductivities: \mathbf{D}_i representing the intracellular conduction tensor and \mathbf{D}_e , which represents the extracellular conduction tensor. \mathbf{D} is the sum of the two conduction tensors. ϕ_e is the extracellular potential, while ϕ is the transmembrane potential. When the two conduction tensors are proportional, some of the divergence operators can be combined, and will yield the mono-domain equation. For the purposes of this dissertation, the constitutive models will concentrate on the mono-domain equations for two reasons. The first reason is that most of the changes in the cell happen within the cell-membrane, and therefore the ionic flux across the membrane is the primary phenomenon that causes changes in voltage. This can also be observed numerically, where the stiffness in solving the bi-domain equations is mainly due to the parabolic mono-domain portion of the equation. The second reason is that for most physiological instances, the proportionality of conductivity is a relatively good assumption; however when simulating defibrillation for example, it is imperative to properly simulate the conduction inhomogeneity between the intracellular and extracellular domains. Fortunately, with the algorithmic framework we will propose in the following section, it will become obvious that the issues between implementing bi-domain and mono-domain models are trivial partially because most electrophysiological constitutive equations regarding ionic flux across the cell-membrane are both dependent upon the transmembrane potential as seen in the source term of (3.1)₁.

While the previous class of equations are deterministic partial differential equations, in the sense that for a given set of boundary conditions and initial conditions, they will always result in the same solution, there are stochastic methods that may result in more random behavior[68]. Typically these models are used to model random

experimental phenomena observed at the subcellular level seen in calcium transient recordings. While these models can be embedded into our source term, it would not be computationally feasible and probably unnecessary to embed this stochastic behavior in our model within each element in our finite element simulations. In physiological conditions, this stochastic behavior would most likely be homogenized over the domain and result in the same behavior as our deterministic models. As the purpose of this thesis is to address macroscale phenomena, deterministic constitutive equations will primarily be used.

The following sections will cover deterministic mono-domain phenomenological and ionic models, and lastly a short section regarding modelling of conduction within the heart.

3.1.1 Phenomenological Models

While there are several different phenomenological models, the two classical electrophysiological models are the FitzHugh-Nagumo[56] and Aliev-Panfilov [4] cell models. Both are relatively simple two-variable models and are relatively easy to understand because phase plots can be determined for the models, and from there some properties of the equations can be determined a priori. The source terms for the FitzHugh-Nagumo model are:

$$f^\phi = c[\phi[\phi - \alpha][1 - \phi] - r] \quad (3.2)$$

$$f^r = \phi - br + a \quad (3.3)$$

The source terms for the Aliev-Panfilov model are:

$$f^\phi = c\phi[\phi - \alpha][1 - \phi] - r\phi \quad (3.4)$$

$$f^r = \left[\gamma + \frac{\mu_1 r}{\mu_2 + \phi} \right] [-r - c\phi[\phi - b - 1]] \quad (3.5)$$

The state variables are the transmembrane potential, ϕ , and the recovery variable, r . The other variables are material parameters that are used to adjust the rate behavior of the two state variables, and generally remain constant temporally, but may vary

spatially.

Historically, phenomenological models were used to reproduce nerve membrane excitation behavior, where an excitation threshold must be reached in order to illicit an action potential [?]. Phenomenological models were also developed to compute the Hodgkin-Huxley (HH) equations for squid giant axon [83] and to better understand the HH equation dynamics through projection to lower order models[56]. The 4 state variable biophysical HH equations were reduced to a two variable model based on whether each state variable elicited a fast or slow response over the duration of excitation. The two phenomenological variables then represent a non-dimensional weighted sum of the original set of 4 state variables. While this reduced the accuracy of exactly representing parts of the action potential, the equations provided a means of studying the excitation threshold behavior, and the oscillating and non-oscillating behavior of different cell types. Over time, the electrophysiological nerve models were modified and adjusted to represent cardiac electrophysiological cells, which have similar excitation threshold properties. The classical Aliev-Paniflov phenomenological model was created to more accurately represent ventricular cell action potentials in comparison to the FitzHugh-Nagumo model.

By looking at the phase plots for a particular set parameters for the FitzHugh-Nagumo (FHN) model (Figure 3.1), we can determine a priori what the action potential profile may look like. The two lines in light blue are the nullclines for the particular rate equations for the two variables. While the phase-plot does not indicate the speed at which a particular set of initial conditions moves along the traced trajectory, the nullclines indicate where the minimum and maximum transmembrane potential, ϕ , and recovery variable, r , occur based on the set of material parameters. Crossing a particular rate equation nullcline indicates a change in direction for that particular rate equations variable. For example, crossing the cubic nullcline for f^ϕ from left to right causes the trajectory of the coupled-p.d.e. to reverse sharply and generate the proper action potential shape. From the phase plot, one can also see that there is a stable limit-cycle, and therefore these material parameters generate stable oscillating action potentials as seen in Figure 3.1.

On the other hand, from observing the phase-plot for the Aliev-Paniflov (AP)

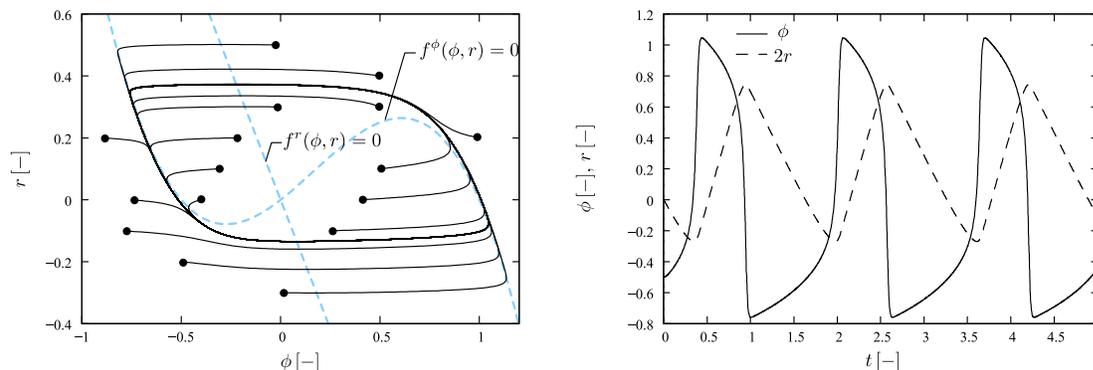


Figure 3.1: The generalized FitzHugh-Nagumo model with $\alpha = -0.5$, $a = 0$, $b = -0.6$, $c = 50$. The phase portrait depicts the trajectories for distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to a stable limiting cycle. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$ (left). Self-oscillatory time plot of the non-dimensional action potential ϕ and the recovery variable r (right).

model (Figure 3.2), it is clear that there is no limit-cycle; thus the Aliev-Panfilov model can be used to properly characterize ventricular cells which are not automatic oscillating cells. The f^ϕ nullcline dictates the peak transmembrane potential and also the necessary excitation potential required to illicit an action potential. The f^r nullclines then dictate the resting potential and the cell relaxation shape. There are several important considerations when determining the suitability of phenomenological models for different simulation scenarios. One of the greatest advantages of using phenomenological models is that by adjusting the nullclines, it is relatively intuitive to adjust the material parameters systematically to reproduce varying material parameter fields that correspond to experimentally measured metrics [?]. For example by inverting the slope of $f^r(\phi, r)$ in the FHN model, the phenomenological oscillating cell model becomes non-oscillating and can be used to represent atrial or ventricular cells (Figure 3.3). Also, the simpler two-variable nature of the source equations and the stable architected nature of the nullclines, means that phenomenological cell models are cheaper than ionic models, and are more likely to be robust and stable over a large set of initial conditions and material parameters. However,

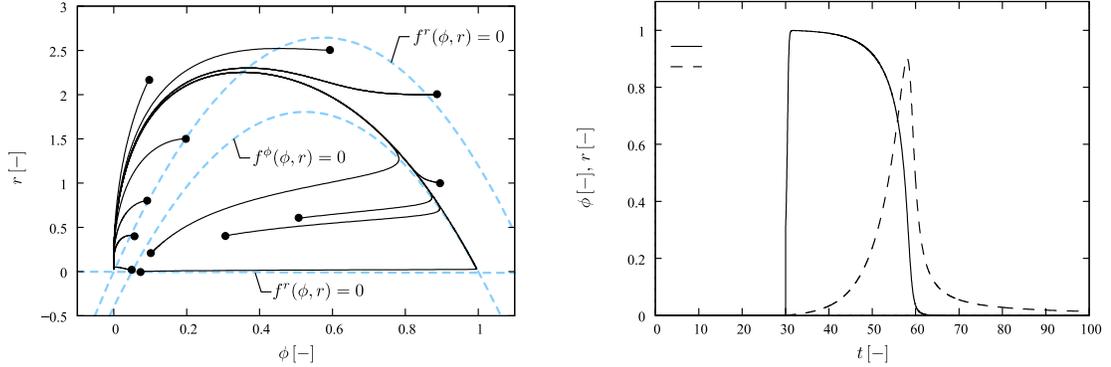


Figure 3.2: The Aliev-Panfilov model with $\alpha = 0.05$, $\gamma = 0.002$, $b = 0.15$, $c = 8$, $\mu_1 = 0.2$, $\mu_2 = 0.3$. The phase portrait trajectories for distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to stable equilibrium point. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$ (left). Non-oscillatory time plot of the non-dimensional action potential ϕ and the recovery variable r is triggered by external stimulation $I=30$ from the steady state $\phi_0=r_0=0$ (right).

the presence of phenomenological parameters and lack of biophysical parameters has some disadvantages. Simpler, lower material parameter models obviously cannot exactly fit experimentally measured data, especially when it is important to model the action potential accurately in certain applications. Naturally, phenomenological parameters can be more difficult to interpret physically, and it may be more difficult, if not impossible, to model physical behavior intuitively through phenomenological state variables and parameters. For example, drug treatments, or ionic concentration modifying scenarios would be more difficult to objectively model using the FHN and AP models introduced above. Likewise, it may be difficult to tune and justify phenomenological electromechanical models, since the electro-chemo-mechanical cascade is inherently affected by the transmembrane voltage and the chemical concentrations available within a given cell. Lastly, phenomenological models cannot model the more detailed and unnatural nor physiological behavior of stem-cell derived experiments, where foreign ion-channels may be introduced to stem cells to provide additional control for experimentalists [1].

Overall, when chemical concentrations and detailed physiological behavior is not

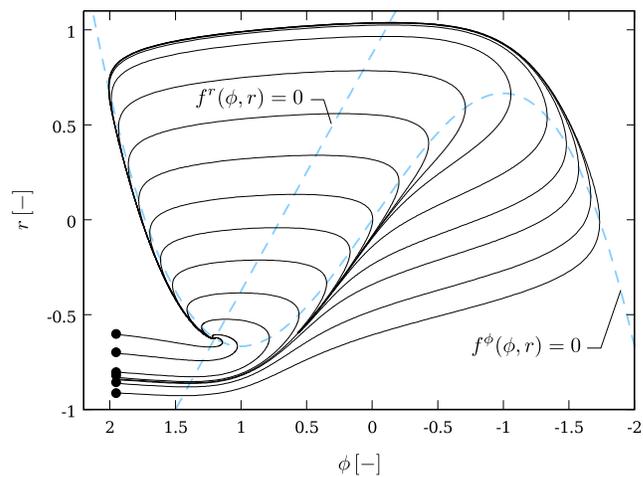


Figure 3.3: Phase portrait of classical FitzHugh-Nagumo model with $a=0.7$, $b=0.8$, $c=3$ (left). Trajectories for the distinct initial values of non-dimensional potential ϕ_0 and recovery variable r_0 (filled circles) converge to steady state. Dashed lines denote nullclines with $f^\phi = 0$ and $f^r = 0$. Circuit diagram of corresponding tunnel-diode nerve model (right).

of concern, FHN and AP-type models are computationally cost-effective, and can provide the necessary accurate electrophysiological behavior for most physiological simulation scenarios that do not interface or depend on extracellular potentials. This, however, may be required in shock treatments that include defibrillation.

3.1.2 Ionic Models

The set of classical Hodgkin and Huxley equations [83] is the first ionic cell model used to represent squid giant axon electrophysiological behavior. In their set of papers in 1952, they established the basis for representing gating variable dynamics with first order rate equations. These assumptions were based on finding the simplest differential equation for representing the time varying nature of the measured normalized current conductances (3.6). They also established the representation of modeling the cell membrane as a capacitor (3.7), representing current as the product of conductance and the difference between the transmembrane and nernst potential of a particular ionic species (3.8). While many of the assumptions made may have

seemed arbitrary at the time, many assumptions have turned out to be correct and correspond well with modern knowledge of these ion channels; i.e. , representing the potassium channel using a quartic polynomial coincides with current markov models of potassium having at least four necessary gates.

$$\dot{g}_{\text{gate}} = \frac{1}{\tau_{\text{gate}}(\phi)} [g_{\text{gate}}^{\infty}(\phi) - g_{\text{gate}}] \quad (3.6)$$

$$\dot{\phi} = f^{\phi}(\phi, g_{\text{gate}}, c_{\text{ion}}) = \frac{-I_{\text{tot}}}{C_{\text{m}}} \quad (3.7)$$

$$I_{\text{ch}} = C_{\text{ch}} g_{\text{gate}}(\phi - \phi_{\text{ion}}) \quad (3.8)$$

$$c_{\text{ion}} \dot{=} = \frac{\sum I_{\text{crt, ion}}}{V_{\text{space}} F} \quad (3.9)$$

Most modern ionic-models are based on the Hodgkin-Huxley equations of the form (3.6 - 3.9). The first equation is the characteristic form of gating variable rate equations. The second equation denotes the rate of change of transmembrane potential, $\dot{\phi}$, with respect to transmembrane currents, I_{tot} , and cell capacitance, C_{m} . The third equation is the characteristic form of ionic currents and denotes how currents are related to conductance and the effective voltage difference. Maximum ion channel conductances are denoted by C_{ch} . These can be extracted by correcting for effects from measured voltage-clamp experiments and also effects from the gating variable, g_{gate} . By constructing a standard I-V curve, one can determine the reversal potential for the particular current, ϕ_{ch} , and determine whether it correlates with the reversal potential for the ion current, ϕ_{ion} . The last equation defines the rate of change in ion concentrations and is important for preserving ion conservation between (3.7) and (3.9). V_{space} is the volume encompassing a particular ion species, and F , is the Faraday's constant. Together they are used to determine the change in ion concentration of a particular species.

The ten Tusscher ventricular model [188] and Dokos Sino-atrial [49] models are more recent ionic models that account for the behavior of several ion channels and gating variables, and both intracellular and extracellular ionic concentrations. Figure 2.2 is a typical diagram of an ionic model where individual ion currents, pumps, and exchangers are highlighted. While the general composition of ionic cell models is

similar, differences in models exist to account for differing action potential profiles and for the differing electrophysiological properties of cells within the heart. For ventricular cells, which are responsible for contraction of the ventricles, modelling the release and uptake of calcium ions is an important detail to include for coupling of electrophysiological models with cardiac mechanics. For pacemaker like cells, such as sino-atrial cells, the hyperpolarizing (funny) current, I_f and the L-type calcium current, I_{CaL} , are particularly important in controlling automatic pacemaker activity.

In many ways, the time constant functions for gating variables, τ_{gate} , and steady state functions, g_{gate}^∞ , are also phenomenological. The physics-based nature of ionic models is mainly due to enforcement that the change in voltage is directly proportional to the instantaneous net ionic current. At the same time the intracellular chemical concentrations change, thus, a physical ionic model should satisfy the condition of ion conservation [86], by ensuring that the total sum transmembrane current, I_{tot} , is equivalent to the change in chemical concentration, c_{ion} , within the cellular membrane. This property is particularly important for reducing drift in numerical calculations, and also when one wants to introduce new currents into existing formulations. New currents should be introduced in a way, such that ion conservation is conserved. An example of this will be shown in Chapter 7 and 8.

Lastly, an important characteristic that relates phenomenological and ionic models is their dependence of gating variables on the transmembrane potential. FitzHugh describes the recovery state variable as "the potential demand of the membrane" where the "transmembrane potential demands a particular state of the membrane." Ionic models, similarly, are generally formulated in a way such that the transmembrane potential directly affects the dynamics, and therefore the steady state of a given gating variable at a particular voltage is also affected. This is partially due to the fact that experimental data was obtained from voltage clamp experiments. However, the implications of the transmembrane voltage as an indirect "controller" of the non-linear coupled system has ramifications ranging from validity of boundary conditions in high-voltage or high-current scenarios, such as defibrillation, to the expected behavior of genetically altered stem cell experiments. This similar characteristic also serves as an interesting segue in the following section where a consistent encompassing

algorithmic framework is introduced.

For the purposes of this dissertation work, the ten Tusscher and Dokos models will be of main interest for several reasons. These models are sufficiently complicated such that benefits of the algorithmic scheme in the following section will be more evident. The two models are representative of their respective cell types: i.e, ventricular, and pacemaker. Lastly, these two models were chosen to better analyze the following algorithmic scheme because they differ in the number of ion concentrations, ion channels, and gating variables. The ten Tusscher model only characterizes intracellular currents, while the Dokos models includes both intracellular and extracellular currents, but models significantly less ionic currents.

3.1.3 Conduction

Currently there are several ways of modelling conduction within the cardiac electrical system. The first is to assume a continuum of cells such that a diffusion-like conduction tensor can be used to propagate electrical signals from one location to another spatially. This seems to be a decent approximation for ventricular cells. In regards to this, the conduction tensor, \mathbf{D} can be framed in terms of isotropic and anisotropic conduction to account for experimentally observed axial preferred conduction directions.

$$\mathbf{D} = \mathbf{D}_{\text{iso}} + \mathbf{D}_{\text{aniso}} = d^{\text{iso}} \mathbf{I} + d^{\text{ani}} \mathbf{f} \otimes \mathbf{f} \quad (3.10)$$

It is still under debate as to how the whole heart excites; however, studies have shown that heart excites very quickly through the His Bundles and Purkinje fiber network. From epicardial activation maps in literature [152], it seems that the heart first activates vertically along the septum, and then from the front around the heart. In addition, t-waves have shown to be important component in analyzing ECGs [58]. The action potential duration of cells is not uniform throughout the heart, and is in fact longer closer to the AV node, and significantly shorter farther away [58]. Following ECG modelling literature [102], we modify the action potential duration behavior of ionic the cells to attain proper heart contraction excitation dynamics by varying the

conductances of three different currents ($I_{\text{cal}}, I_{\text{ks}}, I_{\text{Kr}}$) to attain 5 different levels of action potential duration. It is assumed that changes up to 30% in conductance parameters are a reasonable estimate, and parametric analysis of the effects of the ten Tusscher currents [157] have been used to modify the action potential duration.

While we have chosen the method above, there are alternatives that have been developed to model different parts of conduction. Some groups have chosen to use time-dependent boundary conditions to model the excitation of the heart accurately. [71] This approach was not taken because it is tedious and because it is fixed for certain physiological conditions and for particular cell models. Other groups, have used 3-D line element networks to generate 3-D purkinje networks [189]. For simplicity, the fast bulk diffusion conduction parameters were chosen over a potentially more accurate purkinje network. However, this network has been implemented in generalized models of the ventricles [66, 102] and it should not be difficult to include them into the heart simulations performed in the following sections. Overall, the approach in modeling conduction serves as a simply first-principles representation of the conduction system in the heart and should demonstrate the robustness of the subsequently described algorithmic framework in modeling different physiologically relevant scenarios.

3.2 Mechanics

Due to the electromechanically-coupled nature of contraction, most mechanical models of the heart exhibit passive mechanical characteristics or active mechanical characteristics. Passive mechanical characteristics account for the interaction between vascular and interstitial fluid, and its interaction with the surround extracellular collagen matrix. While, there may be 5-10% volumetric change during passive mechanical tests, heart tissue can generally be modelled as a nearly incompressible material. Active mechanical models, on the other hand, are generally very phenomenological and are calibrated to experimentally measured results, usually with terms related to calcium and transmembrane potential field variables.

3.2.1 Passive Mechanical Behavior Modeling

The anisotropic passive mechanical properties of heart tissue are dependent on myocyte cardiac orientations. Cardiac cells are significantly stiffer along the fiber direction axis in comparison to sheet and sheet normal directions. Collagen within cardiomyocytes is generally thought to be the primary structural property that is responsible for increase in stiffness along the preferred axis [38, 110, 123]. Since collagen exhibits an exponential stress-strain behavior, strain energy density functions of exponential form are appropriate for defining the constitutive stiffness relations. While there are a variety of isotropic [59], transversely isotropic, and orthotropic models, this dissertation work is primarily concerned with using a particular orthotropic model due to its convexity properties and suitability in parameter identification from experimental bi-axial and shear experiments [84]. The strain energy density function is defined below

$$\Psi(I_1, I_{4f}, I_{4s}, I_{8fs}) = \frac{a}{2b} \exp[b(I_1 - 3)] + \sum_{i=f,s} \frac{a_i}{2b_i} \{\exp[b_i(I_{4i} - 1)^2] - 1\} + \frac{a_{fs}}{2b_{fs}} \{\exp[b_{fs}I_{8fs}^2] - 1\} \quad (3.11)$$

in terms of a set of eight material parameters $a, b, a_f, b_f, a_s, b_s, a_{fs}, b_{fs}$ and the invariants I_1, I_{4f}, I_{4s} , and I_{8fs} . The latter are defined through the trace operations in the Lagrangean setting

$$I_1 := \text{tr}(\mathbf{C}) = \mathbf{C} : \mathbf{G}^{-1}, \quad I_{4f} := \mathbf{C} : (\mathbf{f}_0 \otimes \mathbf{f}_0), \quad I_{4s} := \mathbf{C} : (\mathbf{s}_0 \otimes \mathbf{s}_0), \quad I_{8fs} := \mathbf{C} : \text{sym}(\mathbf{f}_0 \otimes \mathbf{s}_0) \quad (3.12)$$

between the right Cauchy-Green tensor \mathbf{C} and the inverse reference metric \mathbf{G}^{-1} , the structural tensors $\mathbf{M}_f := \mathbf{f}_0 \otimes \mathbf{f}_0$, $\mathbf{M}_s := \mathbf{s}_0 \otimes \mathbf{s}_0$, and $\mathbf{M}_{fs} := \text{sym}(\mathbf{f}_0 \otimes \mathbf{s}_0)$, respectively. While it is debated whether or not cardiac tissue is indeed incompressible or closer to a nearly-incompressible material [201], within the context of this dissertation we have chosen to use a nearly-incompressible formulation mainly for numerical stability considerations.

In our nearly incompressible hyperelastic formulation of the orthotropic model, a multiplicative decomposition of the volumetric and isochoric parts of the deformation

were taken, as is commonly done in incompressible material models[57, 169, 125].

$$\mathbf{F} = \bar{\mathbf{F}} \mathbf{F}_{\text{vol}} \quad \text{with} \quad \mathbf{F}_{\text{vol}} := J^{1/3} \mathbf{1} \quad \text{and} \quad \bar{\mathbf{F}} := J^{-1/3} \mathbf{F}, \quad (3.13)$$

implying that $J = \det(\mathbf{F}_{\text{vol}})$ and $\det(\bar{\mathbf{F}}) = 1$. The free energy of passive myocardium can then be additively decomposed into volumetric $U(J)$ and isochoric $\bar{\Psi}$ parts,

$$\Psi = U(J) + \bar{\Psi}(\bar{I}_1, \bar{I}_{4f}, \bar{I}_{4s}, \bar{I}_{8fs}). \quad (3.14)$$

It is the isochoric part of the free energy that is characterized constitutively by the following orthotropic model [65],

$$\bar{\Psi}(\bar{I}_1, \bar{I}_{4f}, \bar{I}_{4s}, \bar{I}_{8fs}) = \frac{a}{2b} \exp[b(\bar{I}_1 - 3)] + \sum_{i=f,s} \frac{a_i}{2b_i} \{\exp[b_i(\bar{I}_{4i} - 1)^2] - 1\} + \frac{a_{fs}}{2b_{fs}} \{\exp[b_{fs} \bar{I}_{8fs}^2] - 1\}. \quad (3.15)$$

While there are several different methods for dealing with numerical issues regarding nearly-incompressible materials, such as penalty, b-bar, and three-field Hu-Washizu mixed-formulations, penalty and three-field techniques were chosen. The details regarding these methods and the actual implementation of the aforementioned orthotropic model are divulged in the following referenced work [65] and in Chapter 10 .

3.2.2 Active Mechanical Behavior Modeling

Currently, most mechanical models of active-contraction in muscle tissue are either biophysical models at the subcellular sarcomeric level or phenomenological models at the macroscopic continuum level. At the subcellular level, several models for myofilament contractile dynamics have been developed and some even include temperature, sarcomeric pre-stretch, and calcium concentrations in determining the level of active stress generated by a filament [153]. On the other hand, phenomenological macroscopic models seem to be able to qualitatively predict a range of macroscopic

physiological parameters [66]. There is an inherent difficulty in obtaining enough data to quantify exactly how the continuum material behavior is dependent on sarcomeric contractile physics, due to the complex structure of cardiac cell orientations within tissue and also due to the gradual adaptation of biological muscle under different conditions. Therefore the general approach to active mechanical characterization in this dissertation will take a phenomenological approach, as the overall goal is to model the macroscopic physiological properties of the heart. However, phenomenological models may be useful in linking sarcomeric dynamics with continuum deformation behavior in cell culture experiments where many properties can be controlled, tested, and measured. Creating robust and stable electromechanical material models of cells will be integral for cell culture and in-vivo experimentalists in analyzing the strongly coupled electrical, chemical, and mechanical properties of cardiac tissue.

For the active Piola stress \mathbf{P}^{act} , we assume that an increase in the intracellular calcium concentration c_{Ca} above a critical level $c_{\text{Ca}}^{\text{crit}}$ induces an active cardiomyocyte contraction F^{act} [34, 69], which is acting along the fiber direction \mathbf{f}_0 [33, ?]. The contractile force F^{act} displays a twitch-type behavior [135], with a smooth off-on transition characterized through the twitch-function ϵ .

$$\begin{aligned}\mathbf{P}^{\text{act}} &= F^{\text{act}} \mathbf{f} \otimes \mathbf{f}_0 \\ \dot{F}^{\text{act}} &= \epsilon \left[\eta \left[c_{\text{Ca}} - c_{\text{Ca}}^{\text{rest}} \right] - F^{\text{act}} \right] \\ \epsilon &= \epsilon_0 + [\epsilon_\infty - \epsilon_0] \exp(-\exp(-\xi [c_{\text{Ca}} - c_{\text{Ca}}^{\text{crit}}]))\end{aligned}\tag{3.16}$$

Here, η controls the saturation of the active contractile force F^{act} , $c_{\text{Ca}}^{\text{rest}}$ is the resting concentration, ϵ_0 and ϵ_∞ are the minimum and maximum values of ϵ , $c_{\text{Ca}}^{\text{crit}}$ is the limit value above which contraction is initiated, and ξ is the transition rate from ϵ_0 to ϵ_∞ at $c_{\text{Ca}}^{\text{crit}}$. To evaluate the iteration matrices (10.19.3) and (10.19.4), we calculate the derivative of the Piola stress \mathbf{P} with respect to the transmembrane potential $\partial_\phi \mathbf{P}$ related to the active stress and with respect to the deformation gradient $\partial_{\mathbf{F}} \mathbf{P}$ related mainly to the passive stress, see Chapter 10 for details.

Chapter 4

Monolithic Algorithmic Framework

While the preceding constitutive models can be integrated in a myriad of ways into an electromechanical model, this chapter follows and expands upon one approach and tries to highlight the benefits and consequences in doing so. The integrative approach taken for this dissertation is to use a global-local variable split, where the macroscopic degrees of freedom are solved globally over the mesh, and the microscopic internal variables are solved locally at the integration point level. The algorithm also imposes conditions necessary for consistency at each iteration between local internal variables and global degrees of freedom for a given nonlinear iterative time-step.

4.1 General Algorithmic Framework

In context of cardiac mechanics, the electrical and mechanical field variables will be considered as global variables. These are the field variables that are almost always measured at the organ and tissue level. While chemical dynamics are often measured in a clinical context, in terms of the “physics” of cardiac mechanics, the chemical components indirectly affect the electrical and mechanical behavior through interactions at the cellular level, and not directly at the macroscopic organ level. However, certain heart conditions arise from differences in ion concentrations spatially in the heart, and to model these disease conditions, the flow of ions from one tissue area to

another may be integral in modeling the mechanics of the disease; thus ion concentrations would then be treated as global variables. Fortunately, for the physiological and diseased conditions covered in this work, it is a reasonable assumption that ion concentrations only have local cellular influence, since ion dynamics mainly occur in the space within the cell and across the cellular membrane. Thus, the phenomenological restitution behavior, active stress activation variables, ion channel gating variables, ion concentrations, and other “local” phenomena will be regarded as local variables.

The algorithmic finite element framework is briefly outlined below in Table 4.1. The framework follows the norm in finite elements. Global degrees of freedom and internal variables are initialized, and the corresponding boundary conditions are prescribed. Linear Lagrangian shape functions are used for the finite element discretization in space, and an implicit first-order backward-Euler finite-difference scheme is used in time. Reduced order-integration was used, as non-linear problems are expensive, and serve to reduce the cost each non-linear Newton-Raphson iteration.

At the first iterative update for global variables at a given time-step, t_{n+1}^0 , a nested local Newton iteration is performed to update the solution of the internal variables at the current iteration keeping the global variables constant, Φ_{n+1}^0 . As the global and local variables are strongly coupled to each other, global consistency is coupled to local consistency conditions via the partial derivative terms as explained in what follows.

In constructing the tangent matrix for the global residual equations, the dependency of the global residuals on local variables is evident in the partial derivatives of the directly coupled terms. For example,

$$R_i^\Phi = f_i^\Phi(\Phi, \mathbf{r}) \triangleq 0 \quad (4.1)$$

$$K_{ij}^\Phi = \partial_{\Phi_j} f_i^\Phi(\Phi, \mathbf{r}) + \partial_{r_k} f_i^\Phi(\Phi, \mathbf{r}) \partial_{\Phi_j} r_k \quad (4.2)$$

The final term in (4.2), $\partial_{\Phi_j} r_k$, can be obtained from utilizing the tangent of the

converged local newton iteration solution at the current global Newton-Raphson iteration.

$$R_i^r = f_i^r(\mathbf{r})|_{\Phi} \triangleq 0 \quad (4.3)$$

$$K_{ij}^{\Phi} = \partial_{r_k} f_i^r(\mathbf{r})|_{\Phi} \quad (4.4)$$

$$d_{\Phi_j} R_i^r = \partial_{\Phi_j} f_i^r(\Phi, \mathbf{r}) + \partial_{r_k} f_i^r(\Phi, \mathbf{r}) \partial_{\Phi_j} r_k \doteq 0 \quad (4.5)$$

$$\partial_{\Phi} \mathbf{r} = -[\partial_{\mathbf{r}} f_i^r(\Phi, \mathbf{r})]^{-1} \partial_{\Phi} f_i^r(\Phi, \mathbf{r}) \quad (4.6)$$

Using this consistent and complete tangent, a global Newton-Raphson iteration is then performed and updated until the convergence criterion has been achieved. It should be noted that in practice, the bandwidth of the local tangent matrices is low and actually banded in the case ionic electrophysiological problems. Currently, coupling between global and local variables is more dominant between the respective electrophysiological and mechanical variables, but cross-electromechanical global-local coupling does and can exist for electromechanical coupled problems weakly through the intracellular calcium concentration and the active contractile force variable. In the following sections, algorithmic details for integrating phenomenological and ionic electrophysiological models into this framework and also coupled electromechanical models will be further discussed. Also, a cursory view of Table 4.1, may lead to the conclusion that the scheme maybe prohibitively expensive at the local iteration point. While this may be a valid assumption generally, the following subsections will attempt to show that for electromechanical models of the heart, the expense is bounded and computation is tractable.

4.2 Phenomenological Electrophysiological formulation

If we examine the case of simple two-variable phenomenological models [?], the global variable is the transmembrane potential, ϕ , and the local internal variable is the restitution variable, r , according to our algorithmic split. Therefore, Table 4.1, can

initialize nodal degrees of freedom (global) $\Phi_J = [\varphi, \phi]_J$
initialize internal variables $\mathbf{r} = [r_1, \dots, r_i]$
global Newton iteration
loop over all elements
loop over all integration points
local Newton iteration
calculate internal variable residuals R_i^r and local iteration matrix $[K_{i,i}^r] = d_{r_i} R_i^r$
update internal variables $\mathbf{r} \leftarrow \mathbf{r}(\Phi, \mathbf{r})$
calculate source term $f^\Phi(\mathbf{r})$ and its linearization $d_\Phi f^\Phi$ (4.6)
calculate element residuals $R_I^{\Phi e}$ and element matrices $K_{IJ}^{\Phi e} = d_{\Phi_j} R_I^{\Phi e}$
calculate global residual R_I^Φ and global iteration matrix $K_{IJ}^\Phi = d_{\Phi_j} R_I^\Phi$
update global variables $\Phi_J \leftarrow \Phi_J - K_{IJ}^{\Phi-1} R_I^\Phi$

Table 4.1: General algorithmic treatment of global-local variable split for cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. Global variables, Φ , are solved at the nodal points of the mesh, while internal variables, r_i , are solved at the integration point level.

be significantly simplified to Table 4.2 . Since the internal variable is simply a scalar quantity at each integration point, the non-linear newton-raphson problem is simply an algebraic update for each newton iteration. No matrix inversion is necessary, and so the general cost of the local Newton problem is fairly cheap.

4.3 Ionic Electrophysiological formulation

Next, we examine the feasibility of implementing more complicated biophysical ionic models within this framework and note possible issues that may arise due to the added complexity. With ionic models, the ionic currents, ionic gates, and ionic concentration dynamics are modeled and are all coupled both directly and indirectly to the

initialize nodal degree of freedom (global) ϕ_J initialize internal variable r
global Newton iteration
loop over all elements
loop over all integration points
local Newton iteration
calculate internal variable residuals R^r and local derivative $d_r R^r$ update internal variables $r \leftarrow r(\phi, r)$
calculate source term $f^\Phi(r)$ and its linearization $d_\phi f^\Phi$ (4.6)
calculate element residuals $R_I^{\Phi e}$ and element matrices $K_{IJ}^{\Phi e} = d_{\phi_J} R_I^{\Phi e}$
calculate global residual R_I^Φ and global iteration matrix $K_{IJ}^\Phi = d_{\phi_J} R_I^\Phi$
update global variables $\phi_J \leftarrow \phi_J - K_{IJ}^{\Phi-1} R_I^\Phi$

Table 4.2: Algorithmic treatment of global-local variable split for the case of phenomenological electrophysiological finite element models. The transmembrane variable, ϕ_J , is solved at the nodal points of the mesh, while internal variable, r_i , are solved at the integration point level.

transmembrane potential and possibly other ionic concentrations. Therefore, the internal Newton-Raphson iteration may be very complex. Fortunately, this is not the case for most models due to the way nature has developed its ionic machinery, as we will see.

For electrophysiological ionic models, the algorithmic strategy can be expanded to Table 4.3, where the global variable is again the transmembrane potential, ϕ_J , and the internal variables are now the gating variables, g_{gate} , and the ion concentrations, c_{ion} . While the general strategy is the same as Table 4.1, extra detail is spent in providing a more optimized algorithm for solving the local Newton-Raphson problem for ionic models. It turns out that there is much that can be utilized to more efficiently integrate ionic models into the proposed framework.

To gain understanding of how complex ionic models work and the numerical issues that they pose, the Ten Tusscher Model [188] will be treated and examined as a typical complex ionic model. An initial concern that rises is the issue of calculating

and inverting the local tangent matrix for a complex ionic model that has many gating variables, g_{gate} , and the ion concentrations, c_{ion} . Therefore, the dependency of the gating variables on ionic concentrations and the transmembrane potential will be studied, and also the inter-dependency of ionic concentration dynamics on each-other.

After a thorough examination of gating variable dependence (Table 4.4), it is clear that while most gating variable differential equations are only dependent upon their own state and the transmembrane potential, ϕ , there several gating variables that are also dependent upon ionic concentrations. The other observation is that the differential equations for the gating variables are linear, if external (non-gating variable) dependencies are considered as constants or parameters at a given time-step or local Newton iteration. This means that at a given time-step, the gating variables only dependent upon ϕ can be updated with one Newton iteration. The other gating variables that are dependent upon ion concentrations, which are then inter-dependent and coupled to other ionic concentrations, however must be updated iteratively. However each update is also a linear differential equation. Thus, we denote gating variables that are independent of chemical concentrations as $g_{\text{gate}}^{\text{I}}$, and those that are dependent upon chemical concentrations or essentially those that must be iteratively updated as $g_{\text{gate}}^{\text{II}}$. This observation will be exploited to more efficiently solve the local Newton-Raphson problem in a subsequent section.

Next, the coupled inter-dependency of ionic concentrations is examined. Dependency of ionic currents to chemical concentrations is shown in Table 4.5, which can be reorganized into Table 4.6. An immediate observation in Table 4.6 is that each ionic residual is at most dependent on three ionic concentration species, and most of the ionic residuals are actually dependent on only two ionic species! In fact, a survey of available ionic models seems to show that ionic currents are generally only directly dependent upon at most two ionic species, and ion concentration dynamics are only dependent on at most three ionic species. However, in certain cases where more ionic species are introduced, i.e. when extracellular ionic concentrations are also modelled, there is a higher degree of inter-dependence among different ionic concentrations for ionic residuals. Almost all models of ionic currents, regardless of the number of ionic species, are only dependent upon two ionic concentrations with possibly multiplicity

of species are also differentiated based on extracellular and intracellular spaces. This is an interesting observation, as it means that in nature, ion channel behavior is dependent upon a relatively low-number, 2, of the available ionic species. Most importantly, it means that the resulting tangent matrix is, in fact, banded and not dense. This has implications then on the added cost and complexity of adding more ionic species dynamics to a model, as solving a tangent matrix problem scales linearly instead of quadratically in the dense case.

For our example ionic model, since R^{Ca_i} is dependent on 3 concentrations and R^{Na} , R^K , $R^{Ca_{sr}}$, one possible and elegant way of constructing the tangent matrix is in the following way.

$$\mathbf{R} = [R^K, R^{Na}, R^{Ca_i}, R^{Ca_{sr}}]^T$$

$$\partial_{\mathbf{c}_{\text{ion}}} \mathbf{R} = \begin{bmatrix} \partial_{K_i} R^K & \partial_{Na_i} R^K & & & \\ 0 & \partial_{Na_i} R^{Na} & \partial_{Ca_i} R^{Na} & & \\ & \partial_{Na_i} R^{Ca_i} & \partial_{Ca_i} R^{Ca_i} & \partial_{Ca_{sr}} R^{Ca_i} & \\ & & \partial_{Ca_i} R^{Ca_{sr}} & \partial_{Ca_{sr}} R^{Ca_{sr}} & \end{bmatrix} \quad (4.7)$$

And we can see that the tangent is indeed a banded tri-diagonal matrix and only 9 of 16 entries in the local tangent are non-zero.

Thus to take advantage of these observations, the local Newton-Raphson problem is solved in the following manner. The gating variables, g_{gate}^I , are solved and updated immediately for the local Newton-Raphson problem. Next, the chemical dependent gating variables, g_{gate}^{II} , are then updated initially, such that initial iterations for the ionic residuals can be obtained. The ionic residuals and concentration-dependent gating variables are iteratively updated and solved, where the ionic residuals require a banded tri-diagonal solver, and g_{gate}^{II} can utilize simple algebraic linear updates. After the ionic residuals converge, the local Newton-Raphson problem has been solved, and coupled terms can be computed and passed onto the global element assembly routine.

4.4 Coupled electromechanical formulation

Fortunately for coupled electromechanical formulations, one can directly apply the general, developed algorithmic implementation, and factor in the phenomenological or ionic electrophysiology algorithmic details. Electromechanical phenomenological formulations [72] can be directly expanded to ionic implementations by making the same modifications in the preceding section. As this algorithmic split is based on those from internal variable formulations of mechanics of plasticity[168], the entire algorithmic framework integrates very naturally with existing finite element mechanics models.

4.5 Algorithmic Splitting Implications

While the algorithmic splitting proposed reduces global degrees of freedom and thus reduces the burden on the finite element solver, other algorithmic schemes, which may possibly be computationally cheaper exist [145]. However, there are several consequences and trade offs in these global-local splitting schemes.

For example, one particular issue with our global-local splitting framework is that the local variables scale according to the number of elements, while the global degrees of freedom scale according to the number of nodes in the mesh. Depending on mesh characteristics, successive refinements of the same mesh may skew the ratio of nodes to elements and cause a shift in the computational cost between the solver and the assembly routines. However, the state variable interpolation method (SVI) [145], makes some unnecessary assumptions in their analysis such that many of their conclusions do not apply to our particular scheme. In particular, authors of the SVI scheme claim that the calculation of residual quantities requires a computationally expensive projection from integration points to the nodes. This is not necessary and does not apply to our scheme. Their implementation stores the values of the state variable at the nodes, but calculates the local Newton problem at the integration point, and then subsequently projects the results back to the nodes. While, this invariably saves storage space, it results in computational cost that outweighs the

benefits of saving storage space. Another flaw with their splitting algorithm is the inherent inability to handle situations where adjacent finite elements have varying mechanical properties. This is due to the projective nature of their SVI method. Lastly, using a nodal representation for the local variables has other computational consequences that add to the difficulty in developing assembly routines that run in parallel on clusters.

Conversely, our algorithmic splitting scheme uses more memory, but is parallelizable, avoids expensive projection computations, and can trivially handle situations where adjacent elements have varying material parameters. This is a consequence of the fact that finite elements are formulated such that calculation of element quantities during assembly is trivially parallelizable; this is not the case for the solution at nodes which is found by solving the global matrix problem. Using an integration point local splitting effectively isolates material properties of other elements away from each individual element. This can then be applied to parallelization of cardiac codes, and will result in relatively easy ports to the GPU in Chapter ???. However, there may be ways of further improving our algorithmic scheme to incorporate the space saving nature of SVI.

4.5.1 Time Stepping Methods

Lastly, there is particular benefit in using an implicit time integration scheme throughout our algorithmic framework. This benefit is that adaptive-time stepping can be used stably to further reduce computation time. FEAP uses a particular simple logarithmic adaptive time stepping scheme. This scheme has been used in conjunction with our framework to reduce the computation time on the CPU and the GPU as will be demonstrated in the following sections.

Table 4.3: Algorithmic treatment of electrochemical coupling in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the chemical unknowns, the two sets of gating variables g_{gate}^I and g_{gate}^{II} and the ion concentrations c_{ion} are introduced locally on the integration point level.

initialize nodal degrees of freedom ϕ_J initialize internal variables $g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}}$
global Newton iteration
loop over all elements
loop over all integration points
update first set of gating variables $g_{\text{gate}}^I \leftarrow g_{\text{gate}}^I + f_{\text{gate}}^I \Delta t$
initialize second set of gating variables $g_{\text{gate}}^{II} \leftarrow g_{\text{gate}}^{II} + f_{\text{gate}}^{II} \Delta t$ initialize ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}})$
local Newton iteration
calculate ion concentration residuals $\mathbf{R}_{\text{ion}}^c$ and local iteration matrix $[\mathbf{K}_{\text{ion ion}}^c] = \mathbf{d}_{c_{\text{ion}}} \mathbf{R}_{\text{ion}}^c$ update ion concentrations $c_{\text{ion}} \leftarrow c_{\text{ion}} - [\mathbf{K}_{\text{ion ion}}^c]^{-1} \mathbf{R}_{\text{ion}}^c$
update second set of gating variables $g_{\text{gate}}^{II} \leftarrow g_{\text{gate}}^{II} + f_{\text{gate}}^{II} \Delta t$ update ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}})$ calculate source term $f^\phi(I_{\text{crt}})$ and its linearization $\mathbf{d}_\phi f^\phi$
calculate element residuals $\mathbf{R}_I^{\phi e}$ and element matrices $\mathbf{K}_{IJ}^{\phi e} = \mathbf{d}_{\phi_J} \mathbf{R}_I^{\phi e}$
calculate global residual \mathbf{R}_I^ϕ and global iteration matrix $\mathbf{K}_{IJ}^\phi = \mathbf{d}_{\phi_J} \mathbf{R}_I^\phi$
update membrane potential $\phi_J \leftarrow \phi_J - \mathbf{K}_{IJ}^{\phi-1} \mathbf{R}_I^\phi$

Table 4.4: Table of the dependence of each gating variable on the transmembrane potential, ϕ , and the ionic concentrations, c_{ion} .

	ϕ	c_{Na_i}	c_{Ca_i}	$c_{Ca_{sr}}$	c_{K_i}
m	x				
h	x				
j	x				
d	x				
f	x				
f_{Ca}			x		
r	x				
s	x				
x_s	x				
x_{r1}	x				
x_{r2}	x				
$x_{K1\infty}$ (no τ)	x				x
g			x		

Table 4.5: Dependency table for each ionic current on each ionic concentration species.

	c_{Na_i}	c_{Ca_i}	$c_{Ca_{sr}}$	c_{K_i}
I_{bCa}		X		
I_{bNa}	X			
I_{CaL}		X		
I_{K1}				X
I_{Kr}				X
I_{Ks}	X			X
I_{leak}		X	X	
I_{Na}	X			
I_{NaCa}	X	X		
I_{NaK}	X			
I_{pCa}		X		
I_{pK}				X
I_{rel}		X	X	
I_{to}				X
I_{up}		X		

Table 4.6: Dependency table of each ionic concentration residual R^i , on ionic concentration species.

		Na_i	Ca_i	Ca_{sr}	K_i	ϕ
R^{Na}	I_{bNa}	X				X
	I_{Na}	X				X
	I_{NaCa}	X	X			X
	I_{NaK}	X				X
R^{Ca_i}	I_{bCa}		X			X
	I_{CaL}		X			X
	I_{leak}		X	X		
	I_{NaCa}	X	X			X
	I_{pCa}		X			
	I_{rel}		X	X		X
	I_{up}		X			
$R^{Ca_{sr}}$	I_{leak}		X	X		
	I_{rel}		X	X		X
	I_{up}		X			
R^K	I_{K1}				X	X
	I_{Kr}				X	X
	I_{Ks}	X			X	X
	I_{NaK}	X				X
	I_{pK}				X	X
	I_{to}				X	X

Chapter 5

Applications in modeling ventricular tissue

We propose a novel, efficient finite element solution technique to simulate the electrochemical response of excitable cardiac tissue. We apply a global-local split in which the membrane potential of the electrical problem is introduced globally as a nodal degree of freedom, while the state variables of the chemical problem are treated locally as internal variables on the integration point level. This particular discretization is efficient and highly modular since different cardiac cell models can be incorporated in a straightforward way through only minor local modifications on the constitutive level. Here, we derive the underlying algorithmic framework for a recently proposed ionic model for human ventricular cardiomyocytes, and demonstrate its integration into an existing nonlinear finite element infrastructure. To ensure unconditional algorithmic stability, we apply an implicit backward Euler scheme to discretize the evolution equations for both the electrical potential and the chemical state variables in time. To increase robustness and guarantee optimal quadratic convergence, we suggest an incremental iterative Newton-Raphson scheme and illustrate the consistent linearization of the weak form of the excitation problem. This particular solution strategy allows us to apply an adaptive time stepping scheme, which automatically generates small time steps during the rapid upstroke, and large time steps during the plateau, the repolarization, and the resting phases. We demonstrate that solving an entire

cardiac cycle for a real patient-specific geometry characterized through a transmembrane potential, four ion concentrations, thirteen gating variables, and fifteen ionic currents requires computation times of less than ten minutes on a standard desktop computer.

5.1 Introduction

Despite intense research over the past decades, cardiovascular disease remains the single most common cause of natural death in developed nations [113, ?]. Sudden cardiac death is estimated to account for approximately half of all these deaths, claiming approximately a thousand lives each day in the United States alone [28]. The high incidence and sudden, unexpected nature of sudden cardiac death, combined with the low success rate of resuscitation, make it a major unsolved problem in clinical cardiology, emergency medicine, and public health [?, ?]. This manuscript is motivated by the vision to create a multi-scale patient-specific computational model of rhythm disorders in the heart to improve our understanding of the basic pathology associated with sudden cardiac death.

Since the famous experiments by Luigi Galvani [?] who impressively demonstrated the electrically stimulated contraction of excised frog leg muscle more than two centuries ago, we have been trying to better understand the electrochemistry of living tissue. We now know that the electrophysiological activity of excitable cells is governed by a delicate balance between electrical and chemical gradients across the cell membrane [18]. These gradients are maintained by means of the membrane's selective permeability with respect to different ions at different points throughout an excitation cycle [?, ?, 143]. In cardiac cells, at rest, the transmembrane potential is approximately -86 mV , meaning the cell's interior is negatively charged with respect to its exterior. Cardiac cells can be excited by an electrical stimulus that generates an initial depolarization across the cell membrane. Once this stimulus exceeds a certain threshold, the transmembrane potential increases rapidly from its resting state of approximately -86 mV to its excited state of $+20\text{ mV}$. After a brief period of partial initial repolarization, we can observe a characteristic plateau of about a fifth of a second before the

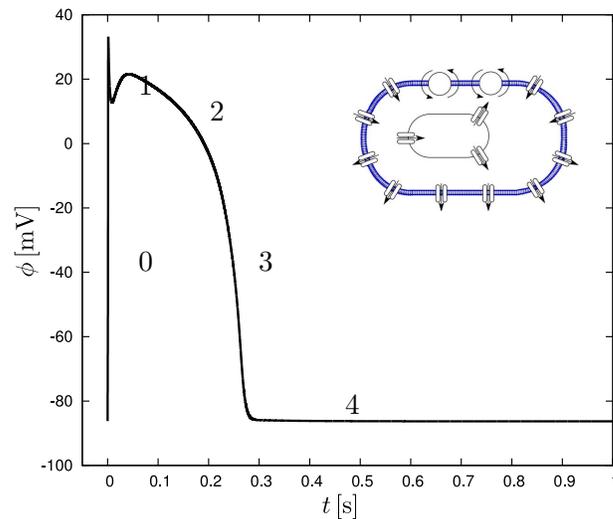


Figure 5.1: Electrochemistry in a human ventricular cardiomyocyte. Time dependent evolution of transmembrane potential ϕ . The characteristic action potential consists of five phases. Phase 0: The rapid upstroke is generated through an influx of positively charged sodium ions through fast sodium channels. Phase 1: Early, partial repolarization is initiated through the efflux of positively charged potassium ions through transient outward channels. Phase 2: During the plateau, the net influx of positively charged calcium ions through L-type sodium channels is balanced by the efflux of positively charged potassium ions through inward rectifier channels, rapid and slow delayed rectifier channels, and transient outward channels. Phase 3: Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. Phase 4: Throughout the interval between end of repolarization and the beginning of the next cycle the cell is at rest.

cell gradually repolarizes to return to its original resting state [147], as illustrated in Figure 10.1.

This characteristic temporal evolution of the transmembrane potential is brought about by the interaction of different ion channels controlling the inward and outward flux of charged sodium, potassium, and calcium ions during the different phases of this excitation cycle. The first model to quantitatively describe the electrophysiological activity of excitable cells was proposed by Hodgkin and Huxley [83] who were awarded the Nobel Prize in Physiology and Medicine for their seminal work on

action potentials in neurons half a century ago. In fact, most currently available cardiac cell models are derived from the classical Hodgkin-Huxley model. A significant conceptual simplification, the celebrated phenomenological two-parameter FitzHugh-Nagumo model [56, 134], was proposed in the early 1960s to allow for a fundamental mathematical analysis of the coupling phenomena between electrical and chemical fields. In the 1970s, two sophisticated new mathematical models for the electrical activity of cardiac cells were introduced, one for cardiac Purkinje fibers [140, 122] and one for mammalian ventricular cardiomyocytes [12]. The latter was calibrated by means of data from voltage-clamp experiments available at that time. Subsequent developments in single-cell and single-channel recording techniques enabled a more accurate control of intracellular and extracellular environments starting in the mid 1980s. These novel experimental techniques paved the way for a rigorous significant refinement of the earlier models for Purkinje fibers [48] and for mammalian ventricular cardiomyocytes [116]. The latter, the celebrated Luo-Rudy model [?, 130, 151], was originally calibrated for guinea pig ventricular cells [117], but soon thereafter adjusted to model human ventricular cardiomyocytes [17, 43, 149], and modified to incorporate intracellular calcium dynamics [68]. Here we will follow its most recent refinement, the ten Tusscher model [188, ?] illustrated in Figure 10.2, which we believe is extremely brilliant and powerful, however, unfortunately computationally demanding in its present explicit finite difference based form. Characterized through four ion concentrations, fifteen ionic currents, and thirteen gating variables, this model captures the essential characteristics of human ventricular cardiomyocytes: It contains the major ionic currents, includes basic intracellular calcium dynamics, and is well-calibrated against experimental data [188, ?, ?]. The goal of this manuscript is therefore to develop an unconditionally stable, efficient, modular, flexible, and easily expandable algorithm for human ventricular cardiomyocytes motivated by the original ten Tusscher model and make it available for efficient whole heart simulations using common, existing finite element infrastructures.

Simulating the electrical activity of the heart is by no means new, and many established research groups have successfully contributed to solving this challenging task

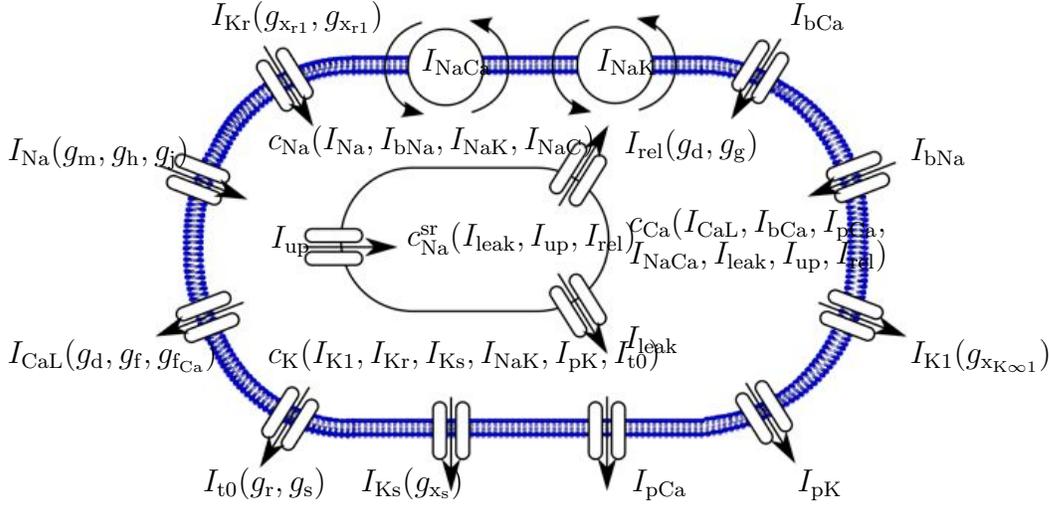


Figure 5.2: Ionic model of a human ventricular cardiomyocyte [12, 68, 116, 188]. In this model, the electrochemical state of the cardiomyocyte is characterized in terms of $n_{\text{ion}} = 4$ ion concentrations, the free intracellular sodium, potassium, and calcium concentrations and the free calcium concentration in the sarcoplasmic reticulum, $c_{\text{ion}} = [c_{\text{Na}}, c_{\text{K}}, c_{\text{Ca}}, c_{\text{Ca}}^{\text{sr}}]$. Ion concentrations are controlled through $n_{\text{crt}} = 15$ ionic currents, $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]$. Their channels are governed by $n_{\text{gate}} = 13$ gating variables $g_{\text{gate}} = [g_{\text{m}}, g_{\text{h}}, g_{\text{j}}, g_{\text{xK1}\infty}, g_{\text{xr1}}, g_{\text{xr2}}, g_{\text{xs}}, g_{\text{r}}, g_{\text{s}}, g_{\text{d}}, g_{\text{f}}, g_{\text{fCa}}, g_{\text{g}}]$ which are functions of the current membrane potential ϕ .

[146, 155, 167, 189]. However, most cardiac excitation models are based on simplifying assumptions to capture the chemical activity on a phenomenological level, similar to the original FitzHugh-Nagumo model [154], as illustrated in various excellent overviews and monographs [40, 96, ?, ?]. As one of the most efficient approaches, the distinguished two-parameter Aliev-Panfilov model [4] seeks to reproduce the major fundamental characteristics of the action potential at minimal computational cost. We have successfully implemented this model in a fully implicit nonlinear finite element framework in the past [72, 73], applied it on patient-specific geometries to extract electrocardiograms [102], applied it in the context of bidomain formulations [44], and coupled it to mechanical contraction in a monolithic whole heart simulation [66]. In this manuscript, rather than using a phenomenological model, we seek to

investigate the potential of ionic models in the context of our previously proposed generic finite element framework [72], embedded in the multipurpose nonlinear finite element program FEAP [?] and its recent parallel version [?]. Within this generic framework, all chemical state variables, in our case the four ion concentrations and the thirteen gating variables, are introduced locally as internal variables on the integration point level.

This manuscript is organized as follows: Section 5.2 briefly summarizes the governing equations of the electrical excitation problem and the chemical ion concentration problem. Section 5.3 then illustrates the computational solution algorithm based on a global node point based solution of the electrical excitation problem combined with a local integration point based solution of the chemical concentration problem. In Section 5.4 we specify the constitutive equations for the particular model problem of a human ventricular cardiomyocyte characterized through thirteen gating variables, fifteen ionic currents, and four ion concentrations which we integrate into the discrete framework in Section 5.5. Section 10.5 documents the features of the proposed algorithm in the context of a single human ventricular cardiomyocyte in Section 5.6.1 and in terms of a real human heart geometry in Section 5.6.2. We close with a final discussion and an outlook with future directions in Section 10.6. Mathematical details about the algorithmic formulation are provided in the Appendix.

5.2 Continuous problem of electrochemistry

In this section, we summarize the generic equations of electrochemical coupling in cardiac tissue characterized through a partial differential equation for the electrical problem and through a system of ordinary differential equations for the chemical problem [?, 131, 151, ?]. We then specify the generic set of equations to represent a particular ionic model of a human ventricular cardiomyocyte [12, 68, 116, 188]. The primary unknown of the electrical problem is the membrane potential ϕ , the unknowns of the chemical problem are the state variables, i.e., the n_{gate} gating variables g_{gate} and the n_{ion} ion concentrations c_{ion} .

5.2.1 Electrical problem - Partial differential equation

The excitation problem is characterized through the spatio-temporal evolution of the membrane potential ϕ in terms of the flux term $\text{div} \mathbf{q}$ and the source term f^ϕ .

$$\dot{\phi} = \text{div} \mathbf{q}(\phi) + f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (5.1)$$

It has become common practice to enhance the initially local equation for cellular excitation by a phenomenological membrane potential flux $\text{div} \mathbf{q}$ with

$$\mathbf{q} = \mathbf{D} \cdot \nabla \phi \quad (5.2)$$

to account for the nonlocal nature of propagating excitation waves. Membrane potential propagation is characterized through the second order diffusion tensor $\mathbf{D} = d_{\text{iso}} \mathbf{I} + d_{\text{ani}} \mathbf{n} \otimes \mathbf{n}$ related to the gap junctions between the cells. The diffusion tensor can account for both isotropic propagation d_{iso} and anisotropic propagation d_{ani} along preferred directions \mathbf{n} . The source term

$$f^\phi = -\sum_{\text{crt}=1}^{n_{\text{crt}}} I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (5.3)$$

is basically directly related to the negative sum of the n_{crt} ionic currents I_{crt} across the cell membrane. Chemoelectrical coupling is introduced through these ionic currents which are parameterized in terms of the gating variables g_{gate} and ion concentrations c_{ion} . The evolution of these chemical state variables will be characterized in detail in the following subsection.

5.2.2 Chemical problem - System of ordinary differential equations

From a mathematical point of view, the chemical problem is defined in terms of two sets of ordinary differential equations, one for the n_{gate} gating variables g_{gate} and one for the n_{ion} ion concentrations c_{ion} [?, 130, 151, ?]. The gating variables essentially characterize the states of the individual ion channels, which can be either open or

closed. It proves convenient to divide the gating variables into two subsets, a first set g_{gate}^I which depends only on the current membrane potential ϕ , and a second set g_{gate}^{II} which depends on both the membrane potential ϕ and the corresponding ion concentration c_{ion} . The gating variables are defined through the following set of ordinary differential equations.

$$\dot{g}_{\text{gate}}^I = f_{\text{gate}}^{gI}(\phi, g_{\text{gate}}^I) = \frac{1}{\tau_{\text{gate}}^I(\phi)} [g_{\text{gate}}^{\infty I}(\phi) - g_{\text{gate}}^I] \quad (5.4)$$

$$\dot{g}_{\text{gate}}^{II} = f_{\text{gate}}^{gII}(\phi, g_{\text{gate}}^{II}, c_{\text{ion}}) = \frac{1}{\tau_{\text{gate}}^{II}(\phi)} [g_{\text{gate}}^{\infty II}(\phi, c_{\text{ion}}) - g_{\text{gate}}^{II}] \quad (5.5)$$

Their evolution is governed by classical Hodgkin-Huxley type equations, each characterized through a steady-state value g_{gate}^{∞} and a time constant τ_{gate} for reaching this steady state, where both are usually exponential functions of the membrane potential ϕ . In addition, the steady state values of the second set $g_{\text{gate}}^{\infty II}$ are also functions of the ion concentration c_{ion} . The relevant ion concentrations in cardiac cells are typically the sodium concentration c_{Na} , the potassium concentration c_{K} , the calcium concentration c_{Ca} , and, in our case, the calcium concentration in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{SR}}$. Collectively, these ion concentrations c_{ion} are defined through a second set of ordinary differential equations.

$$\dot{c}_{\text{ion}} = f_{\text{ion}}^c(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (5.6)$$

Their evolution is driven by the individual righthand sides f_{ion}^c , which represent nothing but the weighted sums of the corresponding individual transmembrane currents I_{crt} . These n_{crt} ionic currents I_{crt}

$$I_{\text{crt}} = I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (5.7)$$

can be expressed in terms of the current potential, the set of gating variables, and the set of ion concentrations. Electrochemical coupling is thus introduced through the

voltage-gated nature of the relevant ion channels which reflects itself in the potential-dependency of the chemical state variables g_{gate} and c_{ion} . The particular cell model illustrated in Figure 10.2, which we will explain in detail in Section 5.4, is characterized in terms of $n_{\text{ion}} = 4$ ion concentrations $c_{\text{ion}} = [c_{\text{Na}}, c_{\text{K}}, c_{\text{Ca}}, c_{\text{Ca}}^{\text{sr}}]$, $n_{\text{crt}} = 15$ ionic currents $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]$, and $n_{\text{gate}} = 13$ gating variables $g_{\text{gate}}^I = [g_{\text{m}}, g_{\text{h}}, g_{\text{j}}, g_{\text{xr1}}, g_{\text{xr2}}, g_{\text{xs}}, g_{\text{r}}, g_{\text{s}}, g_{\text{d}}, g_{\text{f}},]$ and $g_{\text{gate}}^{II} = [g_{\text{xK1}\infty}, g_{\text{fCa}}, g_{\text{g}}]$. It is obvious that the complex, nonlinear coupled system of equations for the membrane potential, the gating variables, and the ion concentrations cannot be solved analytically. In the following section, we will illustrate the discrete problem of electrochemical coupling introducing a consistently linearized fully implicit finite element solution scheme based on a global-local split.

5.3 Discrete problem of electrochemistry

We suggest discretizing the spatio-temporal problem of electrochemical coupling (8.15), (10.24) and (5.6) for the transmembrane potential ϕ , the gating variables g_{gate} , and the intracellular ion concentrations c_{ion} with a finite difference scheme in time and with a finite element scheme in space. Due to the global nature of the membrane potential introduced through the diffusion term $\text{div } \mathbf{q}(\phi)$, we propose a \mathcal{C}^0 -continuous finite element interpolation for the membrane potential ϕ , while a \mathcal{C}^{-1} -continuous interpolation is sufficient for the sets of gating variables g_{gate} and ion concentrations c_{ion} . Accordingly, we introduce the membrane potential as global degree of freedom at each finite element node, whereas the gating variables and ion concentrations are introduced locally on the integration point level. The resulting staggered system is solved with an incremental iterative Newton-Raphson solution procedure based on the consistent linearization of the discrete excitation problem [72, 66, 73, 102]. The use of a fully monolithic implicit solution algorithm allows us to apply an adaptive time stepping procedure, for which the time step size is automatically adjusted in response to the number of Newton iterations towards global equilibrium [?].

5.3.1 Electrical problem - Global discretization on the node point level

Let us first transform the electrical problem (8.15) into its residual format

$$\mathbf{R}^\phi = \dot{\phi} - \operatorname{div}(\mathbf{q}) - f^\phi \doteq 0 \quad \text{in } \mathcal{B} \quad (5.8)$$

which we complement by the corresponding Dirichlet and Neumann boundary conditions $\phi = \bar{\phi}$ on $\partial\mathcal{B}_\phi$ and $\mathbf{q} \cdot \mathbf{n} = \bar{q}$ on $\partial\mathcal{B}_q$. For most physiologically relevant excitation problems, homogeneous Neumann boundary conditions $\mathbf{q} \cdot \mathbf{n} = 0$ are applied on the entire boundary $\partial\mathcal{B}$. As initial conditions, $\phi_0(x) = \phi(x, t_0)$ in \mathcal{B} , we typically set the transmembrane potential to its resting state. The weak form of the electrical residual (8.22) is obtained by the integration over the domain \mathcal{B} , the standard integration by parts, and the inclusion of the Neumann boundary conditions. For the spatial discretization, we discretize the domain of interest \mathcal{B} with n_{el} finite elements \mathcal{B}^e as $\mathcal{B} = \bigcup_{e=1}^{n_{\text{el}}} \mathcal{B}^e$ and apply the standard isoparametric concept to interpolate the trial functions ϕ^h and the test functions $\delta\phi^h$.

$$\delta\phi^h|_{\mathcal{B}^e} = \sum_{i=1}^{n_{\text{en}}} N^i \delta\phi_i \quad \phi^h|_{\mathcal{B}^e} = \sum_{j=1}^{n_{\text{en}}} N^j \phi_j \quad (5.9)$$

Here, N are the standard shape functions on the element level and $i, j = 1, \dots, n_{\text{en}}$ are the n_{en} element nodes. For the temporal discretization, we partition the time interval of interest \mathcal{T} into n_{stp} subintervals $[t^n, t^{n+1}]$ as $\mathcal{T} = \bigcup_{n=0}^{n_{\text{stp}}-1} [t^n, t^{n+1}]$ and apply a standard backward Euler time integration scheme in combination with a finite difference approximation of the first order time derivative $\dot{\phi}$.

$$\dot{\phi} = [\phi - \phi^n] / \Delta t \quad (5.10)$$

Herein, the index $(\circ)^{n+1}$ has been omitted for the sake of clarity, and the common abbreviation $\Delta t := t - t^n > 0$ has been introduced for the current time increment. With the discretizations in space (11.6) and time (8.24), the discrete algorithmic

residual \mathbf{R}_I^ϕ takes the following explicit representation.

$$\mathbf{R}_I^\phi = \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}^e} N^i \frac{\phi - \phi^n}{\Delta t} + \nabla N^i \cdot \mathbf{q} \, dV - \int_{\partial \mathcal{B}_q^e} N^i \bar{q} \, dA - \int_{\mathcal{B}^e} N^i f^\phi \, dV \doteq 0 \quad (5.11)$$

The operator \mathbf{A} symbolizes the assembly of all element contributions at the element nodes $i = 1, \dots, n_{en}$ to the overall residual at the global node points $I = 1, \dots, n_{nd}$. To solve the discrete system of nonlinear equations (8.25), we suggest an incremental iterative Newton Raphson solution technique based on the consistent linearization of the residual which introduces the global iteration matrix \mathbf{K}_{IJ}^ϕ .

$$\mathbf{K}_{IJ}^\phi = d_{\phi_J} \mathbf{R}_I^\phi = \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}^e} N^i \frac{1}{\Delta t} N^j + \nabla N^i \cdot \mathbf{D} \cdot \nabla N^j - N^i d_\phi f^\phi N^j \, dV \quad (5.12)$$

For each incremental iteration, we update the global vector of unknowns $\phi_I \leftarrow \phi_I - \sum_{J=1}^{n_{nd}} \mathbf{K}_{IJ}^{\phi-1} \mathbf{R}_J^\phi$ at all $I = 1, \dots, n_{nd}$ global nodes. In the following subsection, we illustrate the iterative calculation of the source term $f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ and its consistent algorithmic linearization $d_\phi f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ required to evaluate the global residual (8.25) and the global iteration matrix (8.26).

5.3.2 Chemical problem - Local discretization on the integration point level

The chemical problem is characterized through n_{gate} gating variables g_{gate}^I and g_{gate}^{II} , and n_{ion} ion concentrations c_{ion} which we introduce as internal variables to be stored locally on the integration point level. We typically initialize the chemical state variables at t_0 with their resting state values. For their advancement in time, we suggest a finite difference approximation for their temporal discretization,

$$\dot{g}_{\text{gate}}^I = [g_{\text{gate}}^I - g_{\text{gate}}^{In}] / \Delta t \quad \dot{g}_{\text{gate}}^{II} = [g_{\text{gate}}^{II} - g_{\text{gate}}^{II n}] / \Delta t \quad \dot{c}_{\text{ion}} = [c_{\text{ion}} - c_{\text{ion}}^n] / \Delta t \quad (5.13)$$

and apply the classical implicit backward Euler scheme to transform the linear set of gating equations (10.24) into a set of update equations for the gating variables g_{gate}^I

and g_{gate}^{II} at the current time step t .

$$g_{\text{gate}}^I = g_{\text{gate}}^{In} + \frac{1}{\tau_{\text{gate}}^I(\phi)} [g_{\text{gate}}^{\infty I}(\phi) - g_{\text{gate}}^I] \Delta t \quad (5.14)$$

$$g_{\text{gate}}^{II} = g_{\text{gate}}^{II n} + \frac{1}{\tau_{\text{gate}}^{II}(\phi)} [g_{\text{gate}}^{\infty II}(\phi, c_{\text{ion}}) - g_{\text{gate}}^{II}] \Delta t \quad (5.15)$$

Both sets are initialized based on the current membrane potential ϕ . While the first set remains constant throughout the remainder of the constitutive subroutine, the second set is updated iteratively throughout the subsequent local Newton iterations. The gating variables essentially define the n_{crt} ionic currents $I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}})$ which alter the intracellular ion concentrations through the righthand sides f_{ion}^c of equation (5.6). With the help of the finite difference approximation (8.18), the nonlinear set of concentration equations (5.6), which constitutes the core of the chemical problem, is restated in the following residual format.

$$\mathbf{R}_{\text{ion}}^c = c_{\text{ion}} - c_{\text{ion}}^n - f_{\text{ion}}^c(\phi, g_{\text{gate}}, c_{\text{ion}}) \Delta t \doteq 0 \quad (5.16)$$

The discrete algorithmic residual is linearized consistently to yield the $n_{\text{ion}} \times n_{\text{ion}}$ iteration matrix $\mathbf{K}_{\text{ion ion}}^c$ for the local Newton iteration on the integration point level.

$$\mathbf{K}_{\text{ion ion}}^c = \mathbf{d}_{c_{\text{ion}}} \mathbf{R}_{\text{ion}}^c \quad (5.17)$$

At the end of each Newton iteration, we update the set of ion concentrations $c_{\text{ion}} \leftarrow c_{\text{ion}} - [\mathbf{K}_{\text{ion ion}}^c]^{-1} \mathbf{R}_{\text{ion}}^c$, the second set of gate variables $g_{\text{gate}}^{II} \leftarrow g_{\text{gate}}^{II} + f_{\text{gate}}^{gII}(\phi, g_{\text{gate}}, c_{\text{ion}}) \Delta t$ and the set of ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}})$. At convergence, i.e., at chemical equilibrium, we can finally calculate the source term $f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ for the electrical problem (8.25), and its linearization $\mathbf{d}_\phi f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ for the global Newton iteration (8.26). Table 11.1 illustrates the algorithmic solution of the coupled electrochemical problem with its characteristic local-global split. Its local inner loop can be understood as a modern implicit version of the iterative update procedure of the original Rush-Larsen algorithm [160]. Note that, in principle, we could solve for all our internal variables, i.e., for all n_{gate} gating variables and for all n_{ion} ionic concentrations

simultaneously. This would require to invert a $[n_{\text{gate}} + n_{\text{ion}}] \times [n_{\text{gate}} + n_{\text{ion}}]$ iteration matrix, in our case a 17×17 matrix, for each local Newton iteration, at each integration point, during each global iteration step, for each time increment. Because of the particular interdependence of the internal variables, however, we can first update the first set of gating variables g_{gate}^I that only depend on the current membrane potential ϕ , but not on any other internal variables. Then, we calculate the coupled set of ion concentrations c_{ion} which is characterized only through a $n_{\text{ion}} \times n_{\text{ion}}$ iteration matrix, in our case a 4×4 matrix. Last, we update the second set of gating variables g_{gate}^{II} , which then only depends on previously calculated internal variables. Table 11.1 illustrates the local update algorithm tailored to this particular interdependence of internal variables. Overall, this local update is fully implicit.

5.4 Continuous model problem for human ventricular cardiomyocytes

In this section, we will specify the constitutive equations of electrochemistry for an enhanced version of the classical Luo-Rudy model for ventricular cardiomyocytes [116, 117] that incorporates recently proposed modifications [68, 149, 188, ?] as illustrated in Figure 10.2. This model is characterized through $n_{\text{ion}} = 4$ ion concentrations,

$$\dot{c}_{\text{ion}} = \dot{c}_{\text{ion}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad \text{with} \quad c_{\text{ion}} = [c_{\text{Na}}, c_{\text{K}}, c_{\text{Ca}}, c_{\text{Ca}}^{\text{sr}}] \quad (5.18)$$

where c_{Na} , c_{K} , and c_{Ca} are the intracellular sodium, potassium, and calcium concentration, and $c_{\text{Ca}}^{\text{sr}}$ is the calcium concentration in the sarcoplasmic reticulum. Figure 10.2 illustrates the $n_{\text{crt}} = 15$ ionic currents of the model.

$$I_{\text{crt}} = I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (5.19)$$

for $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]$.

In particular, the sodium related currents I_{Na} , I_{bNa} , I_{NaK} , I_{NaCa} induce changes in the intracellular sodium concentration c_{Na} , the potassium related currents I_{K1} , I_{Kr} ,

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

I_{Ks} , I_{NaK} , I_{pK} , I_{t0} induce changes in the intracellular potassium concentration c_K , the calcium related currents I_{CaL} , I_{bCa} , I_{pCa} , I_{NaCa} , I_{leak} , I_{up} , I_{rel} induce changes in the intracellular calcium concentration c_{Ca} , and the calcium related currents I_{leak} , I_{up} , I_{rel} induce changes in the calcium concentration in the sarcoplasmic reticulum c_{Ca}^{sr} , respectively. The states of the channels associated with these currents are gated by $n_{gate} = 13$ gating variables,

$$\begin{aligned} \dot{g}_{gate}^I &= \dot{g}_{gate}^I(\phi, g_{gate}^I) & \text{with} & & g_{gate}^I &= [g_m, g_h, g_j, g_{xr1}, g_{xr2}, g_{xs}, g_r, g_s, g_d, g_f] \\ \dot{g}_{gate}^{II} &= \dot{g}_{gate}^{II}(\phi, g_{gate}^{II}, c_{ion}) & & & g_{gate}^{II} &= [g_{K1\infty}, g_{fCa}, g_g], \end{aligned} \quad (5.20)$$

with g_m , g_h , g_j gating I_{Na} , the fast sodium channel, $g_{K1\infty}$ gating I_{K1} , the inward rectifier channel, g_{xr1} , g_{xr2} gating I_{Kr} , the rapid delayed rectifier channel, g_{xs} gating I_{Ks} , the slow delayed rectifier channel, g_r , g_s gating I_{t0} , the transient outward channel, g_d , g_f , g_{fCa} gating I_{CaL} , the L-type calcium channel, and g_d , g_g gating I_{rel} , the sarcoplasmic reticulum calcium release channel, respectively, see Figure 10.2. For each ion, sodium, potassium, and calcium, we can evaluate the classical Nernst equation,

$$\phi_{ion} = \frac{RT}{z_{ion}F} \log \left(\frac{c_{ion0}}{c_{ion}} \right) \quad \text{with} \quad \phi_{ion} = [\phi_{Na}, \phi_K, \phi_{Ca}] \quad (5.21)$$

to determine the concentration-dependent Nernst or reversal potential ϕ_{ion} , which corresponds to the potential difference across the cell membrane that would be generated by this particular ion if no other ions were present. This implies that at times when the membrane is particularly permeable to a specific ion, its overall membrane potential ϕ tends to approach this ion's equilibrium potential ϕ_{ion} . In the Nernst equation (5.21), $R = 8.3143 \text{ J K}^{-1} \text{ mol}^{-1}$ is the gas constant, $T = 310 \text{ K}$ is the absolute temperature, and $F = 96.4867 \text{ C/mmol}$ is the Faraday constant. The constant z_{ion} is the elementary charge per ion, i.e., $z_{Na}=1$, $z_K=1$, for singly-charged sodium and potassium ions and $z_{Ca}=2$ for doubly-charged calcium ions. The extracellular sodium, potassium, and calcium concentrations are given as $c_{Na0} = 140 \text{ mM}$, $c_{K0} = 5.4 \text{ mM}$, and $c_{Ca0} = 2 \text{ mM}$, respectively, and c_{ion} denotes the corresponding intracellular ion

concentration. In the following subsections, we will specify the individual concentrations, currents, and gating variables for sodium, potassium, and calcium. These will allow us to define the source term f^ϕ for the electrical problem (8.16).

$$f^\phi = -[I_{\text{Na}} + I_{\text{bNa}} + I_{\text{NaK}} + I_{\text{NaCa}} + I_{\text{K1}} + I_{\text{Kr}} + I_{\text{Ks}} + I_{\text{pK}} + I_{\text{t0}} + I_{\text{CaL}} + I_{\text{bCa}} + I_{\text{pCa}}] \quad (5.22)$$

Throughout the remainder of the manuscript, physical units will be used throughout, with time t given in milliseconds, voltage ϕ given in millivolts, ionic currents across the cell membrane given in picoamperes per picofarad, ionic currents across the membrane of the sarcoplasmic reticulum given in millimolar per millisecond, conductances C_{crt} given in nanosiemens per picofarad, and intracellular and extracellular ion concentrations c_{ion} given in millimoles per liter. For the sake of completeness, all material parameters of the human ventricular cardiomyocyte model [68, 116, 149, 188] are summarized in Table 5.2.

5.4.1 Specification of sodium concentration, currents, and gating variables

Sodium plays a crucial role in generating the fast upstroke in the initial phase of the action potential. At rest, the intracellular sodium concentration is approximately $c_{\text{Na}} = 11.6$ mM, which implies that, according to equation (5.21), the sodium equilibrium potential is $\phi_{\text{Na}} = +66.5$ mV. Accordingly, both electrical forces and chemical gradients pull extracellular sodium ions into the cell. The influx of sodium ions is small, however, since at rest, the membrane is relatively impermeable to sodium. Through an external stimulus above a critical threshold value, the fast sodium channels are opened to initiate a rapid inflow of sodium ions associated with the rapid depolarization of the cell membrane. The transmembrane potential increases drastically by more than 100 mV in less than 2 ms, see Figure 10.2. At the end of the upstroke, the cell membrane is positively charged, and the fast sodium channels return to their closed state. In our specific model problem of human ventricular cardiomyocytes, the

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

sodium concentration

$$\dot{c}_{\text{Na}} = -\frac{C}{VF} [I_{\text{Na}} + I_{\text{bNa}} + 3 I_{\text{NaK}} + 3 I_{\text{NaCa}}] \quad (5.23)$$

is evolving in response to the fast sodium current I_{Na} , the background sodium current I_{bNa} , the sodium potassium pump current I_{NaK} , and the sodium calcium exchanger current I_{NaCa} , scaled by the membrane capacitance per unit surface area $C = 185$ pF, the cytoplasmic volume $V = 16404 \mu\text{m}^3$, and the Faraday constant $F = 96.4867$ C/mmol. Note that both the sodium potassium pump and the sodium calcium exchanger operate at a three-to-two ratio as indicated by the scaling factor three. The sodium related currents are defined as follows,

$$\begin{aligned} I_{\text{Na}} &= C_{\text{Na}}^{\text{max}} g_{\text{m}}^3 g_{\text{h}} g_{\text{j}} [\phi - \phi_{\text{Na}}] \\ I_{\text{bNa}} &= C_{\text{bNa}}^{\text{max}} [\phi - \phi_{\text{Na}}] \\ I_{\text{NaK}} &= I_{\text{NaK}}^{\text{max}} [c_{\text{K0}} c_{\text{Na}}] [[c_{\text{Na}} + c_{\text{NaK}}][c_{\text{K0}} + c_{\text{KNa}}] \\ &\quad [1 + 0.1245e^{-0.1\phi F/RT} + 0.0353e^{-\phi F/RT}]^{-1} \\ I_{\text{NaCa}} &= I_{\text{NaCa}}^{\text{max}} [e^{\gamma\phi F/RT} c_{\text{Na}}^3 c_{\text{Ca0}} - e^{(\gamma-1)\phi F/RT} c_{\text{Na0}}^3 c_{\text{Ca}} \gamma_{\text{NaCa}}] \\ &\quad [[c_{\text{NaCa}}^3 + c_{\text{Na0}}^3][c_{\text{CaNa}} + c_{\text{Ca0}}][1 + k_{\text{NaCa}}^{\text{sat}} e^{(\gamma-1)\phi F/RT}]^{-1} \end{aligned} \quad (5.24)$$

where the scaling factors are the maximum fast sodium conductance $C_{\text{Na}}^{\text{max}}=14.838$ nS/pF, the maximum background sodium conductance $C_{\text{bNa}}^{\text{max}} = 0.00029$ nS/pF, the maximum sodium potassium pump current $I_{\text{NaK}}^{\text{max}} = 1.362$ pA/pF, and the maximum sodium calcium exchanger current $I_{\text{NaCa}}^{\text{max}} = 1000$ pA/pF, respectively. The rapid upstroke in the membrane potential is generated by the fast sodium current I_{Na} which is characterized through a three-gate formulation of Beeler-Reuter type [12] in terms of the sodium activation gate g_{m} , the fast sodium inactivation gate g_{h} , and the slow sodium inactivation gate g_{j} . Their evolution is governed by classical Hodgkin-Huxley type equations (10.24) of the format $\dot{g}_{\text{gate}} = [g_{\text{gate}}^{\infty} - g_{\text{gate}}]/\tau_{\text{gate}}$ where g_{gate}^{∞} characterizes the steady state value and τ_{gate} denotes the time constant associated with reaching the steady state. For the sodium activation gate $\dot{g}_{\text{m}} = [g_{\text{m}}^{\infty} - g_{\text{m}}]/\tau_{\text{m}}$, which

initiates the rapid upstroke, they take the following explicit representations.

$$\begin{aligned} g_m^\infty &= [1 + e^{(-56.86-\phi)/9.03}]^{-2} \\ \tau_m &= 0.1 [1 + e^{(-60-\phi)/5}]^{-1} [[1 + e^{(\phi+35)/5}]^{-1} + [1 + e^{(\phi-50)/200}]^{-1}] \end{aligned} \quad (5.25)$$

The kinetics of inactivation are exponential. For the fast sodium inactivation gate $\dot{g}_h = [g_h^\infty - g_h]/\tau_h$, which initiates a fast inactivation of the sodium channel almost instantaneously after the rapid upstroke, the steady state value and the corresponding time constant can be expressed as follows.

$$\begin{aligned} g_h^\infty &= [1 + e^{(\phi+71.55)/7.43}]^{-2} \\ \tau_h &= \begin{cases} 0.1688 [1 + e^{-(\phi+10.66)/11.1}] & \text{if } \phi \geq -40 \\ [0.057 e^{-(\phi+80)/6.8} + 2.7 e^{0.079\phi} + 3.1 \cdot 10^5 e^{0.3485\phi}]^{-1} & \text{if } \phi < -40 \end{cases} \end{aligned} \quad (5.26)$$

For the slow sodium inactivation gate $\dot{g}_j = [g_j^\infty - g_j]/\tau_j$, which gradually inactivates the fast sodium channel over a time span of 100 to 200 ms, these constants take the following form.

$$\begin{aligned} g_j^\infty &= [1 + e^{(\phi+71.55)/7.43}]^{-2} \\ \tau_j &= [\alpha_j + \beta_j]^{-1} \\ \alpha_j &= \begin{cases} 0 & \text{if } \phi \geq -40 \\ \frac{[-2.5428 \cdot 10^4 e^{0.2444\phi} - 6.948 \cdot 10^{-6} e^{-0.04391\phi}][\phi + 37.78]}{[1 + e^{0.311(\phi+79.23)}]} & \text{if } \phi < -40 \end{cases} \\ \beta_j &= \begin{cases} 0.6 e^{0.057\phi} [1 + e^{-0.1(\phi+32)}]^{-1} & \text{if } \phi \geq -40 \\ 0.02424 e^{-0.01052\phi} [1 + e^{-0.1378(\phi+40.14)}]^{-1} & \text{if } \phi < -40 \end{cases} \end{aligned} \quad (5.27)$$

The sodium ions that entered the cell rapidly during the fast upstroke are removed from the cell by the sodium potassium pump I_{NaK} , a metabolic pump that continuously expels sodium ions from the cell interior and pumps in potassium ions. The intracellular sodium concentration is further affected by expulsion of intracellular calcium ions through sodium calcium exchange I_{NaCa} . The additional parameters for the sodium potassium pump current I_{NaK} and for the sodium calcium exchanger current I_{NaCa} are the extracellular sodium, potassium, and calcium concentrations

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

$c_{\text{Na}0} = 140 \text{ mM}$, $c_{\text{K}0} = 5.4 \text{ mM}$, and $c_{\text{Ca}0} = 2 \text{ mM}$, the half saturation constants $c_{\text{CaNa}} = 1.38 \text{ mM}$, $c_{\text{NaCa}} = 87.5 \text{ mM}$, $c_{\text{KNa}} = 1 \text{ mM}$, $c_{\text{NaK}} = 40 \text{ mM}$, the sodium calcium saturation factor $k_{\text{NaCa}}^{\text{sat}} = 0.1$, the outward sodium calcium pump current enhancing factor $\gamma_{\text{NaCa}} = 2.5$, and the voltage dependent sodium calcium parameter $\gamma = 0.35$.

5.4.2 Specification of potassium concentration, currents, and gating variables

Potassium plays an important role in maintaining the appropriate action potential profile in all four phases after the rapid upstroke. At rest, the intracellular potassium concentration is typically about $c_{\text{K}} = 138.3 \text{ mM}$, and the related equilibrium potential would be $\phi_{\text{K}} = -86.6 \text{ mV}$ according to equation (5.21). This value is very close to, but slightly more negative than, the resting potential of $\phi = -86 \text{ mV}$ actually measured in ventricular cardiomyocytes. Unlike for sodium, the electrical force that pulls potassium ions inward is slightly weaker than the chemical force of diffusion pulling potassium ions outward. Accordingly, potassium tends to leave the resting cell. At the end of the rapid upstroke, before the beginning of the plateau, we can observe an early, brief period of limited repolarization governed by the voltage-activated transient outward current I_{t0} . During the following plateau phase, we observe an influx of calcium ions which is balanced by the efflux of an equal amount of positively charged potassium ions, mainly regulated by the rapid and slow delayed rectifier currents I_{Kr} and I_{Ks} . The final repolarization phase can almost exclusive be attributed to potassium ions leaving the cell such that the membrane potential can return to its resting state, see Figure 10.2. In summary, the evolution of the potassium concentration

$$\dot{c}_{\text{K}} = -\frac{C}{VF} [I_{\text{K}1} + I_{\text{Kr}} + I_{\text{Ks}} - 2 I_{\text{NaK}} + I_{\text{pK}} + I_{t0} + I_{\text{stim}}] \quad (5.28)$$

is mainly controlled by four currents, the inward rectifier current $I_{\text{K}1}$, the rapid delayed rectifier current I_{Kr} , the slow delayed rectifier current I_{Ks} , and the transient outward current I_{t0} . Moreover, it is affected by the sodium potassium pump current

I_{NaK} , the plateau potassium current I_{pK} , and the external stimulus current I_{stim} . Currents are scaled by the membrane capacitance per unit surface area $C = 185 \text{ pF}$, the cytoplasmic volume $V = 16404 \mu\text{m}^3$, and the Faraday constant $F = 96.4867 \text{ C/mmol}$. The individual potassium related currents are defined as follows,

$$\begin{aligned}
I_{\text{K1}} &= C_{\text{K1}}^{\text{max}} g_{\text{K1}}^{\infty} [c_{\text{K0}}/5.4]^{1/2} [\phi - \phi_{\text{K}}] \\
I_{\text{Kr}} &= C_{\text{Kr}}^{\text{max}} g_{\text{xr1}} g_{\text{xr2}} [c_{\text{K0}}/5.4]^{1/2} [\phi - \phi_{\text{K}}] \\
I_{\text{Ks}} &= C_{\text{Ks}}^{\text{max}} g_{\text{xs}}^2 [\phi - \phi_{\text{Ks}}] \\
I_{\text{NaK}} &= I_{\text{NaK}}^{\text{max}} [c_{\text{K0}} c_{\text{Na}}] [c_{\text{Na}} + c_{\text{NaK}}] [c_{\text{K0}} + c_{\text{KNa}}] [1 + 0.1245 e^{-0.1\phi F/RT} + 0.0353 e^{-\phi F/RT}]^{-1} \\
I_{\text{pK}} &= C_{\text{pK}}^{\text{max}} [1 + e^{(25-\phi)/5.98}]^{-1} [\phi - \phi_{\text{K}}] \\
I_{\text{t0}} &= C_{\text{t0}}^{\text{max}} g_{\text{r}} g_{\text{s}} [\phi - \phi_{\text{K}}]
\end{aligned} \tag{5.29}$$

where the individual scaling factors are the maximum inward rectifier conductance $C_{\text{K1}}^{\text{max}} = 5.405 \text{ nS/pF}$, the maximum rapid delayed rectifier conductance $C_{\text{Kr}}^{\text{max}} = 0.096 \text{ nS/pF}$, the maximum slow delayed rectifier conductance for epicardial and endocardial cells $C_{\text{Ks, epi}}^{\text{max}} = C_{\text{Ks, endo}}^{\text{max}} = 0.245 \text{ nS/pF}$ and for M cells $C_{\text{Ks, M}}^{\text{max}} = 0.062 \text{ nS/pF}$, the maximum sodium potassium pump current $I_{\text{NaK}}^{\text{max}} = 1.362 \text{ pA/pF}$, the maximum potassium pump conductance $C_{\text{pK}}^{\text{max}} = 0.0146 \text{ nS/pF}$, and the maximum transient outward conductance for epicardial and M cells $C_{\text{t0, epi}}^{\text{max}} = C_{\text{t0, M}}^{\text{max}} = 0.294 \text{ nS/pF}$ and for endocardial cells $C_{\text{t0, endo}}^{\text{max}} = 0.073 \text{ nS/pF}$. The maximum inward rectifier current I_{K1} , which is most active during the later phases of the action potential, depends explicitly on the extracellular potassium concentration $c_{\text{K0}} = 5.4 \text{ mM}$. It is further characterized through the time-independent inward rectification factor g_{K1}^{∞} parameterized in terms of the potential equilibrium potential ϕ_{K} given in equation (5.21).

$$\begin{aligned}
g_{\text{K1}}^{\infty} &= \alpha_{\text{K1}} [\alpha_{\text{K1}} + \beta_{\text{K1}}]^{-1} \\
\text{with} \quad \alpha_{\text{K1}} &= 0.1 [1 + e^{0.06(\phi - \phi_{\text{K}} - 200)}]^{-1} \\
\beta_{\text{K1}} &= [3 e^{0.0002(\phi - \phi_{\text{K}} + 100)} + e^{0.1(\phi - \phi_{\text{K}} - 10)}] [1 + e^{-0.5(\phi - \phi_{\text{K}})}]^{-1}
\end{aligned} \tag{5.30}$$

The action potential plateau is characterized through the influx of charged calcium ions balanced by the efflux of potassium ions. The latter is basically governed by the rapid and slow delayed rectifier current I_{Kr} and I_{Ks} . The channel for the rapid delayed rectifier current I_{Kr} is gated by an activation gate $\dot{g}_{\text{x1}} = [g_{\text{x1}}^{\infty} - g_{\text{x1}}] / \tau_{\text{x1}}$ with

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

the steady state value and time constant given as

$$\begin{aligned} g_{xr1}^{\infty} &= [1 + e^{(-26-\phi)/7}]^{-1} \\ \tau_{xr1} &= 2700 [1 + e^{(-45-\phi)/10}]^{-1} [1 + e^{(\phi+30)/11.5}]^{-1}, \end{aligned} \quad (5.31)$$

and by an inactivation gate $\dot{g}_{x2} = [g_{x2}^{\infty} - g_{x2}]/\tau_{x2}$, with the following steady state value and time constant.

$$\begin{aligned} g_{xr2}^{\infty} &= [1 + e^{(\phi+88)/24}]^{-1} \\ \tau_{xr2} &= 3.36 [1 + e^{(-60-\phi)/20}]^{-1} [1 + e^{(\phi-60)/20}]^{-1}. \end{aligned} \quad (5.32)$$

The channel for the slow delayed rectifier current I_{Ks} is a function of the reversal potential $\phi_{Ks} = RT/F \log ([c_{K0} + p_{KNa} c_{Na0}][c_K + p_{KNa} c_{Na}]^{-1})$ parameterized in terms of its permeability to sodium ions $p_{KNa} = 0.03$. It is gated by an activation gate $\dot{g}_{xs} = [g_{xs}^{\infty} - g_{xs}]/\tau_{xs}$ in terms of the following parameterization.

$$\begin{aligned} g_{xs}^{\infty} &= [1 + e^{(-5-\phi)/14}]^{-1} \\ \tau_{xs} &= 1100 [1 + e^{(-10-\phi)/6}]^{-1/2} [1 + e^{(\phi-60)/20}]^{-1} \end{aligned} \quad (5.33)$$

The transient potassium outward current I_{t0} is responsible for the transition between the rapid upstroke and the plateau phase, where it generates an early short period of limited repolarization. It is gated by a voltage-dependent activation gate g_r with $\dot{g}_r = [g_r^{\infty} - g_r]/\tau_r$ defined through the following steady state value and time constant,

$$\begin{aligned} g_r^{\infty} &= [1 + e^{(20-\phi)/6}]^{-1} \\ \tau_r &= 9.5 e^{-(\phi+40)^2/1800} + 0.8 \end{aligned} \quad (5.34)$$

and by the voltage-dependent inactivation gate g_s with $\dot{g}_s = [g_s^{\infty} - g_s]/\tau_s$ with the steady state value and time constant given as follows.

$$\begin{aligned} g_s^{\infty} &= [1 + e^{(\phi+20)/5}] \\ \tau_s &= 85e^{-(\phi+45)^2/320} + 5[1 + e^{(\phi-20)/5}] + 3 \end{aligned} \left. \vphantom{\begin{aligned} g_s^{\infty} \\ \tau_s \end{aligned}} \right\} \text{epicardium} \\ g_s^{\infty} &= [1 + e^{(\phi+28)/5}] \\ \tau_s &= 1000e^{-(\phi+67)^2/1000} + 8 \end{aligned} \left. \vphantom{\begin{aligned} g_s^{\infty} \\ \tau_s \end{aligned}} \right\} \text{endocardium} \quad (5.35)$$

This voltage dependent potassium inactivation gate displays a significantly different behavior for epicardial and endocardial cells and is therefore characterized differently for the individual cell types. Similar to the previous subsection, we have introduced the extracellular sodium and potassium concentrations $c_{\text{Na}0} = 140 \text{ mM}$ and $c_{\text{K}0} = 5.4 \text{ mM}$, and the half saturation constants $c_{\text{KNa}} = 1 \text{ mM}$ and $c_{\text{NaK}} = 40 \text{ mM}$.

5.4.3 Specification of calcium concentration, currents, and gating variables

Calcium is the key player to translate electrical excitation into mechanical contraction. With a typical intracellular resting concentrations of $c_{\text{Ca}} = 0.08 \mu\text{M}$, its equilibrium potential of $\phi_{\text{Ca}} = +135.3 \text{ mV}$ is much larger than the resting potential. During the plateau of the action potential, calcium ions enter the cell through calcium channels that typically activate and inactivate much more slowly than the fast sodium channels. The influx of positively charged calcium ions through the L-type calcium channel I_{CaL} is balanced by an efflux of positively charged potassium ions. The letter L is meant to indicate the long lasting nature of the inward calcium current. Overall, changes in the intracellular calcium concentration

$$\dot{c}_{\text{Ca}} = \gamma_{\text{Ca}} \left[-\frac{C}{2VF} [I_{\text{CaL}} + I_{\text{bCa}} + I_{\text{pCa}} - 2I_{\text{NaCa}}] + I_{\text{leak}} - I_{\text{up}} + I_{\text{rel}} \right] \quad (5.36)$$

are affected by the L-type calcium current I_{CaL} , the background calcium current I_{bCa} , the plateau calcium current I_{pCa} , and the sodium calcium pump current I_{NaCa} , weighted by the membrane capacitance per unit surface area $C = 185 \text{ pF}$, the cytoplasmic volume $V = 16404 \mu\text{m}^3$, and the Faraday constant $F = 96.4867 \text{ C/mmol}$. In addition, the intracellular calcium concentration is affected by a calcium loss to the sarcoplasmic reticulum characterized through the leakage current I_{leak} , the sarcoplasmic reticulum uptake current I_{up} , and the sarcoplasmic reticulum release current I_{rel} . The

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

individual calcium related currents are defined as follows,

$$\begin{aligned}
I_{\text{CaL}} &= C_{\text{CaL}}^{\text{max}} g_d g_f g_{f\text{Ca}} [4\phi F^2]/[RT] [c_{\text{Ca}} e^{2\phi F/[RT]} - 0.341 c_{\text{Ca0}}] [e^{2\phi F/[RT]} - 1]^{-1} \\
I_{\text{bCa}} &= C_{\text{bCa}}^{\text{max}} [\phi - \phi_{\text{Ca}}] \\
I_{\text{pCa}} &= C_{\text{pCa}}^{\text{max}} c_{\text{Ca}} [c_{\text{pCa}} + c_{\text{Ca}}]^{-1} \\
I_{\text{NaCa}} &= I_{\text{NaCa}}^{\text{max}} [e^{\gamma\phi F/RT} c_{\text{Na}}^3 c_{\text{Ca0}} - e^{(\gamma-1)\phi F/RT} c_{\text{Na0}}^3 c_{\text{Ca}} \gamma_{\text{NaCa}}] \\
&\quad [[c_{\text{NaCa}}^3 + c_{\text{Na0}}^3][c_{\text{CaNa}} + c_{\text{Ca0}}][1 + k_{\text{NaCa}}^{\text{sat}} e^{(\gamma-1)\phi F/RT}]^{-1} \\
I_{\text{leak}} &= I_{\text{leak}}^{\text{max}} [c_{\text{Ca}}^{\text{sr}} - c_{\text{Ca}}] \\
I_{\text{up}} &= I_{\text{up}}^{\text{max}} [1 + c_{\text{up}}^2/c_{\text{Ca}}^2]^{-1} \\
I_{\text{rel}} &= I_{\text{rel}}^{\text{max}} g_d g_g [1 + \gamma_{\text{rel}} c_{\text{Ca}}^{\text{sr}2} [c_{\text{rel}}^2 + c_{\text{Ca}}^{\text{sr}2}]^{-1}]
\end{aligned} \tag{5.37}$$

where the individual scaling factors are the maximum calcium conductance $C_{\text{CaL}}^{\text{max}} = 0.175 \text{ mm}^3 \mu\text{F}^{-1} \text{ s}^{-1}$, the maximum background calcium conductance $C_{\text{bCa}}^{\text{max}} = 0.000592 \text{ nS/pF}$, the maximum plateau calcium conductance $C_{\text{pCa}}^{\text{max}} = 0.825 \text{ nS/pF}$, the maximum sodium calcium pump current $I_{\text{NaCa}}^{\text{max}} = 1000 \text{ pA/pF}$, the maximum leakage current $I_{\text{leak}}^{\text{max}} = 0.08 \text{ s}^{-1}$, the maximum sarcoplasmic reticulum calcium uptake current $I_{\text{up}}^{\text{max}} = 0.000425 \text{ mM/ms}$, and the maximum sarcoplasmic reticulum calcium release current $I_{\text{rel}}^{\text{max}} = 8.232 \text{ mM/s}$. The major calcium channel, the long-lasting L-type calcium channel I_{CaL} , is controlled by the voltage-dependent activation gate $\dot{g}_d = [g_d^\infty - g_d]/\tau_g$ characterized through the following steady state value and time constant

$$\begin{aligned}
g_d^\infty &= [1 + e^{(-5-\phi)/7.5}]^{-1} \\
\tau_d &= [1.4 [1 + e^{(-35-\phi)/13}]^{-1} + 0.25][1.4 [1 + e^{(\phi+5)/5}] + [1 + e^{(50-\phi)/20}],
\end{aligned} \tag{5.38}$$

by the voltage-dependent inactivate gate $\dot{g}_f = [g_f^\infty - g_f]/\tau_f$ characterized through

$$\begin{aligned}
g_f^\infty &= [1 + e^{(\phi+20)/7}]^{-1} \\
\tau_f &= 1125 e^{-(\phi+27)^2/240} + 165 [1 + e^{(25-\phi)/10}]^{-1} + 80,
\end{aligned} \tag{5.39}$$

and by the intracellular calcium dependent inactivation gate $\dot{g}_{\text{fCa}} = [g_{\text{fCa}}^\infty - g_{\text{fCa}}]/\tau_{\text{fCa}}$ characterized through

$$\begin{aligned}
 g_{\text{fCa}}^\infty &= 0.685 [[1 + (c_{\text{Ca}}/0.000325)^8]^{-1} \\
 &\quad + 0.1[1 + e^{(c_{\text{Ca}}-0.0005)/0.0001}]^{-1} \\
 &\quad + 0.2[1 + e^{(c_{\text{Ca}}-0.00075)/0.0008}]^{-1} + 0.23] \\
 \tau_{\text{fCa}} &= \begin{cases} \infty & \text{if } g_{\text{fCa}}^\infty > g_{\text{fCa}} \text{ and } \phi \geq -60\text{mV} \\ 2 \text{ ms} & \text{otherwise.} \end{cases}
 \end{aligned} \tag{5.40}$$

Accordingly, the steady state response g_{fCa}^∞ has a switchlike shape when going from no inactivation to considerable but incomplete inactivation, depending mildly on the calcium concentration c_{Ca} for suprathreshold concentrations. Last, the calcium-induced calcium release current I_{rel} is characterized through the activation gate g_{d} , the same gate that is also activating the L-type calcium channel of I_{CaL} , and through the calcium-dependent inactivation gate $\dot{g}_{\text{g}} = [g_{\text{g}}^\infty - g_{\text{g}}]/\tau_{\text{g}}$ characterized through the following steady state value and time constant.

$$\begin{aligned}
 g_{\text{g}}^\infty &= \begin{cases} [1 + c_{\text{Ca}}^6/0.00035^6]^{-1} & \text{if } c_{\text{Ca}} \leq 0.00035 \\ [1 + c_{\text{Ca}}^{16}/0.00035^{16}]^{-1} & \text{otherwise} \end{cases} \\
 \tau_{\text{g}} &= \begin{cases} \infty & \text{if } g_{\text{g}}^\infty > g_{\text{g}} \text{ and } \phi \geq -60\text{mV} \\ 2 \text{ ms} & \text{otherwise} \end{cases}
 \end{aligned} \tag{5.41}$$

The remaining parameters governing the response of the plateau calcium current I_{pCa} , the calcium uptake current I_{up} , and the sarcoplasmic reticulum calcium release current I_{rel} are the half saturation constants for the plateau calcium concentration $c_{\text{pCa}} = 0.0005 \text{ mM}$, for the sarcoplasmic reticulum calcium uptake $c_{\text{up}} = 0.00025 \text{ mM}$, and for the sarcoplasmic reticulum calcium release $c_{\text{rel}} = 0.25 \text{ mM}$, respectively. The parameter $\gamma_{\text{NaCa}} = 2.5$ has been introduced to enhance the outward nature of the sodium calcium pump current I_{NaCa} . The additional parameter $\gamma_{\text{rel}} = 2$ weighs the relative influence of the sarcoplasmic reticulum calcium concentration on sarcoplasmic reticulum calcium release I_{rel} . Finally, we need to take into account that the total intracellular calcium concentration $c_{\text{Ca}}^{\text{tot}} = c_{\text{Ca}} + c_{\text{Ca}}^{\text{buf}}$ in the cytoplasm is the sum of the

5.4. CONTINUOUS MODEL PROBLEM FOR HUMAN VENTRICULAR CARDIOMYOCYTES

free intracellular calcium concentration c_{Ca} and the buffered calcium concentration $c_{Ca}^{buf} = [c_{Ca} c_{Cabuf}^{tot}][c_{Ca} - c_{Cabuf}]^{-1}$. The definition of the free intracellular calcium concentration in equation (5.36) is therefore weighted by the parameter $\gamma_{Ca} = [1 + [c_{tot} c_{buf}][c_{Ca} + c_{buf}]^{-2}]^{-1}$, where $c_{tot} = 0.15$ mM and $c_{buf} = 0.001$ mM are the total and half saturation cytoplasmic calcium buffer concentrations, respectively.

5.4.4 Specification of sarcoplasmic reticulum calcium concentration, currents, and gating variables

The specification of the sarcoplasmic reticulum calcium concentration

$$\dot{c}_{Ca}^{sr} = \gamma_{Ca}^{sr} \frac{V}{V^{sr}} [-I_{leak} + I_{up} - I_{rel}] \quad (5.42)$$

is now relatively straightforward since it mimics the corresponding loss of intracellular calcium characterized however, now scaled by the ratio between the volume of the cytoplasm $V = 16404 \mu\text{m}^3$ and the volume of the sarcoplasmic reticulum $V^{sr} = 1094 \mu\text{m}^3$. The leakage current I_{leak} , the sarcoplasmic reticulum uptake current I_{up} , and the sarcoplasmic reticulum release current I_{rel} are defined as before.

$$\begin{aligned} I_{leak} &= I_{leak}^{max} [c_{Ca}^{sr} - c_{Ca}] \\ I_{up} &= I_{up}^{max} [1 + c_{up}^2/c_{Ca}^2]^{-1} \\ I_{rel} &= I_{rel}^{max} g_d g_g [1 + \gamma_{rel} c_{Ca}^{sr 2} [c_{rel}^2 + c_{Ca}^{sr 2}]^{-1}] \end{aligned} \quad (5.43)$$

The maximum leakage current $I_{leak}^{max} = 0.08 \text{ s}^{-1}$, the maximum sarcoplasmic reticulum calcium uptake current $I_{up}^{max} = 0.000425 \text{ mM/ms}$, and the maximum sarcoplasmic reticulum calcium release current $I_{rel}^{max} = 8.232 \text{ mM/s}$, the half saturation constants for the calcium uptake $c_{up} = 0.00025 \text{ mM}$, and for the calcium release $c_{rel} = 0.25 \text{ mM}$, and the weighting coefficient $\gamma_{rel} = 2$ have already been introduced in the previous subsection. Similar to the previous subsection, we need to take into account that the total calcium concentration in the sarcoplasmic reticulum $c_{Ca}^{sr tot} = c_{Ca}^{sr} + c_{Ca}^{sr buf}$ is the sum of the free sarcoplasmic reticulum calcium concentration c_{Ca}^{sr} and the buffered sarcoplasmic reticulum calcium concentration $c_{Ca}^{sr buf} = [c_{Ca}^{sr} c_{tot}^{sr}][c_{Ca}^{sr} - c_{buf}^{sr}]^{-1}$. The definition of

the free sarcoplasmic reticulum calcium concentration in equation (5.42) is therefore weighted by the parameter $\gamma_{Ca}^{sr} = [1 + [c_{tot}^{sr} c_{buf}^{sr}][c_{Ca}^{sr} + c_{buf}^{sr}]^{-2}]^{-1}$, where $c_{tot}^{sr} = 10$ mM and $c_{buf}^{sr} = 0.3$ mM are the total and half saturation sarcoplasmic reticulum calcium buffer concentrations, respectively.

5.5 Discrete model problem for human ventricular cardiomyocytes

Finally, we can specify the discrete ion concentration residuals R_{ion} introduced in equation (8.20). For our particular model problem of human ventricular cardiomyocytes we use the individual righthand sides f_{ion}^c defined in equations (5.28), (5.23), (5.36), and (5.42).

$$\begin{aligned}
R_K^c &= c_K - c_K^n + \frac{C}{VF} [I_{K1} + I_{Kr} + I_{Ks} - 2I_{NaK} + I_{pK} + I_{t0} + I_{stim}] & \Delta t &\doteq 0 \\
R_{Na}^c &= c_{Na} - c_{Na}^n + \frac{C}{VF} [I_{Na} + I_{bNa} + 3I_{NaK} + 3I_{NaCa}] & \Delta t &\doteq 0 \\
R_{Ca}^c &= c_{Ca} - c_{Ca}^n + \left[\frac{C}{2VF} [I_{CaL} + I_{bCa} + I_{pCa} - 2I_{NaCa}] - I_{leak} + I_{up} - I_{rel} \right] \gamma_{Ca} & \Delta t &\doteq 0 \\
R_{Ca}^{sr c} &= c_{Ca}^{sr} - c_{Ca}^{sr n} + \frac{V}{V^{sr}} [I_{leak} - I_{up} + I_{rel}] \gamma_{Ca}^{sr} & \Delta t &\doteq 0
\end{aligned} \tag{5.44}$$

Note that for our algorithmic formulation, we have re-arranged the vector of residuals $R_{ion}^c = [R_K^c, R_{Na}^c, R_{Ca}^c, R_{Ca}^{sr c}]$ and the vector of ion concentrations $c_{ion} = [c_K, c_{Na}, c_{Ca}, c_{Ca}^{sr}]$ to obtain a conveniently sparse iteration matrix $K_{ion ion}^c$. According to equation (8.21), this iteration matrix for the local Newton iteration is derived as the linearization of the residual vector R_{ion}^c with respect to the vector of ion concentrations c_{ion} .

$$K_{ion ion}^c = d_{c_{ion}} R_{ion}^c = \begin{bmatrix} d_{c_K} R_K^c & d_{c_{Na}} R_K^c & 0 & 0 \\ 0 & d_{c_{Na}} R_{Na}^c & d_{c_{Ca}} R_{Na}^c & 0 \\ 0 & d_{c_{Na}} R_{Ca}^c & d_{c_{Ca}} R_{Ca}^c & d_{c_{Ca}^{sr}} R_{Ca}^c \\ 0 & 0 & d_{c_{Ca}} R_{Ca}^{sr c} & d_{c_{Ca}^{sr}} R_{Ca}^{sr c} \end{bmatrix} \tag{5.45}$$

At convergence, i.e., at chemical equilibrium, we can finally calculate the source term $f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ of the electrical problem (8.25) according to equation (5.22).

$$f^\phi = -[I_{\text{Na}} + I_{\text{bNa}} + I_{\text{NaK}} + I_{\text{NaCa}} + I_{\text{K1}} + I_{\text{Kr}} + I_{\text{Ks}} + I_{\text{pK}} + I_{\text{t0}} + I_{\text{CaL}} + I_{\text{bCa}} + I_{\text{pCa}}] \quad (5.46)$$

Its linearization $d_\phi f^\phi$ with respect to the membrane potential ϕ

$$\begin{aligned} d_\phi f^\phi = & - [d_\phi I_{\text{Na}} + d_\phi I_{\text{bNa}} + d_\phi I_{\text{NaK}} + d_\phi I_{\text{NaCa}} + d_\phi I_{\text{K1}} + d_\phi I_{\text{Kr}} \\ & + d_\phi I_{\text{Ks}} + d_\phi I_{\text{pK}} + d_\phi I_{\text{t0}} + d_\phi I_{\text{CaL}} + d_\phi I_{\text{bCa}} + d_\phi I_{\text{pCa}}] \end{aligned} \quad (5.47)$$

then enters the iteration matrix for the global Newton iteration (8.26) to ensure optimal quadratic convergence in the proximity of the solution ϕ . The linearizations introduced in equations (5.45) and (5.47) are elaborated in detail in the Appendix.

5.6 Examples

5.6.1 Electrochemistry in a human ventricular cardiomyocyte

To simulate electrochemical coupling in a single epicardial human ventricular cardiomyocyte, we apply the local version of the algorithm described in Table 11.1, ignoring the divergence term of equation (8.22) that has been introduced to model global tissue conductivity. Accordingly, our implementation of the discrete ventricular cell model uses an outer global Newton iteration to solve for the membrane potential ϕ and an inner local Newton iteration to calculate the ion concentrations c_{ion} and the gating variables g_{gate} . We initialize the global membrane potential, $\phi = -86$ mV, and the local ion concentrations, $c_{\text{Na}} = 11.6$ mM, $c_{\text{K}} = 138.3$ mM, and $c_{\text{Ca}} = 0.08$ μ M, mimicking the resting state. For the gating variables, we choose the following initial conditions $g_{\text{m}} = 0$, $g_{\text{h}} = 0.75$, $g_{\text{j}} = 0.75$, $g_{\text{d}} = 0$, $g_{\text{f}} = 1$, $g_{\text{fCa}} = 1$, $g_{\text{r}} = 0$, $g_{\text{s}} = 1$, $g_{\text{xs}} = 0$, $g_{\text{xr1}} = 0$, $g_{\text{xr2}} = 0$, $g_{\text{xK1}\infty} = 0.05$, and $g_{\text{g}} = 1$. Figures 10.3, 10.4, 10.5, and 10.6 represent the electrochemical characteristics for the human ventricular epicardial cardiomyocyte, using the material parameters summarized in Table 5.2, an initial electrical stimulus above the critical threshold, and a discrete time step of $\Delta t = 0.02$ ms,

to match the time step used in the original publication [188]. We have been able to demonstrate though that the time step size could easily be increased by a factor ten without any substantial loss of accuracy. To validate our algorithm against the explicit finite difference results reported in the literature, we reproduce the steady-state profiles g_{gate}^{∞} and the time constant profiles τ_{gate} plotted against the membrane potential ϕ . Figure 10.3 shows the resulting curves for the individual gating variables. Note that the discontinuities in the time constants of τ_h and τ_j which are reported in Figures 10.3c and 10.3e are handled in a piece-wise manner with the partial derivatives and sensitivities calculated for the given membrane potential range. As expected, the steady-state values and time constants coincide perfectly with graphs reported in the original model based on an explicit time integration scheme [188].

Figure 10.4 illustrates the temporal evolution of all thirteen gating variables g_{gate} throughout the duration of a typical action potential. The collection of gating variable profiles nicely illustrates the time sequence of activation and inactivation of the individual ion channels. It also documents which of the gates are slow and fast responding. Figures 10.4 a, b, and c document the activation, the fast inactivation, and the slow inactivation gates g_m , g_h , and g_j for the fast sodium current I_{Na} that governs the rapid upstroke of the action potential ϕ . It is obvious that the inactivation gates g_h and g_j open slightly after the activation gate g_m is closed, with the fast gate g_h responding more rapidly than the slow gate g_j . Figures 10.4 d, e, and f show the activation, the voltage-dependent inactivation, and the intracellular calcium dependent inactivation gates g_d , g_f , and g_{fCa} for the L-type calcium current I_{CaL} that is activated during the action potential upstroke. It is obvious that the long lasting nature of the L-type calcium current can be attributed to the slow response profiles of g_d and g_f . Figures 10.4 g and h illustrate the transient outward channel's activation and inactivation gates g_r and g_s , which manifest themselves in a sharp initial peak in the transient outward current I_{t0} that initiates the short period of early repolarization after initial excitation. Figures 10.4 i, j, k, and l illustrate the slow delayed rectifier gate g_{xs} , the rapid delayed rectifier activation and inactivation gates g_{xr1} and g_{xr2} , and the inward rectification factor g_{K1}^{∞} which collectively determine the potassium concentration profile during the plateau and repolarization phases. Figures 10.4 m

displays the calcium dependent inactivation gate of the sarcoplasmic reticulum release current I_{rel} that characterizes intracellular calcium dynamics through a sharp rapid inactivation towards the end of the repolarization phase.

The evolution of the ionic currents I_{ct} over an excitation cycle is shown in Figure 10.5. The current profiles nicely capture the basic characteristic features of human ventricular cardiomyocytes. The dominance of the fast sodium current I_{Na} in Figure 10.5 a, the transient outward current I_{t0} in Figure 10.5 i, and the L-type calcium current I_{CaL} in Figure 10.5 j is clearly evident. This implies that the sodium, potassium, and calcium concentration profiles primarily depend on these three channels. During the rapid depolarization phase of the cardiomyocyte, we observe a rapid activity of the fast sodium channel I_{Na} . During the following period of partial repolarization, the transient outward current I_{t0} is responsible for a sharp efflux of potassium ions generating the familiar notch in the action potential profile shown in Figure 10.1. The L-type calcium current I_{CaL} is activated rapidly during the depolarization phase and inactivated slowly during the following phases. It is important to note that this calcium current I_{CaL} displays a discontinuity at ϕ equal to zero. Since we need to determine its sensitivities and partial derivatives to guarantee optimal quadratic convergence of our Newton Raphson algorithm, we apply L'Hospital's rule to calculate the algorithmic derivatives in the proximity of this singularity. Altogether, the results in Figure 10.5 correlate well with the reported currents calculated with the explicit time integration scheme reported in the original manuscript [188]. Recall that the ionic current profiles directly feed back into the action potential itself as illustrated in Figure 10.2 and discussed in detail in the introduction.

Lastly, Figure 10.6 documents the evolution of the four ion concentrations throughout a typical action potential cycle. As indicative of the three major currents, i.e., the fast sodium current I_{Na} , the transient outward current I_{t0} , and the L-type calcium current I_{CaL} , the effluxes and influxes directly impact the corresponding ion concentration profiles. The sodium concentration c_{Na} shown in Figure 10.6 a is primarily affected by the fast sodium current I_{Na} initiating a fast intracellular sodium increase to create the rapid upstroke of the action potential. It then decays slowly towards the end of the repolarization phase and increases gradually during the resting phase.

These continuous gradual changes are primarily caused by the sodium potassium pump I_{NaK} and by the sodium calcium exchanger I_{NaCa} . The potassium concentration c_{K} displayed in Figure 10.6 b decreases in a somewhat stepwise fashion regulated by the sequential activation of the transient outward current I_{t0} , the inward rectifier current I_{K1} , and the rapid and slow delayed rectifier currents I_{Kr} and I_{Ks} . At the end of the repolarization phase, we can observe a gradual smooth increase to bring the potassium concentration back to its original value. The calcium concentration c_{Ca} shown in Figure 10.6 c increases rapidly through the opening L-type calcium channel I_{CaL} which is activated slightly after the action potential upstroke. After this sharp increase, the calcium concentration decays smoothly to its original value throughout the remaining phases of the action potential. The intracellular calcium concentration c_{Ca} matches extremely well with the explicit finite difference result [188]. Its profile obviously impacts the intracellular calcium dynamics, and directly affects the calcium concentration in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$ shown in Figure 10.5 d. In summary, the model reproduces the classical characteristics of an initial increase in the sodium concentration followed by an increase in calcium and a decrease in potassium, jointly generating the characteristic plateau. In this model, sodium then experiences a decrease, a minimum, and a gradual increase paired with a potassium increase. Note that despite the drastic changes in the membrane potential from -86 mV to $+20 \text{ mV}$ illustrated in Figure 10.1, the overall changes in the individual ion concentrations remain incredibly small, usually in the order of less than one percent.

5.6.2 Electrochemistry in the human heart

The final example demonstrates the potential of the proposed algorithm in a non-linear finite element analysis of electrochemical coupling using a patient-specific human heart model reconstructed from magnetic resonance images [102]. A tetrahedral heart mesh of 11,347 elements and 3,129 nodes is reconstructed from MRI images. For the global membrane potential, $\phi = -86 \text{ mV}$, and the local ion concentrations, $c_{\text{Na}} = 11.6 \text{ mM}$, $c_{\text{K}} = 138.3 \text{ mM}$, $c_{\text{Ca}} = 0.08 \mu\text{M}$, and $c_{\text{Ca}}^{\text{sr}} = 0.56 \text{ mM}$, we apply initial conditions which mimic the resting state. For the gating variables, we choose the

following initial conditions $g_m = 0, g_h = 0.75, g_j = 0.75, g_d = 0, g_f = 1, g_{fCa} = 1, g_r = 0, g_s = 1, g_{xs} = 0, g_{xr1} = 0, g_{xr2} = 0, g_{xK1\infty} = 0.05$, and $g_g = 1$ similar to the previous single cell example. Moreover, we apply the common assumption of homogeneous Neumann boundary conditions. The heart is excited through the application of an external stimulus in the region of the atrioventricular node in the center of the basal septum. Following the literature, we adopt a time step size of $\Delta t = 0.125$ ms. However, we were able to demonstrate the use of larger time steps, particularly when combined with faster conductivities. For the sake of simplicity, we select an isotropic conductivity $\mathbf{D} = d_{iso}\mathbf{I}$ with $d_{iso} = 0.5 \text{ mm}^2/\text{ms}$. This value is calibrated by means of global electrocardiogram profiles [103], such that the initial excitation of the heart occurs in approximately 30 ms. In the future, we will enhance the model by incorporating an anisotropic conductivity $\mathbf{D} = d_{iso}\mathbf{I} + d_{ani}\mathbf{n} \otimes \mathbf{n}$ with a pronounced signal propagation along preferred directions \mathbf{n} [?], which we are currently calibrating by means of in vitro experiments using microelectrode array recordings [37]. The remaining material parameters which are in agreement with the previous example are listed in Table 5.2. For the sake of simplicity, the entire heart is assumed to be composed of ventricular epicardial cardiomyocytes. Thus, the expected timing and quantitative behavior of different ion concentrations and their currents may not exactly match with what is observed in the actual human heart. However, the incorporation of different cell types is conceptually simple and would require only minor modifications in the finite element input file.

Figures 10.7 and 10.8 illustrate the evolution of the membrane potential ϕ and of the individual ion concentrations c_{Na} , c_K and c_{Ca} during the depolarization and repolarization phases, respectively. Figure 10.7, second row, documents that depolarization is initiated through changes in the intracellular sodium concentration which increases rapidly from 11.60 mM to 11.61 mM within the first 5 ms of the cycle. This increase is associated with a rapid increase in the membrane potential from -86 mV to +20 mV, first row, which, in turn, affects the voltage-gated calcium and potassium channels within the cell membrane. It is primarily through the voltage-gated L-type calcium channel that the intracellular calcium concentration increases from approximately

0.08 μM to 1 μM , fourth row. Sodium follows with a slight time delay of 15 ms decreasing from 138.30 mM to 138.29 mM, third row. After approximately 30 ms, the entire heart is depolarized and the membrane potential has reached its peak value of 20 mV throughout both ventricles. Figure 10.8 displays the repolarization phase characterized through a smooth decrease of the membrane potential back to its initial value of -86 mV after approximately 300 ms, first row. At the same time, the intracellular calcium concentration decreases smoothly back to its resting value of 0.08 μM , fourth row. The intracellular sodium concentration that has initially increased from approximately 11.60 mM to 11.61 mM is now decreasing even below its initial value and reaches a minimum of 11.585 mM after 280 ms, second row. The intracellular potassium concentration reaches its minimum of 138.29 mM at approximately the same time, third row. In the course of time, both sodium and potassium then slowly return to their resting values as their concentrations increase gradually. These results are in excellent qualitative and quantitative agreement with the single cardiomyocyte results documented in Figure 10.6.

Figure 10.9 illustrates the algorithmic performance of the proposed algorithm. The top row shows the non-adaptive time stepping scheme with a fixed time step size of $\Delta t = 0.125$ ms; the bottom row shows the adaptive time stepping scheme with a maximum time step size of $\Delta t^{\text{max}} = 8.0$ ms. Since we apply a Newton Raphson iteration scheme based on the consistent algorithmic linearization of the governing equations, for both time stepping schemes, we typically find convergence within five to six iterations during the upstroke phase, and within three to four iterations during all other phases of the cardiac cycle. Quadratic convergence of global Newton Raphson iteration is confirmed in Table 5.3, which documents representative residuals of the relative error during the five different phases of the cardiac cycle. The total run time of an entire cardiac cycle of $t=1000$ ms, discretized with 8000 time increments of $\Delta t = 0.125$ ms for the non-adaptive scheme, is 3845.74 s on a single core of an i7-950 3.06 GHz desktop with 4GB of memory. Figure 10.9, bottom right, demonstrates that the adaptive time stepping scheme automatically increases the time step size during the plateau phase, between $t=50$ ms and $t=275$ ms, and during the resting phase, after $t=350$ ms. This reduces the number of time increments to 492 and the overall

run time to 395.46 s. Remarkably, when both models use the same fixed time step, the overall run time of our ionic excitation model is only approximately twice as long as the run time of the two-parameter FitzHugh-Nagumo model [56, 134] for which all the information of the chemical problem is lumped into one single phenomenological recovery variable [72, 73].

5.7 Discussion

We have presented a novel finite element based algorithm for electrochemical phenomena in cardiac tissue and demonstrated its potential to simulate cardiac excitation in real patient-specific geometries. In contrast to existing finite difference schemes and collocation methods proposed in the literature, our novel framework is (i) unconditionally stable, (ii) efficient, (iii) highly modular, (iv) geometrically flexible, and (iv) easily expandable.

Unconditional stability is guaranteed by the use of an implicit backward Euler time integration procedure instead of previously proposed explicit time integration schemes [188]. As a result, our time integration procedure is extremely robust [?], in particular in combination with an incremental iterative Newton Raphson solution technique. A comparison of different time-discretization schemes, explicit, semi-implicit, and implicit, in the context of the phenomenological FitzHugh Nagumo model confirmed that implicit electrophysiological models allow for the largest time step size, however, at the prize of having to invert the system matrix at each iteration step of each time increment [53]. That is why many authors prefer to use of operator splitting [151, 175, ?] and semi-implicit schemes, in which the nonlinear reaction term is treated explicitly and the diffusion term is treated implicitly [?, ?].

Efficiency is significantly increased with regard to existing explicit schemes, since we propose a global-local split which only introduces a single global degree of freedom at each finite element node, while all the other state variables are updated locally on the integration point level. In contrast to previous finite element models for electrophysiology, which discretize all unknowns at the finite element nodes [?, 130], we adopt the classical finite element infrastructure of internal variables, which has

proven extremely efficient in materially nonlinear continuum mechanics. Following this well-established approach borrowed from nonlinear mechanics, we solve the nonlinear system of equations by means of two nested Newton-Raphson iterations using an existing finite element framework [?], rather than using an inexact Newton or Krylov subspace method as proposed in the literature [130, ?]. The use of an implicit time integration scheme, which enables larger time steps than existing explicit schemes [53], further enhances the computational efficiency of our algorithm. It allows us to use simple adaptive time stepping schemes which, in the case shown here, reduce the computational time by more than one order of magnitude.

Modularity originates from the particular discretization scheme that treats all unknowns except for the membrane potential as local internal variables on the integration point level. This particular discretization adopts the classical infrastructure of nonlinear finite element programs in continuum mechanics and allows us to recycle a finite element program that was originally designed for structural analysis [?]. Accordingly, the proposed algorithm could be readily integrated into commercial finite element packages by reinterpreting any scalar-valued field, e.g., the temperature field, as the electrical potential field. Algorithmic modifications are restricted exclusively to the constitutive subroutine which would then solve the chemical problem and store the ion concentrations and gating variables as internal variables at each integration point. Moreover, this modular treatment of the chemical problem enables the straightforward combination of different cell models for pacemaker cells, atrial cells, epicardial ventricular cells, and endocardial ventricular cells, allowing for a fully inhomogeneous description of the underlying cardiac microstructure [1].

Geometrical flexibility is the most advantageous feature of finite element techniques when compared to existing finite volume methods or finite difference schemes [131]. Unlike existing schemes which are most powerful on regular grids [24], the proposed finite element based electrochemical model can be applied to arbitrary geometries with arbitrary initial and boundary conditions. It is easily applicable to medical-image based patient-specific geometries [146, ?, ?, 203] as demonstrated in the present manuscript. Here, we have demonstrated the geometric flexibility for a relatively coarse mesh of the heart, which allows us to prototype solutions on single desktop or

laptop computers. We are currently investigating the potential of our algorithm when analyzing finer discretizations of the heart. To this end, we adopt the recent parallel version of FEAP [?], which is a modification of the serial version [?], to interface to the PETSc library system available from Argonne National Laboratories.

Ease of expandability is probably the most crucial advantage of our algorithm. Being finite element based and modular in nature, our approach lays the groundwork for a robust and stable whole heart model of excitation-contraction coupling [66]. Through a straightforward generalization, the proposed excitation algorithm can easily be coupled to cardiac contraction through the additional incorporation of the mechanical deformation field [27, 185]. Also, the incorporation of an additional scalar-valued global unknown for the extracellular potential field is relatively simple, and allows us to extend the proposed formulation to a bidomain model [126, 163, 162, ?]. We have recently undertaken first steps in this direction and have shown that a fully implicit finite element formulation of the bidomain model with two degrees of freedom per finite element node is straightforward within the proposed algorithmic framework [44]. Rather than using phenomenological excitation models which we have successfully applied in the past [72, 73], we are now utilizing a more sophisticated ionic excitation model based on observable phenomena on the molecular scale. Although, from an engineering point of view, the number of material parameters required to characterize all individual ion channel activities might seem tremendous, the parameters of this model are related to well-defined electrochemically observable phenomena. The parameter values are extremely well characterized through a huge body of literature on single-cell and single-channel recordings performed within the past decade, see [68, 116, 149, 188] and the references cited therein. The use of ionic models will allow us, in the future, to elucidate possible arrhythmogenic phenomena on the molecular and cellular levels. For example, we can now explore the correlation between an enhanced activity of the sodium calcium exchanger I_{NaCa} , the reduced activity of the sodium potassium pump I_{NaK} , and prolonged action potential durations typically observable in failing hearts [149]. Along the same lines, we have recently modified our ionic cell model to incorporate a light-activated ion channel, channelrhodopsin, that allows us to activate cardiac cells by photostimulation [1]. This novel technology,

which is known as optogenetics, has gained a tremendous popularity in neuroscience within the past decade, and is currently being adopted in cardiology as well. There is hope that optogenetics could be applied to pace the heart with light, and that ionic computational models might provide further insight into optimal pacing parameters. Using the global finite element approach, we will be able to elaborate correlations between alterations in local action potential profiles and global electric activity through the computation of patient-specific electrocardiograms [22, 103].

Historically, the electrical excitation problem has been solved with finite difference schemes at a high spatial and temporal resolution. After several electrical time steps, the electric potential is mapped onto a coarse grid to solve the mechanical problem with finite element methods, to then map the resulting deformation back to the smaller grid [135]. Unfortunately, spatial mapping errors and temporal energy blow up are inherent to this type of solution procedure. We are currently working on a fully coupled monolithic solution of the electro-chemo-mechanical problem for ionic-excitation-contraction coupling that simultaneously solves for the electrical potential, the chemical ion concentrations, and the mechanical deformation in a unique, robust, and efficient way [66, 65], see also [34, ?]. This framework will allow us to better explain, predict, and prevent rhythm disturbances in the heart. This would have a tremendous potential in the design of novel treatment strategies such as biventricular pacing to prevent heart failure and sudden cardiac death.

Table 5.1: Algorithmic treatment of electrochemical coupling in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the chemical unknowns, the two sets of gating variables g_{gate}^I and g_{gate}^{II} and the ion concentrations c_{ion} are introduced locally on the integration point level.

initialize nodal degrees of freedom ϕ_J initialize internal variables $g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}}$
global Newton iteration
loop over all elements
loop over all integration points
update first set of gating variables $g_{\text{gate}}^I \leftarrow g_{\text{gate}}^I + f_{\text{gate}}^I \Delta t$ (5.25, 5.26, 5.27, 5.31, 5.32, 5.33, 5.34, 5.35, 5.38, 5.39)
initialize second set of gating variables $g_{\text{gate}}^{II} \leftarrow g_{\text{gate}}^{II} + f_{\text{gate}}^{II} \Delta t$ (5.30, 5.40, 5.41)
initialize ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}})$ (5.24, 5.29, 5.37, 5.43)
local Newton iteration
calculate ion concentration residuals R_{ion}^c and local iteration matrix $[K_{\text{ion ion}}^c] = d_{c_{\text{ion}}} R_{\text{ion}}^c$ (5.44, 5.45)
update ion concentrations $c_{\text{ion}} \leftarrow c_{\text{ion}} - [K_{\text{ion ion}}^c]^{-1} R_{\text{ion}}^c$
update second set of gating variables $g_{\text{gate}}^{II} \leftarrow g_{\text{gate}}^{II} + f_{\text{gate}}^{II} \Delta t$ (5.30, 5.40, 5.41)
update ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}^I, g_{\text{gate}}^{II}, c_{\text{ion}})$ (5.24, 5.29, 5.37, 5.43)
calculate source term $f^\phi(I_{\text{crt}})$ and its linearization $d_{\phi} f^\phi$ (5.46, 5.47)
calculate element residuals $R_I^{\phi e}$ and element matrices $K_{IJ}^{\phi e} = d_{\phi_J} R_I^{\phi e}$ (8.25, 8.26)
calculate global residual R_I^ϕ and global iteration matrix $K_{IJ}^\phi = d_{\phi_J} R_I^\phi$ (8.25, 8.26)
update membrane potential $\phi_J \leftarrow \phi_J - K_{IJ}^{\phi -1} R_I^\phi$

Table 5.2: Material parameters of the proposed human ventricular cardiomyocyte model [68, 116, 149, 188].

	Sodium related	Potassium related	Calcium related	Calcium ^{sr} related
Concentrations	$c_{Na} = 140$ mM	$c_{K} = 5.4$ mM	$c_{Ca} = 2$ mM	–
Maximum currents	$I_{Na}^{max} = 1000$ pA/pF $I_{NaCa}^{max} = 1.362$ pA/pF	$I_{K}^{max} = 1000$ pA/pF $I_{KATP}^{max} = 1.362$ pA/pF	$I_{CaL}^{max} = 0.08$ s ⁻¹ $I_{CaT}^{max} = 0.425$ mM/s $I_{CaSR}^{max} = 8.232$ mM/s	$I_{CaSR}^{max} = 0.08$ s ⁻¹ $I_{CaT}^{max} = 0.425$ mM/s $I_{CaSR}^{max} = 8.232$ mM/s
Maximum conductances	$G_{CaNa}^{max} = 14.838$ nS/pF $G_{NaCa}^{max} = 0.00029$ nS/pF	$G_{K}^{max} = 5.405$ nS/pF $G_{KATP}^{max} = 0.0096$ nS/pF $G_{KATP1}^{max} = 0.245$ nS/pF $G_{KATP2}^{max} = 0.245$ nS/pF $G_{K1M}^{max} = 0.062$ nS/pF $G_{K2M}^{max} = 0.0146$ nS/pF $G_{KATP3}^{max} = 0.294$ nS/pF $G_{KATP4}^{max} = 0.073$ nS/pF $G_{K1M}^{max} = 0.294$ nS/pF	$G_{CaL}^{max} = 0.175$ mm ³ /[μF] $G_{CaT}^{max} = 0.000592$ nS/pF $G_{CaSR}^{max} = 0.825$ nS/pF	
Half saturation constants	$c_{CaNa} = 1.38$ mM $c_{NaCa} = 87.50$ mM $c_{Ca} = 1.00$ mM $c_{CaSR} = 40.00$ mM	$c_{K} = 1.00$ mM $c_{KATP} = 40.00$ mM	$c_{CaNa} = 1.38$ mM $c_{NaCa} = 87.50$ mM $c_{Ca} = 0.0005$ mM $c_{CaSR} = 0.00025$ mM $c_{CaT} = 0.25$ mM $c_{CaSR} = 0.001$ mM	$c_{CaSR} = 0.00025$ mM $c_{CaT} = 0.25$ mM $c_{CaSR} = 0.3$ mM
Other parameters	$I_{NaCa}^{int} = 0.10$ $\gamma_{NaCa} = 2.50$ $\gamma = 0.35$	$p_{KATP} = 0.03$	$\gamma_{Ca} = 2$ $c_{Ca} = 0.15$ mM	$\gamma_{CaSR} = 2$ $c_{CaSR} = 10$ mM
Gas constant $R = 8.3143$ J K ⁻¹ mol ⁻¹ Faraday constant $F = 96.4867$ C/mmol		Temperature $T = 310$ K Membrane capacitance $C = 185$ pF		Cytoplasmic volume $V = 16,404$ μm ³ Sarcoplasmic reticulum volume $V^{sr} = 1094$ μm ³

Table 5.3: Algorithmic performance. Characteristic quadratic convergence of global Newton Raphson iteration, illustrated in terms of the representative residuals of the relative error during five different phases of the cardiac cycle.

	phase 0 upstroke [28.5 ms]	phase 1 early repolarization [50 ms]	phase 2 plateau [150 ms]	phase 3 final repolarization [250 ms]	phase 4 resting state [800 ms]
iteration 1	1.0000E+00	1.0000E+00	1.0000E+00	1.0000E+00	1.0000E+00
iteration 2	2.2133E-01	5.1643E-06	3.2777E-06	7.3801E-06	3.9747E-08
iteration 3	4.2979E-03	1.5336E-15	1.0049E-15	1.0146E-15	2.7173E-11
iteration 4	7.7168E-06	–	–	–	–
iteration 5	8.2932E-11	–	–	–	–

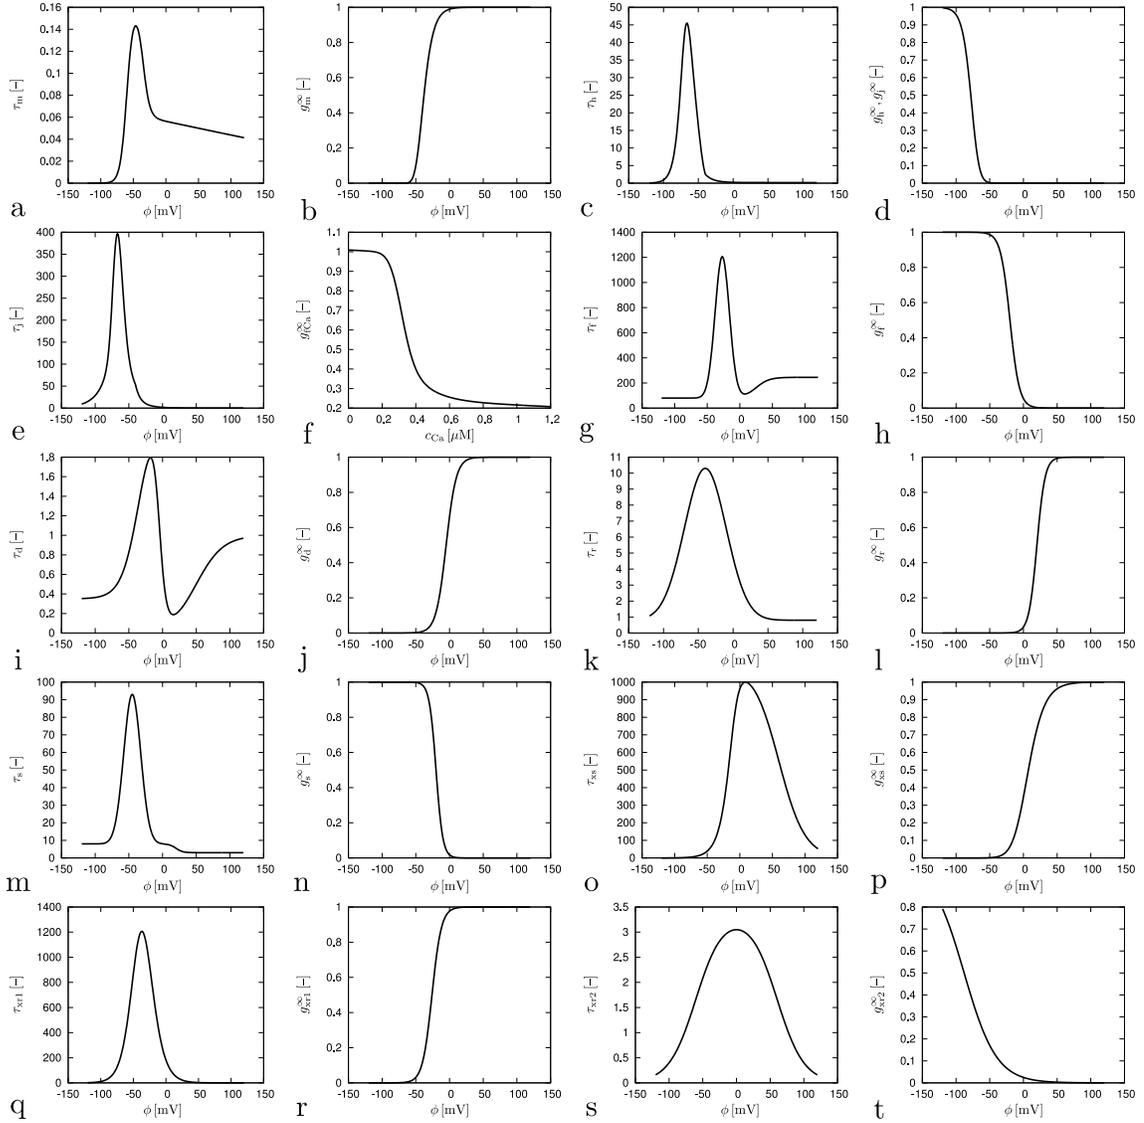


Figure 5.3: Electrochemistry in a human ventricular cardiomyocyte. Voltage dependent evolution of time constants τ_m, τ_h, τ_j and steady state values $g_m^\infty, g_h^\infty, g_j^\infty$ for sodium activation and fast and slow sodium inactivation gates g_m, g_h, g_j . Calcium concentration dependent evolution of steady state value g_{fCa}^∞ for intracellular calcium dependent calcium inactivation gate g_{fCa} . Voltage dependency evolution of time constants τ_d, τ_f and steady state values g_d^∞, g_f^∞ for L-type calcium activation and inactivation gates g_d, g_f . Voltage dependency evolution of time constants τ_r, τ_s and steady state values g_r^∞, g_s^∞ for transient potassium outward activation and inactivation gates g_r, g_s . Voltage dependency evolution of time constants $\tau_{xs}, \tau_{xr1}, \tau_{xr2}$ and steady state values $g_{xs}^\infty, g_{xr1}^\infty, g_{xr2}^\infty$ for slow delayed rectifier gate and rapid delayed rectifier activation and inactivation gates g_{xs}, g_{xr1}, g_{xr2} .

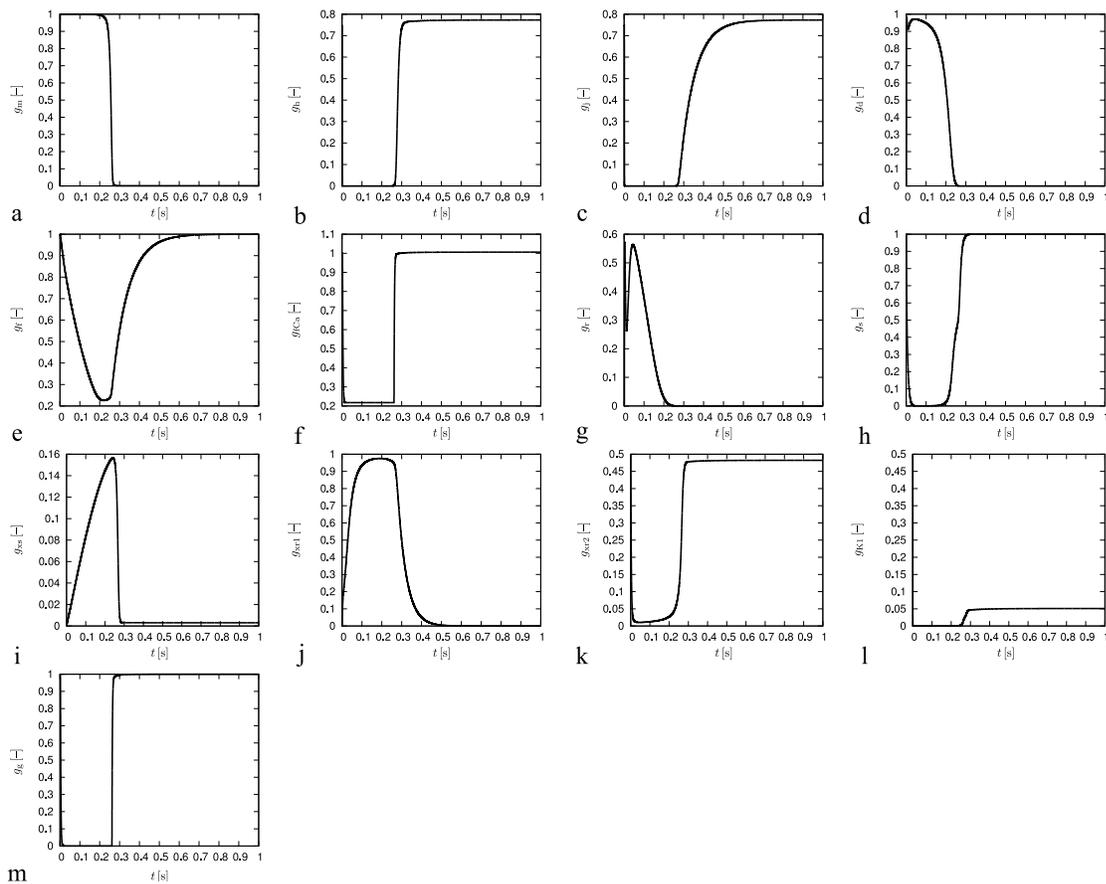


Figure 5.4: Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of sodium activation gate g_m , fast sodium inactivation gate g_h , slow sodium inactivation gate g_j , L-type calcium activation gate g_d , L-type calcium inactivation gate g_f , intracellular calcium dependent calcium inactivation gate g_{fCa} transient outward activation gate g_r , transient outward inactivation gate g_s , slow delayed rectifier gate g_{xs} , rapid delayed rectifier activation gate g_{xr1} , rapid delayed rectifier inactivation gate g_{xr2} , inward rectification factor g_{K1}^∞ , and calcium-dependent inactivation gate g_g .

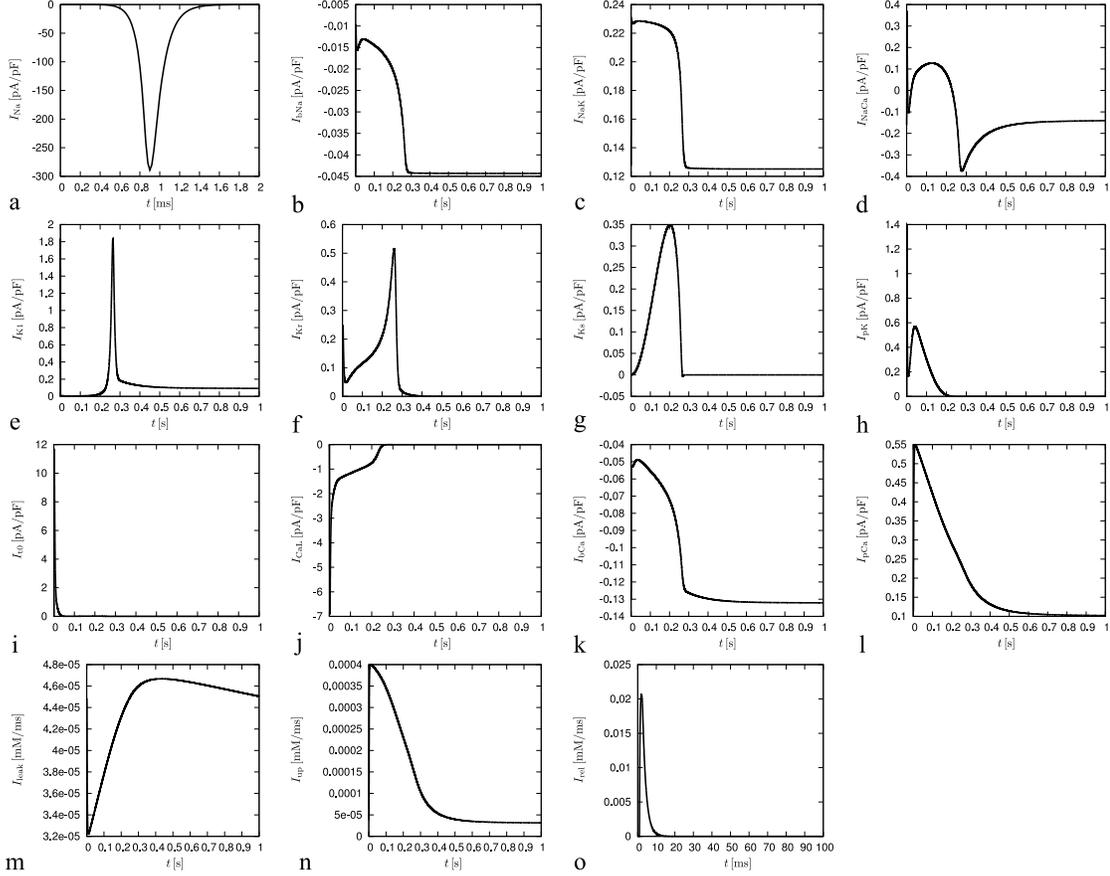


Figure 5.5: Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of the fast sodium current I_{Na} , the background sodium current I_{bNa} , the sodium potassium pump current I_{NaK} , and the sodium calcium exchanger current I_{NaCa} , the inward rectifier current I_{K1} , the rapid delayed rectifier current I_{Kr} , the slow delayed rectifier current I_{Ks} , the plateau potassium current I_{pK} , the transient outward current I_{t0} , the L-type calcium current I_{CaL} , the background calcium current I_{bCa} , the plateau calcium current I_{pCa} , the leakage current I_{leak} , the sarcoplasmic reticulum uptake current I_{up} , and the sarcoplasmic reticulum release current I_{rel} .

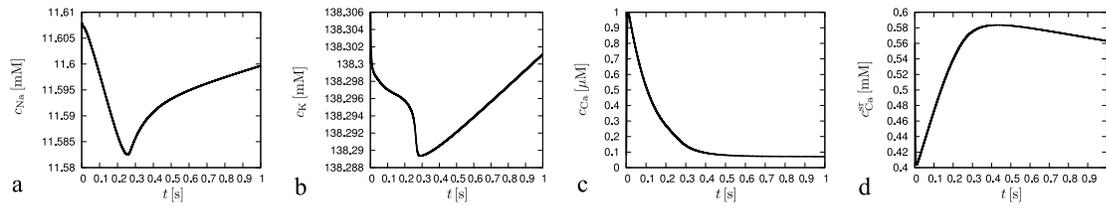


Figure 5.6: Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of intracellular sodium concentration c_{Na} , potassium concentration c_{K} , calcium concentration c_{Ca} , and calcium concentration in the sarcomplastic reticulum $c_{\text{Ca}}^{\text{sr}}$. The sodium concentration increases rapidly from 11.60 mM to 11.61 mM within the first 5 ms to initiate the fast upstroke of the action potential which then, in turn, affects the voltage-gated calcium and potassium channels. Accordingly, the calcium concentration increases quickly to $1.0 \mu\text{M}$ and then decreases gradually back to its resting value of $0.08 \mu\text{M}$. The potassium concentration decreases slowly to 138.29 mM until the beginning of the resting phase at after 0.28s and then gradually returns back to its initial value of 138.30 mM. In this last phase, the sodium concentration which had decreased to 11.585 mM increases gradually to its initial value of 11.60 mM.

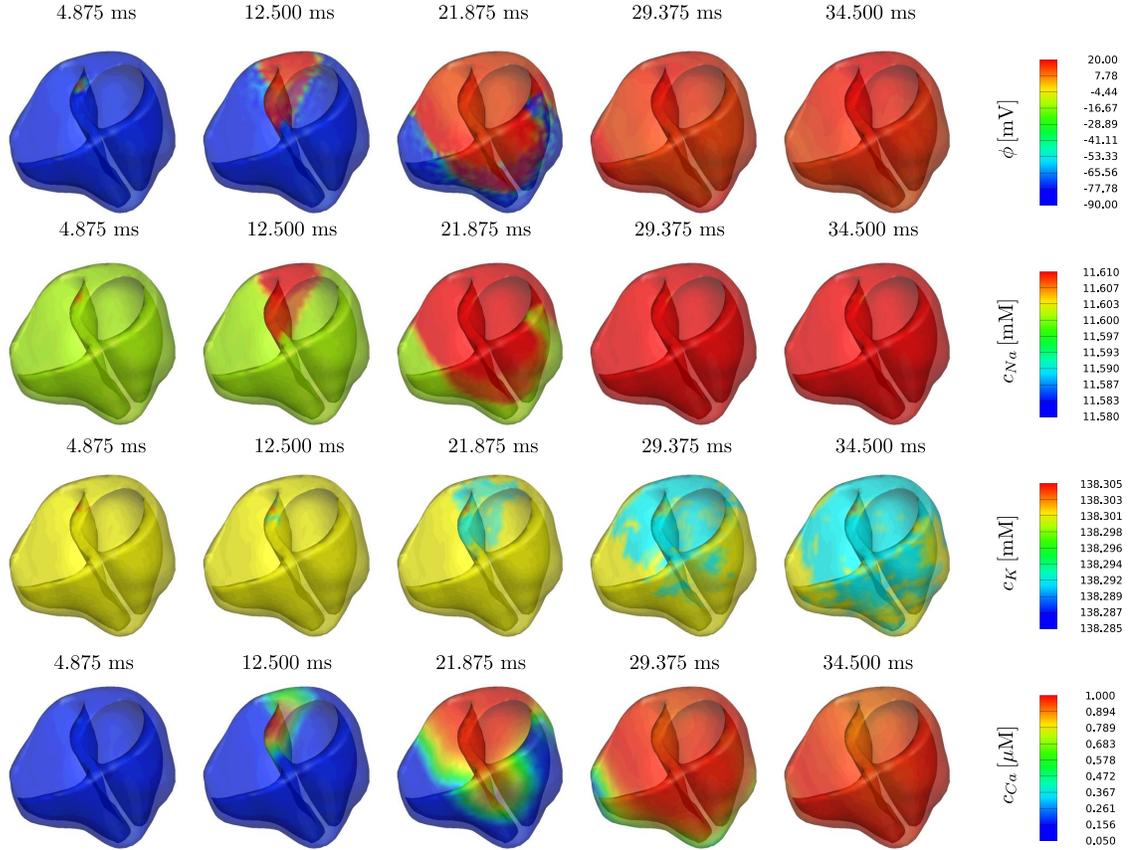


Figure 5.7: Electrochemistry in the human heart. Spatio-temporal evolution of the membrane potential ϕ and the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} during the depolarization phase of the cardiac cycle. Depolarization is initiated through an increase in the intracellular sodium concentration c_{Na} which reflects itself in the rapid depolarization of the cell characterized through an increase in the membrane potential ϕ from -86 mV to $+20$ mV. This affects the voltage-gated potassium and calcium channels and initiates a decrease in the intracellular potassium concentration c_{K} and an increase in the intracellular calcium concentration c_{Ca} . After approximately 30 ms, both ventricles of the heart are fully depolarized.

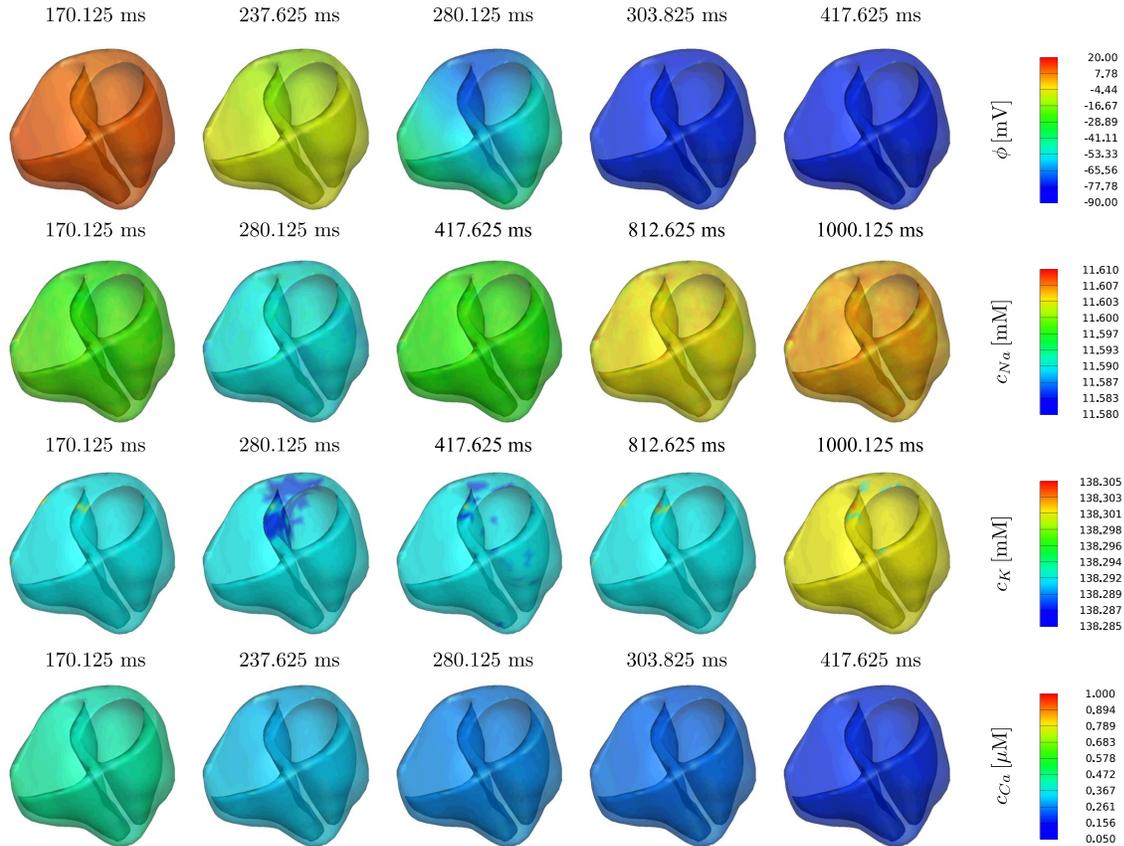


Figure 5.8: Electrochemistry in the human heart. Spatio-temporal evolution of the membrane potential ϕ and the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_{K} , and c_{Ca} during the repolarization phase of the cardiac cycle. Repolarization is characterized through a smooth decrease in the membrane potential ϕ from its excited value of +20 mV back to its resting value of -86 mV. At the same time, the intracellular calcium concentration c_{Ca} decreases smoothly to its resting value. Both sodium c_{Na} and potassium c_{K} respond more slowly and reach minimum concentrations only after 280 ms before increasing gradually back to their initial values at the end of the cycle after 1000 ms.

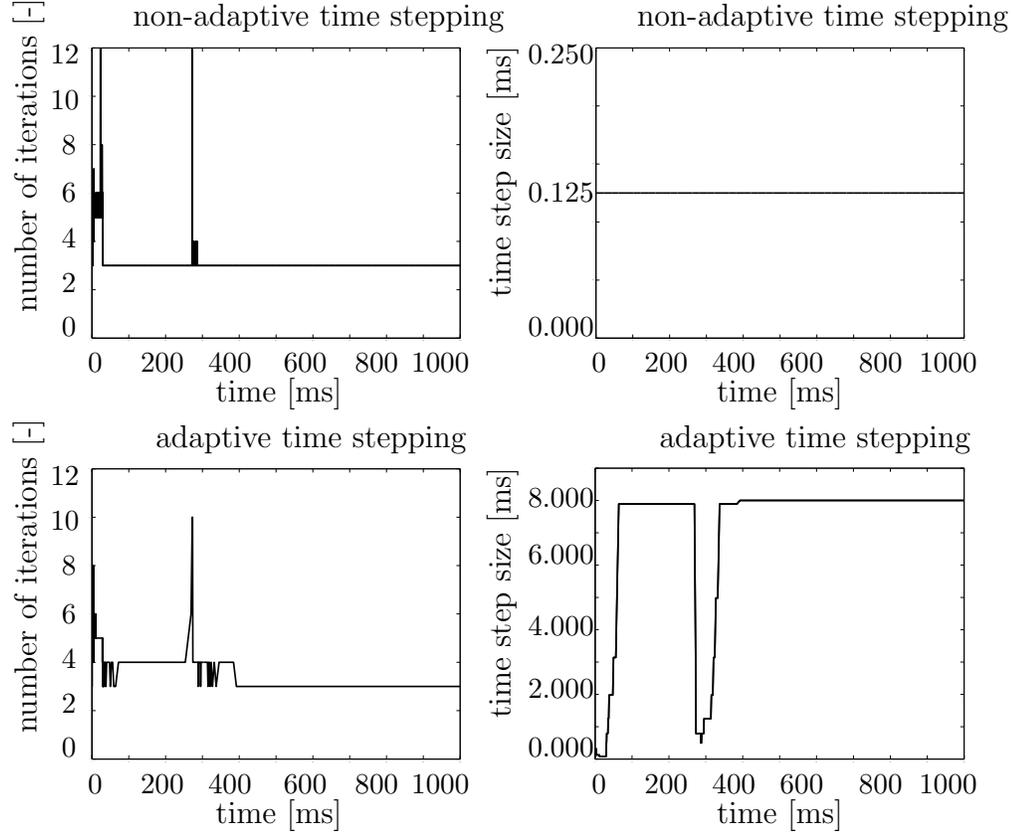


Figure 5.9: Algorithmic performance. Number of iterations and time step size for non-adaptive and adaptive time stepping schemes. For both algorithms, we typically find convergence within five to six Newton Raphson iterations during the upstroke phase, and within three to four iterations during all other phases of the cardiac cycle. This results in a total run time of 3845.74 s for the non-adaptive scheme with a fixed time step size of $\Delta t = 0.125$ ms and 8000 time increments throughout the cardiac cycle of $t=1000$ ms, calculated on a single core of an i7-950 3.06 GHz desktop with 4GB of memory. The adaptive time stepping scheme automatically increases the time step size during the plateau phase, between $t=50$ ms and $t=275$ ms, and during the resting phase, after $t=350$ ms. Adaptive time stepping with a maximum time step size of $\Delta t^{\max} = 8.0$ ms reduces the number of increments to 492, and the overall run time to 395.46 s.

Chapter 6

Applications for cell culture experiments

Cardiac arrhythmias are disturbances of the electrical conduction pattern in the heart with severe clinical implications. The damage of existing cells or the transplantation of foreign cells may disturb functional conduction pathways and increase the risk of arrhythmias. While these conduction disturbances are easily accessible with the human eye, there is no algorithmic method to extract quantitative features that quickly portray the conduction pattern. Here, we show that co-occurrence analysis, a well-established method for feature recognition in texture analysis, provides insightful quantitative information about the uniformity and the homogeneity of an excitation wave. As a first proof-of-principle, we illustrate the potential of co-occurrence analysis by means of conduction patterns of cardiomyocyte-fibroblast co-cultures, generated both *in vitro* and *in silico*. To characterize signal propagation *in vitro*, we perform a conduction analysis of co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts using microelectrode arrays. To characterize signal propagation *in silico*, we establish a conduction analysis of co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts using the finite element method. Our results demonstrate that co-occurrence analysis is a powerful tool to create purity-conduction relationships and to quickly quantify conduction patterns in terms of co-occurrence energy and contrast. We anticipate this first preliminary study to be

a starting point for more sophisticated analyses of different co-culture systems. In particular in view of stem cell therapies, we expect co-occurrence analysis to provide valuable quantitative insight into the integration of foreign cells into a functional host system.

6.1 Introduction

Cardiac arrhythmias are common heart disorders that can lead to severe complications, making them a significant area of research [113]. In the healthy heart, rhythmic contraction is initiated by the generation of an action potential by a group of pacemaker cells, resulting in the depolarization of their neighboring cells, and, consequently, in the formation of an excitation wave [102]. This process provides the basis for the unimpeded smooth propagation of depolarization across the heart [100]. Some cardiac conditions such as ischemia [91, 94] or infarction [184] can lead to changes in tissue structure, associated with the formation of non-conductive fibrous tissue [101]. Non-functional regions may vary in size and shape, but even small areas of damage can lead to severe disruption of the conduction system and initiate cardiac arrhythmias [170].

An emerging cause of arrhythmias is cell injection therapy [32, 193]. This relatively new technique to restore functional tissue within a damaged heart [144] may introduce complications in conduction, since the electrical properties of the newly added cells may not match those of the host tissue [1, 36]. Heterogeneous cell populations, resulting either from impure stem cell derived populations or from diffuse integration within the host, may further increase the risk of arrhythmias rather than improve regeneration and repair [?, 30]. For cell injection therapies to be successful, it is therefore essential to understand the relationship between tissue homogeneity and smooth signal propagation [97]. We hypothesize that a minimum volume fraction of electrically active cells must be present to support the propagation of a uniform, smooth depolarization wave.

To investigate purity-conduction relationships in electrically active tissue, various different in vitro models [60, 29] and in silico models [104, 204] have been proposed to

complement in vivo studies in small and large animals [92, 97]. In contrast to in vivo models, in vitro and in silico models allow us to combine conductive and non-conductive cells under reproducible, well-controlled conditions [127, 158]. They can indicate trends of conduction patterns, and are widely accepted approaches to characterize the electrophysiology of pure cell populations and of co-cultures of different cell types [198]. In view of stem cell therapies, optical mapping of mixed cultures of neonatal rat ventricular myocytes and mesenchymal stem cells has demonstrated the merit of in vitro models for studying arrhythmogenesis [32]. While this study provides general first insight, it requires cytotoxic dyes, and is therefore unsuitable for long-term culture. More importantly, the demonstration of pro-arrhythmic potential presented in this study is not quantitative. In order to study the arrhythmogenic risk associated with stem cell transplantation, there is a clear need for quantitative metrics that assess conduction patterns in a repeatable, efficient, and comparable fashion.

The objective of this study is to establish an easily reproducible method to generate purity-conduction relationships for co-culture systems with different volume fractions. To illustrate the features of this method, we analyze conduction patterns of cardiomyocyte-fibroblast co-cultures, generated both in vitro and in silico. For the in vitro conduction analysis, we study co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts using microelectrode arrays (MEA). For the in silico analysis, we explore co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts using the finite element method (FEM). To quantify conduction patterns and compare the results of the in vitro and in silico analyses, we establish a novel method, the co-occurrence analysis of local activation times (LAT). We illustrate how to create co-occurrence matrices for co-cultures with systematically varying cardiomyocyte-to-fibroblast ratios. To quickly portray the pattern of a depolarization wave, we extract two scalar-valued quantities from the co-occurrence matrix: co-occurrence energy and co-occurrence contrast. In a first proof-of-principle study, we find that both values indicate a significant disturbance of smooth, uniform conduction at cardiomyocyte-to-fibroblast ratios of 80:20. At ratios below the critical threshold of 50:50, we observe a complete loss of the electrical signal.

6.2 Methods

6.2.1 In vitro conduction analysis using MEAs

To characterize the action potential propagation in cardiomyocyte-fibroblast co-cultures in vitro, we acquire activation sequences using microelectrode array recordings.

Cell culture

We use murine atrial tumorigenic HL-1 cardiomyocytes as an electrically active, conductive host system [39]. We culture HL-1 cardiomyocytes in Claycomb media (Sigma-Aldrich, St. Louis, MO), supplemented with 10% fetal bovine serum (Hyclone, Logan, UT), 100 mM norepinephrine (Sigma-Aldrich), 100 units/ml penicillin streptomycin (Invitrogen, Carlsbad, CA), and 4 mM L glutamine (Invitrogen). HL-1 cardiomyocytes undergo mitosis approximately every 18 hours. When cultured to confluency, they are capable of beating spontaneously at a rate of 60-180 beats per minute. To block conduction pathways, we use murine 3T3 fibroblasts as a non-conductive cell type. Our fibroblast media consists of a solution containing 89% DMEM (Invitrogen), 10% FBS (Hyclone), and 1% penicillin streptomycin (Invitrogen). To establish quantitative purity-conduction relationships, we create 21 different co-cultures with a cardiomyocyte-to-fibroblast ratio varying from 100:0 to 0:100 in increments of 5. For each cardiomyocyte-to-fibroblast ratio, we analyze $n=3$ samples. To confirm different cardiomyocytes-to-fibroblast ratios and validate the co-culture method, we transfect our HL-1 cardiomyocytes with a green fluorescent protein (GFP) marker through the lentiviral vector LentiLox PLL3.7 [159]. We plate the samples on a 35 mm Petri dish adhered to the circuit board with our microelectrode array using bio-compatible epoxy (EP42HT, Master Bond; Hackensack, NJ).

Microelectrode array

To record electrical signal propagation, we utilize 36 microelectrodes arrayed in a 6×6 square on a glass substrate [195]. Our microelectrodes are $22 \mu\text{m}$ in diameter and spaced $100 \mu\text{m}$ apart, covering a total area of $500 \times 500 \mu\text{m}^2$ as shown in Figures

10.1A and 10.1B. This setup is capable to monitor the electrical activity of various

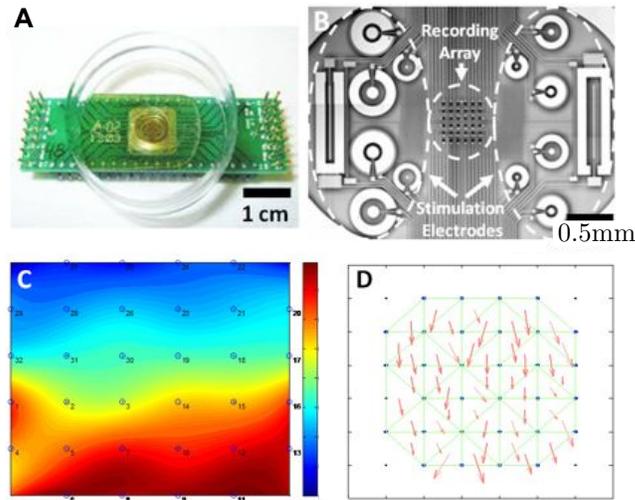


Figure 6.1: Microelectrode array analysis of in vitro activation patterns. **A.** The array of microelectrodes is wire-bonded to a circuit board carrier. A 35 mm Petri dish with an open center is adhered to the circuit board using bio-compatible epoxy. **B.** A 6×6 array of microelectrodes is fabricated on a glass substrate for electrical recording. The electrical propagation pattern can be represented through lateral isochrones maps or as through velocity vector maps. **C.** The lateral isochrones map is an interpolation of the time delay between electrodes, which is displayed by varying shades of color, overlaid on a representation of the 6×6 electrodes. The depolarization wave is initiated in the blue region at an activation time = 0 s, and propagates towards the red region, activation time = 0.04s. **D.** The velocity vector map is calculated by grouping the recording electrodes by three, and solving for the direction and magnitude of the time delay.

different cell types in real time over the course of several days [196]. We acquire MEA data from 32 channels, excluding the four corner electrodes. Using a custom recording system that consisted of a 32 channel amplifier with a two-stage gain of 60 dB, 7 Hz 1st-order high-pass cutoff, and 3 kHz 8th order low-pass cutoff, we process the signals as previously reported [62]. We digitize the analog signals from the amplifier board with 16-bit resolution at 10 kbps using a custom-designed visualization and extraction tool written in Matlab (The MathWorks; Natick, MA) [194].

In vitro conduction analysis

Since each detection electrode of the array is located at a known spatial position, we can easily extract conduction properties within the culture by studying the time delay of action potentials between the different electrodes. We characterize the propagation speed and the direction of propagation by recording the electrical signal at all detection electrodes and analyzing the spatio-temporal variation of the recordings. As a characteristic metric for action potential initiation, we introduce the local activation time (LAT) which we define as the maximum negative slope of the first derivative of the recorded signal.

The propagation of electrical activity is commonly visualized in two ways, lateral isochrones maps [107] and velocity vector maps [36]. To create a lateral isochrones map, we first normalize the 6×6 LATs by subtracting the minimum activation time from each recording. Then, we generate the characteristic activation plot, i.e., the 41×41 pixel image of conduction, by performing a cubic spline interpolation between the LATs. To create a lateral isochrones map, we fit the values of the activation plot to a color-coded contour map. Figure 10.1C displays a representative lateral isochrones map illustrating the initial depolarization in the blue region, and the gradual propagation of the depolarization wave towards the red region.

To create a velocity vector map, we group the detection electrodes in groups of three. For each electrode triplet, we calculate the magnitude and the direction of propagation following a method previously described [11]. Figure 10.1D shows a representative velocity vector map of the electrical activity within our co-culture system.

6.2.2 In silico conduction analysis using FEM

To characterize the action potential propagation in a cardiomyocyte-fibroblast co-culture in silico, we create a finite element model of the microelectrode array system described in Section 6.2.1. In particular, we utilize the classical FitzHugh-Nagumo equations for electroactive cells [56, 134], enhanced by the Aliev-Pafilov modifications [4, 72] for cardiomyocytes.

Mathematical model

Mathematically, we model the excitation of electrically active cells through the action potential ϕ in terms of the flux term, $\text{div } \mathbf{q}$, and the source term, f^ϕ .

$$\dot{\phi} = \text{div } \mathbf{q}(\phi) + f^\phi(\phi, r) \quad (6.1)$$

To characterize the global action potential propagation, we use a phenomenological model for the flux vector \mathbf{q} ,

$$\mathbf{q} = \mathbf{D} \cdot \nabla \phi \quad (6.2)$$

where $\mathbf{D} = d^{\text{iso}} \mathbf{I} + d^{\text{ani}} \mathbf{n} \otimes \mathbf{n}$ is a second order diffusion tensor. \mathbf{D} can be either purely isotropic, with $d^{\text{ani}} = 0$, or possess anisotropic contributions, with $d^{\text{ani}} \neq 0$ [93], along pronounced directions \mathbf{n} [179]. Although in vivo, cardiac conduction is highly anisotropic with faster signal propagation along the cell's long axis, it is common to model randomly oriented in vitro systems like ours as macroscopically isotropic. To study anisotropic conduction in vitro, cells are typically cultured on micropatterned cover slips with pre-defined orientation [156], a method that has recently revealed anisotropy ratios of $d^{\text{ani}} : d^{\text{iso}} = 5.6$ [26]. To characterize the local action potential profile, we introduce the following source term f^ϕ ,

$$f^\phi = c \phi [\phi - \alpha][1 - \phi] - r \phi \quad (6.3)$$

which consists of a cubic polynomial in terms of the action potential, $c \phi [\phi - \alpha][1 - \phi]$, and a coupling term introducing the recovery variable r . Here, we choose the scaling parameter to $c = 8.0$ and the oscillation threshold to $\alpha = 0.05$. While it is common to assume a spatial propagation of the action potential itself, the evolution equation for the recovery variable is usually assumed to be strictly local.

$$\dot{r} = f^r(\phi, r) \quad (6.4)$$

It is governed solely through the source term f^r ,

$$f^r = [\gamma + r \bar{\gamma}(\phi)] [-r - c \phi [\phi - b - 1]] \quad (6.5)$$

in terms of the parameters $\gamma = 0.002$ and $b = 0.150$, and the additional weighting factor $\gamma + r \bar{\gamma}(\phi)$ with $\bar{\gamma}(\phi) = \mu_1 / [\mu_2 + \phi]$. This term allows us to phenomenologically tune the restitution curve through the parameters $\mu_1 = 0.02$ and $\mu_2 = 0.03$ [4].

Computational model

Computationally, we discretize the coupled spatio-temporal system of equations (8.15) and (6.4) with finite differences in time and with finite elements in space [?, 197]. We introduce the action potential ϕ as global degree of freedom at each finite element node, while we evaluate the recovery variable r locally on the integration point level [73, 102].

Globally, we transform the non-linear excitation problem (8.15) into its residual format $\mathbf{R}^\phi = \dot{\phi} - \text{div}(\mathbf{q}) - f^\phi \doteq 0$ in \mathcal{B} , which we complement by the corresponding Dirichlet and Neumann boundary conditions $\phi = \bar{\phi}$ on $\partial\mathcal{B}_\phi$ and $\mathbf{q} \cdot \mathbf{n} = \bar{q}$ on $\partial\mathcal{B}_q$. Here, on the entire boundary $\partial\mathcal{B}$, we assume homogeneous Neumann boundary conditions, $\mathbf{q} \cdot \mathbf{n} = 0$. We obtain the weak form of the residual \mathbf{R}^ϕ by an integration over the domain \mathcal{B} , the standard intergration by parts, and the inclusion of the Neumann boundary conditions. For the spatial discretization, we discretize the domain of interest \mathcal{B} with n_{el} finite elements \mathcal{B}_e as $\mathcal{B} = \bigcup_{e=1}^{n_{\text{el}}} \mathcal{B}^e$ and apply the standard isoparametric concept to interpolate the trial and test functions ϕ and $\delta\phi$.

$$\delta\phi = \sum_{i=1}^{n_{\text{el}}} N^i \delta\phi_i \quad \phi = \sum_{j=1}^{n_{\text{el}}} N^j \phi_j \quad (6.6)$$

For the linear quadrilateral elements used in the sequel, N are the standard bi-linear shape functions and $i, j = 1, \dots, n_{\text{el}}$ are the $n_{\text{el}} = 4$ element nodes. For the temporal discretization, we partition the time interval of interest \mathcal{T} into n_{stp} subintervals $[t_n, t_{n+1}]$ as $\mathcal{T} = \bigcup_{n=0}^{n_{\text{stp}}-1} [t_n, t_{n+1}]$ and apply a standard backward Euler time integration scheme in combination with a finite difference approximation of the first order

time derivative $\dot{\phi}$.

$$\dot{\phi} = [\phi - \phi_n] / \Delta t \quad (6.7)$$

Here and from now on, we omit the index $(\circ)_{n+1}$ for the sake of brevity, and use the common abbreviation $\Delta t := t - t_n > 0$ for the current time increment. With the discretizations in space (11.6) and time (8.24), we can introduce the discrete algorithmic residual,

$$\mathbf{R}_I^\phi = \mathbf{A} \int N^i \frac{\phi - \phi_n}{\Delta t} + \nabla N^i \cdot \mathbf{q} - N^i f^\phi \, dV \doteq 0 \quad (6.8)$$

where the operator \mathbf{A} symbolizes the assembly of all local element residuals at the element nodes i to the global residual at the global nodes I . To solve the discrete system of non-linear equations (6.8), we apply an incremental iterative Newton Raphson solution technique. It is based on the consistent linearization of the residual $\mathbf{K}_{IJ}^\phi = \partial_{\phi_J} \mathbf{R}_I^\phi$, introducing the global iteration matrix,

$$\mathbf{K}_{IJ}^\phi = \mathbf{A} \int_{\mathcal{B}^e} N^i \frac{1}{\Delta t} N^j + \nabla N^i \cdot \mathbf{D} \cdot \nabla N^j - N^i d_\phi f^\phi N^j \, dV \quad (6.9)$$

which defines the update of the global vector of unknowns $\phi_I \leftarrow \phi_I - \sum_J \mathbf{K}_{IJ}^{\phi^{-1}} \mathbf{R}_J^\phi$ at all global nodes I .

Locally, to discretize the recovery variable r in time, we apply a standard finite difference approximation.

$$\dot{r} = [r - r_n] / \Delta t \quad (6.10)$$

Using an implicit Euler backward scheme, we rephrase the discrete residual statement of the recovery equation (6.4) in the following form.

$$\mathbf{R}^r = r - r_n - [[\gamma + r \bar{\gamma}] [-r - c \phi [\phi - b - 1]]] \Delta t \doteq 0 \quad (6.11)$$

Its consistent linearization $\mathbf{K}^r = \partial_r \mathbf{R}^r$ with

$$\mathbf{K}^r = 1 + [\gamma + \bar{\gamma} [2r + c \phi [\phi - b - 1]]] \Delta t \quad (6.12)$$

defines the iteration scheme for the incremental update of the recovery variable $r \leftarrow r - \mathbf{K}^{-1} \mathbf{R}^r$ on the integration point level. At local equilibrium, we compute the source term for $f^\phi = c\phi[\phi - \alpha][1 - \phi] - r\phi$ the global residual of the excitation problem (6.8) and its linearization $d_\phi f^\phi = \partial_\phi f^\phi + \partial_r f^\phi d_\phi r$ for the global Newton iteration (10.19).

In silico conduction analysis

Motivated by the in vitro model of Section 6.2.1, we create an in silico finite element model of the co-culture system of cardiomyocytes and fibroblasts. We represent the culture system through a two-dimensional flat sheet discretized with 1225 bi-linear finite elements of $20 \mu\text{m} \times 20 \mu\text{m}$ each, covering a total area of $700 \times 700 \mu\text{m}^2$. We model cardiomyocytes as electrically active and conductive as described in Sections 6.2.2 and 6.2.2. For the sake of simplicity, we model fibroblasts as non-conductive, keeping in mind that, in reality, they might be actively engaged in the conduction system, however, at a much lower inherent conductivity than cardiomyocytes [198]. To identify the effect of cell morphology on signal propagation, we conduct two sets of simulations. For the first set, labeled as 1:1 aspect ratio, we assume that cardiomyocytes and fibroblasts are of equal size. For the second set, labeled as a 1:2 aspect ratio, we assume that fibroblasts are twice as long as cardiomyocytes.

To establish quantitative purity-conduction relationships, we create 21 different co-culture models with a cardiomyocyte-to-fibroblast ratio varying from 100:0 to 0:100 in increments of 5. For each cardiomyocyte-to-fibroblast ratio, we analyze $n=5$ realizations. Similar to a technique reported in the literature [204], for each realization, we assign each of the 1225 elements a randomly selected cell type, either cardiomyocyte or fibroblast.

To initiate signal propagation throughout the culture system, we select one random cardiomyocyte within the culture and stimulate it with an external electrical signal. In analogy to the in vitro setup, we record the electrical signal at 6×6 virtual recording sites spaced $100 \mu\text{m}$ apart, covering an area of $500 \times 500 \mu\text{m}^2$ located in the center of our sample. Similar to the in vitro conduction analysis described in Section 6.2.1, we then generate the characteristic activation plot, i.e., the 41×41 pixel image of

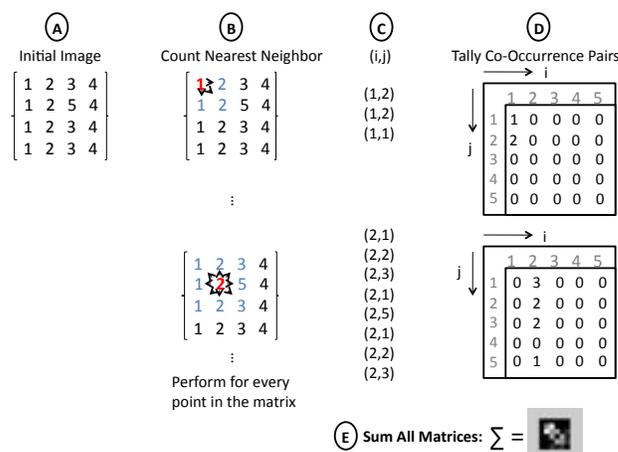


Figure 6.2: Algorithm to calculate the co-occurrence matrix based on counting the number of incidents that each local activation time occurs in an activation plot. Within a given activation plot, **A**, each entry is paired with its nearest neighbor, **B**, and noted, **C**. The values within each pair correspond to the coordinates of the co-occurrence matrix, and the number of occurrences is tallied, **D**. Steps **B-D** are repeated for each entry in the original matrix in **A** such that the final co-occurrence matrix is the sum of all these matrices.

conduction, by performing a cubic spline interpolation between the virtually recorded 6×6 LATs.

6.2.3 Co-occurrence matrix

In Section 6.2.1 we have illustrated how to extract lateral isochrones maps and velocity vector maps from our MEA recordings. Although both mappings provide insight into the qualitative nature of the overall conduction pattern, neither of them is capable to characterize disturbed conduction in heterogeneous co-culture systems through quickly accessible scalar-valued metrics. Here, we suggest a novel evaluation technique, co-occurrence analysis, to quantify the uniformity and homogeneity of an electrical conduction pattern [76]. Originally designed for feature recognition in texture analysis [95], co-occurrence analysis can be adapted to post-process activation plots obtained from MEA recordings or FEM simulations. In our case, the underlying activation plot consists of 41×41 discrete values, generated from an interpolation of

the in vitro and in silico recorded 6×6 LATs. By its definition, the corresponding co-occurrence matrix

$$C_{ij} = \sum_{r=1}^n \sum_{s=1}^m \begin{cases} 1 & \text{if } I_{rs} = i \text{ and } I_{r\pm 1, s\pm 1} = j \\ 0 & \text{otherwise} \end{cases} \quad (6.13)$$

is a correlation matrix, which quantifies how often two neighboring entries occur in the $n \times m$ activation plot I_{rs} , one with a LAT of i , the other one with a LAT of j . For example, the co-occurrence matrix component $C_{i=2, j=1} = 8$ indicates that a LAT of $i = 2$ is found next to a LAT of $j = 1$ a total of 8 times in the discrete activation plot.

Since the co-occurrence matrix is a transform based on LATs, its dimension $i \times j$ depends on the number of different activation times found within the activation plot. When constructing the co-occurrence matrix, we can define its sensitivity based on the distance between pairs of LATs [76]. In equation (6.13), we have chosen a distance of ± 1 to capture all pairs of nearest neighbors, as illustrated in Figure 10.2. Since LAT pairs do not account for directionality, the co-occurrence matrix is always square and always symmetric, $C_{ij} = C_{ji}$. In addition, the total number of co-occurrence pairs within activation plots of the same size is equivalent, here it is always 41×41 . We therefore normalize the co-occurrence matrix by dividing each entry by the total number of occurrence pairs, here 1681, thus making the co-occurrence matrix a collection of probabilities of occurrences. Throughout this manuscript, we use Matlab (The Mathworks, Natick, MA) to calculate the co-occurrence matrices for the in vitro and in silico generated activation plots from Sections 6.2.1 and 6.2.2, respectively. The underlying algorithm is illustrated in Figure 10.2.

Figure 10.3 displays representative activation plots and the corresponding co-occurrence matrices of five characteristic examples: linear gradient, sine wave, criss cross pattern, spiral wave, and random noise, from left to right. It is important to remember that the population of the co-occurrence matrix is not representative of the geometric excitation pattern itself. Rather, it represents the local distribution of pixel pairs. Uniform conduction patterns, e.g., the linear gradient pattern, exhibit an activation plot that changes steadily as the depolarization wave propagates. This

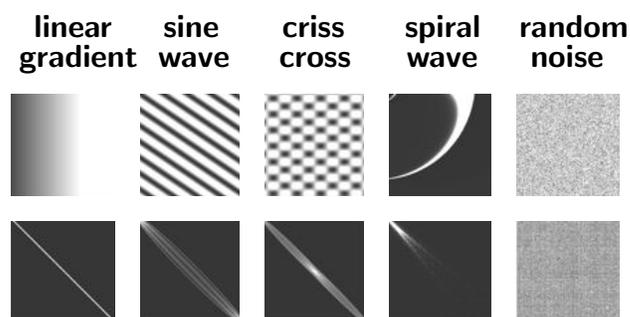


Figure 6.3: Examples of characteristic activation plots (top) and their co-occurrence matrices (bottom). Activation plots consist of 300×300 pixels and exhibit grayscale intensities from 0 to 255. Smooth images are characterized through co-occurrence matrices that are clustered around the diagonal (left). The more the image is disordered, the farther away from the diagonal the co-occurrence matrix is populated (right).

level of regularity creates a co-occurrence matrix with values heavily concentrated along the diagonal. Non-uniform conduction patterns, e.g., the spiral wave pattern, exhibit large grayscale differences between neighboring nodes, yielding a co-occurrence matrix with values farther away from the diagonal. The co-occurrence matrix can further indicate whether the depolarization wave is evenly propagated. A disproportionate number of early activation times results in a cluster of occurrences in the upper left corner, e.g., for the spiral wave pattern, while a disproportionate number of late activation times results in a cluster of occurrences in the lower right corner. The co-occurrence matrix is therefore a measure of uniformity and homogeneity.

6.2.4 Co-occurrence energy and contrast

Two characteristic features of the co-occurrence matrix, co-occurrence energy and co-occurrence contrast, are capable to quickly portray the status of a depolarization wave [77, 205]. The energy of a co-occurrence matrix, which is sometimes also referred to as the angular second moment,

$$\text{energy} = \sum_{i=1}^n \sum_{j=1}^m C_{ij}^2 \quad (6.14)$$

is the sum of the squares of each matrix entry, here evaluated for an $n \times m$ co-occurrence matrix C_{ij} . In our case, $n = m$ is the number of different activation times found within the activation plot. The co-occurrence energy is a measure for the *uniformity* of the co-occurrence matrix [45]. Accordingly, it is lowest when all entries are equal. While the energy is a useful metric to access whether there are many particular occurrences, it fails to quantify whether values are clustered along the diagonal or not. The contrast of a co-occurrence matrix

$$\text{contrast} = \sum_{i=1}^n \sum_{j=1}^m [i - j]^2 C_{ij} \quad (6.15)$$

is a metric that characterizes whether a large proportion of co-occurrence pairs is located along the diagonal. This is achieved by assigning a larger weight, $[i - j]^2$, to entries farther away from the diagonal. The co-occurrence contrast is a measure for the *homogeneity* of the spatial distribution of local features [45]. Accordingly, larger contrast values indicate a more heterogeneous conduction.

Table 6.1: Examples of characteristic activation plots from Figure 10.3 and their corresponding co-occurrence energy and contrast. The energy is highest for images with sharp fronts like the spiral wave and lowest for uniform images like the random noise. The contrast is highest for images with large spatial variations like the random noise and lowest for homogeneous images like the linear gradient.

	linear gradient	sine wave	criss cross	spiral wave	random noise
energy	893.5	2978.1	595.9	134210	8.6
contrast	0.5	155.5	780.4	240.3	7806.8

Table 6.1 illustrates the co-occurrence energy and contrast for the characteristic activation plots shown in Figure 10.3. The energy is highest for images with sharp fronts like the spiral wave. It is lowest for uniform images like the random noise. The contrast is highest for images with large spatial variations like the random noise. It is lowest for homogeneous images like the linear gradient. In the following section,

we illustrate the potential of the co-occurrence matrix and its energy and contrast to quantify conduction patterns of heterogeneous cardiomyocyte-fibroblast co-cultures for both the *in vitro* and *in silico* co-culture models created in Sections 6.2.1 and 6.2.2.

6.3 Results

6.3.1 *In vitro* conduction analysis using MEAs

Figure 10.4 shows a representative set of micrographs for varying cardiomyocyte-to-fibroblast ratios, here illustrated for 0:100, 60:40, 70:30, 80:20, 90:10, and 100:0. Cardiomyocytes and fibroblasts can be clearly distinguished by their characteristic morphologies. While murine HL-1 cardiomyocytes, dominant towards the lower right corner, take a more spherical shape, murine 3T3 fibroblasts, dominant towards the upper left corner, are characterized through a spindle-type morphology.

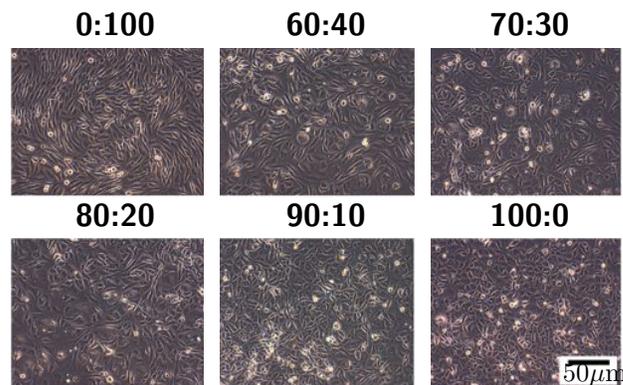


Figure 6.4: Optical micrographs of *in vitro* co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically increase the cardiomyocyte-to-fibroblast ratio from 0:100 to 100:0 in increments of 5. As the cardiomyocyte-to-fibroblast ratio increases, here shown for 0:100, 60:40, 70:30, 80:20, 90:10, and 100:0, the co-culture system becomes increasingly dominated by circular-shaped cardiomyocytes (bottom right) and less dominated by spindle-shaped fibroblasts (top left).

Figure 10.5 documents the fluorescence expression of co-cultures of murine HL-1 cardiomyocytes and murine 3T3 fibroblasts, where HL-1 cardiomyocytes are transfected with a GFP marker through a lentiviral vector. Qualitatively, fluorescent imaging with 450 nm excitation and 515 nm emission wavelengths reveals decreased GFP expression with decreasing cardiomyocyte-to-fibroblast ratios, here illustrated for ratios of 100:0, 90:10, 70:30, and 50:50. Although the pure population of cardiomyocytes is confluent, GFP expression is not ubiquitous throughout the culture of transfected cells. We therefore normalize the fluorescence intensity with respect to the pure cardiomyocyte population displayed in the upper left corner. Quantitatively, GFP expression decreases approximately linearly with decreasing cardiomyocyte-to-fibroblast ratios, shown in terms of averages and standard deviations for groups of $n=3$, in the bottom graph of Figure 10.5. Figure 10.6 summarizes the action potential profiles for co-cultures of different cardiomyocyte-to-fibroblast ratios for all 32 electrodes on the microelectrode arrays. When decreasing the cardiomyocyte-to-fibroblast ratio, here shown for 100:0, 80:20, 70:30, and 50:50, the smoothness of the conduction pattern clearly decreases. While the 100:0 plain cardiomyocyte culture, shown on the left, displays a homogeneous conduction pattern, the 80:20 and 70:30 co-cultures, shown in the middle, demonstrate an intermediate behavior between smooth conduction and isolated spontaneous beating. The 50:50 co-culture, shown on the right, displays spontaneous beating, indicated through isolated signals in specific channels, with no sign of homogenous conduction.

Figure 10.7 shows two representative velocity vector maps and the percent active area for co-cultures of murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. While pure 100:0 cultures of cardiomyocytes, shown on the left, support continuous almost unidirectional propagation throughout the system, heterogeneous 80:20 populations, shown on the right, exhibit a non-smooth conduction pattern impeded by non-conductive fibroblasts.

By systematically decreasing the cardiomyocyte-to-fibroblast ratio in increments of 5, and exploring $n=3$ groups for each ratio, we identify the 50:50 ratio to be as the lowest threshold necessary to ensure signal detection, as demonstrated in the bottom graph of Figure 10.7. In two out of three MEAs plated with this ratio, we record

action potentials in only two channels, whereas the third MEA in the group does not display any action potentials at all. Recorded beats are asynchronous between the different channels. This indicates that the 50:50 ratio, while sufficient for islets of activity of about 1mm^2 each, does not contain enough cardiomyocytes to support continuous conduction throughout the entire sample.

Beginning with a 70:30 cardiomyocyte-to-fibroblast ratio, samples display similar asynchronous patterns of spontaneous beating between 12 ± 2 channels. Activation patterns typically group into two beating patches of approximately 6mm^2 each. One of the 80:20 ratios demonstrates a small depolarization wavefront throughout the heterogeneous population, with one large synchronously beating patch that averages 20 ± 3 active electrodes, approximately corresponding to an area of 20mm^2 .

6.3.2 In silico conduction analysis using FEM

Figure 10.8 illustrates the lateral isochrones maps simulated for cardiomyocyte-to-fibroblast ratios of 100:0, 90:10, 70:30, 50:50, 30:70, and 10:90 using our finite element model. Similar to the lateral isochrones map in Figure 10.1C, early stage depolarization is shown in blue, while late depolarization is shown in red. When systematically decreasing the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5, and evaluating $n=5$ realizations of statistically random cell distributions for each set, we observe that continuous electrical conduction paths do not form below a critical cardiomyocyte-to-fibroblast threshold of 50:50. The top row of Figure 10.8 shows the results for the 1:1 aspect ratio simulation, for which cardiomyocytes and fibroblasts are of equal size. The bottom row illustrates the results for the 1:2 aspect ratio simulation, for which fibroblasts are twice as long as cardiomyocytes. The conduction pattern is slightly smoother for the equal aspect ratio, which is particularly apparent close to the critical threshold ratio of 50:50.

6.3.3 Co-occurrence matrix

Figure 10.9 illustrates representative activation plots, which indicate the local activation times as a depolarization wave travels from the white regions to the dark regions.

While pure cultures of cardiomyocytes, shown on the left, support gradual, smooth activation patterns, heterogeneous co-cultures, shown in the middle and on the right, exhibit non-smooth activation patterns, resulting in a more diffuse signal propagation. The bottom row of Figure 10.9 documents the corresponding co-occurrence matrices. With decreasing cardiomyocyte-to-fibroblast ratios, shown from left to right, the co-occurrence matrices become less focused along their diagonals and more populated in off-diagonal entries. This population of the co-occurrence matrices agrees nicely with the examples of characteristic co-occurrence matrices documented in Figure 10.3.

6.3.4 Co-occurrence energy and contrast

Figure 6.10 displays the co-occurrence energy and contrast for co-cultures of murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. For systems with a cardiomyocyte-to-fibroblast ratio of less than 80:20, we observe that less than half of the electrodes of the MEA record an electrical signal. Analyses of samples below this threshold are rather based on interpolated values than on actual data, and are therefore not included in the co-occurrence analysis. Accordingly, in Figure 6.10, we only show results above this critical threshold of 80:20. We display energy and contrast values for each co-culture group in terms of averages and standard deviations for groups of $n=3$, normalized to the average value of the 100:0 group. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating a severely disrupted electrical conduction system. Figure 6.10 demonstrates an increase of energy and contrast with a decreasing cardiomyocyte-to-fibroblast ratio. Since the energy is the sum of squared values of the whole co-occurrence matrix, it increases considerably in the presence of a large number of a particular pair of LATs. In this case, areas on the array that do not exhibit electrical activity result in activation plots with regions that are all the same value. A high energy value therefore indicates inconsistencies in a co-culture, here caused by the lack of conduction. The contrast increases steadily as the cardiomyocyte-to-fibroblast ratio decreases, but not as drastically as the energy. Quantifying the deviation from the diagonal of a co-occurrence matrix, large contrast values indicate increased disruption of conduction.

Figure 6.11 illustrates the co-occurrence energy and contrast extracted from our finite element simulations. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating a severely disrupted electrical conduction system. Simulations based on the square fibroblast model, 1:1 aspect ratio, do not exhibit a significant change in energy from the pure cardiomyocyte population down to a cardiomyocyte-to-fibroblast ratio of 45:55, while simulations based on the rectangular fibroblast model, 1:2 aspect ratio, exhibit this change at the 75:25 ratio. Significant changes in contrast begin at the 80:20 ratio for the square fibroblast model, while they already begin at the 95:5 ratio for rectangular fibroblast model. As the relative number of fibroblasts increases, the energy begins to increase due to large patches of non-conductive regions, which lead to large quantities of similar co-occurrence pairs. With increasing numbers of co-occurrence pairs, fewer variations are observed, which leads to the decrease in contrast. For square fibroblasts, this decrease takes place at ratios of 40:60, for rectangular fibroblasts, it takes places slightly earlier, at ratios of 60:40.

6.4 Discussion

In this study, we establish an easily reproducible method to generate purity-conduction relationships for co-culture systems with different volume fractions. To characterize conduction patterns in vitro, we utilize a custom-designed MEA platform to analyze conduction in co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. For our two-dimensional culture system, we observe that fibroblast concentrations greater than 10% severely disrupt the electrical conduction, as shown in the bottom graph of Figure 10.7. The increase in co-occurrence energy reported in Figure 6.10, top, which indicates a loss of uniformity, is in excellent agreement with this observation. The increase co-occurrence contrast reported in Figure 6.10, bottom, which indicates a loss of homogeneity, is in general agreement with these findings; yet the contrast increases more gradually and almost linearly. The population of the co-occurrence matrix shown in Figure 10.9 reveals that heterogeneous samples exhibit co-occurrence matrices that are not evenly populated along the diagonal. This

implies that there is an uneven distribution of time delays between the individual electrodes. This breakdown of smooth conduction is already visible for cardiomyocyte-to-fibroblast ratios of 90:10, illustrated in Figure 10.9, middle, and in Figure 6.10. While local electrical signals can be recorded at relatively low cardiomyocyte-to-fibroblast ratios, the global conduction pattern is extremely sensitive to the presence of small numbers of fibroblasts. This observation is in good qualitative agreement with findings reported in the literature [127]. Since this study is meant as a first proof-of-principle, we accept the inherent limitation that we only acquire signals at 6×6 recording sites, from which we interpolate a discrete 41×41 activation plot for the co-occurrence analysis. However, it remains to be shown, how the purity-conduction relationships of our two-dimensional model system scale up to three dimensions [101]. We anticipate that trends will be similar, while the critical threshold ratio to maintain smooth conduction might be lower in the three-dimensional in vivo setting [97].

To characterize conduction patterns in silico, we use finite element models of co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts. Although we could potentially record our electrical signal at a much higher spatial resolution, we mimic the in vitro analysis and establish 6×6 virtual recording sites, from which we interpolate the discrete 41×41 activation plot. Qualitatively, the observed trends of the FEM-based in silico conduction analysis displayed in Figure 10.8 correlate well with the MEA-based in vitro conduction analysis shown in Figure 10.9. Quantitatively, the first set of computations, shown in the top row of Figure 10.8, predicts a lower critical threshold to maintain smooth conduction using equal-sized cells. Compared to the critical 80:20 ratio from the in vitro conduction analysis, the first in silico analysis remains conductive well past the 50:50 ratio. We hypothesize that the equal-cell-size model is overly simplistic because it does not capture the morphological differences between spindle shaped 3T3 fibroblasts and round HL-1 cardiomyocytes. Accordingly, we refine the computational model and perform a second set of simulations now using elongated fibroblasts, shown in the bottom row of Figure 10.8. With this morphological refinement, the critical thresholds to maintain conduction increases to approximately 80:20 as illustrated in Figure 6.11, and agrees nicely with in vitro model. Despite these promising preliminary results,

our computational model has two major limitations. First, we model fibroblasts as non-conductive, although experimental studies suggest that fibroblasts might play a more important role in cardiac electrophysiology than simply acting as passive electrical insulators [60, 158]. To further elaborate this issue, we are currently analyzing the impact of slightly conducting fibroblasts in analogy to a recently proposed finite-difference based conductivity analysis [204]. Second, for the sake of simplicity, we have applied a pure monodomain formulation for the co-culture system [72]. We have recently developed a bidomain formulation for cardiac tissue [?], which we will utilize in the future to better adjust the computational readouts to the electrical recordings [174].

Co-occurrence analysis allows us to quantify and compare our *in vitro* and *in silico* created conduction patterns. Co-occurrence energy and contrast are known to quickly monitor uniformity and homogeneity of a texture [76]. Both features display significant differences between smooth conduction, *i.e.*, in cultures with 100:0 and 95:5 cardiomyocyte-to-fibroblast ratios, and disrupted conduction, *i.e.*, in cultures with lower cardiomyocyte ratios. Co-occurrence analysis has significant relevance towards understanding the arrhythmogenic risk of stem cell transplantation in the heart. This procedure remains controversial because the mechanisms behind the electrical integration of transplanted cells into the host system remain unclear [97]. Even the type of graft cells to transplant is still under considerable debate. In addition to using cells from a cardiac lineage, there is evidence that cells outside this lineage, such as mesenchymal stem cells [119] or skeletal myoblasts [124], may support cardiac tissue regeneration. The complexity of live models has made it difficult to truly understand how host tissue reacts electrically to transplanted cells. Therefore, further work is necessary to probe how the electrical environment affects stem cell integration. Until electrical integration takes place, transplanted cells may not play any functional role [?].

6.5 Conclusion

Various cardiac diseases and treatment options change the local volume fraction of conducting cells. This may alter existing conduction pathways and increase the risk of arrhythmias. Co-culture systems are illustrative model systems to systematically investigate the mismatch in conductivity of different cell types and its impact on arrhythmogenesis. From microelectrode array recordings, disturbances in conduction are easily detectable with the human eye; however, there is no automated method to quantify conduction patterns algorithmically in terms of easily accessible characteristic metrics. Here, we demonstrate that co-occurrence analysis, a method that was originally developed for feature recognition in texture analysis, can quickly portray the status of an excitation wave, and provide insightful information about conduction disturbances. To illustrate the potential of this technique, we calculate co-occurrence matrices, energy, and contrast using experimentally and computationally created conduction patterns of different cardiomyocyte-fibroblast co-cultures. We show that the population of the co-occurrence matrix is correlated to the smoothness of signal propagation. By systematically reducing the cardiomyocyte-to-fibroblast ratio, we demonstrate that co-occurrence energy is a measure for pattern uniformity, while co-occurrence contrast is an indicator for pattern homogeneity. Overall, co-occurrence analysis is an easily reproducible method to create purity-conduction relationships for co-culture systems with different volume fractions. In particular, in view of stem cell therapies, we expect co-occurrence analysis to provide valuable quantitative insight into the integration of foreign cells into a functional host system.

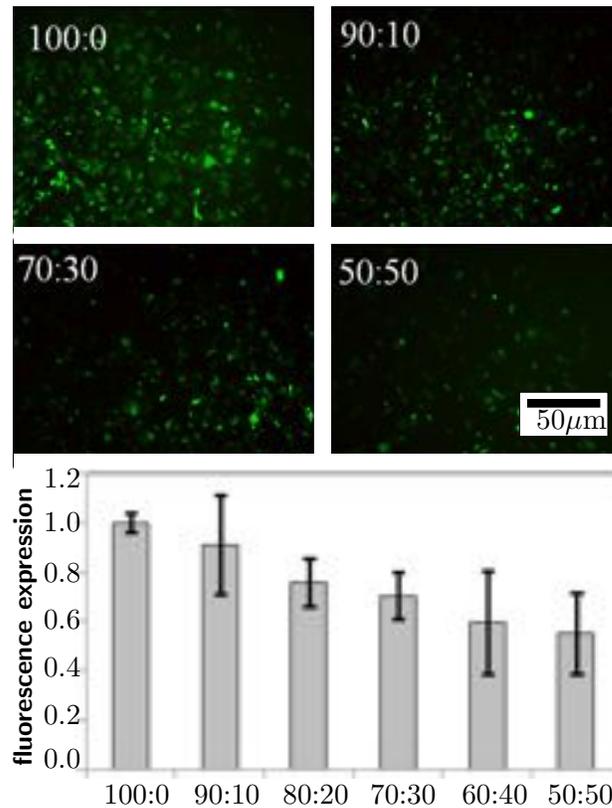


Figure 6.5: Fluorescence expression of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts, where HL-1 cardiomyocytes are transfected with green fluorescence protein (GFP) through a lentiviral vector. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Qualitatively, GFP expression decreases with decreasing cardiomyocyte-to-fibroblast ratio, here shown for 100:0, 90:10, 70:30, and 50:50 (top). Quantitatively, GFP expression decreases approximately linearly with decreasing cardiomyocyte-to-fibroblast ratios, shown in terms of averages and standard deviations for groups of $n=3$ (bottom).

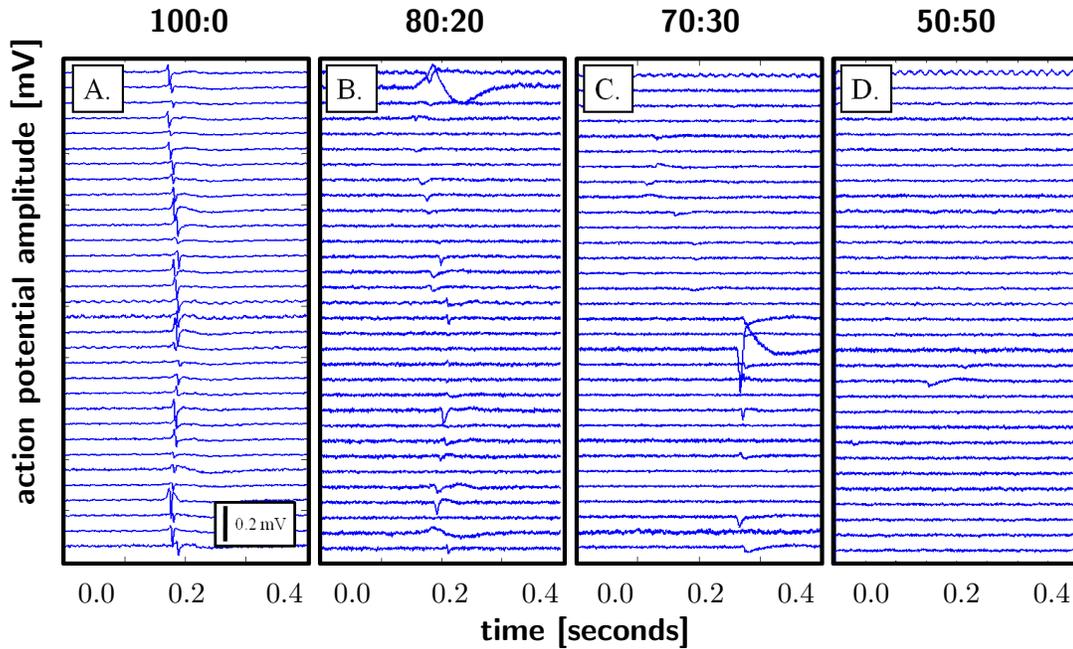


Figure 6.6: Action potential profiles of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts, displayed for all 32 electrodes on the microelectrode array. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. The smoothness of the conduction pattern decreases with decreasing cardiomyocyte-to-fibroblast ratio, here shown for 100:0, 80:20, 70:30, and 50:50. The 100:0 plain cardiomyocyte culture displays a homogeneous conduction pattern (left). The 80:20 and 70:30 co-cultures demonstrate an intermediate behavior between smooth conduction and isolated spontaneous beating (middle). The 50:50 co-culture displays spontaneous beating, indicated through isolated signals in specific channels, with no sign of homogenous conduction (right).

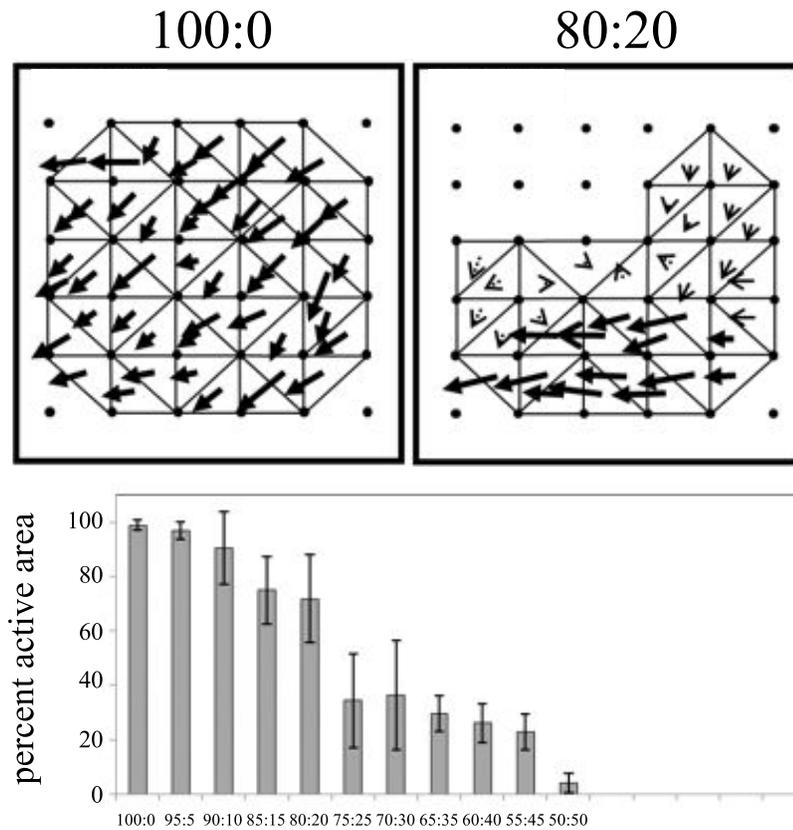


Figure 6.7: Velocity vector maps (top) and percent active area (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. The active area fraction decreases with decreasing cardiomyocyte-to-fibroblast ratio, as the electrical activity becomes less synchronized. Pure 100:0 cultures of cardiomyocytes (left) support continuous, almost unidirectional, propagation throughout the homogeneous population. Heterogeneous 80:20 populations (right) exhibit a non-smooth conduction pattern impeded by non-conductive fibroblasts. The active area fraction (bottom) reveals that action potentials can only be recorded above a critical 50:50 ratio threshold.

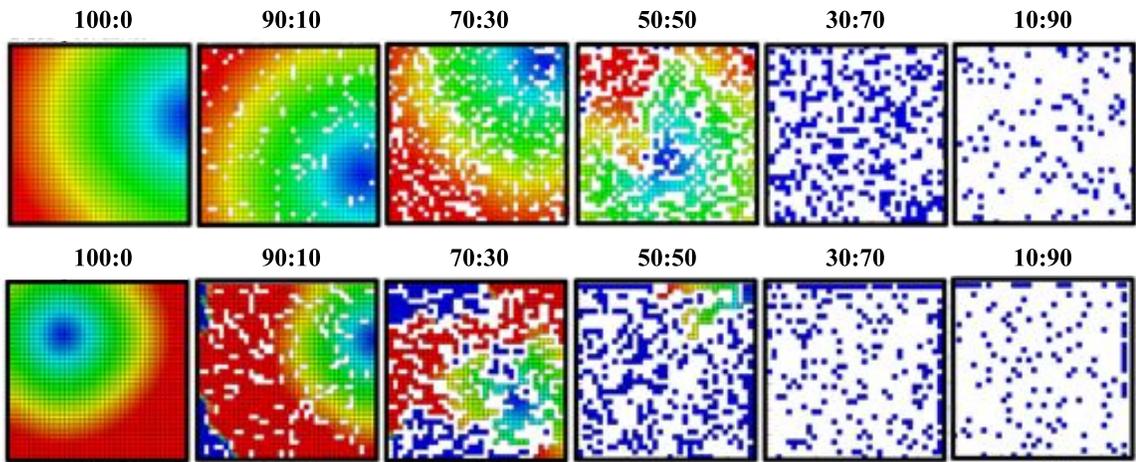


Figure 6.8: Lateral isochrones map of in silico co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts, simulated with our finite element model. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5, here shown for 100:0, 90:10, 70:30, 50:50, 30:70, and 10:90. For the first set of simulations, 1:1 aspect ratio (top row), we model cardiomyocytes and fibroblasts to be of equal size. For the second set of simulations, 1:2 aspect ratio (bottom row), we model fibroblasts twice as long as cardiomyocytes. Continuous electrical conduction paths do not form below a critical cardiomyocyte-to-fibroblast threshold of 50:50. Similar to the lateral isochrones map in Figure 10.1C, early stage depolarization is shown in blue, while late depolarization is shown in red.

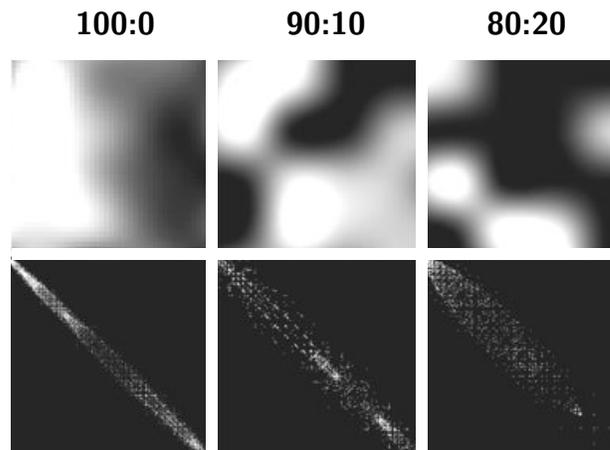


Figure 6.9: Isolated activation plots (top) and co-occurrence matrices (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Representative activation plots display the local activation times as a depolarization wave travels from the white regions to the dark regions. Pure 100:0 cardiomyocyte cultures (left) support gradual, smooth activation patterns, indicated through a single unique propagation front. More heterogeneous 90:10 (middle) and 80:20 (right) populations exhibit non-smooth activation patterns, resulting in a more diffuse signal propagation. With decreasing cardiomyocyte-to-fibroblast ratios, the co-occurrence matrices become less focused along their diagonals, and more populated in off-diagonal regions.

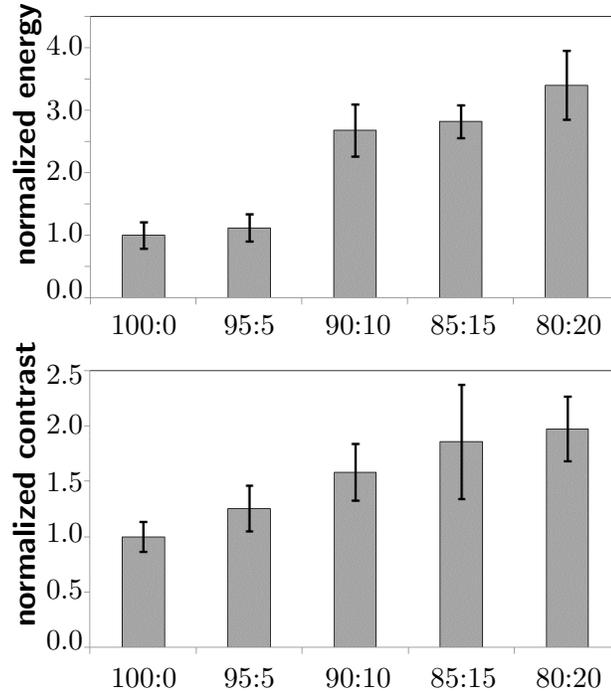


Figure 6.10: Co-occurrence energy (top) and co-occurrence contrast (bottom) of in vitro co-cultured murine HL-1 cardiomyocytes and murine 3T3 fibroblasts. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. We exclude results below the critical cardiomyocyte-to-fibroblast threshold of 80:20, for which less than half of the microelectrodes record an electrical signal. Energy and contrast values for each co-culture group are displayed in terms of averages and standard deviations for groups of $n=3$, normalized to the average value of the 100:0 group. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating the gradual disruption of the electrical conduction system.

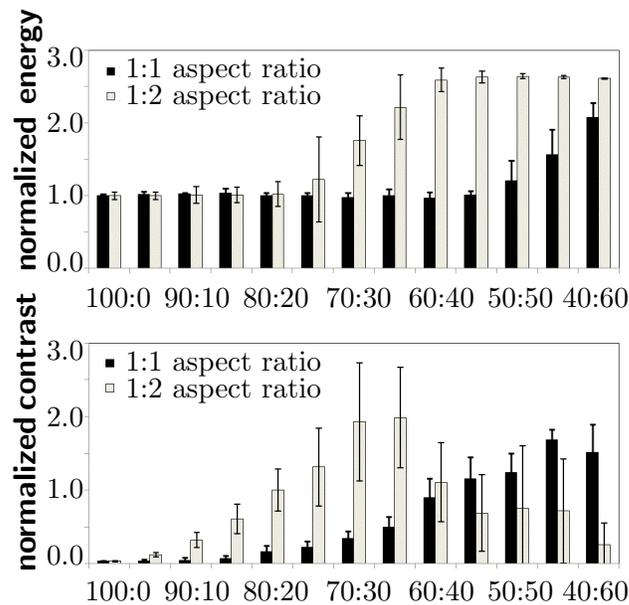


Figure 6.11: Co-occurrence energy (top) and co-occurrence contrast (bottom) of in silico co-cultured electrically active, conductive cardiomyocytes and non-conductive fibroblasts, simulated with our finite element model. We systematically decrease the cardiomyocyte-to-fibroblast ratio from 100:0 to 0:100 in increments of 5. Energy and contrast values for each co-culture group are displayed in terms of averages and standard deviations for groups of $n=5$, normalized to the average value of the 100:0 group. Both co-occurrence energy and contrast increase with decreasing cardiomyocyte-to-fibroblast ratios, indicating the gradual disruption of the electrical conduction system.

Chapter 7

Applications for optogenetically modified tissue cultures

The ability to stimulate mammalian cells with light has significantly changed our understanding of electrically excitable tissues in health and disease, paving the way towards various novel therapeutic applications. Here, we demonstrate the potential of optogenetic control in cardiac cells using a hybrid experimental/computational technique. Experimentally, we introduced channelrhodopsin-2 into undifferentiated human embryonic stem cells via a lentiviral vector, and sorted and expanded the genetically engineered cells. Via directed differentiation, we created channelrhodopsin-expressing cardiomyocytes, which we subjected to optical stimulation. To quantify the impact of photostimulation, we assessed electrical, biochemical, and mechanical signals using patch clamping, multielectrode array recordings, and video microscopy. Computationally, we introduced channelrhodopsin-2 into a classic autorhythmic cardiac cell model via an additional photocurrent governed by a light-sensitive gating variable. Upon optical stimulation, the channel opens and allows sodium ions to enter the cell, inducing a fast upstroke of the transmembrane potential. We calibrated the channelrhodopsin-expressing cell model using single action potential readings for different photostimulation amplitudes, pulse widths, and frequencies. To illustrate the potential of the proposed approach, we virtually injected channelrhodopsin-expressing cells into different locations of a human heart, and explored its activation sequences

upon optical stimulation. Our experimentally calibrated computational tool box allows us to virtually probe landscapes of process parameters, and identify optimal photostimulation sequences towards pacing hearts with light.

INTRODUCTION

For more than 40 years, biologists have studied microorganisms that produce proteins to directly regulate the flow of charged ions across their plasma membrane in response to light. The first identified protein of this kind was the light-gated ion pump bacteriorhodopsin, transporting positively charged hydrogen ions across the cell membrane [142]. While bacteriorhodopsin naturally acts as an on-switch for electrically active cells, the second reported light-gated ion pump, halorhodopsin, transports negatively charged chloride ions, thereby acting as an off-switch [120]. The first reported light-gated ionic channel illustrated in Figure 10.1, channelrhodopsin, was only identified thirty years later [132, 133], but has since then revolutionized neuroscience.

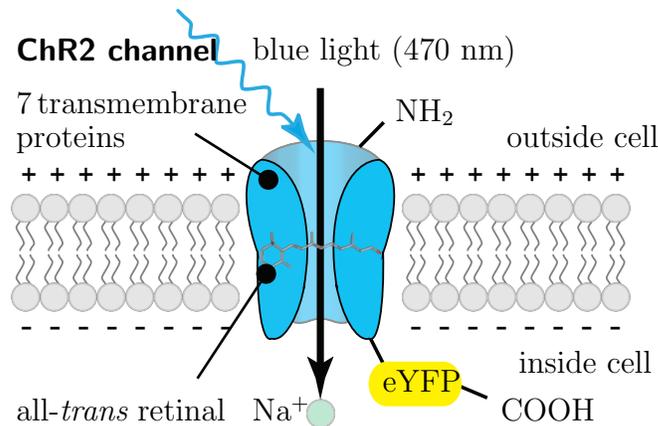


Figure 7.1: Channelrhodopsin-2 (ChR2) is a light-gated cation channel native to the green alga *Chlamydomonas reinhardtii*. It consists of seven transmembrane proteins and absorbs light through its interaction with retinal. Here, we induce channelrhodopsin coupled to yellow fluorescent protein (eYFP) into undifferentiated human embryonic stem cells via a lentiviral vector and differentiate these cells into cardiomyocytes.

Since the early nineties, we have known that phototaxis and photophobic responses

in the green alga *Chlamydomonas reinhardtii* are mediated by rhodopsins with a microbial-type all-*trans* retinal chromophore [78, 108]. The photochemical isomerization of this all-*trans* retinal to 13-*cis* retinal is illustrated in Figure 10.2. It occurs at peak absorption wavelengths of 470 nm, opening the channel to sodium, potassium, and calcium cations in response to blue light. In the dark, the covalently bound retinal spontaneously relaxes to all-*trans*, providing closure of the channel and regeneration of the chromophore.

A breakthrough enabling technology was reported in 2005, when the light-sensitive target was first introduced genetically using engineered viruses [23, 202], a technique that is now known as optogenetics [46, 80]. Since then, optical tools for controlling the electrical activity of neurons have rapidly evolved, and are now gaining widespread use in neuronal research and medicine [111, 176]. While initial applications of optogenetics have been restricted exclusively to the neuronal system, optogenetic tools have now advanced to a level of maturity, where they can confidently be applied to other cells and organs [46]. Natural first candidates of choice are stem cells, and electrically active glial cells, muscle cells, and cardiac cells [1, 5].

The objective of this study is to demonstrate the potential of optogenetic control of the cardiac system using a hybrid experimental/computational technique. We demonstrate that channelrhodopsin-2 (ChR2) can be expressed stably and safely in human embryonic stem cells (hESC), which can then be differentiated into cardiomyocytes (hESC^{ChR2}-CM). Upon photostimulation, ChR2 opens rapidly and allows sodium ions to enter the cell, inducing a defined transmembrane potential, commonly known as the action potential. We illustrate how ChR2 can be introduced into a computational autorhythmic cell model via an additional photocurrent governed by a light-sensitive gating variable to simulate this effect. The calibrated cell model is capable of reliably reproducing photostimulation amplitudes, pulse widths, and frequencies from single action potential readings. Using a custom designed finite element model, we virtually inject our calibrated model cells into different locations of a human heart to illustrate the potential of the proposed approach towards pacing hearts with light.

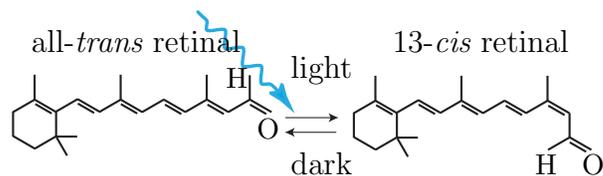


Figure 7.2: Channelrhodopsin-2 (ChR2) is activated by photoisomerization of all-*trans* retinal to 13-*cis* retinal at wavelengths of 470 nm. After photoisomerization, the covalently bound retinal spontaneously relaxes to all-*trans* in the dark, providing closure of the ion channel and regeneration of the chromophore.

MATERIALS AND METHODS

All experiments, methods, and protocols for this study were approved by the Stanford University Stem Cell Research Oversight (SCRO) committee.

Opsin sources and lentiviral vector

The channelrhodopsin-2 (ChR2) variant described here was optimized for mammalian expression by truncating the native sequence from 2241 bp to 933 bp, by changing the native histidine (H) codon (CAC) to the arginine (R) codon (CGC) coding for protein residue 134 of ChR2 (H134R), and by changing its gene's codon usage to conform to human codon usage distribution [67, 203]. The lentiviral vector pLenti-EF1 α -ChR2-eYFP-WPRE (pLECYT) was constructed as previously described [23, 203]. The pLET plasmid contains the ubiquitously expressed elongation factor 1- α (EF1 α), to obtain high levels of ChR2-eYFP expression in a mammalian system [1, 202]. All constructs have been fully sequenced previously for accuracy of cloning [23, 67]. High-titer lentivirus was produced using a second generation lentiviral system by co-transfection of 293FT cells (Invitrogen, Carlsbad, CA), the pLECYT viral vector described above, pCMVR Δ 8.74 (containing GAG and POL), pMD2.G (containing VSVg), and calcium phosphate as previously described [165, 203].

Stem cell culture and differentiation

Channelrhodopsin-expressing human embryonic stem cells (hESC^{ChR2+}) were grown as monolayers [1] on hESC-qualified Matrigel (BD Biosciences, San Jose, CA) and maintained

in the pluripotent state through daily feeding with mTeSR1 media (StemCell Technologies, Vancouver, Canada) [115], supplemented with 1x penicillin / streptomycin (Invitrogen, Carlsbad, CA). Cardiomyocyte differentiation was usually begun 2-5 days after initially seeding hESC^{ChR2+} on Matrigel (BD Biosciences). At this time, the cells were transferred to RPMI-1640 media supplemented with B27, 1x non-essential amino acids, 1x penicillin/streptomycin, and 0.1 mM β -mercaptoethanol (all Invitrogen) and our differentiation method was begun using aspects of other methods previously described [105, 200]. On the first day of differentiation, Day 0, RPMI media with 50 ng/mL of Activin A (R&D Systems, Minneapolis, MN) was added to each well. On the subsequent day, Day 1, 5 ng/mL of BMP-4 (R&D Systems) was added to each well. On Day 3, fresh RPMI media was added to each well and was replaced every 48 hours until Day 11, when the cells were transferred to a DMEM (Invitrogen) media supplemented with 5% FBS (Invitrogen), 1x non-essential amino acids, 1x penicillin / streptomycin, and 0.1 mM beta-mercaptoethanol. This DMEM media was then replaced approximately every 48 hours. Cardiomyocytes generally began spontaneously beating between Days 9 and 20.

Fluorescence-activated cell sorting (FACS)

Fluorescence-activated cell sorting (FACS) was performed with a BD FACSAria instrument (BD Biosciences) equipped with BD FACSDiva 6.0 software. Up to 1×10^6 cells transduced with the ChR2-eYFP lentivirus were sorted. Sorted cells were then resuspended in mTeSR1 media and replated on Matrigel-coated wells. After 2-3 days in culture, eYFP signal was confirmed via fluorescence microscopy. Analysis of FACS data was performed offline with FlowJo 7.6.1 software (Tree Star, Ashland, OR).

Polymerase chain reaction (PCR)

For undifferentiated hESC, PCR primers with gene product length for the following genes were used: GAPDH (152 bp), Oct-4 (169 bp), Nanog (154 bp), a region within eYFP (187 bp), a region spanning eYFP-ChR2 (197 bp), and a region within ChR2 (174 bp). Total RNA was isolated and RNA yield was then quantified using a Quant-iT kit (Invitrogen) and Qubit fluorometer (Invitrogen) per the manufacturer's instructions. For cDNA synthesis, 1 μ g total RNA, random hexamers, annealing buffer, 2x First-Strand Reaction Mix (Invitrogen) and SuperScript III/RNase OUT Enzyme Mix (Invitrogen), and RNAase/DNase-free

water were combined then incubated per the manufacturer instructions. For PCR amplification, AccuPrime Pfx SuperMix (Invitrogen), custom primers for pluripotency markers as described above, and cDNA were combined. Non-template control (NTC) reactions were prepared by substituting cDNA with distilled water. Samples were transferred to a thermal cycler and the following cycling program was used: i) initial denaturation at 95°C for 2 min; ii) 30 cycles of 95°C, 30 sec; 60°C, 30 sec; 68°C, 1 min; iii) final extension at 68°C for 5 min. Finally, PCR products, a 100 bp ladder (Invitrogen), and NTC were loaded in separate wells of a 2% agarose E-gel with SYBR-Safe (Invitrogen) and run for 30 min. Bands were then visualized with an E-gel iBase blue light transilluminator (Invitrogen).

Immunocytochemistry (ICC)

Human embryonic stem cell derived cardiomyocytes (hESC-CM) were labeled with primary antibodies for the cardiac markers α -actinin (Sigma, IgG, 1:500) and TnI (Millipore, IgG, 1:200). The secondary antibody used was goat anti-mouse IgG-Alexa 594 (Invitrogen, 1:1000) for both α -actinin and TnI. Cells were counterstained with DAPI (Sigma) for 10 minutes. An AxioObserver Z1 (Carl Zeiss, Göttingen, Germany) inverted microscope was used to visualize hESC-CMs. The Zeiss microscope was equipped with a Lambda DG-4 300 W Xenon light source (Sutter Instruments, Novato, CA), an ORCA-ER CCD camera (Hamamatsu, Bridgewater, NJ), and AxioVison 4.7 software (Zeiss).

Optical stimulation

Optical stimulation was delivered to hESC-CM via a Lambda DG-4 300 W Xenon light source or with a 470 nm LED at 7 mW/mm² (Thorlabs, Newton, NJ). For multielectrode array (MEA) electrophysiology, optical stimulation consisted of a monophasic waveform with peak amplitude of 0, 33, 75, or 100% of maximum power (10 mW/mm² for 40x objective), pulse width of 100 ms, and frequency of 0.5, 1.0, or 1.5 Hz. For whole-cell patch clamp (PC) electrophysiology, optical stimulation consisted of a monophasic waveform with peak amplitude of 0, 12.5, 25, 50, or 100% of maximum power (10 mW/mm² for 40x objective) and a pulse width of 1000 ms. Optical power delivered to cells at each microscope objective was measured with a digital power meter (Thorlabs) at the focal plane of the objective.

Multielectrode array (MEA) electrophysiology

Multielectrode arrays (MEA) with sixty $30\ \mu\text{m}$ titanium nitride (TiN) electrodes equally spaced $200\ \mu\text{m}$ apart, with indium tin oxide (ITO) leads, and with an internal reference (Thin MEA 200/30 iR ITO, Multi Channel Systems, MCS GmbH, Reutlingen, Germany) were coated with $25\ \mu\text{g}/\text{mL}$ fibronectin (Sigma). Desired cardiomyocyte (CM) colonies were then manually dissected off their plates, transferred to the MEAs, and allowed to attach. A single MEA containing cells and Tyrode's solution (Sigma) was then placed in the amplifier (MEA 1060-Inv-BC, MCS) for recordings. Signals were acquired at 1kHz from a USB-6225 M Series DAQ (NI, Austin, TX). Videos of contracting CM were captured at 30 fps for a duration of 1-30 s with a Retiga 2000R color cooled camera (QImaging). The MEA amplifier was configured with MEA Select 1.1.0 software (MCS) and electrical and video signals were acquired and controlled with a custom program created with LabVIEW 2009 (NI). To visualize hESC-CM contractions, a custom edge detection algorithm in LabVIEW 2009 (NI) was used to detect rising and falling edge locations along a grayscale profile generated from a user-defined region of interest based on a user-defined threshold value.

Patch clamp (PC) electrophysiology

hESC-CM were recorded by means of whole-cell patch clamp (PC), using an Axon Multi-clamp 700B amplifier (Molecular Devices, Sunnyvale, CA), an Axon Digidata 1440A data acquisition system (Molecular Devices), and pClamp 10 software (Molecular Devices) as previously described [67, 129]. Cells were visualized and optically stimulated on an Olympus upright microscope equipped with a 470 nm LED (Thorlabs) and EXFO X-Cite halogen light source (Lumen Dynamics, Ontario, Canada) through a 40X/0.8 NA water immersion objective. When using the halogen light source coupled to a shutter (VCM-D1, Uniblitz, Rochester, NY), an excitation filter of HQ470/40, dichroic Q495LP (Chroma) was used for delivering blue light for ChR2 activation (10, 5, 2.5, $1.25\ \text{mW}/\text{mm}^2$). eYFP was visualized with a standard eYFP filter set (excitation 500/20, dichroic 515LP, emission 535/30; Chroma). Borosilicate glass (Sutter Instruments) pipette resistance ranged from 3-6 $\text{M}\Omega$. Whole-cell PC recordings were performed as previously described [67]; intracellular solution: 129 mM K-gluconate, 10 mM HEPES, 10 mM KCl, 4 mM MgATP, 0.3 mM NaGTP, titrated to pH 7.2; extracellular Tyrode: 125 mM NaCl, 2 mM KCl, 3 mM CaCl_2 , 1 mM MgCl_2 , 30 mM glucose, and 25 mM HEPES, titrated to pH 7.3. For voltage clamp recordings cells

were held at -70mV . All experiments were performed at room temperature, $22\text{-}25^\circ\text{C}$. Fluorescent cells were patched immersed in Tyrode solution containing $5\text{-}10\ \mu\text{M}$ blebbistatin to eliminate contractility while preserving electrical activity [54]. Patch clamp data was analyzed using Clampfit 10.2 (Molecular Devices).

Mathematical model of channelrhodopsin photocycle

To model the channelrhodopsin photocycle, we adopt a three-state photocycle model [79, 139], which is characterized through an open, a closed but still recovering, and a fully closed state [133], as illustrated in Figure 10.3. Upon photo absorption, molecules which are in

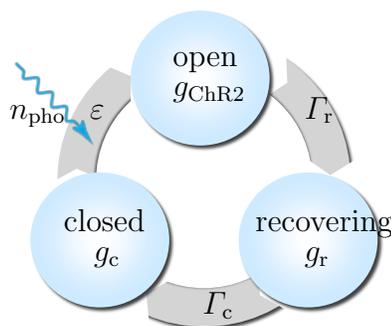


Figure 7.3: Three-state model for the channelrhodopsin photocycle. Upon photo absorption, molecules in the closed state g_c undergo a fast transition into the open state g_{ChR2} . After for some time, molecules spontaneously turn into the recovering state g_r where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After a recovery period, the molecules finally return to the closed state g_c , ready to undergo a new photocycle when subjected to light.

the closed state g_c undergo a fast transition into the open state g_{ChR2} . After being open for some time, molecules spontaneously transition into the recovering state g_r , where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After a recovery period, the molecules finally return to the closed state g_c , ready to undergo a new photocycle when exposed to light [138]. Figure 10.3 suggests the following first order model for the channelrhodopsin photocycle,

$$\begin{aligned}
 \dot{g}_{\text{ChR2}} &= \varepsilon n_{\text{pho}} g_c & - & \Gamma_r g_{\text{ChR2}} \\
 \dot{g}_r &= \Gamma_r g_{\text{ChR2}} & - & \Gamma_c g_r \\
 \dot{g}_c &= \Gamma_c g_r & - & \varepsilon n_{\text{pho}} g_c
 \end{aligned} \tag{7.1}$$

where Γ_r and Γ_c are the rates of recovery and full closure, ε is the quantum efficiency of the channelrhodopsin system, and n_{pho} is the number of photons hitting the cell per second. Herein, g_{ChR2} , g_r , and g_c define the fraction of molecules in the open, recovering, and closed states, scaled such that they sum up to one, $g_{\text{ChR2}} + g_r + g_c = 1$. This implies that the photocycle system (8.1) can be characterized through two independent variables, e.g., the fraction of molecules in the open and in the recovering states.

$$\begin{aligned} \dot{g}_{\text{ChR2}} &= \varepsilon n_{\text{pho}} - [\varepsilon n_{\text{pho}} + \Gamma_r] g_{\text{ChR2}} - \varepsilon n_{\text{pho}} g_r \\ \dot{g}_r &= \Gamma_r g_{\text{ChR2}} - \Gamma_c g_r \end{aligned} \quad (7.2)$$

We identify the state g_{ChR2} as the channelrhodopsin gating variable and integrate it into a

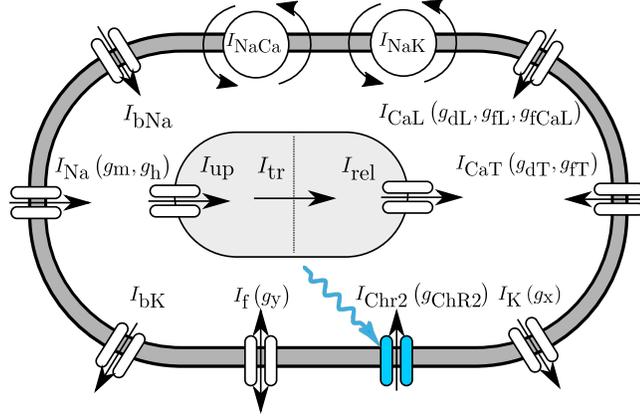


Figure 7.4: Ionic model of genetically engineered light sensitive cardiac cell. The electrochemical state of the cell is characterized in terms of $n_{\text{ion}} = 8$ ion concentrations, $c_{\text{ion}} = [c_{\text{Na}}^e, c_{\text{K}}^e, c_{\text{Ca}}^e, c_{\text{Na}}^i, c_{\text{K}}^i, c_{\text{Ca}}^i, c_{\text{Ca}}^{\text{up}}, c_{\text{Ca}}^{\text{rel}}]$, the extracellular and intracellular sodium, potassium, and calcium concentrations, and the sarcoplasmic reticulum calcium uptake and release. Ion concentrations are controlled through $n_{\text{crt}} = 12$ ionic currents, $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{f}}, I_{\text{NaCa}}, I_{\text{NaK}}, I_{\text{ChR2}}, I_{\text{K}}, I_{\text{bK}}, I_{\text{CaL}}, I_{\text{CaT}}, I_{\text{up}}, I_{\text{rel}}]$, where the baseline autorhythmic cell model [49] has been enhanced with the channelrhodopsin current I_{ChR2} , here shown in blue. The channels are governed by $n_{\text{gate}} = 10$ gating variables $g_{\text{gate}} = [g_m, g_h, g_{\text{ChR2}}, g_y, g_x, g_{\text{dL}}, g_{\text{fL}}, g_{\text{fCa}}, g_{\text{dT}}, g_{\text{fT}}]$ which may be functions of the current membrane potential ϕ .

well-defined autorhythmic cell model [49] characterized through $n_{\text{gate}} = 10$ gating variables in total as illustrated in Figure 10.4.

$$g_{\text{gate}} = [g_m, g_h, g_{\text{ChR2}}, g_y, g_x, g_{\text{dL}}, g_{\text{fL}}, g_{\text{fCa}}, g_{\text{dT}}, g_{\text{fT}}] \quad (7.3)$$

In particular, these are the fast sodium channel activation gate g_m , the fast sodium channel inactivation gate g_h , the channelrhodopsin activation gate g_{ChR2} , the hyperpolarization activated inward current activation gate g_y , the delayed rectifier current activation gate g_x , the long-lasting calcium channel activation gate g_{dL} , the voltage dependent long-lasting calcium channel inactivation gate g_{fL} , the calcium dependent long-lasting calcium channel inactivation gate g_{fCa} , the transient calcium channel activation gate g_{dT} , and the transient calcium channel inactivation gate g_{fT} . The gating variables are parameterized in terms of the transmembrane potential ϕ , the ionic concentrations c_{ion} , and the gating variables g_{gate} themselves. Their evolution is governed by classic Hodgkin-Huxley type equations,

$$\dot{g}_{\text{gate}} = \frac{1}{\tau_{\text{gate}}(\phi)} [g_{\text{gate}}^{\infty}(\phi, c_{\text{ion}}) - g_{\text{gate}}] \quad (7.4)$$

each characterized through a steady state value g_{gate}^{∞} and a time constant τ_{gate} for reaching this steady state [197]. Both are usually exponential functions of the transmembrane potential ϕ .

Mathematical model of ionic currents

The channelrhodopsin gating variable g_{ChR2} introduced in the previous section governs the channelrhodopsin current I_{ChR2} , for which we make the following ansatz.

$$I_{\text{ChR2}} = C_{\text{ChR2}} g_{\text{ChR2}} [\phi - \phi_{\text{ChR2}}] \quad (7.5)$$

Here, C_{ChR2} is the channelrhodopsin conductance and ϕ_{ChR2} is the reversal potential of channelrhodopsin, see supplemental material. We integrate the channelrhodopsin current I_{ChR2} into an autorhythmic cardiac cell model [47, 49], defined through a total of $n_{\text{crt}} = 12$ ionic currents,

$$I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{f}}, I_{\text{NaCa}}, I_{\text{NaK}}, I_{\text{ChR2}}, I_{\text{K}}, I_{\text{bK}}, I_{\text{CaL}}, I_{\text{CaT}}, I_{\text{up}}, I_{\text{rel}}] \quad (7.6)$$

which are, mathematically speaking, functions of the transmembrane potential ϕ , the individual gating variables g_{gate} , and the ion concentrations c_{ion} .

$$I_{\text{crt}} = I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (7.7)$$

In particular, these are eight transmembrane currents, the fast sodium current I_{Na} , the background sodium current I_{bNa} , the hyperpolarization activated sodium and potassium currents I_f , the sodium calcium exchanger current I_{NaCa} , the sodium potassium pump current I_{NaK} , the channelrhodopsin sodium current I_{ChR2} , the delayed rectifying potassium current I_{K} , the background potassium current I_{bK} , the long-lasting L-type calcium current I_{CaL} , the transient T-type calcium current I_{CaT} , and two sarcoplasmic reticulum currents, the calcium uptake I_{up} and the calcium release I_{rel} .

Mathematical model of ionic concentrations

From a chemical point of view, light induces a channelrhodopsin current I_{ChR2} , which directly impacts the intracellular sodium concentrations c_{Na}^i ,

$$\dot{c}_{\text{Na}}^i = -\frac{1}{FV^i} [I_{\text{Na}} + I_{\text{bNa}} + I_{\text{fNa}} + 3I_{\text{NaCa}} + 3I_{\text{NaK}} + I_{\text{KNa}} + I_{\text{ChR2}}] \quad (7.8)$$

where V^i is the cytosolic volume and F is the Faraday constant. The sodium concentration will directly, and indirectly through the resulting changes in the transmembrane potential ϕ , affect all other ionic concentrations in the cell. The biochemistry of our cell model is characterized through $n_{\text{ion}} = 8$ ion concentrations,

$$c_{\text{ion}} = [c_{\text{Na}}^e, c_{\text{K}}^e, c_{\text{Ca}}^e, c_{\text{Na}}^i, c_{\text{K}}^i, c_{\text{Ca}}^i, c_{\text{Ca}}^{\text{up}}, c_{\text{Ca}}^{\text{rel}}] \quad (7.9)$$

defined in terms of their evolution equations,

$$\dot{c}_{\text{ion}} = \dot{c}_{\text{ion}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (7.10)$$

which are parameterized in terms of the transmembrane potential ϕ , the gating variables g_{gate} , and the ion concentrations c_{ion} themselves. In particular, c_{Na}^e , c_{K}^e , and c_{Ca}^e are the extracellular sodium, potassium, and calcium concentrations, c_{Na}^i , c_{K}^i , and c_{Ca}^i are the intracellular sodium, potassium, and calcium concentrations, and $c_{\text{Ca}}^{\text{up}}$ and $c_{\text{Ca}}^{\text{rel}}$ are the sarcoplasmic reticulum calcium uptake and release, see supplemental material for details about the cell model and its material parameters.

Mathematical model of action potential propagation

From an electrical point of view, light induces a channelrhodopsin current I_{ChR2} , which directly impacts the action potential propagation ϕ in the heart. Unlike the local ion concentrations for single cells, however, the action potential is a global field variable [72]. Its spatio-temporal evolution

$$\dot{\phi} = f^{\phi}(\phi, g_{\text{gate}}, c_{\text{ion}}) + \text{div } \mathbf{q}(\phi) \quad (7.11)$$

is driven by a local source term f^{ϕ} at the single cell level, and by a global flux term, $\text{div } \mathbf{q}$, the divergence of the propagation vector \mathbf{q} at the organ level. The local source term

$$\begin{aligned} f^{\phi} = -\frac{1}{C} [& I_{\text{Na}} + I_{\text{bNa}} + I_{\text{f}} + I_{\text{NaCa}} + I_{\text{NaK}} \\ & + I_{\text{ChR2}} + I_{\text{K}} + I_{\text{bK}} + I_{\text{CaL}} + I_{\text{CaT}}] \end{aligned} \quad (7.12)$$

is directly related to the negative sum of all transmembrane currents scaled by the individual cell membrane capacitance per unit surface area C . To account for the nonlocal nature of propagating excitation waves in the heart, we introduce the propagation vector

$$\mathbf{q} = \mathbf{D} \cdot \nabla \phi \quad (7.13)$$

through the second order diffusion tensor \mathbf{D} scaling the gradient of the action potential field $\nabla \phi$, see supplemental material for the finite element based solution of the action potential propagation problem [?, 72, 197].

RESULTS

Figure 10.5 demonstrates our ability to stably transduce undifferentiated hESC with a ChR2-eYFP lentiviral vector. The resulting hESC^{ChR2+} remain pluripotent and can differentiate into hESC^{ChR2+}-CM. Figure 10.5a, the PCR shows that hESC^{ChR2+} express the pluripotent Oct-4 gene (169 bp, lane 4) and Nanog gene (154 bp, lane 5), indicated through the solid blue box. In addition, amplification within the ChR2 gene (174 bp, lane 6), across the ChR2-eYFP gene (197 bp, lane 7), and within the eYFP gene (187bp, lane 8), confirms stable transduction of the ChR2-eYFP lentivirus in

undifferentiated hESC^{ChR2+}, indicated through the solid yellow box. A ladder (100 bp, lane 1) confirms the predicted sizes of PCR products. Non-template control (lane 2) and GAPDH (152 bp, lane 3) serve as negative and positive controls, respectively. Figure 10.5b illustrates that pluripotent hESC^{ChR2+} stain positive for alkaline phosphatase shown in red. Figure 10.5c demonstrates the hESC^{ChR2+} positive eYFP signal upon fluorescence microscopy shown in green. Figure 10.5d confirms that hESC^{ChR2+}-CM have positive TnI signals, shown in red, consistent with a CM phenotype. DAPI staining, shown in blue, demonstrates the position of nuclei. Figure 10.5e illustrates transmission electron microscopy of hESC^{ChR2+}-CM showing sarcomeres with characteristic z-lines and mitochondria. Figure 10.5f shows the light microscopy of three hESC^{ChR2+}-CM colonies, indicated through dashed white circles, on a multielectrode array.

Figure 10.6 demonstrates the experimental and computational sensitivity of hESC^{ChR2}-CM with respect to different light intensities. With light on, the photocurrent I_{ChR2} increases rapidly, peaks, and decays towards a characteristic plateau value. With light off, the photocurrent I_{ChR2} drops rapidly and decays to zero. Figure 10.6, top, illustrates the varying light intensities from 12.5% to 25%, 50%, and 100% for which we measure the photocurrent I_{ChR2} upon whole cell voltage clamp. Figure 10.6, middle, demonstrates the experimentally measured photocurrent I_{ChR2} , which increases with increasing light intensity. Figure 10.6, bottom, shows the calibrated computational cell model which captures the characteristic light sensitivity, displaying increased photocurrents I_{ChR2} with increased light intensity, see supplemental material.

Figure 10.7 demonstrates the experimental and computational sensitivity of hESC^{ChR2}-CM with respect to different stimulation frequencies. Blue lines illustrate the applied light stimulation at 100% intensity, at 100 ms pulse width. Black and red lines display the experimentally measured electrical field potentials and mechanical contractions. Green lines display the computationally predicted electrical transmembrane potentials. We probe the cells with light stimulation at 0.5 Hz, top, 1.0 Hz, middle, and 1.5 Hz, bottom. Electrical and mechanical signals during light stimulation, shown in the center, are significantly different from pre- and post-stimulation signals at all frequencies, shown at the beginning and end of the readings. The computational

hESC^{ChR2}-CM model excellently captures the electrical signal both during light stimulation, and pre- and post-stimulation. The two initial double spikes in the green curves of the computational model paced at 0.5 Hz are caused by an interference of the photostimulation with the cell's natural frequency of 0.7 Hz. This interference is only present when cells are paced below their natural frequency, and does not occur during photostimulation at 1.0 Hz and 1.5 Hz.

Figure 10.8 illustrates the potential of the proposed technology to virtually pace a human heart with light. The finite element model of the heart created from magnetic resonance images consists of 3,129 nodes and 11,347 three-dimensional tetrahedral elements [102], see supplemental material. Figure 10.8, top, shows the effects of atrioventricular (AV) node photostimulation initiated through hESC^{ChR2}-CM, which are virtually injected into the basal region of the septum, while all other regions are modeled as standard ventricular CM. For AV node pacing, the depolarization wave is initiated at the AV node, travels down the septum, and then activates the left and right ventricles. Figure 10.8, bottom, shows the effects of bi-ventricular (bi-VP) photostimulation, initiated through hESC^{ChR2}-CM, which are virtually injected into the lateral wall of the left and right ventricles, while all other regions are modeled as standard ventricular CM. For bi-VP, the depolarization wave is initiated at the lateral left and right ventricular walls, travels along the ventricles, and then activates the apex and the septum. The color code indicates the magnitude of the transmembrane potential varying from -90mV, shown in blue, to +20mV, shown in red.

DISCUSSION

In Figure 10.5, we have shown that ChR2 can be expressed stably and safely in hESC-CM to drive CM depolarization via photostimulation. Using a lentiviral vector [23, 203], we have introduced ChR2 coupled to YFP into undifferentiated hESC. After confirming expression via immunocytochemistry, we have expanded these ChR2-expressing cells and demonstrated their pluripotency using PCR. Via directed differentiation, we have created hESC^{ChR2}-CM, which we have then subjected to optical

stimulation. We have successfully recorded their electrical, biochemical, and mechanical signals using patch clamping, MEA recordings, and video microscopy. These data have allowed us to calibrate our computational hESC^{ChR2}-CM model.

In Figure 10.6, we have demonstrated the sensitivity of both experimental and computational photocurrents with respect to the stimulating light intensity [138]. Both graphs illustrate the characteristic rapid increase to the intensity-dependent peak current, followed by a slower decrease towards the intensity-dependent asymptotic plateau value [89, 187].

While the translation of optogenetic techniques into clinical practice may still have many technical hurdles to pass, the technology itself can already serve as a valuable research tool in cardiac electrophysiology [1]. Traditional tools based on the simultaneous use of electrical stimulation and electrical recording typically suffer from unavoidable artifacts [109]. The inherent orthogonality of optical and electrical techniques allows us to significantly reduce these spurious errors when using optical stimulation combined with electrical recordings, as shown in Figure 10.7.

A tremendous potential of opsin-based systems for optical manipulation lies in their inherent ability to not only turn cells on using channelrhodopsin as a blue-light-gated ion channel transporting positively charged cations along their concentration gradients [46], but also to turn cells off using halorhodopsin as a yellow-light-driving ion pump transporting negatively charged chloride ions against their concentration gradients [80]. A promising first study in zebrafish has shown that a combination of channelrhodopsin and halorhodopsin allows for optically controlling heart rate, reversing cardiac conduction, and inducing disease-like arrhythmias [5].

Computational modeling allows us to predict the response of living cells, both in isolation and in interaction with their environment. In Figure 10.8, we have demonstrated our very first proof-of-principle, using finite element modeling in an attempt to bridge the scales from cells to systems [?]. Finite element based models allow us to combine virtually any cell type [?] on virtually any geometry [193]. Here, they allow us to predict the activation sequences in the human heart for different pacing sites [?].

Pacing hearts *in silico* by means of photostimulation will allow us to virtually probe

different scenarios [150, 73] towards our ultimate goal of light pacing hearts *in vivo* [5]. In contrast to electrical pacemakers, light delivery is minimally invasive, genetically targeted, and temporally precise. Most importantly, light can be delivered at a distance. Unlike pacing leads for electrical stimulation, which are known to have a high failure rate due to mechanical fatigue [99], the light source for optical stimulation does not have to sit directly on the continuously moving heart muscle. Light pacing might therefore be an attractive remote, less invasive, and more durable alternative to current electrical pacing leads [191].

Limitations

This manuscript presents our first attempts towards pacing the heart with light using a multiscale approach. On the photocycle level, we have adapted a classic three state model for the ChR2 photocycle, which was initially proposed for ChR2-expressing neurons [133]. This conceptually elegant model has allowed us to reliably reproduce the characteristic features of a light-evoked response in ChR2-expressing hESC-CM. However, the extension to a four [139] or five [8] state model, which are known to reproduce the bi-exponential decay of the light-off current more accurately [79], is conceptually straightforward. In terms of photostimulation, we have applied pulse widths of 100 and 1000 ms. These relatively long durations of channel opening may increase sodium overload and adversely influence action potential profiles. Since the rapid upstroke of the action potential takes place in the first 10 ms [49], the stimulation pulse width could potentially be decreased to 10 ms or less [1], which would favorably limit light exposure time, see supplemental material. Novel developments in ultrafast optogenetics suggest using the rationally engineered channelrhodopsin ChETA, which provides high fidelity optical control of spiking at high frequencies and eliminates plateau potentials during continued stimulation [70].

On the ion channel level, we have assumed that the ChR2 current is driven exclusively by concentration gradients in the sodium concentration [139]. Accordingly, we have modeled ChR2 to be selectively permeable to sodium ions only. While ChR2 is known to be a general cation channel [176], its effects on potassium and calcium

remain poorly characterized [15]. However, making the channel permeable to other monovalent and divalent cations [133], would require only modular changes in the mathematical model.

On the cellular level, due to the lack of mathematical models for stem cell-derived cardiac cells, we have adopted a widely used and well characterized model for mature cardiac cells [47, 49]. A recent study confirmed that ChR2 expresses its characteristic features independent of the particular expression system [15]. Therefore, we have modified the mature cell model via an additional photocurrent governed by a light-sensitive gating variable. A thorough identification of the individual channel characteristics of hESC-CM and their quantitative comparison with mature CM and hESC^{ChR2+}-CM remain to be addressed to fully validate our conceptual approach [129].

On the whole heart level, the simulation of the light-paced heart is admittedly relatively simplistic. Although our algorithm can, in principle, handle arbitrary mixtures of different cell types [?], here, we have assumed that the injected cells are pure hESC^{ChR2+}-CM. We have adopted a common procedure to model cell injection [193], which does not address additional obstacles commonly associated with cell delivery such as cell migration away from the injection site or cell survival in the myocardial wall. At this stage, our model also fails to appropriately predict the effects of light scattering by tissue and absorption by blood. Light scattering is a general barrier to in vivo translation, and efforts are underway by our groups and others to red-shift the opsins and maximize the efficiency of light delivery through tissues [16, 67].

CONCLUSION

At an unprecedented temporal and spatial precision, optogenetic tools now enable us to manipulate electrically active cells. This study capitalizes on recent developments in optics and genetics, supplemented by novel technologies in stem cell biology, electrophysiology, and computational mechanics. It documents our first attempts to

introduce a light-sensitive ion channel in human embryonic stem cell derived cardiomyocytes, with the ultimate goal to control the cardiac system by means of photostimulation. Unlike traditional electrical stimulation, optogenetics allows us to precisely control the selective permeability of the plasma membrane, its conductivity with respect to different ions, its sensitivity to light of different wavelengths, and the spatio-temporal evolution of different opening and closing profiles. Given this incredible freedom, we need to establish an economical strategy to optimize the matrix of input variables. Predictive computational models allow us to virtually probe landscapes of process parameters and identify optimal photostimulation sequences in various different tissues and organs. Here, as a very first proof-of-principle, we virtually inject photosensitive cells into different locations of a human heart model, and pace the heart with light. We believe that this concept will be widely applicable to systematically manipulate electrically active cells and, ultimately, support the design of novel therapies for various types of neuronal, musculoskeletal, pancreatic, and cardiac disorders such as depression, schizophrenia, cerebral palsy, paralysis, diabetes, pain syndromes, and cardiac arrhythmias.

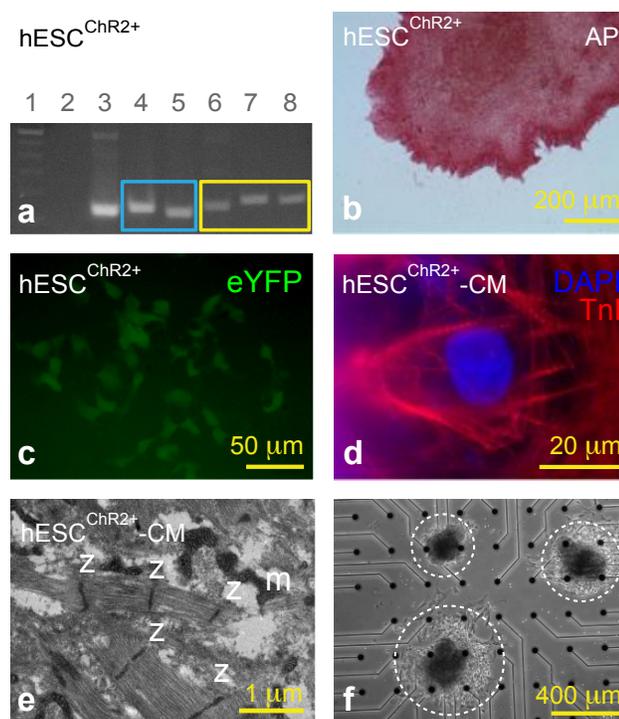


Figure 7.5: Undifferentiated human embryonic stem cells (hESC) stably transduced with a ChR2-eYFP lentiviral vector (hESC^{ChR2+}) remain pluripotent and can differentiate into cardiomyocytes (hESC^{ChR2+}-CM). **a.** PCR shows that hESC^{ChR2+} express the pluripotent Oct-4 gene (169 bp, lane 4) and Nanog gene (154 bp, lane 5), solid blue box. In addition, amplification within the ChR2 gene (174 bp, lane 6), across the ChR2-eYFP gene (197 bp, lane 7), and within the eYFP gene (187bp, lane 8), confirms stable transduction of the ChR2-eYFP lentivirus in undifferentiated hESC^{ChR2+}, solid yellow box. A ladder (100 bp, lane 1) confirms the predicted sizes of PCR products. Non-template control (lane 2) and GAPDH (152 bp, lane 3) serve as negative and positive controls, respectively. **b.** Pluripotent hESC^{ChR2+} stain positive for alkaline phosphatase (red). **c.** Fluorescence microscopy shows hESC^{ChR2+} positive eYFP signal (green). **d.** hESC^{ChR2+}-CM have positive TnI signals (red), consistent with a CM phenotype. DAPI staining (blue) demonstrates the position of nuclei. **e.** Transmission electron microscopy shows sarcomeres with associated z-lines (z) and mitochondria (m) in hESC^{ChR2+}-CM. **f.** Light microscopy shows three hESC^{ChR2+}-CM colonies (dashed white circles) on a multielectrode array.

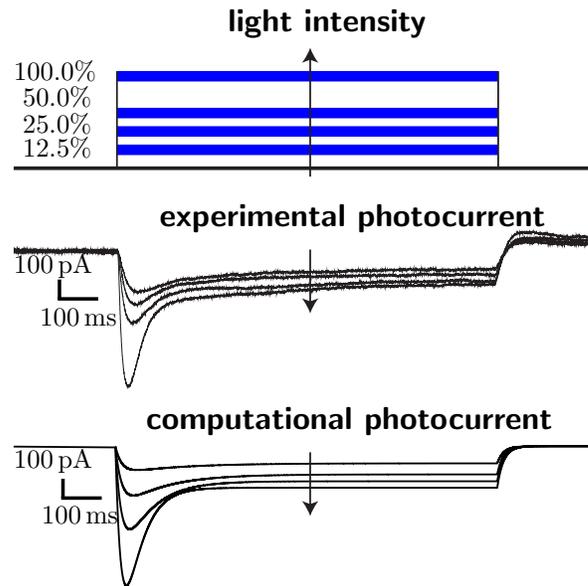


Figure 7.6: Experimental and computational sensitivity of $\text{hESC}^{\text{ChR2}}\text{-CM}$ with respect to light intensity. With the light turned on, the photocurrent I_{ChR2} increases rapidly, peaks, and decays towards a characteristic plateau value. With the light turned off, the photocurrent I_{ChR2} drops rapidly and decays to zero. Light intensity is varied from 12.5% to 25%, 50%, and 100% (top). Whole cell voltage clamp reveals an increased photocurrent I_{ChR2} as the light intensity increases. The computational $\text{hESC}^{\text{ChR2}}\text{-CM}$ model captures the light sensitivity and displays increased photocurrents I_{ChR2} with increased light intensity (bottom).

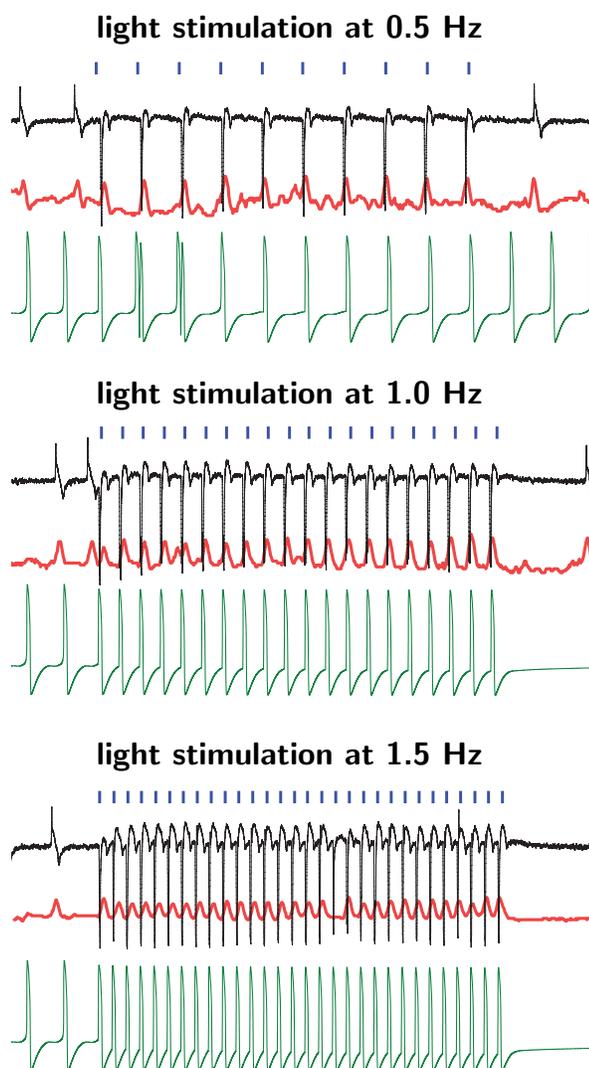


Figure 7.7: Experimental and computational sensitivity of $\text{hESC}^{\text{ChR2}}\text{-CM}$ with respect to stimulation frequency. Light stimulation (blue) evokes field potentials (black) which translate into mechanical contractions (red). Light stimulation at 100% intensity is performed at 0.5 Hz (top), 1.0 Hz (middle), and 1.5 Hz (bottom). Evoked signals during light stimulation (center) are markedly different from pre- and post-stimulation signals at all frequencies (left and right). The computational $\text{hESC}^{\text{ChR2}}\text{-CM}$ model (green) captures the electrical signal at all frequencies, both during light stimulation (center), and pre and post stimulation (left and right).

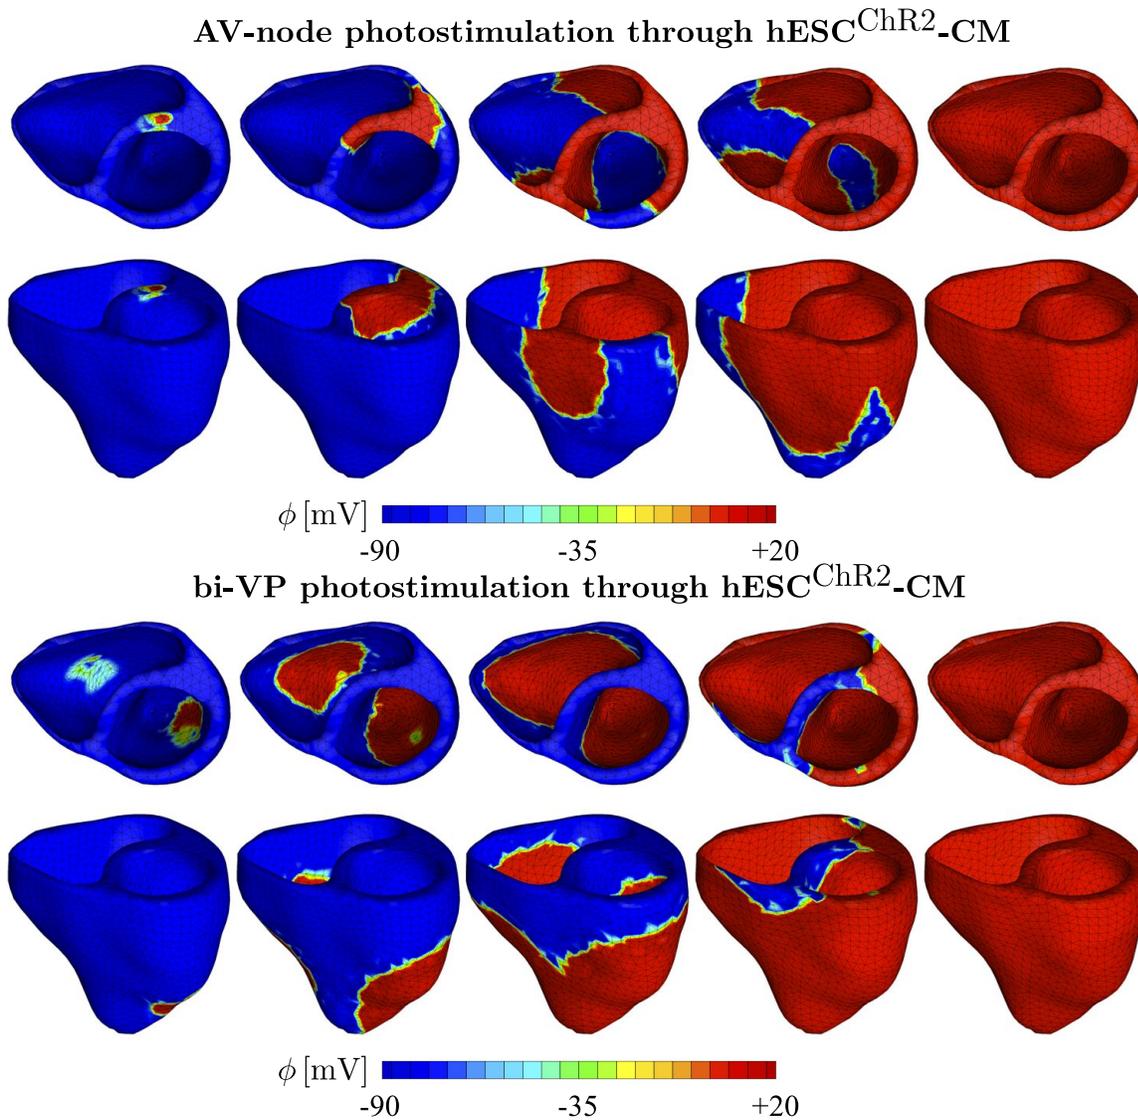


Figure 7.8: Virtual activation sequences of light-paced hearts. Atrioventricular (AV) node photostimulation (top) is initiated through hESC^{ChR2}-CM, virtually injected into the basal region of the septum, while all other regions are modeled is standard ventricular CM. A depolarization wave forms at the AV node, travels down the septum, and activates the left and right ventricles. Bi-ventricular (bi-VP) photostimulation (bottom) is initiated through hESC^{ChR2}-CM, virtually injected into the lateral wall of the left and right ventricles, while all other regions are modeled is standard ventricular CM. A depolarization wave forms at the lateral left and right ventricular walls, travels along the ventricles, and activates the apex and the septum. The color code indicates the magnitude of the transmembrane potential ϕ varying from -90mV (blue) to +20mV (red).

Chapter 8

Applications in virtual treatments using Chr2 cells

Electrical stimulation is currently the gold standard treatment for heart rhythm disorders. However, electrical pacing is associated with technical limitations and unavoidable potential complications. Recent developments now enable the stimulation of mammalian cells with light using a novel technology known as optogenetics. The optical stimulation of genetically engineered cells has significantly changed our understanding of electrically excitable tissues, paving the way towards controlling heart rhythm disorders by means of photostimulation. Controlling these disorders, in turn, restores coordinated force generation to avoid sudden cardiac death. Here, we report a novel continuum framework for the photoelectrochemistry of living systems that allows us to decipher the mechanisms by which this technology regulates the electrical and mechanical function of the heart. Using a modular multiscale approach, we introduce a non-selective cation channel, channelrhodopsin-2, into a conventional cardiac muscle cell model via an additional photocurrent governed by a light-sensitive gating variable. Upon optical stimulation, this channel opens and allows sodium ions to enter the cell, inducing electrical activation. In side-by-side comparisons with conventional heart muscle cells, we show that photostimulation directly increases the sodium concentration, which indirectly decreases the potassium concentration in the

cell, while all other characteristics of the cell remain virtually unchanged. We integrate our model cells into a continuum model for excitable tissue using a nonlinear parabolic second order partial differential equation, which we discretize in time using finite differences and in space using finite elements. To illustrate the potential of this computational model, we virtually inject our photosensitive cells into different locations of a human heart, and explore its activation sequences upon photostimulation. Our computational optogenetics tool box allows us to virtually probe landscapes of process parameters, and to identify optimal photostimulation sequences with the goal to pace human hearts with light and, ultimately, to restore mechanical function.

8.1 Motivation

The human heart propels over 7000 liters of blood through the body daily, beating more than 40 million times a year. It is a remarkably efficient, durable, and reliable mechanical pump, precisely regulated by spatially and temporally varying electrical and chemical fields. Disturbed conduction and uncoordinated electrical signals can induce abnormal heart rhythms, which may critically reduce mechanical function [?, 185]. A delicate balance between the electrical and chemical gradients across the cell membrane governs the electrophysiological activity of excitable cardiac cells. To maintain these gradients, the cell membrane is selectively permeable with respect to different ions at different points in time [?]. In a cardiac muscle cell, for example, the potential difference across the cell membrane is approximately -80 mV at rest, meaning the cell is negatively charged. It can be excited by an external electrical stimulus that pushes its transmembrane potential beyond a critical threshold value. This triggers the influx of sodium ions, generating a rapid depolarization, and the transmembrane potential increases to +20 mV. Over time, the coordinated interplay between sodium, potassium, and calcium ions brings the transmembrane potential smoothly back to its resting state [96].

Since the late 1950s, permanent cardiac pacemakers have been used successfully to correct rhythm disorders and restore smoothly propagating waves of electrical excitation [?]. Pacemakers deliver electrical stimuli via implanted pacing leads that are

in contact with the heart muscle. Although permanent pacemaker implantation is a standard procedure today, pacing lead placement continues to be a subject of great controversy [?]. Choosing the optimal pacing site is becoming more challenging [150], in particular in view of novel stimulation strategies such as single chamber, dual chamber, and biventricular pacing. A potential complication of lead placement is cardiac perforation, which can have severe clinical consequences including pericardial effusion, cardiac tamponade, pneumothorax, and death [?]. However, to date, there is no conceptual alternative to electrical pacing that avoids direct contact stimulus delivery [99].

For more than 40 years, biologists have studied microorganisms that produce proteins to directly regulate the flow of charged ions across their plasma membrane in response to light [142]. The first identified protein of this kind was the light-gated ion pump bacteriorhodopsin, transporting positively charged hydrogen ions across the cell membrane [142]. While bacteriorhodopsin naturally acts as an on-switch for electrically active cells, the second reported light-gated ion pump, halorhodopsin, transports negatively charged chloride ions, thereby acting as an off-switch [120]. Despite intense research [?], the first reported light-gated ionic channel, channelrhodopsin, was only identified ten years ago [132], but has since then revolutionized neuroscience. Channelrhodopsin consists of seven transmembrane proteins embedded in the lipid bilayer [?] and absorbs blue light through its interaction with retinal, see Figure 10.1. Since the early nineties, we have known that phototaxis and photophobic responses in the green alga *Chlamydomonas reinhardtii* are mediated by rhodopsins with a microbial-type all-*trans* retinal chromophore [78]. The photochemical isomerization of this all-*trans* retinal to 13-*cis* retinal is illustrated in Figure 8.2. It occurs at peak absorption wavelengths of 470 nm, opening channelrhodopsin non-specifically to sodium, potassium, and calcium cations in response to blue light. In the dark, the covalently bound retinal spontaneously relaxes to all-*trans*, providing closure of the channel and regeneration of the chromophore. In 2005, channelrhodopsin was first introduced genetically using engineered viruses [23], a technique that is now widely known as optogenetics [46]. Since then, optical tools to control the electrical activity of nerve cells have rapidly evolved [176], and are now gaining widespread use in neuronal research

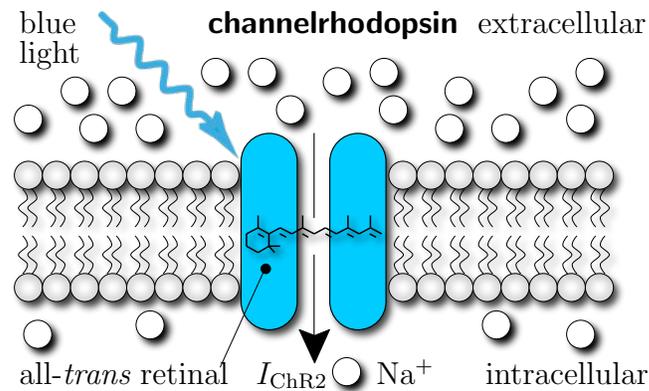


Figure 8.1: Channelrhodopsin-2 is a light-gated cation channel native to the green alga *Chlamydomonas reinhardtii*. It consists of seven transmembrane proteins and absorbs blue light through its interaction with retinal. Photoisomerization of retinal opens the channel to sodium ions, which have a higher concentration outside than inside the cell. To make our cells responsive to light, and allow sodium concentrations to equilibrate, we induce channelrhodopsin into a conventional cardiac muscle cell model.

and medicine [80]. While initial applications of optogenetics have been restricted exclusively to the neuronal system, optogenetic tools have now advanced to a level of maturity where they can confidently be applied to other cells and organs. Natural first candidates are stem cells, electrically active glial cells, muscle cells, and cardiac cells [5, ?]. Since light can be delivered at a distance, optical pacing has the potential to become a more precise and less invasive alternative to existing electrical pacing strategies [1].

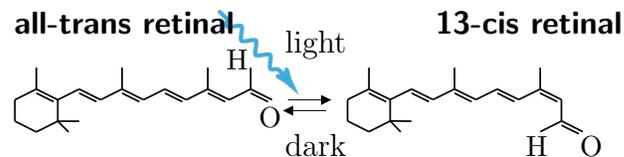


Figure 8.2: Channelrhodopsin-2 is activated by photoisomerization of all-*trans* retinal to 13-*cis* retinal at wavelengths of 470 nm. After photoisomerization, the covalently bound retinal spontaneously relaxes to all-*trans* in the dark, providing closure of the ion channel and regeneration of the chromophore.

The objective of this manuscript is to establish a novel continuum model for the

photoelectrochemistry of living systems that will allow us to virtually explore the potential of optogenetic pacing. The rationale for creating such a model is that it will enable patient-specific predictions of ion channel dynamics, ionic concentrations, and action potential profiles across the heart, which are outside the reach of experimental measurements in humans. Figure 8.3 illustrates the underlying approach in which the different physical fields interact across the different scales: On the molecular level, optical stimulation opens the cation channel channelrhodopsin initiating a photocurrent. On the subcellular level, this photocurrent increases the chemical concentration of sodium ions inside the cell. On the cellular level, concentration changes evoke changes in the electrical potential and excite the cell. On the tissue level, changes in the electrical potential propagates across the system in the form of smooth excitation waves.

In Section 8.2, we summarize the mathematical model for the channelrhodopsin photocycle and illustrate how channelrhodopsin can be introduced into a conventional cardiac cell model via an additional photocurrent governed by a light-sensitive gating variable. We characterize this additional photocurrent in Section 8.3, and demonstrate its impact on the intracellular ion concentrations in Section 8.4, and on the transmembrane potential ϕ . In a similar fashion, we compare this model to a conventional

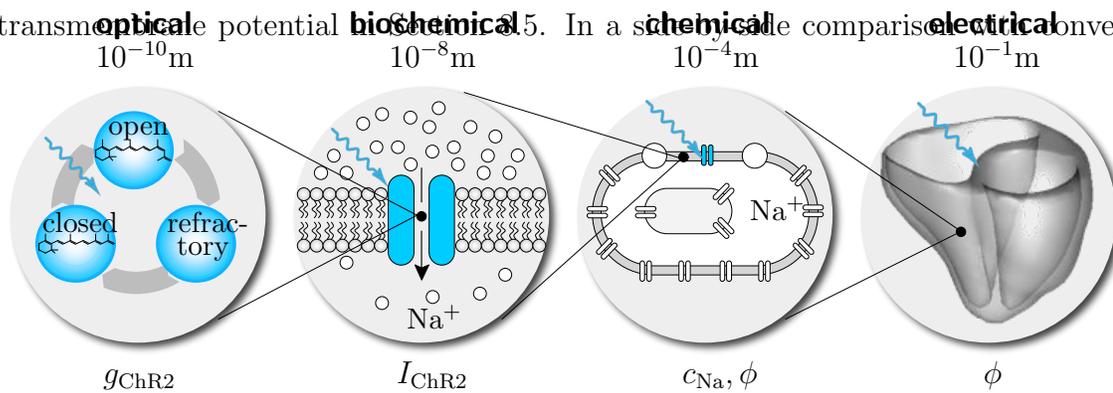


Figure 8.3: Multiscale model for the photoelectrochemistry of living systems. Optical stimulation opens the cation channel channelrhodopsin g_{ChR2} . This initiates a photocurrent I_{ChR2} increasing the chemical concentration of sodium ions c_{Na} inside the cell. Concentration changes evoke changes in the electrical potential ϕ , which propagates across the tissue system in the form of smooth excitation waves.

tional cardiac muscle cells, we perform virtual case studies to demonstrate the features

of genetically engineered light-sensitive cells in terms of gating variables, ionic currents, ion concentrations, and transmembrane potentials. Since our baseline model for conventional cardiac cells is well-characterized and fairly standard, we summarize its equations in the Appendix. We proceed by embedding these single cells into a novel continuum framework for photoelectrochemistry to simulate the spatial and temporal variation of electrical and chemical fields across the human heart. We briefly summarize the finite element discretization of the governing equations in Section 8.6. In Section 8.7, we virtually inject our light-sensitive cells into different locations of a human heart and explore its response to photostimulation. In Section 10.6, we conclude by discussing computational optogenetics and the potential of turning light into force in the context of pacing hearts with light.

8.2 Mathematical model of channelrhodopsin photocycle

To model the light sensitivity of channelrhodopsin, we adopt the kinetics of a three-state photocycle model [79, 139], which consists of a closed, photosensitive, excitable state g_{closed} , an open state g_{ChR2} , and a closed, desensitized, absolutely refractory state g_{refrac} [133], as illustrated in Figure 8.4. Upon photo absorption, molecules which are in the closed state g_{closed} undergo a fast transition into the open state g_{ChR2} . After being open for some time, molecules spontaneously transition into the absolutely refractory state g_{refrac} , where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After the refractory period, the molecules finally return to the closed state g_{closed} , ready to undergo a new photocycle when exposed to light [138]. Figure 8.4 suggests the following first order model for the channelrhodopsin photocycle,

$$\begin{aligned}
 \dot{g}_{\text{ChR2}} &= \varepsilon n_{\text{pho}} g_{\text{closed}} - \Gamma_{\text{r}} g_{\text{ChR2}} \\
 \dot{g}_{\text{refrac}} &= \Gamma_{\text{r}} g_{\text{ChR2}} - \Gamma_{\text{c}} g_{\text{refrac}} \\
 \dot{g}_{\text{closed}} &= \Gamma_{\text{c}} g_{\text{refrac}} - \varepsilon n_{\text{pho}} g_{\text{closed}}
 \end{aligned} \tag{8.1}$$

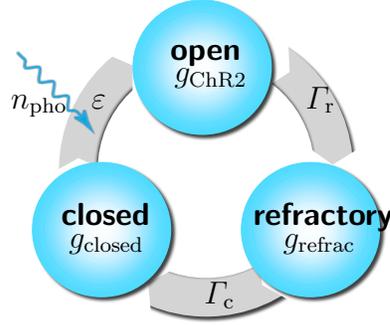


Figure 8.4: Three-state model for the channelrhodopsin photocycle. Upon photo absorption, molecules in the closed state g_{closed} undergo a fast transition into the open state g_{ChR2} . The molecules spontaneously turn into the refractory state g_{refrac} where the ion channels are closed, but the molecules are not yet ready to photoswitch again. After the refractory period, the molecules return to the closed state g_{closed} , ready to undergo a new photocycle when subjected to light [133, 138].

where Γ_r and Γ_c are the rates of recovery and full closure, ε is the quantum efficiency of the channelrhodopsin system, and n_{pho} is the number of photons hitting the cell per second. Here, we choose $\Gamma_r = 0.06129$, $\Gamma_c = 0.00422$, and $\varepsilon n_{\text{pho}} = 0.018$. Herein, g_{ChR2} , g_{refrac} , and g_{closed} define the fraction of molecules in the open, refractory, and closed states. Initially, all molecules are in the closed state, i.e., $g_{\text{closed}} = 1$, $g_{\text{ChR2}} = 0$, and $g_{\text{refrac}} = 0$. Since all three fractions sum up to one, i.e., $g_{\text{ChR2}} + g_{\text{refrac}} + g_{\text{closed}} = 1$, the photocycle system (8.1) can be characterized through two independent variables [89], e.g., the fraction of molecules in the open and in the refractory states.

$$\begin{aligned} \dot{g}_{\text{ChR2}} &= \varepsilon n_{\text{pho}} - [\varepsilon n_{\text{pho}} + \Gamma_r] g_{\text{ChR2}} - \varepsilon n_{\text{pho}} g_{\text{refrac}} \\ \dot{g}_{\text{refrac}} &= \Gamma_r g_{\text{ChR2}} - \Gamma_c g_{\text{refrac}} \end{aligned} \quad (8.2)$$

We identify the state g_{ChR2} as the channelrhodopsin gating variable and integrate it into a well-characterized ventricular cell model [188, 197] characterized through a total of $n_{\text{gate}} = 14$ gating variables as illustrated in Figure 8.5.

$$g_{\text{gate}} = [g_m, g_h, g_j, g_{\text{ChR2}}, g_{\text{xr1}}, g_{\text{xr2}}, g_{\text{xs}}, g_r, g_s, g_d, g_f, g_{\text{K1}\infty}, g_{\text{fCa}}, g_g], \quad (8.3)$$

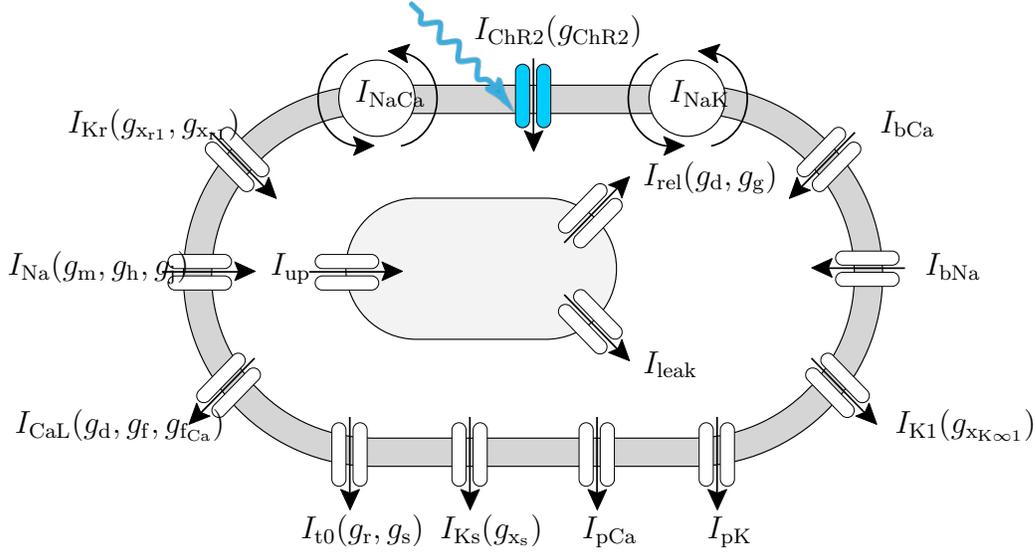


Figure 8.5: Genetically engineered light sensitive cardiac cell. The electrophysiology of the cell is characterized in terms of $n_{\text{ion}} = 4$ ion concentrations, the intracellular sodium, potassium, and calcium concentrations and the calcium concentration in the sarcoplasmic reticulum. Ion concentrations are controlled through $n_{\text{crt}} = 16$ ionic currents, where we have enhanced the conventional cell model [188, 197] with the channelrhodopsin photocurrent I_{ChR2} , here shown in blue [1]. The channels are governed by $n_{\text{gate}} = 14$ gating variables, where we have added the channelrhodopsin gating variable g_{ChR2} to characterize the cell's response to photostimulation.

In particular, these are the sodium activation gate g_m , the fast and slow sodium inactivation gates g_h and g_j , the channelrhodopsin activation gate g_{ChR2} , the rapid delayed rectifier activation and inactivation gates g_{xr1} and g_{xr2} , the slow delayed rectifier activation gate g_{xs} , the transient potassium activation and inactivation gates g_r and g_s , the long-lasting L-type calcium channel activation and inactivation gates g_d , g_f , and g_{fCa} , and the calcium release activation gate g_g . The gating variables are parameterized in terms of the transmembrane potential ϕ , the ionic concentrations c_{ion} , and the gating variables g_{gate} themselves. Their evolution is governed by classic Hodgkin-Huxley type equations,

$$\dot{g}_{\text{gate}} = \frac{1}{\tau_{\text{gate}}(\phi)} [g_{\text{gate}}^{\infty}(\phi, c_{\text{ion}}) - g_{\text{gate}}] \quad (8.4)$$

each characterized through a steady state value g_{gate}^{∞} and a time constant τ_{gate} for reaching this steady state [197]. Both are usually exponential functions of the transmembrane potential ϕ , see Appendix for details. Figure 8.6 illustrates the temporal evolution of all $n_{\text{gate}} = 14$ gating variables for the conventional cell, electrically stimulated for 1 ms, dashed black lines, and for the genetically engineered cell, light stimulated for 30 ms, solid blue lines. Overall, the gating variables of the conventional cell and the genetically engineered cell display a similar behavior. The g_{ChR2} gate is, of course, only active for the photosensitive cell.

8.3 Mathematical model of ionic currents

The channelrhodopsin gating variable g_{ChR2} introduced in the previous section governs the channelrhodopsin photocurrent I_{ChR2} , for which we make the following ansatz.

$$I_{\text{ChR2}} = C_{\text{ChR2}} g_{\text{ChR2}} [\phi - \phi_{\text{ChR2}}] \quad (8.5)$$

This implies that in our current model, the effects of conduction C_{ChR2} and photostimulation g_{ChR2} on the photocurrent I_{ChR2} are separable. This assumption is sufficient to accurately capture most phenomena reported in the literature [1, 89]. Some more advanced four state models suggest the use of rate variables with dependence on light intensity to finetune of the photocurrent [139]. However, because of the lack of experimental data, it is currently unclear whether these electrical potential effects are truly significant. In our model, the channelrhodopsin conductance

$$C_{\text{ChR2}} = \frac{1}{g_{\text{ChR2}}^{\infty} [\phi_{\text{clamp}} - \phi_{\text{ChR2}}]} I_{\text{ChR2}}^{\infty} \quad (8.6)$$

is expressed in terms of a quadratic polynomial, which we have identified to

$$I_{\text{ChR2}}^{\infty} / [\phi_{\text{clamp}} - \phi_{\text{ChR2}}] = 0.000266\phi^2 - 0.003658\phi + 0.498819$$

using a least squares fit. With a plateau value g_{Chr2}^∞ according to [138]

$$g_{\text{Chr2}}^\infty = \frac{\varepsilon n_{\text{pho}} \Gamma_c}{\varepsilon n_{\text{pho}} [\Gamma_r + \Gamma_c] + \Gamma_r \Gamma_c} \quad (8.7)$$

i.e., in our case $g_{\text{Chr2}}^\infty = 0.05283$, the conductance becomes $C_{\text{Chr2}} = [0.05\phi^2 - 0.0692\phi + 9.442]$. Moreover, ϕ_{Chr2} is the reversal potential of channelrhodopsin,

$$\phi_{\text{Chr2}} = \phi_{\text{Na}} - \tilde{\phi}_{\text{Na}} \quad (8.8)$$

which we approximate as the difference of the concentration-dependent reversal potential for sodium ϕ_{Na} and the experimental reversal potential $\tilde{\phi}_{\text{Na}}$.

$$\phi_{\text{Na}} = \frac{RT}{zF} \log \left(\frac{c_{\text{Na}}^e}{c_{\text{Na}}^i} \right) \quad \text{and} \quad \tilde{\phi}_{\text{Na}} = \frac{RT}{zF} \log \left(\frac{\tilde{c}_{\text{Na}}^e}{\tilde{c}_{\text{Na}}^i} \right) + \tilde{\phi}_{\text{Na},0} \quad (8.9)$$

Here, we choose $c_{\text{Na}}^e = 153$ mM [89] and $c_{\text{Na}}^i = 11.6$ mM, calibrated to match the experimental reversal potential of channelrhodopsin, and $\tilde{c}_{\text{Na}}^e = 140$ mM, $\tilde{c}_{\text{Na}}^i = 11.6$ mM, and $\tilde{\phi}_{\text{Na},0} = 1.79$ mV. Accordingly, $\phi_{\text{Chr2}} = [RT]/[zF] [\log (c_{\text{Na}}^e/c_{\text{Na}}^i) - \log (\tilde{c}_{\text{Na}}^e/\tilde{c}_{\text{Na}}^i) + \tilde{\phi}_{\text{Na},0}]$. In contrast to our previous approach, where we have photostimulated an atrial nodal cell [1], we now integrate the channelrhodopsin photocurrent I_{Chr2} into our conventional ventricular cell model [188, 197], defined through a total of $n_{\text{crt}} = 16$ ionic currents, see Figure 8.5.

$$I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{Chr2}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]. \quad (8.10)$$

Mathematically speaking, these currents are functions of the transmembrane potential ϕ , the individual gating variables g_{gate} , and the ion concentrations c_{ion} .

$$I_{\text{crt}} = I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \quad (8.11)$$

In particular, these are the fast sodium current I_{Na} , the background sodium current I_{bNa} , the sodium potassium pump current I_{NaK} , the sodium calcium exchanger current I_{NaCa} , the channelrhodopsin photocurrent I_{Chr2} the inward rectifier current I_{K1} , the

rapid delayed rectifier current I_{Kr} , the slow delayed rectifier current I_{Ks} , the plateau potassium current I_{pK} , the transient outward current I_{t0} , the L-type calcium current I_{CaL} , the background calcium current I_{bCa} , the plateau calcium current I_{pCa} , the leakage current I_{leak} , and sarcoplasmic reticulum uptake and release currents I_{up} and I_{rel} , see Appendix for details. Figure 8.7 illustrates the temporal evolution of all $n_{crt} = 16$ currents for the conventional cell, electrically stimulated for 1 ms, dashed black lines, and for the genetically engineered cell, light stimulated for 30 ms, solid blue lines. Except for the fast sodium current I_{Na} , the currents of the conventional cell and the genetically engineered cell display a relatively similar behavior. Overshoots in the potassium related currents are less pronounced for the photosensitive cell. The channelrhodopsin photocurrent I_{ChR2} is, of course, only active for the photosensitive cell.

8.4 Mathematical model of ionic concentrations

From a chemical point of view, light induces a channelrhodopsin photocurrent I_{ChR2} , which directly impacts the intracellular sodium concentration c_{Na} ,

$$\dot{c}_{Na} = -\frac{1}{FV} [I_{Na} + I_{bNa} + 3I_{NaCa} + 3I_{NaK} + I_{ChR2}] \quad (8.12)$$

where V is the cytosolic volume and F is the Faraday constant. The sodium concentration will directly, and indirectly through the resulting changes in the transmembrane potential ϕ , affect all other ionic concentrations in the cell. The biochemistry of our cell model is characterized through $n_{ion} = 4$ ion concentrations,

$$c_{ion} = [c_{Na}, c_K, c_{Ca}, c_{Ca}^{sr}] \quad (8.13)$$

where c_{Na} , c_K , and c_{Ca} are the intracellular sodium, potassium, and calcium concentrations, and c_{Ca}^{sr} is the calcium concentration in the sarcoplasmic reticulum. The concentrations obey evolution equations of the following format,

$$\dot{c}_{ion} = f_{ion}(\phi, g_{gate}, c_{ion}) \quad (8.14)$$

which are parameterized in terms of the transmembrane potential ϕ , the gating variables g_{gate} , and the ion concentrations c_{ion} themselves, see Appendix for details. Figure 8.8 illustrates the temporal evolution of all $n_{\text{gate}} = 14$ gating variables for the conventional cell, electrically stimulated for 1 ms, dashed black lines, and for the genetically engineered cell, light stimulated for 30 ms, solid blue lines. While the calcium concentrations inside the cell c_{Ca} and in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$ remain unaffected by photostimulation, the intracellular sodium concentration c_{Na} increases with light. Through the sodium potassium pump current I_{NaK} , this has a direct impact on the intracellular potassium concentration c_{K} , which decreases upon photostimulation.

8.5 Mathematical model of action potential propagation

From an electrical point of view, light induces a channelrhodopsin current I_{ChR2} , which directly impacts the sodium concentration c_{Na} , and thus, the action potential propagation ϕ in the heart. Unlike the local ion concentrations for single cells, however, the action potential ϕ is a global field variable [72]. It propagates from cell to cell via gap junction channels [?]. We can characterize its spatio-temporal evolution through a nonlinear parabolic second order partial differential equation

$$\dot{\phi} = f^{\phi}(\phi, g_{\text{gate}}, c_{\text{ion}}) + \text{div } \mathbf{q}(\phi) \quad (8.15)$$

driven by a nonlinear local source term f^{ϕ} at the single cell level, and by a linear global flux term, $\text{div } \mathbf{q}$, the divergence of the propagation vector \mathbf{q} at the organ level. We identify the local source term

$$f^{\phi} = -\frac{1}{C} [I_{\text{Na}} + I_{\text{bNa}} + I_{\text{NaK}} + I_{\text{NaCa}} + I_{\text{ChR2}} + I_{\text{K1}} + I_{\text{Kr}} + I_{\text{Ks}} + I_{\text{pK}} + I_{\text{t0}} + I_{\text{CaL}} + I_{\text{bCa}} + I_{\text{pCa}}] \quad (8.16)$$

as the negative sum of all transmembrane currents scaled by the cell membrane capacitance per unit surface area C . To account for the nonlocal nature of propagating

excitation waves in the heart, we introduce the propagation vector

$$\mathbf{q} = \mathbf{D} \cdot \nabla \phi \quad (8.17)$$

through the second order diffusion tensor \mathbf{D} scaling the gradient of the action potential field $\nabla \phi$. Figure 8.9 illustrates the five phases of the action potential ϕ for the conventional cell, electrically stimulated for 1 ms, dashed black lines, and for the genetically engineered cell, light stimulated for 30 ms, solid blue lines. We can clearly see the impact of photostimulation, which requires the entire stimulation period of 30 ms to push the cell beyond its excitation threshold. For visualization purposes, we have shifted the transmembrane potential of the electrical stimulation, dashed lines, by 34 ms, which agrees qualitatively with the time delay of 19.7 ± 3.4 ms reported in the literature [?]. Beyond the stimulation threshold, both cells display an almost identical electrical response with the characteristic rapid upstroke, early, partial repolarization, plateau, final repolarization, and rest phases. Again, this agrees nicely with experimental findings [?].

8.6 Computational model of light-activated cells

To solve the spatio-temporal problem of electrochemical coupling (10.24), (8.14), and (8.15) for the gating variables g_{gate} , the intracellular ion concentrations c_{ion} , and the transmembrane potential ϕ , we apply a finite difference scheme in time and a finite element scheme in space. Due to the global nature of the membrane potential introduced through the diffusion term $\text{div } \mathbf{q}(\phi)$, we suggest a \mathcal{C}^0 -continuous finite element interpolation for the membrane potential ϕ , while a \mathcal{C}^{-1} -continuous interpolation is sufficient for the sets of gating variables g_{gate} and ion concentrations c_{ion} [72, ?]. Accordingly, we introduce the membrane potential as global degree of freedom at each finite element node, and the gating variables and ion concentrations locally on the integration point level [197]. We solve the resulting staggered system with an incremental iterative Newton-Raphson solution procedure based on the consistent linearization of the discrete excitation problem [?, ?, 73].

8.6.1 Chemical problem - Local discretization on the integration point level

We typically initialize the chemical state variables at t_0 with their resting state values. Our goal is to determine the chemical state variables, i.e., the $n_{\text{gate}} = 14$ gating variables g_{gate} and the $n_{\text{ion}} = 4$ ion concentrations c_{ion} of the current time step t for a given action potential ϕ at t , and known chemical state variables from the previous time step t^n . For the temporal discretization, we partition the time interval of interest \mathcal{T} into n_{stp} subintervals $[t^n, t^{n+1}]$ as $\mathcal{T} = \bigcup_{n=0}^{n_{\text{stp}}-1} [t^n, t^{n+1}]$ and apply a standard backward Euler time integration scheme in combination with a finite difference approximation of the first order time derivatives $\dot{g}_{\text{gate}}\phi$ and \dot{c}_{ion} .

$$\dot{g}_{\text{gate}} = [g_{\text{gate}} - g_{\text{gate}}^n] / \Delta t \quad \text{and} \quad \dot{c}_{\text{ion}} = [c_{\text{ion}} - c_{\text{ion}}^n] / \Delta t \quad (8.18)$$

Here and from now on, we omit the index $(\circ)^{n+1}$ for the sake of clarity, and introduce the common abbreviation $\Delta t := t - t^n > 0$ for the current time increment. This allows us to transform the linear set of gating equations (10.24) into a set of update equations for the gating variables g_{gate} at the current time step t .

$$g_{\text{gate}} = g_{\text{gate}}^n + \frac{1}{\tau_{\text{gate}}(\phi)} [g_{\text{gate}}^\infty(\phi) - g_{\text{gate}}] \Delta t \quad (8.19)$$

The gating variables define the $n_{\text{crt}} = 16$ ionic currents $I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}})$, which alter the $n_{\text{ion}} = 4$ intracellular ion concentrations c_{ion} . With the help of the finite difference approximation (8.18), we reformulate the nonlinear set of concentration equations (8.14) in the following residual format.

$$\mathbf{R}_{\text{ion}}^c = c_{\text{ion}} - c_{\text{ion}}^n - f_{\text{ion}}^c(\phi, g_{\text{gate}}, c_{\text{ion}}) \Delta t \doteq 0 \quad (8.20)$$

To obtain the $n_{\text{ion}} \times n_{\text{ion}}$ iteration matrix $\mathbf{K}_{\text{ion ion}}^c$ for the local Newton iteration on the integration point level, we linearize the discrete algorithmic residual.

$$\mathbf{K}_{\text{ion}}^c = d_{c_{\text{ion}}} \mathbf{R}_{\text{ion}}^c = 1 - d_{c_{\text{ion}}} f_{\text{ion}}^c \quad (8.21)$$

At the end of each Newton iteration, we update the set of ion concentrations $c_{\text{ion}} \leftarrow c_{\text{ion}} - [\mathbf{K}_{\text{ion}}^c]^{-1} \mathbf{R}_{\text{ion}}^c$, the set of gating variables $g_{\text{gate}} \leftarrow g_{\text{gate}} + [g_{\text{gate}}^\infty(\phi) - g_{\text{gate}}] \Delta t / \tau$ and the set of ionic currents $I_{\text{crt}} \leftarrow I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}})$, and continue to iterate locally until we achieve convergence. This local inner loop can be understood as a modern implicit version of the iterative update procedure of the original Rush-Larsen algorithm [160].

8.6.2 Electrical problem - Global discretization on the node point level

To evaluate the action potential ϕ on the node point level, we transform the electrical problem (8.15) into its residual format

$$\mathbf{R}^\phi = \dot{\phi} - \text{div}(\mathbf{q}) - f^\phi \doteq 0 \quad (8.22)$$

in the domain of interest \mathcal{B} . We complement this initial boundary value problem with the corresponding Dirichlet and Neumann boundary conditions $\phi = \bar{\phi}$ on $\partial\mathcal{B}_\phi$ and $\mathbf{q} \cdot \mathbf{n} = \bar{q}$ on $\partial\mathcal{B}_q$, typically using homogeneous Neumann boundary conditions $\mathbf{q} \cdot \mathbf{n} = 0$ on the entire boundary $\partial\mathcal{B}$. As initial conditions, $\phi_0(x) = \phi(x, t_0)$ in \mathcal{B} , we simply set the transmembrane potential to its resting state. We obtain the weak form of the electrical residual (8.22) through the integration over the domain \mathcal{B} , the standard integration by parts, and the inclusion of the Neumann boundary conditions. For the spatial discretization, we discretize the domain of interest \mathcal{B} with n_{el} finite elements \mathcal{B}^e as $\mathcal{B} = \bigcup_{e=1}^{n_{\text{el}}} \mathcal{B}^e$ and apply the standard isoparametric concept to interpolate the trial functions ϕ and the test functions $\delta\phi$.

$$\delta\phi = \sum_{i=1}^{n_{\text{en}}} N^i \delta\phi_i \quad \phi = \sum_{j=1}^{n_{\text{en}}} N^j \phi_j \quad (8.23)$$

Here, N are the standard shape functions on the element level and $i, j = 1, \dots, n_{\text{en}}$ are the n_{en} element nodes. For the temporal discretization, we suggest a finite difference approximation of the first order time derivative $\dot{\phi}$,

$$\dot{\phi} = [\phi - \phi^n] / \Delta t \quad (8.24)$$

in combination with a standard backward Euler time integration scheme. Using the discretizations in space (11.6) and time (8.24), we can introduce the discrete algorithmic residual \mathbf{R}_I^ϕ .

$$\mathbf{R}_I^\phi = \mathbf{A}_{e=1}^{n_{\text{el}}} \int_{\mathcal{B}^e} N^i \frac{\phi - \phi^n}{\Delta t} + \nabla N^i \cdot \mathbf{q} \, dV - \int_{\partial \mathcal{B}_q^e} N^i \bar{q} \, dA - \int_{\mathcal{B}^e} N^i f^\phi \, dV \doteq 0 \quad (8.25)$$

The operator \mathbf{A} symbolizes the assembly of all element contributions at the element nodes $i = 1, \dots, n_{\text{en}}$ to the overall residual at the global node points $I = 1, \dots, n_{\text{nd}}$. To solve the discrete system of nonlinear equations (8.25), we suggest an incremental iterative Newton Raphson solution technique based on the consistent linearization of the residual, which introduces the global iteration matrix \mathbf{K}_{IJ}^ϕ .

$$\mathbf{K}_{IJ}^\phi = d_{\phi_J} \mathbf{R}_I^\phi = \mathbf{A}_{e=1}^{n_{\text{el}}} \int_{\mathcal{B}^e} N^i \frac{1}{\Delta t} N^j + \nabla N^i \cdot \mathbf{D} \cdot \nabla N^j - N^i d_{\phi} f^\phi N^j \, dV \quad (8.26)$$

For each incremental iteration, we update the global vector of unknowns $\phi_I \leftarrow \phi_I - \sum_{J=1}^{n_{\text{nd}}} \mathbf{K}_{IJ}^{\phi-1} \mathbf{R}_J^\phi$ at all $I = 1, \dots, n_{\text{nd}}$ global nodes. At convergence of the local Newton iteration, i.e., at chemical equilibrium, we can evaluate the source term $f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ for the electrical problem (8.25), and its linearization $d_{\phi} f^\phi(\phi, g_{\text{gate}}, c_{\text{ion}})$ for the global Newton iteration (8.26). The use of a fully monolithic implicit solution algorithm allows us to apply an adaptive time stepping procedure, for which the time step size is automatically adjusted in response to the number of Newton iterations towards global equilibrium [197]. In principle, the above equations for \mathbf{R}_I^ϕ and \mathbf{K}_{IJ}^ϕ are generically portable and can interface with any commercially available finite element package. Here, we embed the proposed algorithm in the general multipurpose nonlinear finite element program FEAP [?].

8.7 Computational model of a human heart

To illustrate the features of the proposed algorithm, we model the photostimulation of a human heart stimulated at three different pacing sites. We create a patient-specific heart model from magnetic resonance images taken at different depths along

the heart's long axis, see Figure 8.10. On these two-dimensional slices, we use semi-automated image processing tools to section the regions of the cardiac muscle. From the raw and noisy grayscale images, we create monochrome images with sharply defined boundaries using thresholding and binary masking. From the resulting black and white slices, we generate three preliminary triangular surface meshes, one for the outer surface and one for the inner surface of each ventricle. Finally, we convert the surface meshes into a tetrahedral volume mesh consisting of 3,129 nodes and 11,347 tetrahedral elements [102], see Figure 8.10.

We virtually inject our genetically engineered, light sensitive cardiac cells into different regions of the heart [?], and model all other cells as conventional cardiac muscle cells [188, 197]. The former group of cells possesses the additional channelrhodopsin current I_{ChR2} , while the latter does not. Following the literature, we model cell injection by modifying the material properties of the myocardial wall at the injection sites [193]. Here, these regions span a volume of approximately 0.02cm^3 , which is roughly the size of the $0.6 \times 0.3 \times 0.1\text{cm}^3$ large atrioventricular node in humans, corresponding to approximately 10^7 cells [?]. To initiate electrical activation, we photostimulate the light sensitive cells for a period of 30 ms with blue light, then turn off the light, and record the action potential propagation across the heart.

8.7.1 Atrioventricular node pacing through photostimulation

Figure 8.11 illustrates the activation profile in response to atrioventricular node pacing through photostimulation. To reproduce the native activation sequence of the human heart, we virtually inject photosensitive cells into the basal septal region, the location of the atrioventricular node, and pace them with blue light. A depolarization wave forms in the region of the atrioventricular node, travels down the septum, and activates the left and right ventricles. When comparing the temporal evolution of the transmembrane potential ϕ , top row, with the different ion concentrations c_{Na} , c_{K} , and c_{Ca} , lower rows, we observe that the sodium concentration changes first, with a sharp front, followed by changes in the potassium and calcium concentrations, with smoother fronts. These observations are in agreement with the temporal evolution

of the local ion concentrations illustrated in Figure 8.8. After approximately 50 ms, the entire heart is depolarized. The transmembrane potential ϕ has changed from approximately -80 mV to +20 mV.

8.7.2 Apical pacing through photostimulation

Figure 8.12 shows the response of the heart to apical pacing through photostimulation. The right ventricular apex has traditionally been the most prominent pacing site, providing a stable lead position associated with few complications. To predict the activation sequence in response to apical pacing, we virtually inject photosensitive cells into the apex, stimulate them with light for 30 ms, turn off the light, and monitor the electrical and chemical fields. As the transmembrane potential profile indicates, a depolarization wave forms at the apex, travels up towards the base, and activates the septum and both ventricles simultaneously. Again, we observe the rapid increase in the intracellular sodium concentration c_{Na} causing the heart to depolarize. The transmembrane potential changes from -80 mV to +20 mV within the first 50 ms of the cycle, until the heart is completely depolarized. The intracellular calcium concentration c_{Ca} follows smoothly, while the intracellular potassium concentration c_{K} changes only in the later stages.

8.7.3 Biventricular pacing through photostimulation

Figure 8.13 displays the activation of the heart, photostimulated through bi-ventricular pacing. Biventricular pacing is a common pacing strategy in patients with heart failure for which both ventricles do not contract synchronously. To simulate biventricular pacing, we virtually inject photosensitive cells into the lateral walls of the left and right ventricles and pace both locations simultaneously. In response to photostimulation, two depolarization waves form, one at each pacing site. The waves travel along the ventricles to finally activate the apex and the septum. Again, during depolarization, the transmembrane potential profile ϕ closely follows the intracellular sodium concentration c_{Na} , which triggers the rapid upstroke of the action potential, while the intracellular calcium concentration c_{Ca} follows slightly later with a smoother

propagation front.

8.8 Discussion

8.8.1 Comparison with the literature

Within the past decade, optogenetics has become a tremendously active field of research [46], however, little is known about its impact on cardiac cells, neither in isolation, nor in the cardiac system. Since the channelrhodopsin photocycle itself is relatively well established from intense studies of the neuronal system [138, 139], we simply adopt the kinetics of a well-characterized three-state photocycle model [79, 133]. Creating a cellular model system that expresses Channelrhodopsin-2 allows us to validate its photocurrent using whole cell patch-clamp [?] or multielectrode array recordings [?]. Using both techniques, we have recently demonstrated that our cell model is capable of reliably reproducing experimentally measured photostimulation amplitudes, pulse widths, and frequencies from single cell recordings [1]. In this manuscript, for the first time, we quantify the impact of photostimulation in side-by-side comparisons between conventional cells, electrically stimulated for 1 ms, and genetically engineered cells, optically stimulated for 30 ms, see Figures 8.6 - 8.9. The gating variables in Figures 8.6 can be understood as internal variables, which cannot be measured experimentally. The ionic currents in Figure 8.7 can be extracted experimentally from a series of voltage and current clamp recordings using drugs to selectively block the activity of different combinations of channels. Although theoretically possible, for optogenetically engineered cells, typically only the channelrhodopsin photocurrent is measured experimentally. Its indirect impact on all other currents, however, has not been characterized systematically to date. The ion concentrations in Figure 8.8 could potentially be measured experimentally, however, for optogenetically engineered cells, this has not been done to date. Unfortunately, only three groups worldwide have reported the use of optogenetics to manipulate the cardiac system, one in zebrafish [5], one in transgenic mice [?], and one in stem cell derived cardiomyocytes [1]. The last study by our own group is the only one that

has demonstrated the potential of optogenetics in human heart cells. This underlines the importance of computational models, which provide easy access to ion channel dynamics and ionic concentrations, both in healthy and diseased conditions.

The transmembrane potential in Figure 8.9 could potentially be extracted from patch clamp or microelectrode array recordings. Here, for the selected light intensity, we find 30 ms to be the minimum light exposure time to evoke a standard action potential. For photostimulation just slightly above this threshold, we identify the time delay between electrical stimulation and photostimulation to be 34 ms, which agrees nicely with the values of 19.7 ± 3.4 ms reported in the literature [?]. Although not explicitly reported here, in agreement with the literature, we observe a decrease in activation delay for increased light intensities [?]. Beyond the stimulation threshold, however, electrically and optically activated cells display an almost identical electrophysiology with the characteristic rapid upstroke, early, partial repolarization, a pronounced plateau, a final repolarization, and a return to the resting phase, see Figures 8.9. Again, this agrees nicely with experimental findings, which report that channelrhodopsin expression in murine embryonic stem cells is not associated with additional leak currents, or changes in the resting membrane potential, membrane resistance, and action potential duration [?].

A first qualitative study in zebrafish has demonstrated that optogenetic tools can, in principle, be used to control the cardiac system *in vivo* [5]. In that study, optical perturbation of genetically engineered pacemaker cells was entirely reversible. This agrees nicely with our optically stimulated transmembrane potential returning smoothly to its resting state, see Figure 8.9. The study also showed that only few genetically engineered cells are needed to deliver a sufficient optical stimulus in zebrafish hearts [5]. In our human heart, we virtually inject approximately 10^7 cells covering a volume of 0.02 cm^3 , which is roughly the size of the atrioventricular node in humans, see Figure 8.10. In agreement with an *in vivo* study in transgenic mice [?], we observe that smaller cell volumes require either increased light intensities or longer stimulation times. Similar to the reported activation delays of 12.2 ± 3.7 ms and 9.3 ± 2.2 ms upon atrial and ventricular photostimulation in mice [?], we observe an activation delay of 19.5 ms when comparing our light activated heart in Figure

8.11 with previous simulations of native activation sequences [102, 185] and electrical pacing [197].

Overall, because of the inherent frequency mismatch in rodents and humans, the complete *in vivo* validation of our model remains challenging, and can only be addressed satisfactorily by means of large animal models. This is, however, a generic limitation of the clinical translational character of cardiac optogenetics in general. Large animal models would also allow us to quantify the impact of optogenetical stimulation on cardiac function *in vivo*. In a previous study, we have demonstrated how light can be turned into force using optogenetic tools *in vitro* [1]. A related study by another group reported light-induced mechanical strains of up to 4.5% [?]. We are currently integrating our continuum model for the photoelectrochemistry of living systems into a model for excitation-contraction coupling [73]. This will allow us to characterize the mechanically-mediated ionic response of strain-gated ion channels [?, ?], and, vice versa, the calcium-mediated mechanical response of the heart [?, ?], and its impact on muscle contraction and cardiac function [?, ?].

8.8.2 Conclusion

At an unprecedented temporal and spatial precision, optogenetic tools enable us to manipulate electrically active cells with light. This study documents our first attempts to model genetically engineered light-sensitive heart muscle cells, with the ultimate goal to control the cardiac system by mean of photostimulation. In contrast to electrical pacemakers, light delivery is minimally invasive, genetically targeted, and temporally precise. Most importantly, light can be delivered at a distance. Unlike pacing leads for electrical stimulation, which are known to have a high perforation risk, the light source for optical stimulation does not need to be in direct contact with the moving heart muscle. Light pacing might therefore be an attractive remote, less invasive, and more durable alternative to conventional electrical pacing. Unlike traditional electrical stimulation, optogenetic stimulation allows us to precisely control the selective permeability of the plasma membrane, its conductivity with respect to different ions, its sensitivity to light of different wavelengths, and the spatio-temporal

evolution of different opening and closing profiles. Given this incredible freedom, there is a pressing need to establish economic strategies to optimize the matrix of input variables. Predictive continuum models for photoelectrochemistry allow us to virtually probe landscapes of process parameters and identify optimal photostimulation sequences in various different tissues and organs. Discretized with finite element tools, they enable the combination of virtually any cell type on virtually any geometry. Our time-adaptive algorithm is highly efficient, conceptually modular, and easily portable to any commercial nonlinear finite element platform. While the photoelectrochemical behavior of living tissue is almost impossible to characterize *in vivo*, and little is known about the spatio-temporal distribution of sodium, potassium, and calcium ions, computational tools like ours allow us to predict action potential profiles and ionic concentrations across the heart *in silico* within the order of minutes. We believe that computational optogenetics will be widely applicable to predict the response of other genetically engineered, electrically-active cells and, ultimately, support the design of novel therapies for various types of neuronal, musculoskeletal, pancreatic, and cardiac disorders such as depression, schizophrenia, cerebral palsy, paralysis, diabetes, pain syndromes, and cardiac arrhythmias.

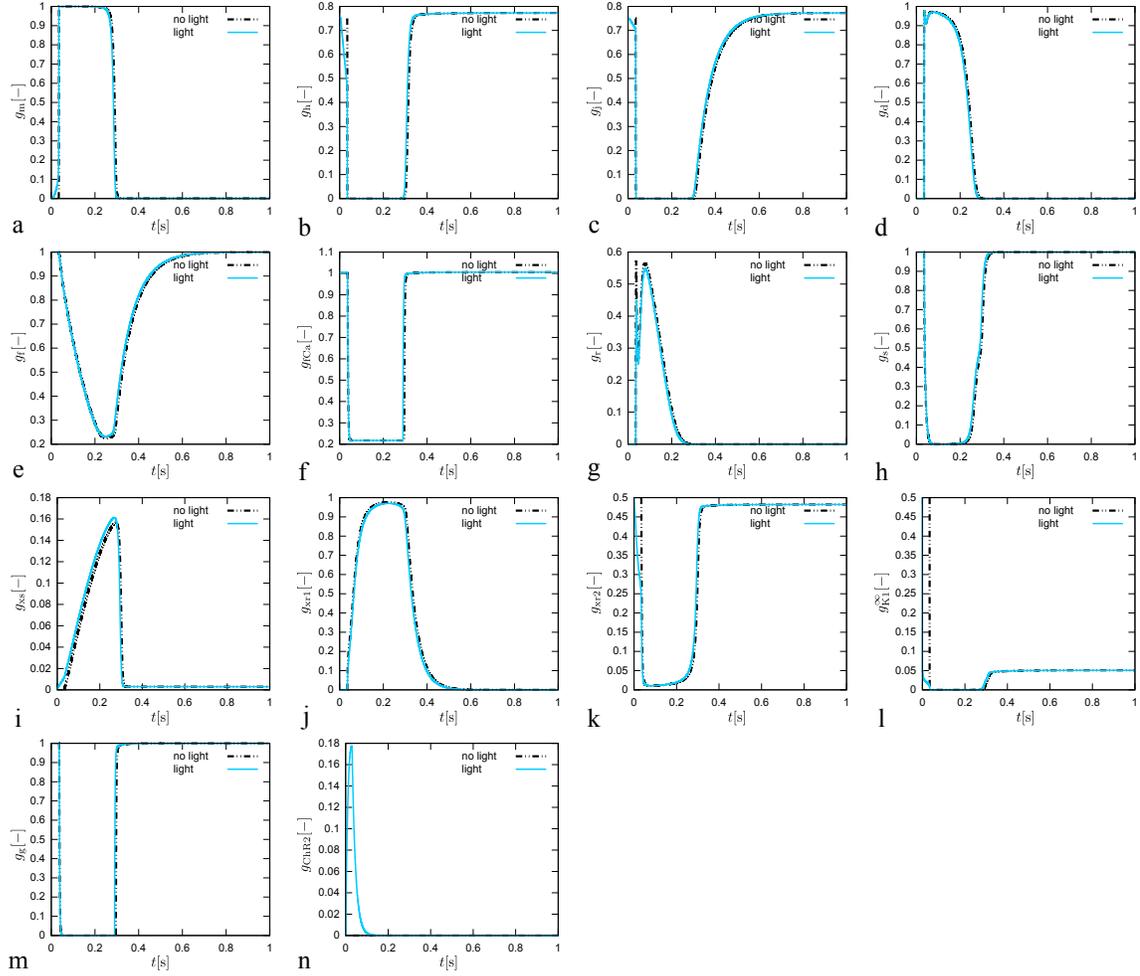


Figure 8.6: Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of sodium activation gate g_m , fast sodium inactivation gate g_h , slow sodium inactivation gate g_j , L-type calcium activation gate g_d , L-type calcium inactivation gate g_f , intracellular calcium dependent calcium inactivation gate g_{TCa} transient outward activation gate g_r , transient outward inactivation gate g_s , slow delayed rectifier gate g_{xs} , rapid delayed rectifier activation gate g_{xr1} , rapid delayed rectifier inactivation gate g_{xr2} , inward rectification factor g_{K1}^∞ , calcium-dependent inactivation gate g_g , and channelrhodospin activation gate g_{ChR2} . The gating dynamics for the electrically stimulated cell have been delayed by 34ms for the purposes of comparison against the optically stimulated cell.

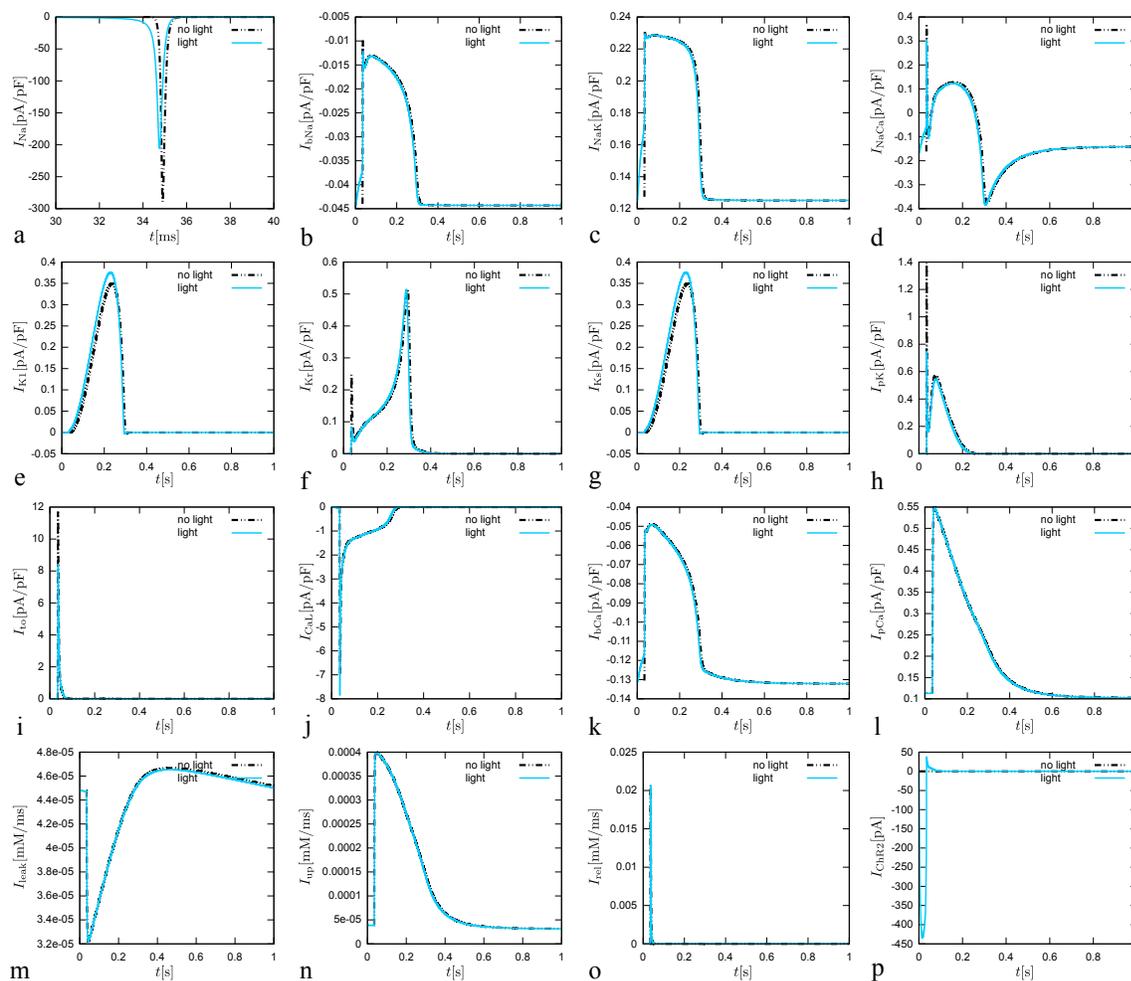


Figure 8.7: Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of fast sodium current I_{Na} , background sodium current I_{bNa} , sodium potassium pump current I_{NaK} , sodium calcium exchanger current I_{NaCa} , inward rectifier current I_{K1} , rapid delayed rectifier current I_{Kr} , slow delayed rectifier current I_{Ks} , plateau potassium current I_{pK} , transient outward current I_{t0} , L-type calcium current I_{CaL} , background calcium current I_{bCa} , plateau calcium current I_{pCa} , leakage current I_{leak} , sarcoplasmic reticulum uptake current I_{up} , sarcoplasmic reticulum release current I_{rel} , and channelrhodopsin current I_{ChR2} . The current dynamics for the electrically stimulated cell have been delayed by 34 ms for the purposes of comparison against the optically stimulated cell.

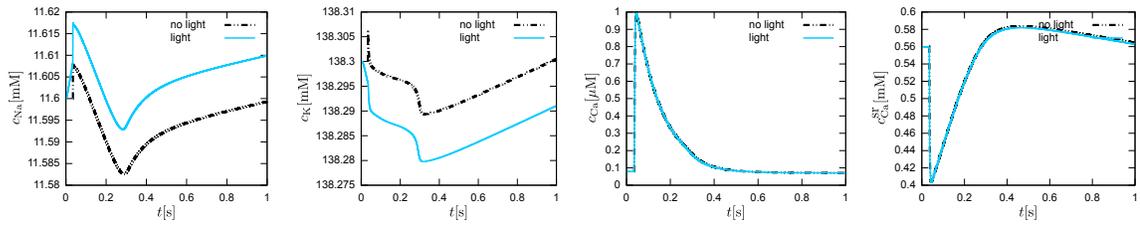


Figure 8.8: Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of intracellular sodium c_{Na} , potassium c_K , calcium c_{Ca} concentrations, and calcium concentration c_{Ca}^{SR} in the sarcomplastic reticulum. The chemical concentration dynamics for the electrically stimulated cell have been delayed by 34 ms for the purposes of comparison against the optically stimulated cell.

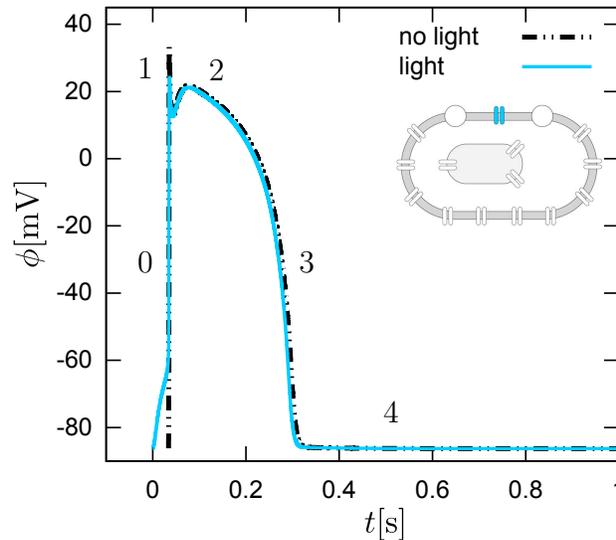


Figure 8.9: Genetically engineered light sensitive cardiac cell stimulated conventionally with an electric field, dashed lines, and optically with light, solid lines. Temporal evolution of transmembrane potential ϕ . The characteristic action potential consists of five phases. Phase 0: The rapid upstroke is generated through an influx of positively charged sodium ions. Phase 1: Early, partial repolarization is initiated through the efflux of positively charged potassium ions. Phase 2: During the plateau, the net influx of positively charged calcium ions is balanced by the efflux of positively charged potassium ions. Phase 3: Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. Phase 4: Throughout the interval between end of repolarization and the beginning of the next cycle the cell is at rest. The transmembrane potential for the electrically stimulated cell has been delayed by 34 ms for the purposes of comparison against the optically stimulated cell. This delay agrees nicely with the time delay of activation of 19.7 ± 3.4 ms reported in the literature [?].

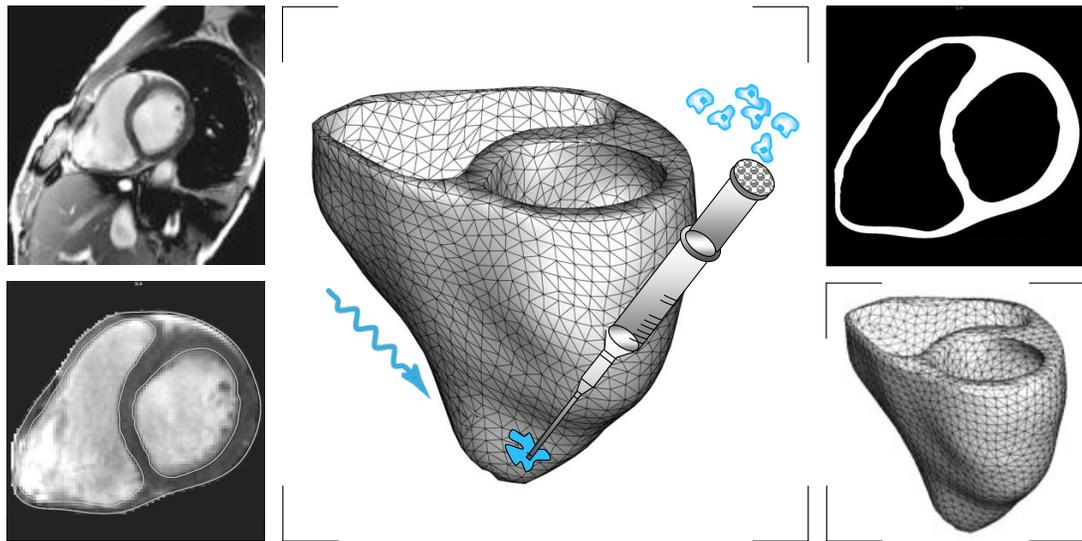


Figure 8.10: Virtual injection of genetically engineered light sensitive cardiac cells into a human heart. Magnetic resonance imaging generates a sequence of two-dimensional images at different depths (top, left). We segment cardiac muscle tissue semi-manually using standard image processing techniques (bottom, left). Thresholding and binary masking convert the raw grayscale images to monochrome images with sharply defined boundaries (top, right). From these slices, we create a preliminary triangular surface mesh and converted it into the final tetrahedral volume mesh consisting of 3,129 nodes and 11,347 tetrahedral elements (bottom, right). Last, we virtually inject photosensitive cells into different regions of the heart and stimulated with light (middle).

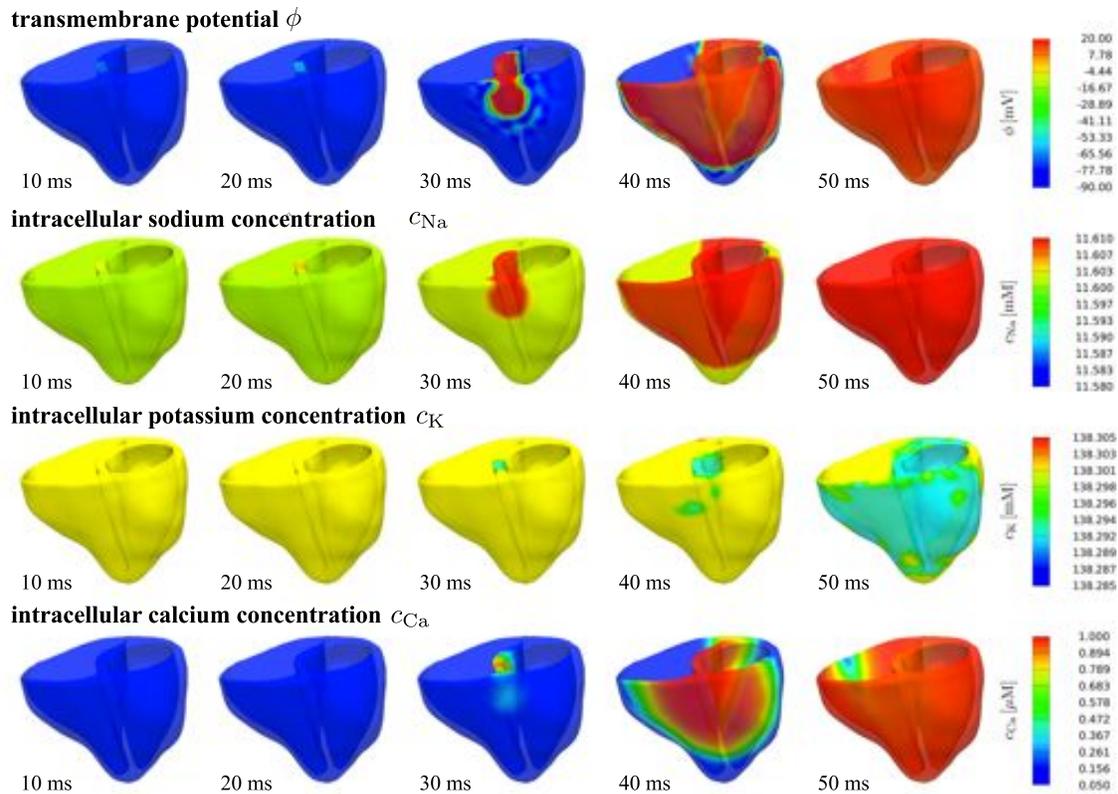


Figure 8.11: Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for atrioventricular node paced heart. Photosensitive cells are virtually injected into the basal region of the septum, while all other regions are modeled as conventional cardiac muscle cells. A depolarization wave forms at the atrioventricular node, travels down the septum, and activates the left and right ventricles.

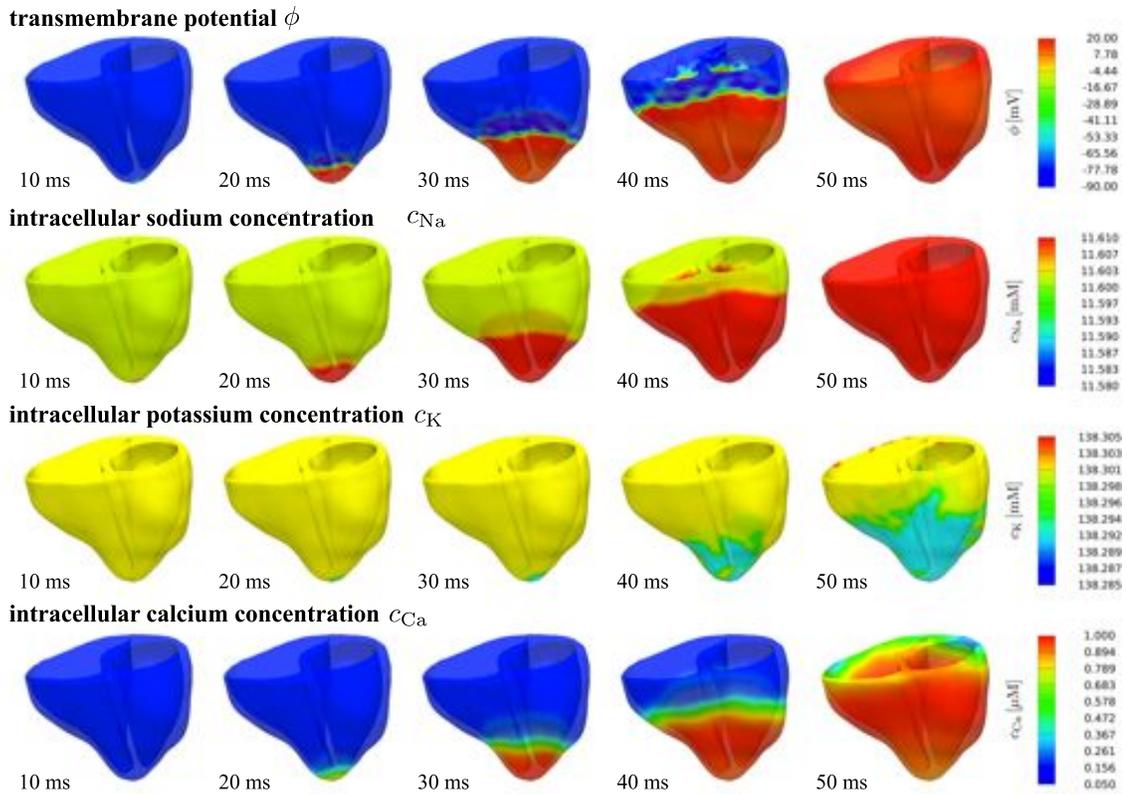


Figure 8.12: Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for apically paced heart. Photosensitive cells are virtually injected into the apex, while all other regions are modeled as conventional cardiac muscle cells. A depolarization wave forms at the apex, travels up, and activates the septum and both ventricles simultaneously.

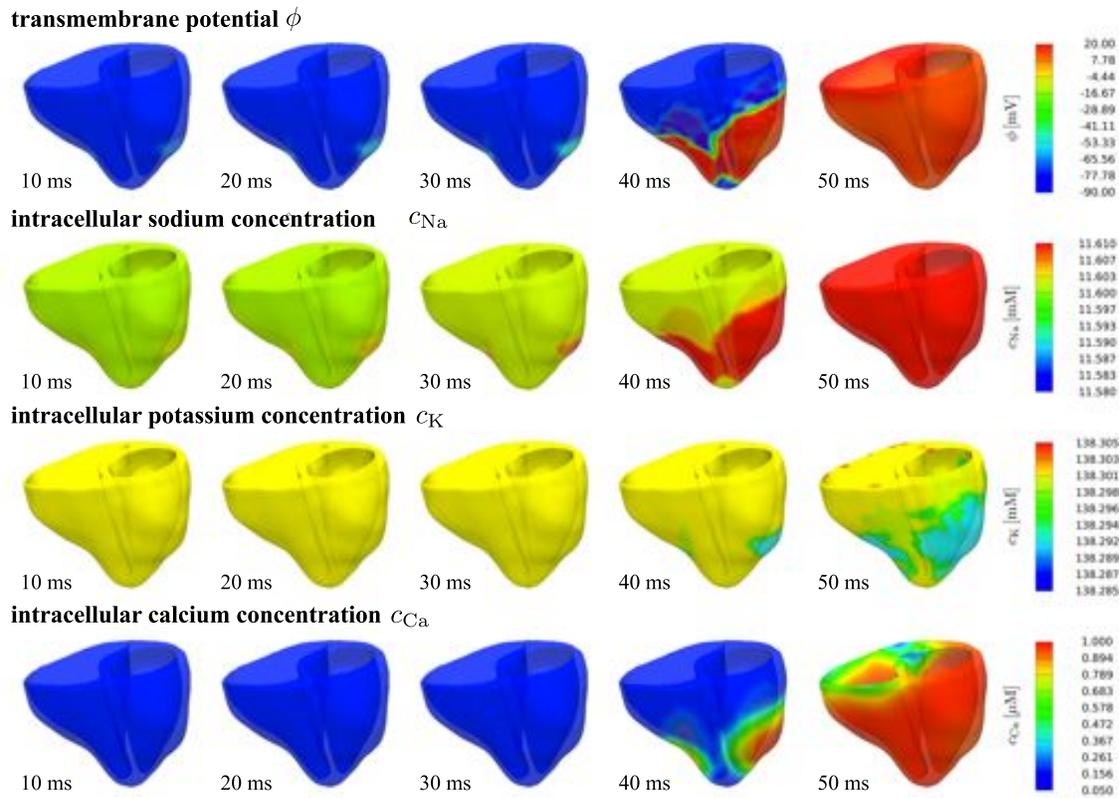


Figure 8.13: Photostimulation of a human heart. Spatio-temporal evolution of transmembrane potential ϕ , intracellular sodium c_{Na} , potassium c_K , and calcium c_{Ca} concentrations for bi-ventricularly paced heart. Photosensitive cells are virtually injected into the lateral walls of the left and right ventricles, while all other regions are modeled as conventional cardiac muscle cells. Two depolarization waves form in the lateral left and right ventricular walls to travel along the ventricles and activate the apex and the septum.

Chapter 9

Designating fibers in the heart

Up until this point, we have assumed that fiber and sheet directions within the heart are already defined for a given heart geometry. However, this is generally not the case. While it is possible to measure fiber directions with a variety of medical imaging modalities, they are not performed often during the course of clinical protocols. Luckily the general fiber orientation distribution of the heart is known, and this chapter proposes a convenient geometrical-based algorithm that utilizes common finite element tools to generate fiber and sheet directions for arbitrary heart meshes algorithmically.

9.1 Introduction

The structural architecture of the heart is one of the most important factors to healthy pump function [?]. While many geometric factors such as shape, thickness, and size play a critical role in disease diagnosis, the correct alignment of cardiac muscle fibers is crucial for proper choreography of cardiac muscle contraction [?]. In fact, many cardiac disorders or traumatic events induce cardiac fiber misalignment, which may cyclically reduce cardiac function and lead to further disorganization [?]. For example, in response to cardiac infarctions, residual fibers are secreted to structurally stabilize the heart. However, these fibers lose organization and may actually form nodes of orthogonal myofiber intersection or contact that impede proper cardiac function

[171]. Other pathologies such as myocardial fibrosis are also known to modify organized fiber orientation, ultimately impairing both electrical and mechanical function [164]. Therefore, it is imperative for computational models of the heart to include proper fiber orientation distributions to accurately predict the electrical and mechanical function of the heart in health and disease.

Fortunately, the overall structural alignment of the human heart is well characterized in the literature [173], where fibers have been reported to vary transmurally from approximately -70° in the epicardium, the outer wall, to $+80^\circ$ in the endocardium, the inner wall [6, 61]. However, along the septal wall, the right endocardium displays a fiber orientation of approximately -70° . Although there is a wide agreement on these fiber orientations in the epicardial and endocardial walls, there is no generally accepted concept to assign these fiber orientations to realistic anatomical computer models of the human heart. Simplified attempts represent the two ventricles as nested, truncated ellipsoids and allow for an analytical characterization of the fiber orientation distribution [66]. Other techniques rely on complex projections and use least squares fitting to generate a continuum description of fiber orientations [181]. Recent attempts have adopted a geometrically-based approach to mitigate some of the difficulties in determining approximate fiber orientations in irregular geometries [177].

When heart fiber orientations were first characterized, the goal was to identify an analytical description for a single ellipsoid representing the left ventricle [137]. Initially, Hammer projections served to map a real heart geometry onto an idealized prolate spheroidal ellipsoid to create a mathematical model of the fiber directions in the heart. A more recent variation of this technique has been applied to diffusion tensor image interpolation [181]. The irregular geometry is first projected to a prolate spheroid coordinate system, and then smoothed using an averaging kernel and least squares approach to generate an interpolation in the idealized domain. In addition, some groups use least squares techniques to fit the experimentally measured fiber directions to an approximated geometrical mesh [42, 181]. The fiber directions from a least squares approach can then be interpolated onto element centroids, nodes, or integration points of the finite element mesh by using the corresponding finite element shape functions. For more complex interpolating schemes, different weights can be

assigned to the individual interpolants. Most recently, a geometrical-based approach was proposed to provide sketch-based methods for generating layered fiber structures for the heart [177].

Unfortunately, there are several essential drawbacks in using ellipsoid, projected, or least squares techniques. For example, some difficulties may arise when trying to determine proper projections for mapping a non-ellipsoidal geometry onto an idealized ellipsoidal representation of the heart. With the relative accessibility of different three-dimensional imaging modalities today, we can attain patient-specific heart geometries relatively easily [103]. Generally these geometries do not look like ellipsoids, and there may be inherent singularities when mapping each ventricle onto an idealized ellipsoid, in particular in the septal region and close to the apex [182]. Of all existing techniques, least squares techniques might therefore be best suited to fit experimentally measured data. However, they may require some fine tuning to create sufficiently continuous fiber orientations. In addition, they typically require a sufficiently complete set of experimental data.

Overall, the challenges above mainly add complexity to designating fiber orientation maps throughout the heart. Here we propose a novel fiber interpolation concept, which significantly reduces the complexity of generating fiber orientation distributions within arbitrary geometries. In addition, it allows for more flexibility with regards to mesh refinement and is easily accessible and intuitive to the field of computational biomechanics. The underlying approach is adopted from feature-based interpolation techniques [199] used in computer graphics [55], and is conceptually similar to sketch-based methods [177]. However, our approach is finite-element based and provides a straightforward refinement to properly describe both ventricles of the heart, while maintaining a continuous fiber orientation distribution. The key to reducing complexity, ironically, is to reformulate the orientation of the heart from an explicit representation of measured or mathematically parametrized orientations to a more implicit one. This eliminates having to make difficult and subjective assumptions by simply calculating the fiber orientation as the solution to a boundary value problem. While the concept is inherently applicable to different types of boundary value problems, here we use the Poisson equation as an interpolation method to solve

for our designated fiber angle orientation distribution.

This paper is organized into four additional sections. First, in Section 9.2, we introduce the Poisson interpolation, the generation of local cardiac coordinate systems, and the creation of Dirichlet boundary conditions, which serve as the basis for the fiber interpolation. Then, in Section 9.3, we illustrate the concept of the Poisson interpolation on an idealized geometry to quantitatively compare it to an analytical fiber orientation and on a patient-specific geometry to demonstrate the generality of the underlying scheme. We also perform two sensitivity analyses to identify the amount of information required to generate smooth fiber orientation maps. This mimics the potential to interpolate fiber orientations distributions using experimentally acquired data. Last, in Section 9.4, we discuss the merits and limitations of our algorithm and highlight future possible avenues.

9.2 Methods

9.2.1 Poisson interpolation of scalar-valued feature θ

The key idea of our fiber orientation algorithm is to generate a smooth coordinate-free fiber interpolation using algorithms from computer graphics motivated by feature-based interpolation. In essence, we solve the following Poisson equation for the scalar-valued feature θ ,

$$\operatorname{div}(\mathbf{K} \cdot \nabla \theta) = 0 \quad \text{in } \mathcal{B} \quad (9.1)$$

for given Dirichlet boundary conditions

$$\theta = \bar{\theta} \quad \text{on } \partial \mathcal{B}_\theta. \quad (9.2)$$

For isotropic constant diffusion coefficients, $\mathbf{K} = K \mathbf{I}$, the Poisson equation reduces to the homogeneous Laplacian, $\Delta \theta = 0$, generating a smooth linear interpolation of the feature θ across the cardiac domain \mathcal{B} . To solve the underlying partial differential equation (9.1), we transform it into its weak form, integrate it over the domain \mathcal{B} , discretize it with linear tetrahedral elements, and solve the resulting system with

the corresponding Dirichlet boundary conditions (9.2). This implies that we solve the fiber interpolation (9.1) in a weak sense, while we enforce the set of user-defined constraints (9.2) in a strong sense. For the finite element solution, we can adopt any standard Finite Element program for linear diffusion or thermal problems, and interpret the scalar-valued feature θ as the primary unknown.

9.2.2 Local cardiac coordinate system $\{\mathbf{n}, \mathbf{c}, \mathbf{z}\}$

To establish a local cardiac coordinate system $\{\mathbf{n}, \mathbf{c}, \mathbf{z}\}$, we identify the normal \mathbf{n} , circumferential \mathbf{c} , and longitudinal \mathbf{z} directions at each node [186]. First, we assign the vector \mathbf{z} pointing along the long axis of the heart. Second, we calculate the nodal normals \mathbf{n} using an area-weighted face-normal averaging algorithm commonly used in computer graphics [121]. We calculate the facet normals to each surface triangle by taking the cross product of two edge vectors of the corresponding facet. We then calculate the nodal normal \mathbf{n} as average over all facet normals connected at the particular node. Since the magnitude of the cross product is proportional to the area of the corresponding facet, our nodal normal is automatically area averaged. Third, we calculate the circumferential direction \mathbf{c} as the weighted cross product of the longitudinal and normal vectors \mathbf{z} and \mathbf{n} ,

$$\mathbf{c} = \phi [\mathbf{z} \times \mathbf{n}], \quad (9.3)$$

To ensure a proper circumferential orientation we weight the cross product with the orientation index ϕ ,

$$\phi = \bar{\phi} = \begin{cases} +1 & \text{on } \partial\mathcal{B}_{\phi^+} \\ -1 & \text{on } \partial\mathcal{B}_{\phi^-} \end{cases}, \quad (9.4)$$

which is positive on the epicardium and the right septal endocardium $\partial\mathcal{B}_{\phi^+}$ and negative on the left endocardium and right free wall endocardium $\partial\mathcal{B}_{\phi^-}$.

9.2.3 Nodal fiber and sheet orientations \mathbf{f} and \mathbf{s}

We can now assign the fiber and sheet orientations \mathbf{f} and \mathbf{s} . In particular, we assume that fibers \mathbf{f} at the inner and outer surfaces of the heart lie within the tangent plane to the corresponding surface. To begin, we calculate the outward pointing normal \mathbf{n}_{cz} with respect to the local cz -plane,

$$\mathbf{n}_{cz} = \mathbf{c} \times \mathbf{z}. \quad (9.5)$$

Although we could, in principle, assign any given sheet angle β to define the sheet direction, here, without loss of generality, we select the sheet normal \mathbf{s} to be aligned with the outward facing nodal normal,

$$\mathbf{s} = \text{sign}(\mathbf{n}_{cz} \cdot \mathbf{n}) \mathbf{n}. \quad (9.6)$$

To specify the fiber angle with respect to the horizontal plane, we assume that all fibers within the epicardium and right septal endocardium are inclined with -70° and all fibers within the left endocardium and right free wall endocardium are inclined with 80° [137],

$$\alpha = \begin{cases} -70^\circ & \text{on } \partial\mathcal{B}_{\phi^+} \\ +80^\circ & \text{on } \partial\mathcal{B}_{\phi^-} \end{cases}. \quad (9.7)$$

We then construct the projection \mathbf{p} of the fiber direction \mathbf{f} on the cz -plane corresponding to the fiber orientation angle α with respect to the circumferential direction \mathbf{c} ,

$$\mathbf{p} = \text{proj}(\mathbf{f}) = \cos(\alpha) \mathbf{c} + \sin(\alpha) \mathbf{z}. \quad (9.8)$$

Finally, we rotate the fiber direction \mathbf{p} from the cz -plane back into the sheet plane tangent to the surface,

$$\mathbf{f} = [-\mathbf{p} \cdot \mathbf{s}] \mathbf{n}_{cz} + [\mathbf{n}_{cz} \cdot \mathbf{s}] \mathbf{p}. \quad (9.9)$$

At this point, we have assigned fiber and sheet directions \mathbf{f} and \mathbf{s} to each surface node on $\partial\mathcal{B}_{\phi+}$ and $\partial\mathcal{B}_{\phi-}$, which we can then normalize for the sake of convenience. To obtain smooth fiber distributions, we interpolate \mathbf{f} and \mathbf{s} into the cardiac domain \mathcal{B} by interpreting each vector component as scalar-valued feature θ and solve the corresponding Poisson problem (9.1) a total of six times. Table 9.1 summarizes the fiber and sheet interpolation algorithm for a given heart mesh and given fiber angles α on the cardiac surfaces. **Remark:** If we preferred the sheet directions to be

Table 9.1: Algorithm for creating smoothly varying fiber and sheet orientations \mathbf{f} and \mathbf{s} within a human heart, for a given mesh and prescribed fiber angles α on the surfaces of the heart.

Identify epicardial and endocardial surfaces $\partial\mathcal{B}_{\phi+}$ and $\partial\mathcal{B}_{\phi-}$	
For all nodes on $\partial\mathcal{B}_{\phi+}$ and $\partial\mathcal{B}_{\phi-}$	
Assign local cardiac coordinate system $\{\mathbf{n}, \mathbf{c}, \mathbf{z}\}$	(9.3)
Assign sheet plane normal \mathbf{s}	(9.6)
Assign fiber angles α with respect to horizontal circumferential direction \mathbf{c}	(9.7)
Calculate projection \mathbf{p} of the fiber direction	(9.8)
Calculate fiber direction \mathbf{f} from rotation of projection into sheet plane	(9.9)
Solve Poisson equation for each component of \mathbf{f} and \mathbf{s} interpreted as feature θ	
(9.1)	
using \mathbf{f} and \mathbf{s} on $\partial\mathcal{B}_{\phi+}$ and $\partial\mathcal{B}_{\phi-}$ as Dirichlet boundary conditions	(9.2)
Collect and normalize the three components of \mathbf{f} and \mathbf{s} at each node	
[Optional] Correct sheet normals such that \mathbf{f} and $\tilde{\mathbf{s}}$ form an orthonormal basis	
(9.10)	
[Optional] Interpolate \mathbf{f} and \mathbf{s} to integration points or element centroids	

orthonormal, we could first interpolate the normal directions \mathbf{n} using the Poisson interpolation. Then, we would create sheet normals $\tilde{\mathbf{s}}$ orthogonal to the fiber direction \mathbf{f} .

$$\tilde{\mathbf{s}} = [\mathbf{f} \times \mathbf{s}] \times \mathbf{f} \quad (9.10)$$

We recommend reprojecting the sheet normal directions \mathbf{s} instead of the fiber directions \mathbf{f} because of the no-contact assumption.

Remark: Our algorithm is designed to interpolate sparse vector fields across a given domain. The user-defined constraints, which serve as Dirichlet boundary conditions, do not necessarily have to be prescribed at the epicardial and endocardial surfaces $\partial\mathcal{B}_{\phi^+}$ and $\partial\mathcal{B}_{\phi^-}$, but could potentially also be prescribed anywhere inside the domain.

9.3 Results

To illustrate our algorithm, we analyze four benchmark problems. First, we use an idealized geometry compatible with prolate spheroid coordinate systems to compare our method against an analytically assigned fiber orientation. Second, we utilize a human heart mesh to demonstrate how our algorithm works on a real patient-specific geometry. Third, we analyze the sensitivity of the fiber distribution with respect to randomly distributed boundary conditions. Fourth, we perform a second sensitivity analysis, however now with respect to uniformly distributed boundary conditions.

9.3.1 Generic bi-ventricular heart model

Figure 9.1 shows the Poisson interpolation providing a smooth continuum description across the predefined boundary sets. We observe a spiral-like fiber orientation in the left endocardium, which contributes to the characteristic torsional motion in the left ventricle.

Figure 9.2 shows the fibers and sheet plane normals at slices towards the top, middle, and bottom of the ventricles. Even the fiber orientation around the apex, a challenging region for fiber assignment, seems reasonable and smooth. In summary, our feature-based Poisson interpolation agrees excellently with the analytically assigned fiber orientations of the generic bi-ventricular heart model [66].

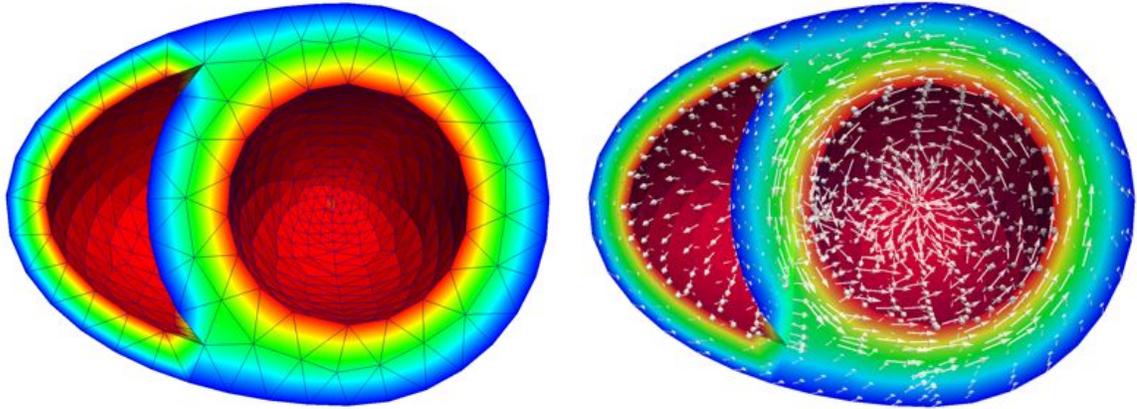


Figure 9.1: (Left) Poisson interpolation results of fiber orientation angle with respect to circumferential fiber. Epicardial and right ventricular septal surfaces are assigned as -70° , while the endocardial surfaces are assigned as 80° as boundary conditions. (Right) Resulting interpolated fiber directions throughout the heart.

9.3.2 Patient-specific human heart model

Figure 9.3 shows the results for the patient-specific heart model generated from magnetic resonance images [103]. Similar to the generic idealized heart, the Poisson interpolation provides a smooth continuum description across the predefined boundary sets.

Figure 9.4 illustrates the resulting fiber and sheet interpolations. The solution of the Poisson problem creates a smooth interpolation through the thickness of the myocardium for both fiber directions and sheet plane normals. Again, we can observe the characteristic torsion-inducing fiber orientation in both ventricles. Like the generic model, the fiber distribution near the apex is smooth and continuous. In summary, the algorithm is capable of creating smooth fiber and sheet interpolations on arbitrary patient-specific geometries, even in geometrically challenging regions such as the apex.

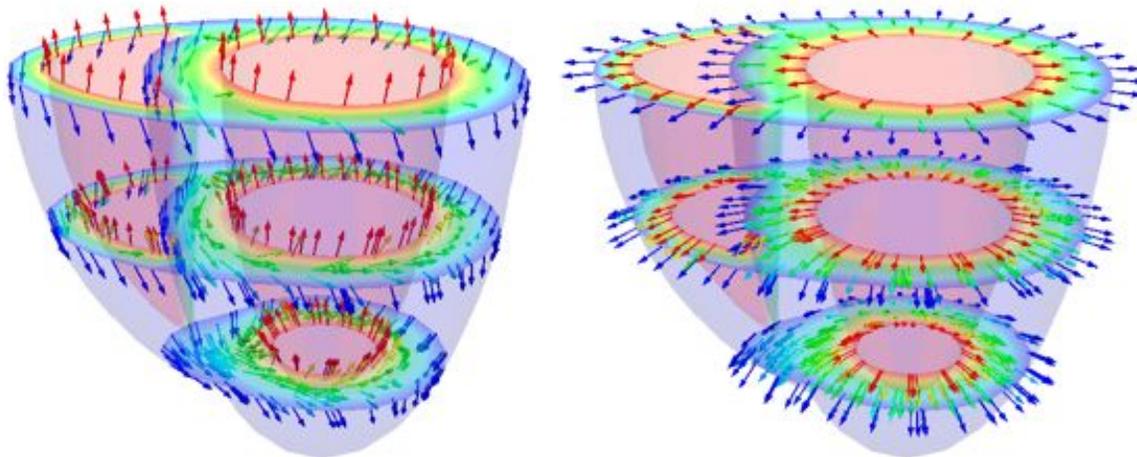


Figure 9.2: (Left) Fiber directions at various slices in the heart are shown. (Right) The normal directions at the corresponding slices are shown. The colors at each point of the cross-sectional slices correspond to the fiber angle orientation with blue representing -70° and red representing $+80^\circ$.

9.3.3 Sensitivity analysis for randomly distributed boundary conditions

Since our fiber interpolation algorithm is framed as a boundary value problem with user-defined constraints, we next perform numerical tests to determine the relative sensitivity of the overall fiber distribution with respect to changes in the boundary conditions. In our first set of tests, we select random subsets of our boundary conditions ranging from 10% to 100% of all boundary conditions at the surface nodes. These tests serve to determine the amount of error in angle distributions we might expect from a random, and possibly poor spatial uniformity of measurements. We define the angle error e^\angle as

$$e^\angle = |\arccos(\mathbf{f} \cdot \mathbf{f}^{\text{int}})|, \quad (9.11)$$

where \mathbf{f} and \mathbf{f}^{int} are the exact and interpolated fiber orientation.

Figure 9.5 illustrates representative fiber interpolation errors from 10%, top left, to 90%, bottom right, of all possible boundary conditions used. Since we have chosen

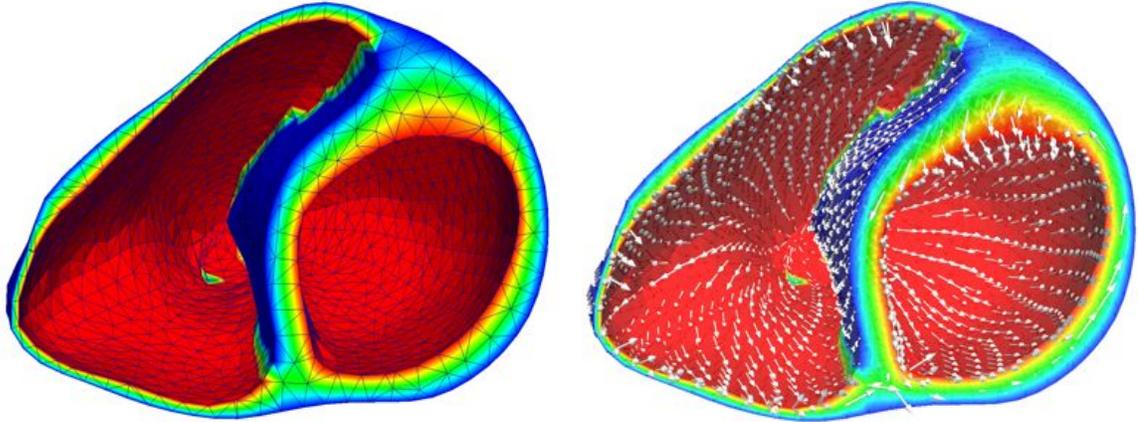


Figure 9.3: (Left) Poisson interpolation results of fiber orientation angle with respect to circumferential fiber. Epicardial and right ventricular septal surfaces are assigned as -70° , while the endocardial surfaces are assigned as 80° as boundary conditions. (Right) Resulting interpolated fiber directions throughout the heart.

the random subsets of the boundary conditions randomly for each ratio, we expect the angle error to be more or less randomly distributed across the mesh. As the subset of assigned boundary condition grows from 10% to 90%, the angle error decreases. Figure 9.6 displays the average angle error for varying random boundary condition ratios from 10% to 100%. As a rule of thumb, when using randomly distributed boundary conditions, to obtain an angle error of 10% or less on average, we have to prescribe at least 70% of the surfaces nodes as boundary conditions.

9.3.4 Sensitivity analysis for uniformly distributed boundary conditions

Last, since real fiber data extracted from histology or diffusion tensor imaging are typically acquired in a spatially uniform manner, we performed a mesh refinement test. We generated a coarse patient-specific mesh of 4,563 nodes [103] and refined it recursively twice using edge bi-section subdivision. For each of the three meshes, we performed a Poisson interpolation to generate fiber orientations and assigned boundary conditions only to the original 4,563 nodes. We then calculated the angle error

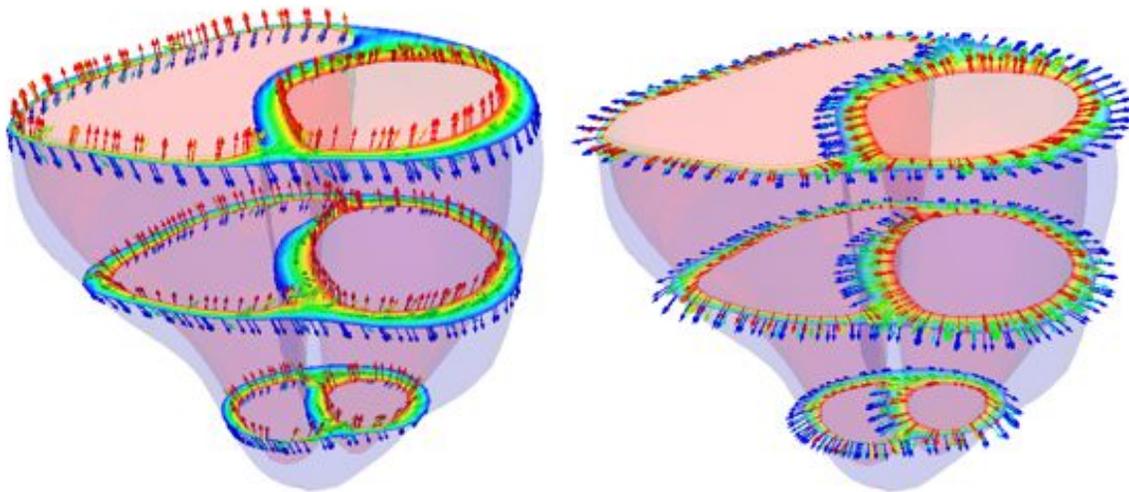


Figure 9.4: (Left) Fiber directions at various slices in the heart are shown. (Right) The normal directions at the corresponding slices are shown.

(9.11) between the two subdivided meshes averaged over all nodes. Table 9.3.4 summarizes the generating parameters and mesh data.

Table 9.2: Table of parameters for random and subdivision examples.

Surfaces	Fiber angle [°]	
Epicardium, right septal wall, $\partial\mathcal{B}_{\phi^+}$	-70	
Left endocardium, right free wall, $\partial\mathcal{B}_{\phi^-}$	+80	
Subdivision	# of surface nodes	# of total nodes
Initial mesh	3,839	4,563
1st Subdivision	15,350	28,639
2nd Subdivision	61,394	198,041
Average angle error	9.965°	

Figure 9.7 illustrates the results of the mesh refinement study. With spatially uniform subsets of the boundary conditions, in contrast to the previous problem, we observe that the error is less than 10% for only 6.25% of the surface nodes prescribed as boundary conditions.

Figure 9.8 illustrates the fiber angle error in first subdivision mesh compared to second subdivision mesh. When using uniformly distributed boundary conditions, most

fibers display an error of 20% or less. This indicates that our algorithm works best when interpolating uniformly distributed features across the heart.

9.4 Discussion

We believe that the results from the previous section demonstrate the advantages of a geometrical approach, compared to a mathematical approach, when generating fiber orientation maps for computational human heart models. The proposed method is able to reproduce fiber orientations for given fiber angles on an idealized geometry [137]. As the analytical fiber distribution elegantly enforces a linear interpolation across the thickness of the ellipsoidal ventricle [66], a Poisson interpolation with uniform properties on the same geometry generates identical results. However, by using our geometrical approach, we can avoid singularities in the septum and at the apex inherent to the mathematical approach by directly controlling the fiber orientation when \mathbf{n} and \mathbf{z} are perfectly alligned.

Our geometrical approach allows us to easily transfer the elegance of the mathematical approach to arbitrary patient-specific geometries [103]. For finite element models of anisotropic electrical conduction [37, 44] or anisotropic cardiac mechanics [65, ?], our geometrical approach is convenient since its results can immediately feed into a finite element analysis. The Poisson interpolation is easy, fast, and robust at various mesh granularities.

Here we have chosen linear tetrahedral elements since they are closely related to concepts of discrete exterior calculus used in computer graphics to interpolate surface features on arbitrary geometries [55]. In addition, tetrahedral elements are widely used and easy to understand. However, the proposed concept is easily expandable to hexahedral elements. In fact, it is even compatible with cubic Hermitian elements [137], which already contain the necessary information regarding surface normals, sheet direction, and tangent direction.

Another advantage of solving the Poisson interpolation with finite element solvers is

that the finite element method ensures the best approximating solution if the governing differential equation is linear. This has some interesting implications. For example, it means that the method presented can also properly account for non-constant value boundary conditions, which may occur if more detailed fiber orientation measurements are available. Given experimental measurements, our algorithm returns the best approximating solution that matches the fiber orientation for different mesh refinement levels. To further fine-tune control of the fiber orientation profile, the finite element method allows us to easily integrate additional constraints.

Previously, least squares methods have been proposed to generate fiber orientation distributions [42]. However, there are fundamental differences between the finite element method and the least squares methods. The least squares approach will try to find the solution that most closely resembles the measured fiber orientation [181]. As such, it strongly depends on the character of the interpolating shape functions and on the underlying discretization scheme. The finite element method, however, enforces user-defined fiber orientations at the specific nodes. It satisfies the Dirichlet boundary conditions in a strong sense, and determines the best approximating fiber orientation at all other nodes. However, least squares methods and this finite element method are not exclusive approaches and can be used in tandem when necessary. For example least squares methods can be used to smoothen out the distribution on surfaces, which our finite element method can then interpolate through the thickness. Perhaps the most exciting and most useful application of our feature-based Poisson interpolation compared to analytical and least squares approaches is the ability to fill in gaps when there is lack of experimental data. Since it is known that the heart has an organized fiber orientation and it may be significantly easier to measure the surfaces of tissue rather than the inner layers, our method can provide a reasonable approximation using a combination of observable results and a more algorithmically-methodical approach. Even if we only measure a relatively sparse, yet spatially uniform, set of fiber directions, either histologically [52] or from diffusion tensor imaging [182], fiber diffusion can fill in the gaps and create a reasonable and smooth solution as shown in Figure 9.7. In addition, we can easily integrate experimental data at non-nodal points, e.g., from magnetic resonance imaging [20], using constraint boundary conditions.

The proposed method provides a variety of natural extensions and generalizations that may be of interest to the biomechanics community. For example, we could easily apply the underlying concept to other muscular tissues such as skeletal muscle [19, 114], or to collagenous structures in other biological tissues. If we assume that the transitioning nature between two surfaces in a body is not linear, we could vary the diffusion tensor at varying points within the body to match the description accurately. While the heart's transmural fiber orientation transition is almost linear [52], by varying the transmural or in-plane conductivity, we could potentially approximate nonlinear transmural and regional fiber variations. Likewise, the fiber orientation distribution can be more carefully adjusted and tuned by possibly using inequality constraints, which might be more physiological. Another possible extension would be using noisy source terms with the finite element model to account for biological variability. Along a similar vein, fiber interpolation can be used as a tool to generate smoothly varying fiber fields with random defects to represent diseased tissue conditions [164], similar to the bottom row of Figure 9.6. We could generate and control the amount of disorganization within the fiber structure, and virtually investigate the impact of different clinical therapies on electrical and mechanical function in cases where the fiber orientation is globally irregular or locally distorted.

9.5 Conclusion

We have presented a novel Poisson interpolation-based approach to create fiber orientation maps in real patient-specific heart geometries. The algorithm makes creative use of existing and common linear finite element programs for diffusion or thermal problems. Adopted from computer graphics, the underlying concept is based on diffusing a defined subset of fiber orientations throughout the domain of interest to create smoothly interpolating fields. In the finite element setting, this subset is simply treated as Dirichlet boundary condition and enforced in a strong sense. Here we have shown the application of the Poisson interpolation scheme to create fiber orientation maps across the heart, but the general concept can easily be applied to other arbitrarily shaped biological tissues with known organized fiber orientations.

We have shown that the algorithm does generally work for both arbitrarily and uniformly distributed input data. However, our interpolation errors were significantly lower for uniform distributions. In summary, the proposed method is computationally efficient, robust, and easy to solve. It is immediately transformative to generate smoothly interpolated fields in experimental and clinical settings, in which data sets are sparse but uniformly distributed.

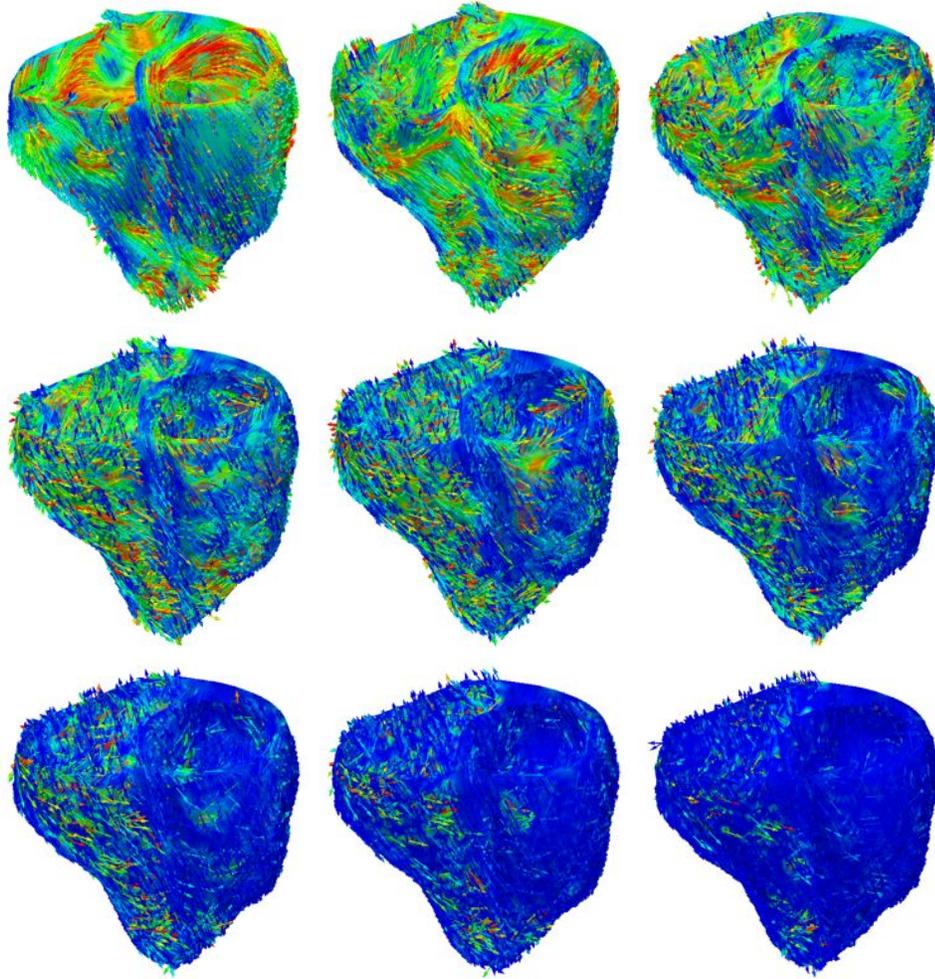


Figure 9.5: Representative images/fiber of angle fiber error given random subsets of boundary conditions for the fiber interpolation algorithm. Fibers highlighted in blue have no error (0°) whereas red fibers are closer to perpendicular (90°). Increasing sizes of random subsets are chosen and are shown starting from 10% (top left), 20% (top center), 30% (top right), 40% (left), 50% (center), 60% (right), 70% (bottom left), 80% (bottom center), 90% (bottom right).

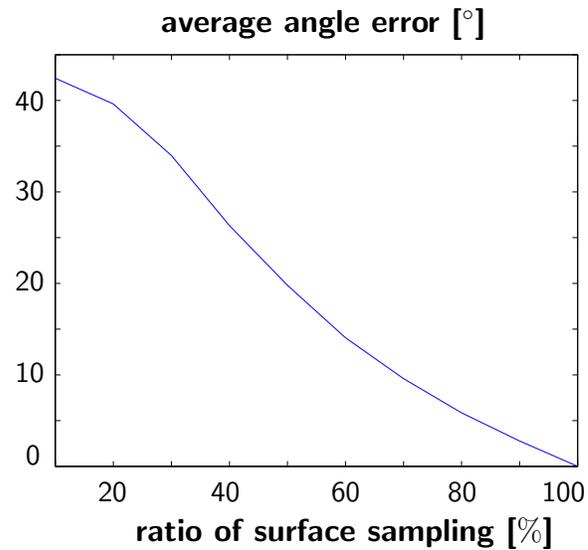


Figure 9.6: Average angle error when 10% to 100% of the surface nodes used as boundary conditions to generate fiber orientation distribution.

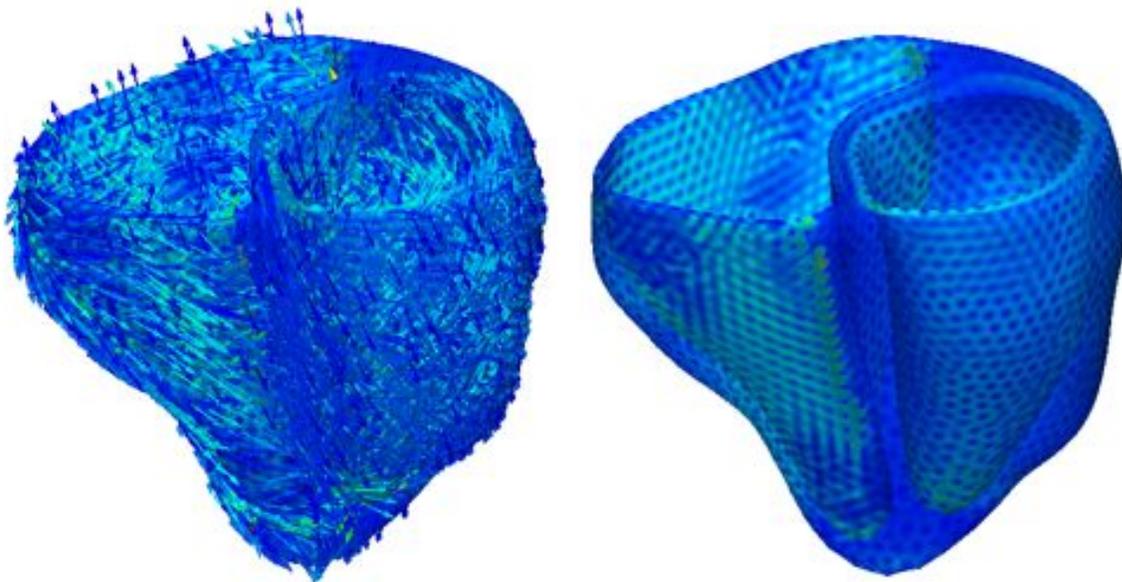


Figure 9.7: (left) Fiber direction color coded by angle error. (right) Fiber angle error over the heart without fibers. Small hexagonal patterns are formed, because the nodes at the coarsest mesh are used as the boundary condition set, and therefore the new nodes on the surfaces are not part of this boundary condition set.

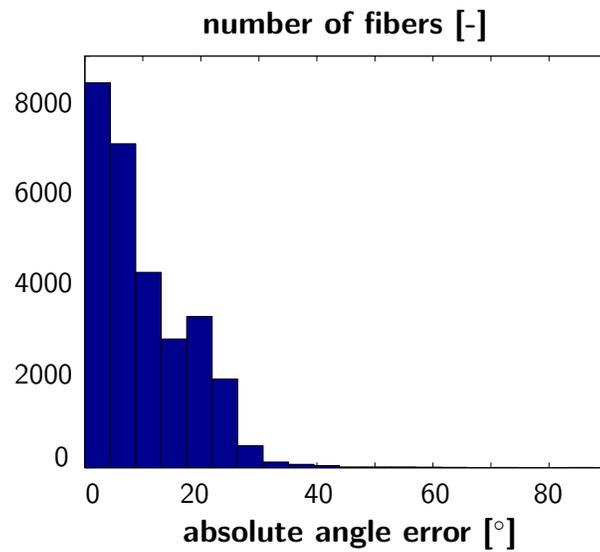


Figure 9.8: Histogram of fiber angle error in first subdivision mesh compared to second subdivision mesh. The histogram is organized into bins of 5° error.

Chapter 10

Electromechanically-coupled Applications

Computational modeling of the human heart allows us to predict how chemical, electrical, and mechanical fields interact throughout a cardiac cycle. Pharmacologic treatment of cardiac disease has advanced significantly over the past two decades, yet it remains unclear how the local biochemistry of individual heart cells translates into global cardiac function. Here we propose a novel, unified strategy to simulate excitable biological systems across three biological scales; from the molecular level via the cellular level to the organ level. To discretize the governing chemical, electrical, and mechanical equations in space, we propose a monolithic finite element scheme. We apply a global-local split in which the deformation of the mechanical problem and the transmembrane potential and electrical problems are introduced globally as a nodal degrees of freedom, while the state variables of the chemical problem are treated locally as internal variables on the integration point level. This particular discretization scheme is highly efficient and inherently modular, since it allows us to combine various different cell types through only minor local modifications on the constitutive level. To ensure unconditional algorithmic stability, we apply an implicit backward Euler finite difference scheme to discretize the resulting system in time. To increase algorithmic robustness and guarantee optimal quadratic convergence, we suggest an incremental iterative Newton-Raphson scheme. The proposed

algorithm allows us to simulate the interaction of chemical, electrical, and mechanical fields during a representative cardiac cycle on a real patient-specific geometry, robust and stable, with calculation times on the order of four days on a standard desktop computer.

10.1 Introduction

Pharmacological treatment has opened new avenues for managing various types of cardiac disease. On a daily basis, cardiologists now prescribe antiarrhythmic agents to control heart rhythm disorders such as atrial fibrillation, atrial flutter, ventricular tachycardiac, and ventricular fibrillation [?]. While the pharmacological control of the electrical activity of the heart is reasonably well understood, the pharmacological manipulation of the mechanical activity of the heart remains severely understudied. This is an important problem in heart failure [?], a disease associated with an annual health care cost of more than \$30 billion in the United States alone [87]. To understand how a new drug affects the interaction between chemical, electrical, and mechanical fields, systematic drug testing is of incredible clinical importance. Not surprisingly, it covers a huge market ranging from single cell testing using patch clamp electrophysiology [1, 75], via cell culture testing using microelectroarray recordings [37, 180], to large animal experiments [90, 185]. While the pharmacological manipulation of chemo-electro-mechanical coupling is relatively well understood on the single cell level [18, 27], little is known about how this knowledge translates into clinically relevant function on the organ level [143]. This knowledge gap presents a tremendous opportunity for quantitative, predictive computational modeling. Most importantly, the nature of coupling between the underlying chemical, electrical, and mechanical fields is ideally tailored for finite element simulations, a circumstance that has been largely overlooked until today.

The first model to quantitatively characterize the electrical activity of excitable cells was the Nobel-price winning Hodgkin-Huxley model introduced more than half a century ago [83]. Initially designed for nerve cells [56, 134], the model was soon adopted for other cell types, such as pacemaker cells [140], purkinje fiber cells [122], atrial

cells [43], and ventricular cells [12, 116, 188] in the heart. Originally proposed for single cells, these approaches were generalized to multiple cells, tissues, and organs by adding a phenomenological flux term to characterize the propagation of the excitation wave. Traditionally, simulations of propagating electrical signals were dominated by biophysicists and electrical engineerings [4, 96]. Their models were based on simple straightforward algorithms, discretized in space using finite differences, discretized in time using explicit time stepping schemes [40]. To compensate for the lack of sophistication in algorithmic design, these initial models generally use a high spatial and temporal resolution, small grids and small time steps. Not surprisingly, these initial methods are extremely expensive from a computational point of view [106].

Within the past decade, physiological function has become a key focus in cardiac simulations, paving the way for mechanical models and finite element methods [34, 145, 154]. However, progress was dampened by the finite difference nature of existing algorithms, making it virtually impossible to integrate mechanical deformation, in particular in the context of finite strains. The first generation of electro-mechanical heart models combined previously established finite difference based electrical algorithms with finite element based mechanical algorithms [98]. Most versions of these models are coupled unidirectionally, i.e., the algorithm first calculates the electrical field and then uses it as an input to calculate the mechanical field [71]. The advantage of this approach is that it allows us to combine different spatial and temporal resolutions for both fields. For loosely coupled problems, these algorithms typically perform sufficiently well [183], although we can not really quantify the loss of information and the possible energy blow-up associated with the explicit discretization of the coupling terms. For strongly coupled problems, these algorithms require an extremely fine spatial and temporal resolution, especially during the rapid upstroke phase when all fields undergo rapid changes. To eliminate potential algorithmic instabilities, revised versions of these models are coupled bidirectionally, i.e., they iterate between electrical and mechanical fields. It is not surprising that those algorithms, which integrate more information about the nature of coupling upfront have enhanced stability and performance properties [136].

Here we challenge existing excitation-contraction algorithms and propose a second

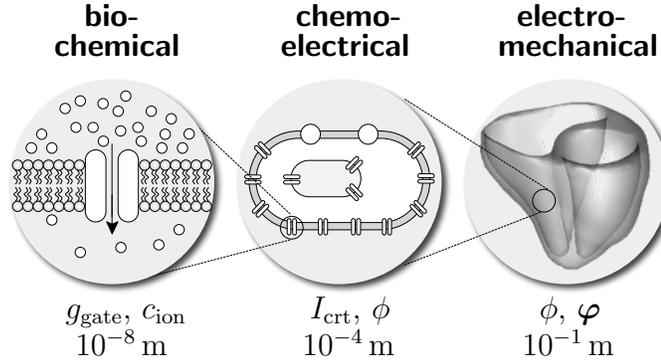


Figure 10.1: Multiscale model of the human heart. At the molecular level, gating variables g_{gate} and ion concentrations c_{ion} characterize the bio-chemical response. At the cellular level, ionic currents I_{crt} and the transmembrane potential ϕ characterize the chemo-electrical response. At the organ level, the propagation of the electrical potential ϕ and the deformation φ characterize the electro-mechanical response.

generation of chemo-electro-mechanical heart models, algorithmically redesigned from scratch. We propose a novel unified algorithm, which is entirely finite element based, fully coupled, implicitly time-integrated, and consistently linearized. This allows us to use existing finite element infrastructures, such as simple, ad hoc adaptive time stepping schemes [?]. In designing our new algorithm, we take advantage of the multiscale nature of the underlying problem illustrated in Figure 10.1, and discretize all chemical unknowns locally on the integration point level and all electrical and mechanical unknowns globally on the node point level [72, 151, 197]. In Section 10.2, we summarize the underlying kinematic, balance, and constitutive equations for chemo-electro-mechanical problems. In Section 10.3, we then illustrate their temporal and spatial discretizations. We introduce the global system of equations, which we solve using an incremental iterative Newton Raphson strategy. In Section 10.4, we specify the constitutive equations for the electrical and mechanical source and flux terms, which naturally introduce the coupling between the underlying chemical, electrical, and mechanical fields. In Section 10.5, we illustrate the features of the proposed model, first locally, at the single cell level, and then globally, at the whole heart level. We conclude with a discussion and a brief outlook in Section 10.6.

10.2 Continuous problem of chemo-electro-mechanics

In this section, we summarize the generic continuous equations of chemo-electro-mechanical coupling characterized through a set of partial differential equations for the electrical and mechanical problems and through a system of ordinary differential equations for the chemical problem. The primary unknowns of the electrical and mechanical problems are the transmembrane potential ϕ and the deformation φ . The unknowns of the chemical problem are the local state variables which we collectively summarize in the vector \mathbf{q} . For simple two-parameter models, \mathbf{q} would only contain a single variable, the phenomenological recovery variable r . For more sophisticated ionic models, \mathbf{q} contains a set of gating variables g_{gate} and a set of ion concentrations c_{ion} , which, at any point in time, characterize the local ionic currents I_{crt} .

10.2.1 Kinematic equations

To characterize the kinematic state of the body under consideration, we introduce the nonlinear deformation φ that maps particles from the undeformed reference configuration \mathcal{B}_0 at time t_0 to the deformed current configuration \mathcal{B}_t at time $t \in \mathbb{R}$,

$$\mathbf{x} = \varphi(\mathbf{X}, t) : \mathcal{B}_0 \times \mathbb{R} \rightarrow \mathcal{B}_t. \quad (10.1)$$

In what follows, $\{\dot{\circ}\} = d_t\{\circ\}|_X$ and $\nabla\{\circ\} = d_X\{\circ\}|_t$ denote the material time derivative and the material gradient of a quantity $\{\circ\}$. Accordingly, $\text{Div}\{\circ\} = \nabla\{\circ\} : \mathbf{I}$ denotes the material divergence, where \mathbf{I} is the second-order identity tensor. With these definitions, we can introduce the deformation gradient \mathbf{F} as the linear tangent map from the material tangent space $T\mathcal{B}_0$ to the spatial tangent space $T\mathcal{B}_t$,

$$\mathbf{F} = \nabla\varphi(\mathbf{X}, t) : T\mathcal{B}_0 \rightarrow T\mathcal{B}_t. \quad (10.2)$$

We will utilize the right Cauchy-Green deformation tensor and its inverse

$$\mathbf{C} = \mathbf{F}^t \cdot \mathbf{F} \quad \text{and} \quad \mathbf{C}^{-1} = \mathbf{F}^{-1} \cdot \mathbf{F}^{-t} \quad (10.3)$$

to define relevant strain measures. In particular, we introduce the Jacobian J and the trace I_1 as characteristic isotropic invariants,

$$J = \det(\mathbf{F}) \quad I_1 = \mathbf{C} : \mathbf{I} \quad (10.4)$$

and I_{ff} , I_{ss} , and I_{fs} as characteristic anisotropic invariants,

$$I_{\text{ff}} = \mathbf{C} : [\mathbf{f}_0 \otimes \mathbf{f}_0] \quad I_{\text{ss}} = \mathbf{C} : [\mathbf{s}_0 \otimes \mathbf{s}_0] \quad I_{\text{fs}} = \mathbf{C} : [\mathbf{f}_0 \otimes \mathbf{s}_0]^{\text{sym}}. \quad (10.5)$$

Here, I_{ff} and I_{ss} are the stretches squared along the myocardial fiber and sheet directions, which we denote by \mathbf{f}_0 and \mathbf{s}_0 in the reference configuration \mathcal{B}_0 and by $\mathbf{f} = \mathbf{F} \cdot \mathbf{f}_0$ and $\mathbf{s} = \mathbf{F} \cdot \mathbf{s}_0$ in the current configuration \mathcal{B}_t .

10.2.2 Balance equations

The balance equation of the electrical problem balances the rate of change of the transmembrane potential ϕ with the divergence of the electrical flux, $\text{Div } \mathbf{Q}$, and the electrical source, F^ϕ ,

$$\dot{\phi} = \text{Div } \mathbf{Q} + F^\phi. \quad (10.6)$$

The balance equation of the mechanical problem balances the rate of change of the linear momentum with the divergence of the momentum flux, $\text{Div } \mathbf{P}$, and the momentum source, \mathbf{F}^φ ,

$$\mathbf{0} = \text{Div } \mathbf{P} + \mathbf{F}^\varphi. \quad (10.7)$$

In the quasi-static case considered here, the rate of change of the linear momentum vanishes identically.

10.2.3 Constitutive equations

The electrical and mechanical problems (10.6) and (10.7) are coupled constitutively through the corresponding flux and source terms. The electrical flux \mathbf{Q} is typically introduced phenomenologically and characterizes the propagation speed of the electrical signal. It is usually proportional to the potential gradient $\nabla\phi$ and can potentially be

coupled to the mechanical problem through the deformation gradient $\nabla\boldsymbol{\varphi}$ to account for stretch-induced changes in the propagation speed,

$$\mathbf{Q} = \mathbf{Q}(\nabla\phi, \nabla\boldsymbol{\varphi}). \quad (10.8)$$

The electrical source F^ϕ characterizes the electrophysiology of the individual cells on the local level. Through voltage-gated ion channels, F^ϕ depends on the electrical potential ϕ . Through possible stretch-activated ion-channels, F^ϕ may depend on the deformation gradient $\nabla\boldsymbol{\varphi}$. Through the cell's biochemistry, F^ϕ also depends on the set of internal variables collectively summarized in the vector \mathbf{q} , which, in our case, contains a set of gating variables g_{gate} and a set of ion concentrations c_{ion} [197],

$$F^\phi = F^\phi(\phi, \nabla\boldsymbol{\varphi}, \mathbf{q}). \quad (10.9)$$

The momentum flux \mathbf{P} is simply the Piola stress, which we can additively decompose into a passive and an active part. The passive stress \mathbf{P}^{pas} depends on the deformation gradient $\nabla\boldsymbol{\varphi}$ and characterizes the passive myocardium. The active stress \mathbf{P}^{act} either depends on the electrical potential ϕ [66] or on the set of internal variables \mathbf{q} [183], as proposed here, and introduces coupling to the electro-chemical problem. The two-field nature of the Piola stress introduces an additional dependence on the deformation gradient $\nabla\boldsymbol{\varphi}$,

$$\mathbf{P} = \mathbf{P}^{\text{pas}}(\nabla\boldsymbol{\varphi}) + \mathbf{P}^{\text{act}}(\nabla\boldsymbol{\varphi}, \mathbf{q}). \quad (10.10)$$

The momentum source \mathbf{F}^φ characterizes volume forces such as gravity, which we assume to be negligibly small in the subsequent analyses,

$$\mathbf{F}^\varphi = \mathbf{0}. \quad (10.11)$$

We will now illustrate the computational solution of the coupled chemo-electro-mechanical problem using a weighted-residual based finite element approach.

10.3 Discrete problem of chemo-electro-mechanics

To discretize the continuous chemo-electro-mechanical problem, we rephrase the electrical and mechanical balance equations (10.6) and (10.7) in their residual formats,

$$\begin{aligned} \mathbf{R}^\phi &= \dot{\phi} - \text{Div } \mathbf{Q} - F^\phi \doteq 0 & \text{in } \mathcal{B}_0 \\ \mathbf{R}^\varphi &= - \text{Div } \mathbf{P} - \mathbf{F}^\varphi \doteq \mathbf{0} & \text{in } \mathcal{B}_0, \end{aligned} \quad (10.12)$$

where both are valid in the entire domain \mathcal{B}_0 . We then partition the boundary $\partial\mathcal{B}_0$ into disjoint parts $\partial\mathcal{B}_0^\phi$ and $\partial\mathcal{B}_0^Q$ for the electrical problem and equivalently into $\partial\mathcal{B}_0^\varphi$ and $\partial\mathcal{B}_0^P$ for the mechanical problem and prescribe the corresponding Dirichlet and Neumann boundary conditions,

$$\begin{aligned} \phi &= \bar{\phi} \text{ on } \partial\mathcal{B}_0^\phi & \mathbf{Q} \cdot \mathbf{N} &= \bar{T}^Q \text{ on } \partial\mathcal{B}_0^Q \\ \varphi &= \bar{\varphi} \text{ on } \partial\mathcal{B}_0^\varphi & \mathbf{P} \cdot \mathbf{N} &= \bar{T}^P \text{ on } \partial\mathcal{B}_0^P, \end{aligned} \quad (10.13)$$

where \mathbf{N} denotes the outward normal to $\partial\mathcal{B}_0$. To derive the weak forms of the electrical and mechanical problems \mathbf{G}^ϕ and \mathbf{G}^φ , we integrate the the residual statements (10.12) over the domain \mathcal{B}_0 , multiply both with the scalar- and vector-valued test functions $\delta\phi$ in $H_1^0(\mathcal{B}_0)$ and $\delta\varphi$ in $H_1^0(\mathcal{B}_0)$, integrate them by parts, and apply the corresponding Neumann boundary conditions (10.13.2) and (10.13.4),

$$\begin{aligned} \mathbf{G}^\phi &= \int_{\mathcal{B}_0} \delta\phi \dot{\phi} \, dV + \int_{\mathcal{B}_0} \nabla\delta\phi \cdot \mathbf{Q} \, dV \\ &\quad - \int_{\partial\mathcal{B}_0^Q} \delta\phi \bar{T}^Q \, dA - \int_{\mathcal{B}_0} \delta\phi F^\phi \, dV = 0 \quad \forall \delta\phi \\ \mathbf{G}^\varphi &= \int_{\mathcal{B}_0} \nabla\delta\varphi : \mathbf{P} \, dV \\ &\quad - \int_{\partial\mathcal{B}_0^P} \delta\varphi \cdot \bar{T}^P \, dA - \int_{\mathcal{B}_0} \delta\varphi \cdot \mathbf{F}^\varphi \, dV = 0 \quad \forall \delta\varphi. \end{aligned} \quad (10.14)$$

10.3.1 Temporal discretization

For the temporal discretiation, we partition the time interval of interest \mathcal{T} into n_{step} subintervals $[t_n, t]$ as $\mathcal{T} = \bigcup_{n=0}^{n_{\text{step}}-1} [t_n, t]$ and focus on a typical time slab $[t_n, t]$. Here and from now on we omit the index $n+1$ associated with the current time step. We assume, that the primary unknowns ϕ_n and φ_n and all derivable flux terms,

source terms, and state variables are known at the beginning of the current interval. To approximate the material time derivative of the transmembrane potential ϕ , we apply a first order finite difference scheme,

$$\dot{\phi} = [\phi - \phi_n] / \Delta t, \quad (10.15)$$

where $\Delta t := t - t_n > 0$ denotes the current time increment. To solve for the unknowns ϕ and φ , we then apply a classical backward Euler time integration scheme and evaluate the discrete set of governing equations (10.14) at the current time point t .

10.3.2 Spatial discretization

For the spatial discretization, we apply a \mathcal{C}^0 -continuous interpolation of the transmembrane potential ϕ and of the deformation φ and introduce both ϕ and φ as global degrees of freedom at the node point level. We partition the domain of interest \mathcal{B}_0 into n_{el} elements \mathcal{B}_0^e as $\mathcal{B}_0 = \bigcup_{e=1}^{n_{\text{el}}} \mathcal{B}_0^e$. Using the isoparametric concept, we interpolate the trial functions $\phi^h, \varphi^h \in H_1(\mathcal{B}_0)$ on the element level with the same basis function N^ϕ and N^φ as the element geometry. Using the Bubnov-Galerkin approach, we interpolate the test functions $\delta\phi^h, \delta\varphi^h \in H_1^0(\mathcal{B}_0)$ on the element level with the same basis function N^ϕ and N^φ as the trial functions,

$$\begin{aligned} \delta\phi^h|_{\mathcal{B}_0^e} &= \sum_{i=1}^{n_{e\phi}} N_i^\phi \delta\phi_i & \phi^h|_{\mathcal{B}_0^e} &= \sum_{k=1}^{n_{e\phi}} N_k^\phi \phi_k \\ \delta\varphi^h|_{\mathcal{B}_0^e} &= \sum_{j=1}^{n_{e\varphi}} N_j^\varphi \delta\varphi_j & \varphi^h|_{\mathcal{B}_0^e} &= \sum_{l=1}^{n_{e\varphi}} N_l^\varphi \varphi_l. \end{aligned} \quad (10.16)$$

We then rephrase the residuals of the electrical and the mechanical problem (10.12) in their discrete forms,

$$\begin{aligned} \mathbf{R}_I^\phi &= \mathbf{A} \int_{\mathcal{B}_0^e} N_i^\phi \frac{1}{\Delta t} [\phi - \phi_n] dV_e + \int_{\mathcal{B}_0^e} \nabla N_i^\phi \cdot \mathbf{Q} dV_e \\ &\quad - \int_{\partial\mathcal{B}_0^Q} N_i^\phi \bar{T}^Q dA_e - \int_{\mathcal{B}_0^e} N_i^\phi F^\phi dV_e = 0 \\ \mathbf{R}_J^\varphi &= \mathbf{A} \int_{\mathcal{B}_0^e} N_j^\varphi \frac{1}{\Delta t} [\varphi - \varphi_n] dV_e + \int_{\mathcal{B}_0^e} \nabla N_j^\varphi \cdot \mathbf{P} dV_e \\ &\quad - \int_{\partial\mathcal{B}_0^P} N_j^\varphi \bar{T}^P dA_e - \int_{\mathcal{B}_0^e} N_j^\varphi \mathbf{F}^\varphi dV_e = \mathbf{0}. \end{aligned} \quad (10.17)$$

Here the operator \mathbf{A} symbolizes the assembly of all element contributions at the local electrical and mechanical element nodes $i = 1, \dots, n_{e\phi}$ and $j = 1, \dots, n_{e\varphi}$ to the overall residuals at the global electrical and mechanical nodes $I = 1, \dots, n_{n\phi}$ and $J = 1, \dots, n_{n\varphi}$.

10.3.3 Linearization

To solve the resulting coupled nonlinear system of equations (10.17), we propose a monolithic incremental iterative Newton-Raphson solution strategy based on consistent linearization of the governing equations,

$$\begin{aligned} \mathbf{R}_{I\mathbf{k}+1}^\phi &= \mathbf{R}_I^\phi + \sum_K \mathbf{K}_{IK}^{\phi\phi} d\phi_K + \sum_L \mathbf{K}_{IL}^{\phi\varphi} \cdot d\varphi_L \doteq 0 \\ \mathbf{R}_{J\mathbf{k}+1}^\varphi &= \mathbf{R}_J^\varphi + \sum_K \mathbf{K}_{JK}^{\varphi\phi} d\phi_K + \sum_L \mathbf{K}_{JL}^{\varphi\varphi} \cdot d\varphi_L \doteq \mathbf{0}. \end{aligned} \quad (10.18)$$

in terms of the following iteration matrices,

$$\begin{aligned} \mathbf{K}_{IK}^{\phi\phi} &= \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} N_i^\phi \frac{1}{\Delta t} N_k^\phi - N_i^\phi \partial_\phi F^\phi N_k^\phi dV_e \\ &\quad + \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} \nabla N_i^\phi \cdot \partial_{\nabla\phi} \mathbf{Q} \cdot \nabla N_k^\phi dV_e \\ \mathbf{K}_{IL}^{\phi\varphi} &= \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} N_i^\phi \cdot \partial_{\mathbf{F}} F^\phi \cdot \nabla N_l^\varphi dV_e \\ &\quad + \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} \nabla N_i^\phi \cdot \partial_{\mathbf{F}} \mathbf{Q} \cdot \nabla N_l^\varphi dV_e \\ \mathbf{K}_{JK}^{\varphi\phi} &= \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} \nabla N_j^\varphi \cdot \partial_\phi \mathbf{P} N_k^\phi dV_e \\ \mathbf{K}_{JL}^{\varphi\varphi} &= \mathbf{A}_{e=1}^{n_{el}} \int_{\mathcal{B}_0^e} \nabla N_j^\varphi \cdot \partial_{\mathbf{F}} \mathbf{P} \cdot \nabla N_l^\varphi dV_e. \end{aligned} \quad (10.19)$$

The solution of the system of equations (10.18) renders the iterative update for the increments of the global unknowns $\phi_I \leftarrow \phi_I + d\phi_I$ and $\varphi_J \leftarrow \varphi_J + d\varphi_J$.

10.4 Model problem of chemo-electro-mechanics

10.4.1 Electrical flux

The electrical flux \mathbf{Q} in equation (10.17.1) is typically assumed to depend on both the potential gradient $\nabla\phi$ and the deformation gradient $\nabla\varphi$. We now specify this

dependency to be multiplicative. In analogy to Fick's law of diffusion and Fourier's law of heat transfer, we assume that the electrical flux \mathbf{Q} is proportional to the gradient of the potential $\nabla\phi$,

$$\mathbf{Q} = \mathbf{D} \cdot \nabla\phi \quad \text{with} \quad \mathbf{D} = d^{\text{iso}}\mathbf{C}^{-1} + d^{\text{ani}}\mathbf{f}_0 \otimes \mathbf{f}_0. \quad (10.20)$$

The second order diffusion tensor \mathbf{D} can account for both isotropic propagation d^{iso} and anisotropic propagation d^{ani} along preferred directions \mathbf{f}_0 . Stretch-induced changes in the propagation speed are incorporated indirectly through the inverse left Cauchy-Green tensor \mathbf{C}^{-1} motivated by the assumption of a spatial rather than material isotropy, for which the isotropic term would simply scale with the second order identity tensor \mathbf{I} . To evaluate the iteration matrices (10.19.1) and (10.19.2), we calculate the sensitivity of the electrical flux \mathbf{Q} with respect to the electrical gradient $\partial_{\nabla\phi}\mathbf{Q}$ and with respect to the deformation gradient $\partial_{\mathbf{F}}\mathbf{Q}$, where the former is nothing but the diffusion tensor \mathbf{D} and the latter reflects the above-discussed stretch-induced change in the propagation speed, see [66] for details.

10.4.2 Electrical source

The electrical source F^ϕ in equation (10.17.1) is a result of the local electrophysiology on the cellular level. As such, it is a function of the electrical potential ϕ , the deformation gradient $\nabla\boldsymbol{\varphi}$, and a set of internal variables \mathbf{q} , which characterize the electrochemical behavior of the cell. For the simplest possible models, \mathbf{q} only contains a single variable, the phenomenological recovery variable r [72, ?]. For the particular ventricular cardiomyocytes we consider here [116, 117, 188], \mathbf{q} contains a total of 17 variables, i.e., $n_{\text{gate}} = 13$ gating variables $g_{\text{gate}} = [g_m, g_h, g_j, g_{xr1}, g_{xr2}, g_{xs}, g_r, g_s, g_d, g_f, g_{xK1\infty}, g_{fCa}, g_g]$ and $n_{\text{ion}} = 4$ ion concentrations $c_{\text{ion}} = [c_{\text{Na}}, c_{\text{K}}, c_{\text{Ca}}, c_{\text{Ca}}^{\text{st}}]$. These state variables define $n_{\text{crt}} = 15$ ionic currents $I_{\text{crt}} = [I_{\text{Na}}, I_{\text{bNa}}, I_{\text{NaK}}, I_{\text{NaCa}}, I_{\text{K1}}, I_{\text{Kr}}, I_{\text{Ks}}, I_{\text{pK}}, I_{\text{t0}}, I_{\text{CaL}}, I_{\text{bCa}}, I_{\text{pCa}}, I_{\text{leak}}, I_{\text{up}}, I_{\text{rel}}]$, as illustrated in Figure 10.2. The electrical source F^ϕ is directly related to the negative sum of all these currents I_{crt}

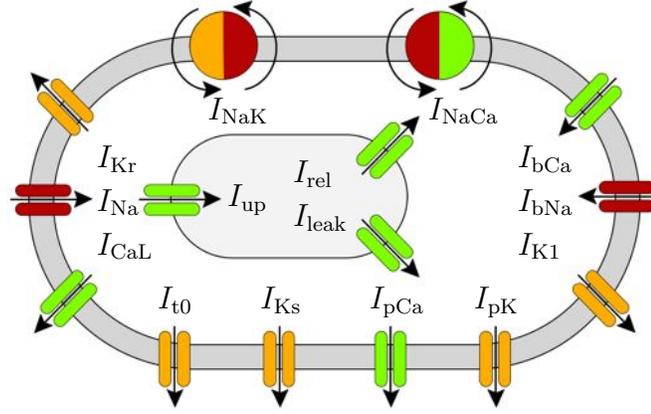


Figure 10.2: Human ventricular cardiomyocyte model with 15 ionic currents resulting from ten transmembrane channels, one exchanger, and one pump. Three additional currents characterize ionic changes inside the sarcoplasmic reticulum, shown in grey. Sodium currents are indicated in red, potassium currents in orange, and calcium currents in green.

across the cell membrane due to the outward positive convection established in experiments,

$$F^\phi = - [I_{Na} + I_{bNa} + I_{NaK} + I_{NaCa} + I_{K1} + I_{Kr} + I_{Ks} + I_{pK} + I_{t0} + I_{CaL} + I_{bCa} + I_{pCa}]. \quad (10.21)$$

Here, I_{Na} is the fast sodium current, I_{bNa} is the background sodium current, I_{NaK} is the sodium potassium pump current, I_{NaCa} is the sodium calcium exchanger current, I_{K1} is the inward rectifier current, I_{Kr} and I_{Ks} are the rapid and slow delayed rectifier currents, I_{pK} is the plateau potassium current, I_{CaL} is the long-lasting L-type calcium current, I_{bCa} is the background calcium current, and I_{pCa} is the plateau calcium current. In addition, we also have three intracellular currents, I_{leak} is the leakage current, I_{up} is the sarcoplasmic reticulum uptake current, and I_{rel} is the sarcoplasmic reticulum release current. For our particular cell model, none of the channels are mechanically gated, i.e., all currents are independent of the deformation gradient $\nabla\varphi$. Instead, all channels are voltage gated and their currents depend on the electrical potential ϕ . In addition, the currents depend on the set of internal variables \mathbf{q} consisting of

the chemical state variables, i.e., the gating variables g_{gate} and the ion concentrations c_{ion} . We can characterize all ionic currents through generic equations of the following generic form,

$$I_{\text{crt}} = I_{\text{crt}}(\phi, g_{\text{gate}}, c_{\text{ion}}), \quad (10.22)$$

which we specify in detail in the Appendix. From a mathematical point of view, the chemical problem is defined in terms of two sets of state variables, the n_{gate} gating variables g_{gate} and the n_{ion} ion concentrations c_{ion} . Both are governed through ordinary differential equations depending on the transmembrane potential ϕ , on the gating variables g_{gate} , and on the ion concentrations c_{ion} ,

$$\begin{aligned} \dot{g}_{\text{gate}} &= f_{\text{gate}}(\phi, g_{\text{gate}}, c_{\text{ion}}) \\ \dot{c}_{\text{ion}} &= f_{\text{ion}}(\phi, g_{\text{gate}}, c_{\text{ion}}). \end{aligned} \quad (10.23)$$

The gating variables g_{gate} characterize the states of the individual ion channels, either open or closed. They are defined through a set of ordinary differential equations of Hodgkin-Huxley type,

$$\dot{g}_{\text{gate}} = [g_{\text{gate}}^{\infty}(\phi, c_{\text{ion}}) - g_{\text{gate}}] / \tau_{\text{gate}}(\phi), \quad (10.24)$$

which we specify in detail in the Appendix. Here, g_{gate}^{∞} is a steady-state value and τ_{gate} is the time constant for reaching this steady state. Both are usually exponential functions of the membrane potential ϕ . The ion concentrations inside the cell c_{ion} change in response to the transmembrane currents I_{crt} . For our particular cardiomyocyte model, the relevant ion concentrations are the sodium concentration c_{Na} , the potassium concentration c_{K} , the calcium concentration c_{Ca} , and the calcium concentration in the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$. Collectively, these ion concentrations c_{ion}

are defined through a set of ordinary differential equations,

$$\begin{aligned}
\dot{c}_{\text{Na}} &= -\frac{C}{VF} [I_{\text{Na}} + I_{\text{bNa}} + 3I_{\text{NaK}} + 3I_{\text{NaCa}}] \\
\dot{c}_{\text{K}} &= -\frac{C}{VF} [I_{\text{Kl}} + I_{\text{Kr}} + I_{\text{Ks}} - 2I_{\text{NaK}} + I_{\text{pK}} + I_{\text{t0}} + I_{\text{stim}}] \\
\dot{c}_{\text{Ca}} &= -\frac{C}{2VF} [I_{\text{CaL}} + I_{\text{bCa}} + I_{\text{pCa}} - 2I_{\text{NaCa}}] \gamma_{\text{Ca}} \\
&\quad + [I_{\text{leak}} - I_{\text{up}} + I_{\text{rel}}] \gamma_{\text{Ca}} \\
\dot{c}_{\text{Ca}}^{\text{sr}} &= +\frac{V}{V^{\text{sr}}} [I_{\text{up}} - I_{\text{leak}} - I_{\text{rel}}] \gamma_{\text{Ca}}^{\text{sr}}.
\end{aligned} \tag{10.25}$$

Here C is the membrane capacitance per unit surface area, V is the cytoplasmic volume, V^{sr} is the volume of the sarcoplasmic reticulum, F is the Faraday constant, and γ_{Ca} and $\gamma_{\text{Ca}}^{\text{sr}}$ are scaling coefficients. While the electrical and mechanical problems are global in nature, the chemical problem remains strictly local. When using a finite element discretization, this allows us to store the chemical state variables g_{gate} and c_{ion} locally as internal variables on the integration point level. It is obvious that their complex, nonlinear coupled system of ordinary differential equations (10.24) and (10.25) cannot be solved analytically. Here, we apply a numerical solution using an implicit Euler backward time stepping scheme embedded in a local Newton iteration as proposed in [197]. To evaluate the iteration matrices (10.19.1) and (10.19.2), we calculate the derivative of the electrical source F^ϕ with respect to the transmembrane potential $\partial_\phi F^\phi$ related to voltage-gated ion channels and with respect to the deformation gradient $\partial_{\mathbf{F}} F^\phi$ related to stretch-activated ion channels, see [197] for details. For our particular cell model, in the absence of stretch-activated ion channels, the second term vanishes identically.

10.4.3 Momentum flux

The momentum flux \mathbf{P} in equation (10.17.2) depends on both the electrical potential ϕ and the deformation gradient $\boldsymbol{\varphi}$. We adopt the common assumption to decompose the overall stress additively into a passive mechanically-induced part \mathbf{P}^{pas} and an active electrically-induced part \mathbf{P}^{act} , such that $\mathbf{P} = \mathbf{P}^{\text{pas}} + \mathbf{P}^{\text{act}}$. For the passive Piola stress, we adopt an orthotropic model [65, 84]

Table 10.1: Chemo-electrical material parameters of human ventricular

	sodium related	potassium related	calcium related	calcium ^{sr} related	
concentrations	$c_{\text{Na}0} = 140 \text{ mM}$	$c_{\text{K}0} = 5.4 \text{ mM}$	$c_{\text{Ca}0} = 2 \text{ mM}$	-	
maximum currents	$I_{\text{NaCa}}^{\text{max}} = 1000 \text{ pA/pF}$ $I_{\text{NaK}}^{\text{max}} = 1.362 \text{ pA/pF}$	$I_{\text{NaK}}^{\text{max}} = 1.362 \text{ pA/pF}$	$I_{\text{NaCa}}^{\text{max}} = 1000 \text{ pA/pF}$ $I_{\text{leak}}^{\text{max}} = 0.08 \text{ s}^{-1}$ $I_{\text{up}}^{\text{max}} = 0.425 \text{ mM/s}$ $I_{\text{rel}}^{\text{max}} = 8.232 \text{ mM/s}$	$I_{\text{leak}}^{\text{max}} = 0.08 \text{ s}^{-1}$ $I_{\text{up}}^{\text{max}} = 0.425 \text{ mM/s}$ $I_{\text{rel}}^{\text{max}} = 8.232 \text{ mM/s}$	
maximum conductances	$C_{\text{Na}}^{\text{max}} = 14.838 \text{ nS/pF}$ $C_{\text{bNa}}^{\text{max}} = 0.00029 \text{ nS/pF}$	$C_{\text{K}1}^{\text{max}} = 5.405 \text{ nS/pF}$ $C_{\text{K}r}^{\text{max}} = 0.0096 \text{ nS/pF}$ $C_{\text{Ks, epi}}^{\text{max}} = 0.245 \text{ nS/pF}$ $C_{\text{Ks, endo}}^{\text{max}} = 0.245 \text{ nS/pF}$ $C_{\text{Ks, M}}^{\text{max}} = 0.062 \text{ nS/pF}$ $C_{\text{pK}}^{\text{max}} = 0.0146 \text{ nS/pF}$ $C_{\text{t}0, \text{epi}}^{\text{max}} = 0.294 \text{ nS/pF}$ $C_{\text{t}0, \text{endo}}^{\text{max}} = 0.073 \text{ nS/pF}$ $C_{\text{t}0, \text{M}}^{\text{max}} = 0.294 \text{ nS/pF}$	$C_{\text{CaL}}^{\text{max}} = 0.175 \text{ mm}^3/[\mu\text{F}]$ $C_{\text{bCa}}^{\text{max}} = 0.000592 \text{ nS/pF}$ $C_{\text{pCa}}^{\text{max}} = 0.825 \text{ pA/pF}$		
half saturation constants	$c_{\text{CaNa}} = 1.38 \text{ mM}$ $c_{\text{NaCa}} = 87.50 \text{ mM}$ $c_{\text{KNa}} = 1.00 \text{ mM}$ $c_{\text{NaK}} = 40.00 \text{ mM}$	$c_{\text{KNa}} = 1.00 \text{ mM}$ $c_{\text{NaK}} = 40.00 \text{ mM}$	$c_{\text{CaNa}} = 1.38 \text{ mM}$ $c_{\text{NaCa}} = 87.50 \text{ mM}$ $c_{\text{pCa}} = 0.0005 \text{ mM}$ $c_{\text{up}} = 0.00025 \text{ mM}$ $c_{\text{rel}} = 0.25 \text{ mM}$ $c_{\text{buf}} = 0.001 \text{ mM}$	$c_{\text{up}} = 0.00025 \text{ mM}$ $c_{\text{rel}} = 0.25 \text{ mM}$ $c_{\text{buf}}^{\text{sr}} = 0.3 \text{ mM}$	
other parameters	$k_{\text{NaCa}}^{\text{sat}} = 0.10$ $\gamma_{\text{NaCa}} = 2.50$ $\gamma = 0.35$	$p_{\text{KNa}} = 0.03$	$\gamma_{\text{rel}} = 2$ $c_{\text{tot}} = 0.15 \text{ mM}$	$\gamma_{\text{rel}} = 2$ $c_{\text{tot}}^{\text{sr}} = 10 \text{ mM}$	
gas constant	$R = 8.3143 \text{ J K}^{-1} \text{ mol}^{-1}$	temperature	$T = 310 \text{ K}$	cytoplasmic volume	$V = 16404 \mu\text{m}^3$
Faraday constant	$F = 96.4867 \text{ C/mmol}$	membrane capacitance	$C = 185 \text{ pF}$	sarcoplasmic reticulum volume	$V^{\text{sr}} = 1094 \mu\text{m}^3$

$$\begin{aligned}
\mathbf{P}^{\text{pas}} = & \kappa [J - 1] \mathbf{F}^{-t} \\
& + a \exp(b [I_1 - 3]) \mathbf{F} \\
& + 2 a_{\text{ff}} [I_{\text{ff}} - 1] \exp(b_{\text{ff}} [I_{\text{ff}} - 1]^2) \mathbf{f} \otimes \mathbf{f}_0 \\
& + 2 a_{\text{ss}} [I_{\text{ss}} - 1] \exp(b_{\text{ss}} [I_{\text{ss}} - 1]^2) \mathbf{s} \otimes \mathbf{s}_0 \\
& + a_{\text{fs}} I_{\text{fs}} \exp(b_{\text{fs}} I_{\text{fs}}^2) \mathbf{f} \otimes \mathbf{s}_0 \\
& + a_{\text{fs}} I_{\text{fs}} \exp(b_{\text{fs}} I_{\text{fs}}^2) \mathbf{s} \otimes \mathbf{f}_0.
\end{aligned} \tag{10.26}$$

parameterized in terms of the isotropic invariants J and I_1 , and the anisotropic invariants I_{ff} , I_{ss} , and I_{fs} , weighted by the bulk modulus κ and the four sets of parameters a and b . For the active Piola stress \mathbf{P}^{act} , we assume that an increase in the intracellular calcium concentration c_{Ca} above a critical level $c_{\text{Ca}}^{\text{crit}}$ induces an active cardiomyocyte contraction F^{act} [34, 69], which is acting along the fiber direction \mathbf{f}_0 [33, 66]. The contractile force F^{act} displays a twitch-type behavior [135], with a smooth off-on

transition characterized through the twitch-function ϵ .

$$\begin{aligned} \mathbf{P}^{\text{act}} &= F^{\text{act}} \mathbf{f} \otimes \mathbf{f}_0 \\ \dot{F}^{\text{act}} &= \epsilon [\eta [c_{\text{Ca}} - c_{\text{Ca}}^{\text{rest}}] - F^{\text{act}}] \\ \epsilon &= \epsilon_0 + [\epsilon_\infty - \epsilon_0] \exp(-\exp(-\xi [c_{\text{Ca}} - c_{\text{Ca}}^{\text{crit}}])) \end{aligned} \quad (10.27)$$

Here, η controls the saturation of the active contractile force F^{act} , $c_{\text{Ca}}^{\text{rest}}$ is the resting concentration, ϵ_0 and ϵ_∞ are the minimum and maximum values of ϵ , $c_{\text{Ca}}^{\text{crit}}$ is the limit value above which contraction is initiated, and ξ is the transition rate from ϵ_0 to ϵ_∞ at $c_{\text{Ca}}^{\text{crit}}$ [66]. To evaluate the iteration matrices (10.19.3) and (10.19.4), we calculate the derivative of the Piola stress \mathbf{P} with respect to the transmembrane potential $\partial_\phi \mathbf{P}$ related to the active stress and with respect to the deformation gradient $\partial_{\mathbf{F}} \mathbf{P}$ related mainly to the passive stress, see [65] for details.

10.5 Examples

10.5.1 Chemo-electro-mechanical coupling in a single cell

To illustrate the local features of our chemo-electrical-mechanical model, we simulate the electrophysiology of an epicardial human ventricular cardiomyocyte throughout a representative excitation cycle. For the chemical parameters, we use the values summarized in Table 10.1. For the electro-mechanical coupling parameters, we choose the saturation of cardiomyocyte contraction to $\eta = 12.5 \text{ kPa}/\mu\text{M}$, the resting concentration of calcium to $c_{\text{Ca}}^{\text{rest}} = 0.05 \mu\text{M}$, the minimum and maximum values scaling fiber contraction to $\epsilon_0 = 0.1/\text{ms}$ and $\epsilon_\infty = 1.0/\text{ms}$, the critical calcium concentration above which contraction is initiated to $c_{\text{Ca}}^{\text{crit}} = 0.8 \mu\text{M}$, and the transition rate to $\xi = 4.0/\mu\text{M}$. We initialize the global membrane potential with $\phi = -86 \text{ mV}$, and the local ion concentrations with $c_{\text{Na}} = 11.6 \text{ mM}$, $c_{\text{K}} = 138.3 \text{ mM}$, and $c_{\text{Ca}} = 0.08 \mu\text{M}$, mimicking the resting state. For the gating variables, we choose the following initial conditions $g_{\text{m}} = 0$, $g_{\text{h}} = 0.75$, $g_{\text{j}} = 0.75$, $g_{\text{d}} = 0$, $g_{\text{f}} = 1$, $g_{\text{fCa}} = 1$, $g_{\text{r}} = 0$, $g_{\text{s}} = 1$, $g_{\text{xs}} = 0$, $g_{\text{xr1}} = 0$, $g_{\text{xr2}} = 0$, $g_{\text{xK1}\infty} = 0.05$, and $g_{\text{g}} = 1$. To initiate a characteristic action potential, we apply an initial electrical stimulus slightly above the critical stimulation

threshold [197].

Figure 10.3, top left, illustrates the evolution of the transmembrane potential ϕ . In cardiac cells at rest, the transmembrane potential is -86 mV, which implies that the intracellular domain is negatively charged in comparison to the extracellular domain. The application of an external stimulus generates an initial depolarization across the cell membrane. Once the stimulus exceeds the critical threshold, the transmembrane potential increases rapidly from its resting state of -86 mV via an overshoot of $+38$ mV to its excited state of $+20$ mV. After a brief period of partial initial repolarization, the transmembrane potential experiences a characteristic plateau of 0.2 ms, before the cell gradually repolarizes to return to its initial resting state.

Figure 10.3, top right, illustrates the evolution of the intracellular sodium concentration c_{Na} , which rises sharply at the beginning of the cycle to create the rapid upstroke of the transmembrane potential. The sodium concentration then decays slowly towards the end of the repolarization phase and increases gradually during the resting phase to return to its initial value. Figure 10.3, bottom left, illustrates the evolution of the intracellular potassium concentration c_{K} . After a rapid increase, c_{K} decreases in a stepwise fashion, regulated by the sequential activation of the individual potassium channels. At the end of the repolarization phase, c_{K} increases gradually to smoothly return to its initial value. Figure 10.3, bottom right, illustrates the evolution of the intracellular calcium concentration c_{Ca} . Slightly after the upstroke of the transmembrane potential, the calcium concentration increases to its peak value and then decays smoothly to its original value throughout the remaining phases of the cycle. In the following section, we will demonstrate how an increase in the intracellular calcium concentration can initiate mechanical contraction. In summary, the model reproduces all characteristic features of human ventricular cardiomyocytes [188, ?, 197]: an initial increase in sodium to create a rapid upstroke in the transmembrane potential, a combined decrease in potassium and increase in calcium to generate the characteristic plateau, and an increase in potassium during the recovery phase to bring the cell back to its resting state. Despite drastic changes in the membrane potential from -86 mV to $+20$ mV, changes in the individual ion concentrations remain remarkably small, typically in the order of less than one percent.

Figure 10.4 illustrates the evolution of the active contractile force F^{act} throughout an excitation cycle. The rapid increase in the intracellular calcium concentration c_{Ca} initiates a rapid increase in the active force. After reaching its peak value, the force gradually returns to zero.

10.5.2 Chemo-electro-mechanical coupling in the human heart

To illustrate the global features of our chemo-electrical-mechanical model, we simulate excitation-contraction coupling in a human heart throughout a representative cardiac cycle. We reconstruct a patient-specific human heart model from magnetic resonance images [102]. Figure see 10.5, left, illustrates the finite element discretization consisting of 46,896 linear tetrahedral elements, 13,831 nodes, and 55,324 degrees of freedom. To account for the characteristic microstructure of the heart, we assign locally varying fiber vectors \mathbf{f}_0 and sheet vectors \mathbf{s}_0 from a feature-based interpolation using a Lagrangian interpolant. Figure 10.5, right, illustrates the fiber distribution across the left and right ventricles. Fiber directions vary gradually from -70° in the epicardium, the outer wall shown in blue, to $+80^\circ$ in the endocardium, the inner wall shown in red. Sheet directions are outward-pointing with respect to the epicardial surface.

Similar to the single cell example in Section 10.5.1, we apply initial conditions which mimic the resting state, with a global membrane potential of $\phi = -86 \text{ mV}$, and the local ion concentrations of $c_{\text{Na}} = 11.6 \text{ mM}$, $c_{\text{K}} = 138.3 \text{ mM}$, $c_{\text{Ca}} = 0.08 \mu\text{M}$, and $c_{\text{Ca}}^{\text{sr}} = 0.56 \text{ mM}$, and gating variables of $g_{\text{m}} = 0$, $g_{\text{h}} = 0.75$, $g_{\text{j}} = 0.75$, $g_{\text{d}} = 0$, $g_{\text{f}} = 1$, $g_{\text{fCa}} = 1$, $g_{\text{r}} = 0$, $g_{\text{s}} = 1$, $g_{\text{xs}} = 0$, $g_{\text{xr1}} = 0$, $g_{\text{xr2}} = 0$, $g_{\text{xK1}\infty} = 0.05$, and $g_{\text{g}} = 1$. Tables 10.1 summarizes the chemo-electrical parameters which are similar to single cell example in Section 10.5.1. To account for regionally varying action potential durations, we divide the heart in five regions, basal septum, apical septum, apex, mid-ventricular wall, and lateral ventricular wall [?]. We systematically increase the bulk ion channel conductances $C_{\text{Kr}}^{\text{max}}$, $C_{\text{Ks}}^{\text{max}}$, and $C_{\text{CaL}}^{\text{max}}$ from upper septum to lateral wall by $\pm 30\%$. The

Table 10.2: Electro-mechanical material parameters of human cardiac tissue.

	electrical parameters
isotropic conduction	$d^{iso} = 5 \text{ mm}^2/\text{ms}$ [66]
anisotropic conduction	$d^{ani} = 10 \text{ mm}^2/\text{ms}$ [66]
	mechanical parameters
isotropic bulk	$\kappa = 100 \text{ kPa}$ [65]
isotropic myocardium	$a = 0.496 \text{ kPa}, b = 7.209$ [65]
anisotropic myocardium	$a_{ff} = 15.193 \text{ kPa}, b_{ff} = 20.417$ [65]
	$a_{ss} = 3.283 \text{ kPa}, b_{ss} = 11.176$ [65]
	$a_{fs} = 0.662 \text{ kPa}, b_{fs} = 9.466$ [65]
	electro-mechanical parameters
saturation of contraction	$\eta = 12.5 \text{ kPa}/\mu\text{M}$ [66]
resting calcium concentration	$c_{Ca}^{rest} = 0.05 \mu\text{M}$ [197]
critical calcium concentration	$c_{Ca}^{crit} = 0.8 \mu\text{M}$
minimum activation	$\epsilon_0 = 0.1 / \text{ms}$ [66]
maximum activation	$\epsilon_\infty = 1.0 / \text{ms}$ [66]
transition rate	$\xi = 4.00 / \mu\text{M}$ [66]

electro-mechanical coupling parameters are the saturation of cardiomyocyte contraction $\eta = 12.5 \text{ kPa}/\mu\text{M}$, the resting concentration of calcium $c_{Ca}^{rest} = 0.05 \mu\text{M}$, the minimum and maximum values scaling fiber contraction $\epsilon_0 = 0.1/\text{ms}$ and $\epsilon_\infty = 1.0/\text{ms}$, the critical calcium concentration above which contraction is initiated $c_{Ca}^{crit} = 0.8 \mu\text{M}$, and the transition rate $\xi = 4.0/\mu\text{M}$. These are the same values as in the single cell example in Section 10.5.1, which have been calibrated such that the maximum fiber contraction λ_{ff} is approximately 15% [185]. The electrical parameters are the isotropic and anisotropic conductivities $d^{iso} = 5 \text{ mm}^2/\text{ms}$ and $d^{ani} = 10 \text{ mm}^2/\text{ms}$. The mechanical parameters are the isotropic elastic bulk modulus $\kappa = 100.0 \text{ kPa}$, the isotropic elastic tissue parameters $a = 0.496 \text{ kPa}$ and $b = 7.209$, the anisotropic elastic parameters, $a_{ff} = 15.193 \text{ kPa}$ and $b_{ff} = 20.417$, $a_{ss} = 3.283 \text{ kPa}$ and $b_{ss} = 11.176$, and $a_{fs} = 0.662 \text{ kPa}$ and $b_{fs} = 9.466$, which we have identified [65] using simple shear experiments from the literature [50]. Table 10.2 summarizes the electro-mechanical parameters.

For the electrical problem, we apply the common assumption of homogeneous Neumann boundary conditions. For the mechanical problem, we apply homogeneous

Dirichlet boundary conditions throughout the basal plane. We excite the heart through an external stimulus in the region of the atrioventricular node located in the center of the basal septum. We apply an adaptive time stepping scheme, for which we select the optimal number of iterations to four. For a larger number of iterations, the adaptive scheme decreases the time step size; for a smaller number of iterations, the adaptive scheme increases the time step size.

Figures 10.6 and 10.7 illustrate the evolution of the fiber contraction λ_{ff} , of the transmembrane potential ϕ , and of the individual ion concentrations c_{Na} , c_{K} , c_{Ca} , and $c_{\text{Ca}}^{\text{sr}}$ during the depolarization and repolarization phases, respectively. Figure 10.6 shows how depolarization is initiated through changes in the intracellular sodium concentration c_{Na} , which increases rapidly within the first milliseconds of the cardiac cycle, third row. This increase is associated with a rapid increase in the membrane potential ϕ , second row, which, in turn, affects the voltage-gated calcium and potassium channels within the cell membrane. The intracellular calcium concentration c_{Ca} increases, fifth row. The intracellular potassium concentration c_{K} follows with a slight time delay of 15 ms, fourth row. The intracellular calcium concentration c_{Ca} increases further as calcium is released from the sarcoplasmic reticulum $c_{\text{Ca}}^{\text{sr}}$, sixth row. The increase in the intracellular calcium concentration directly initiates cardiomyocyte contraction λ_{ff} , first row. Contraction varies regionally and transmurally with maximum values of $\lambda_{\text{ff}} = -10\%$. As the heart contracts, the apex moves markedly upward towards the fixed base, columns four and five. After 50 ms, the heart is entirely depolarized. The transmembrane potential ϕ has reached its peak value of 20 mV throughout both ventricles, and the heart is maximally contracted.

Figure 10.7 displays the repolarization phase characterized through a smooth decrease of the transmembrane potential ϕ and the mechanical contraction λ_{ff} back to their resting values, first and second row. Decrease in mechanical contraction is caused by a gradual decrease of the intracellular calcium c_{Ca} concentration back to its resting value, fifth row. The sarcoplasmic reticulum takes up the intracellular calcium, and $c_{\text{Ca}}^{\text{sr}}$ returns back to its resting value, sixth row. At the same time, the intracellular sodium concentration c_{Na} , which has initially increased, now dips even below its initial value and reaches a minimum after 260 ms, third row. The intracellular potassium

concentration c_K reaches its minimum approximately at the same time, fourth row. In the course of time, both sodium and potassium then slowly return to their resting values as their concentrations increase gradually. The temporal evolution of the mechanical, electrical, and chemical fields is in excellent qualitatively and quantitatively agreement with the single cardiomyocyte transients documented in Figures 10.3 and 10.4.

Figure 10.8 illustrates the global performance of the heart in dry pumping in terms of two characteristic clinical metrics of cardiac function, apical lift δ and of ventricular torsion ϑ . Figure 10.8, left, shows the apical lift, i.e., the vertical movement of the apex along the heart's long axis towards the fixed base. Shortly after the onset of excitation, the apex lifts rapidly towards the base moving upward by approximately 14mm. Figure 10.8, right, shows the ventricular torsion, i.e., the rotation of two marked locations in the lateral left ventricular wall, at approximately 1/3 and 2/3 height, around the heart's long axis. Shortly after the onset of excitation, the heart undergoes a rapid twist, rotating clockwise by approximately 8° and 17° , with the amount of torsion increasing from the fixed base to the free apex. Both apical lift and ventricular torsion then decrease gradually as the heart returns to its original position.

Figure 10.9 demonstrates the performance of our fully implicit monolithic finite-element based algorithm. Figure 10.9, left shows the variation of the time step size and Figure 10.9, right, shows the corresponding number of Newton iterations within the adaptive time stepping scheme. The algorithm typically convergences within four Newton Raphson iterations. For more required iterations, the adaptive algorithm automatically decreases the time step size, for example, during the rapid upstroke phase before $t = 0.05$ s and during the repolarization phase between $t = 0.25$ s and $t = 0.32$ s. For less required iterations, the adaptive algorithm automatically increases the time step size, for example during the plateau phase, between $t = 0.05$ s and $t = 0.25$ s and during the resting phase after $t = 0.32$ s. We do not observe stability issues, which we attribute to the implicit nature of the underlying time integration scheme. The simulation run of an entire cardiac cycle finishes after a total number of time increments of 1,288. The overall run time is 51.97 hours, calculated on a single core

of an i7-950 3.06 GHz desktop with 12GB of memory.

10.6 Discussion

We have presented a unified, fully coupled finite element formulation for chemo-electro-mechanical phenomena in living biological systems and demonstrated its potential to simulate excitation-contraction coupling in a patient-specific human heart. The novel aspect of this work is that all chemical, electrical, and mechanical fields are solved monolithically using an implicit time integration scheme, consistently linearized, embedded in a Newton-Raphson solution strategy. In contrast to most existing algorithms, the proposed discretization scheme is unconditionally stable, computationally efficient, highly modular, geometrically flexible, and easily expandable. *Unconditional stability* is guaranteed by using a fully coupled, implicitly integrated, consistently linearized finite element approach. Existing algorithms are typically based on sequential, staggered solution techniques [71, 183] and utilize explicit time marching schemes [98, 106]. They are inherently unstable and limited in time step size, which might make them less robust and less efficient. Especially during the rapid upstroke phase, steep spatial and temporal gradients in the unknown fields might initiate spurious instabilities when using explicit time stepping schemes [136]. To avoid these potential limitations, we have applied an implicit backward Euler time integration scheme [66, 197]. We have shown that this scheme is capable of handling sharp chemical, electrical, and mechanical profiles associated with rapid changes in the local and global unknowns. Since our algorithm follows the classical layout of nonlinear finite element schemes, we can utilize readily available adaptive time stepping schemes at no extra cost or effort [?]. We have demonstrated that a simple, ad-hoc, iteration-counter based time adaptive scheme automatically decreases the time step size during phases with steep temporal gradients and, conversely, increases the time size when all unknowns evolve smoothly.

Efficiency is not only increased by using time adaptive schemes, but also by using a classical finite-element specific global-local split [72, 151]. While most existing algorithms discretize all unknowns globally, we only introduce four global degrees of

freedom at the node point level, i.e., the vector-valued mechanical deformation and the scalar-valued electrical potential. We introduce, update, and store all other state variables locally on the integration point level [73, 145, 161], i.e., the thirteen chemical gating variables and the four ion concentrations for our particular cardiomyocyte model. Accordingly, our global system matrix remains small and efficiently to invert during the solution procedure.

Modularity originates from the nature of the underlying finite element discretization, which introduces all cell-specific unknowns as local internal variables on the integration point level. This allows us to modularly integrate the proposed algorithm into any commercial finite element package that can handle a coupled nonlinear system of vector- and scalar-valued governing equations [?]. The simplest strategy would be to use an existing thermo-mechanical element formulation and re-interpret the temperature field as the transmembrane potential. Algorithmic modifications are then restricted exclusively to the constitutive subroutine, in which we would solve the chemical problem and store the ion concentrations and gating variables as internal variables at each integration point [197]. Another natural benefit of using finite element schemes is that the modular treatment of the constitutive equations allows us to combine arbitrary cell types, e.g., epicardial and endocardial ventricular cells [116, 188], purkinje fiber cells [122], atrial cells [43], and pacemaker cells [140], to effortlessly account for microstructural inhomogeneities.

Geometrical flexibility is probably the most advantageous feature of finite element techniques when compared to finite difference schemes or finite volume methods. Unlike existing schemes which are most powerful on regular grids [24], the proposed algorithm can be applied to arbitrary geometries with arbitrary initial and boundary conditions [167]. Finite element algorithms can easily handle medical-image based patient-specific geometries [20, 102, 203]. In a simple pre-processing step, we could even utilize finite element algorithms to create fiber orientations on arbitrary patient-specific meshes using Lagrangian feature-based interpolation as illustrated in Figure 10.5. The key advantage of finite element algorithms, however, is that they allow us to simulate finite deformations throughout the cardiac cycle in a straightforward and natural way [66, 183].

Ease of expandability is attributed to the fact that we use a single unified discretization technique. Being finite-element based and transparent in nature, our approach lays the groundwork for a robust and stable whole heart model of excitation-contraction coupling. Through the incorporation of an additional scalar-valued global unknown, e.g., to characterize the extracellular potential field [?, 44], we could easily expand the proposed formulation into a bidomain formulation [126, 163, 162]. Through the incorporation of additional gating variables as local unknowns, e.g., to characterize the optical manipulation of cardiac cells [1], we could easily expand the proposed formulation into a photo-electro-chemical formulation [?].

In summary, we strongly believe that there are compelling reasons to consider the use of fully coupled, implicitly integrated, consistently linearized discretization strategies that enjoy the advantages inherent to finite element schemes. Initially, it may seem tedious to transition existing algorithms into a single unified algorithmic framework. However, we are convinced that these efforts will pay off when it comes to truly predicting the impact of pharmacological, interventional, and surgical treatment options to systematically manipulate chemical, electrical, and mechanical fields in the human heart.

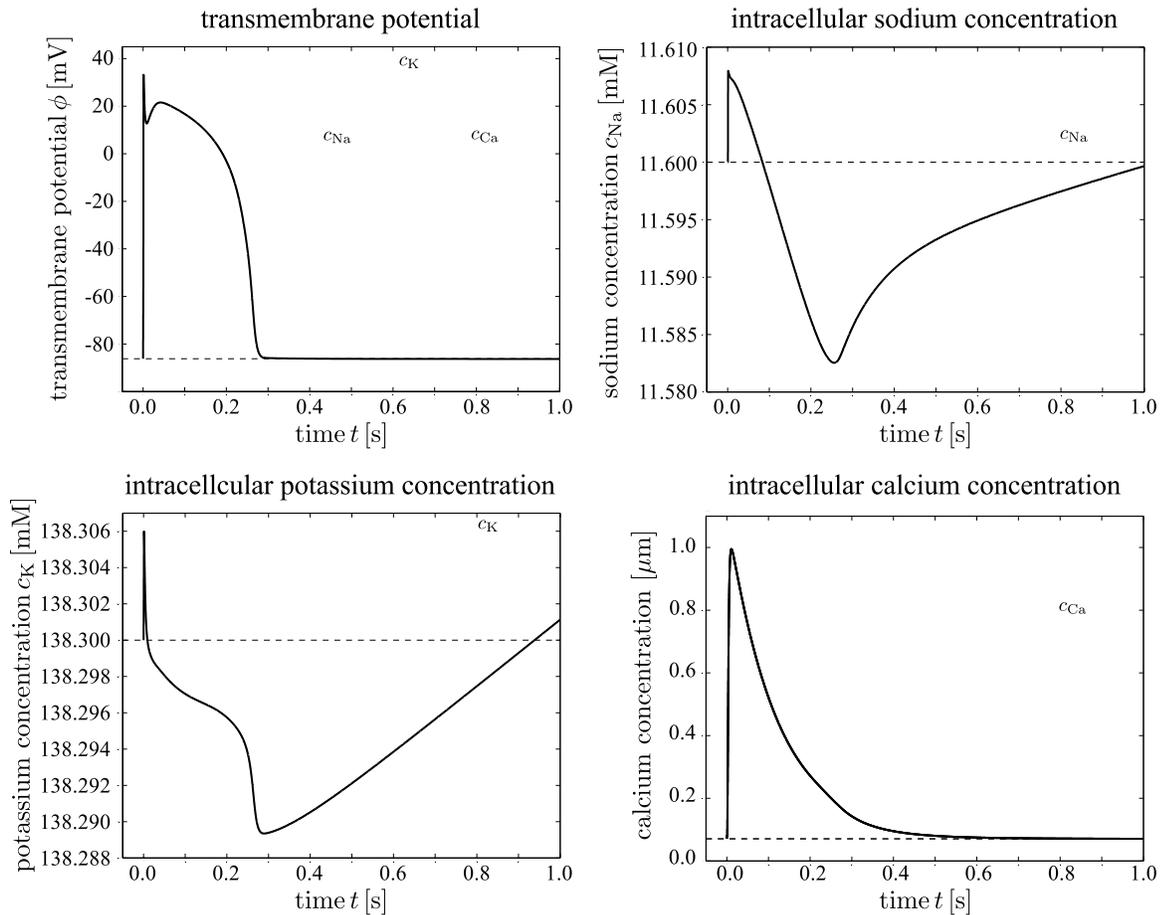


Figure 10.3: Electrochemistry in a human ventricular cardiomyocyte. Temporal evolution of the transmembrane potential ϕ and of the intracellular sodium potassium and calcium concentrations c_{Ca} , c_K , and c_{Ca} . The influx of positively charged sodium ions generates a rapid upstroke in the transmembrane potential. At peak, the efflux of positively charged potassium ions initiates an early, partial repolarization. During the plateau, the influx of positively charged calcium ions balances the efflux of positively charged potassium ions. Final repolarization begins when the efflux of potassium ions exceeds the influx of calcium ions. The cell is at rest throughout the interval between the end of repolarization and the beginning of the next cycle.

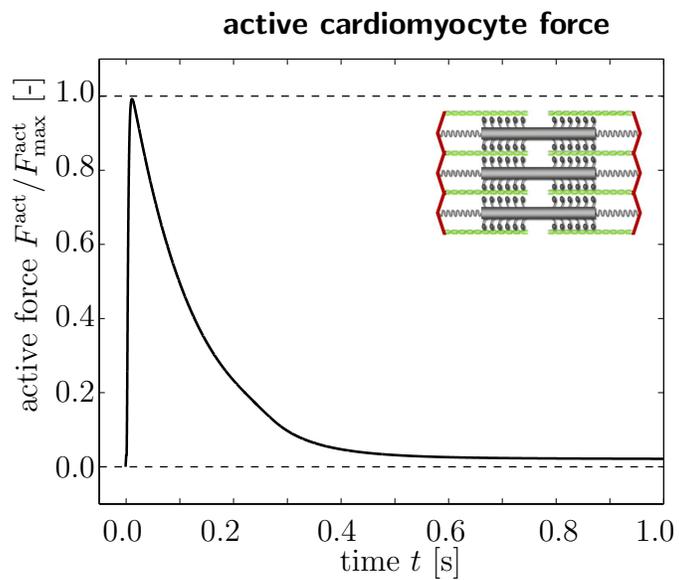


Figure 10.4: Mechanical contraction in a human ventricular cardiomyocyte. Temporal evolution of the active force F^{act} throughout an excitation cycle. The rapid increase in the intracellular calcium concentration c_{Ca} initiates a rapid increase in the active force. After reaching its peak value, the force gradually returns to zero.

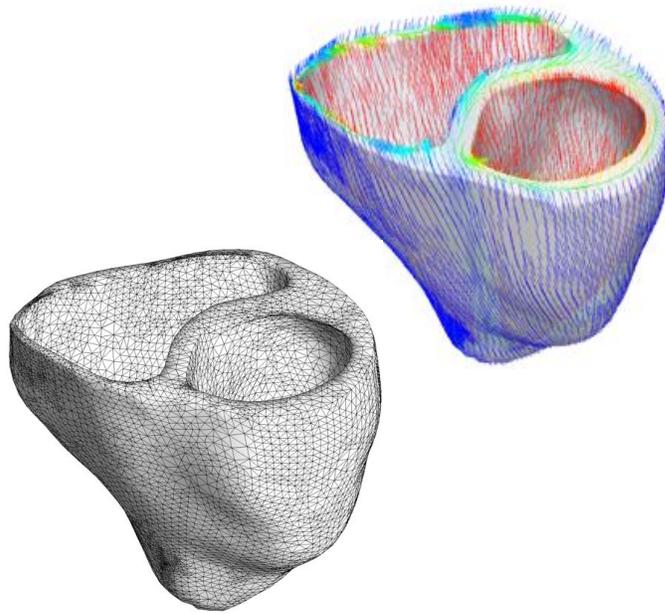


Figure 10.5: Human heart model. The finite element discretization is generated from magnetic resonance images, left. The mesh consists of 46,896 linear tetrahedral elements, 13,831 nodes, and 55,324 degrees of freedom. The fiber orientation is created from feature interpolation using a Lagrangian interpolant, right. Fiber directions vary gradually from -70° in the epicardium, the outer wall shown in blue, to $+80^\circ$ in the endocardium, the inner wall shown in red.

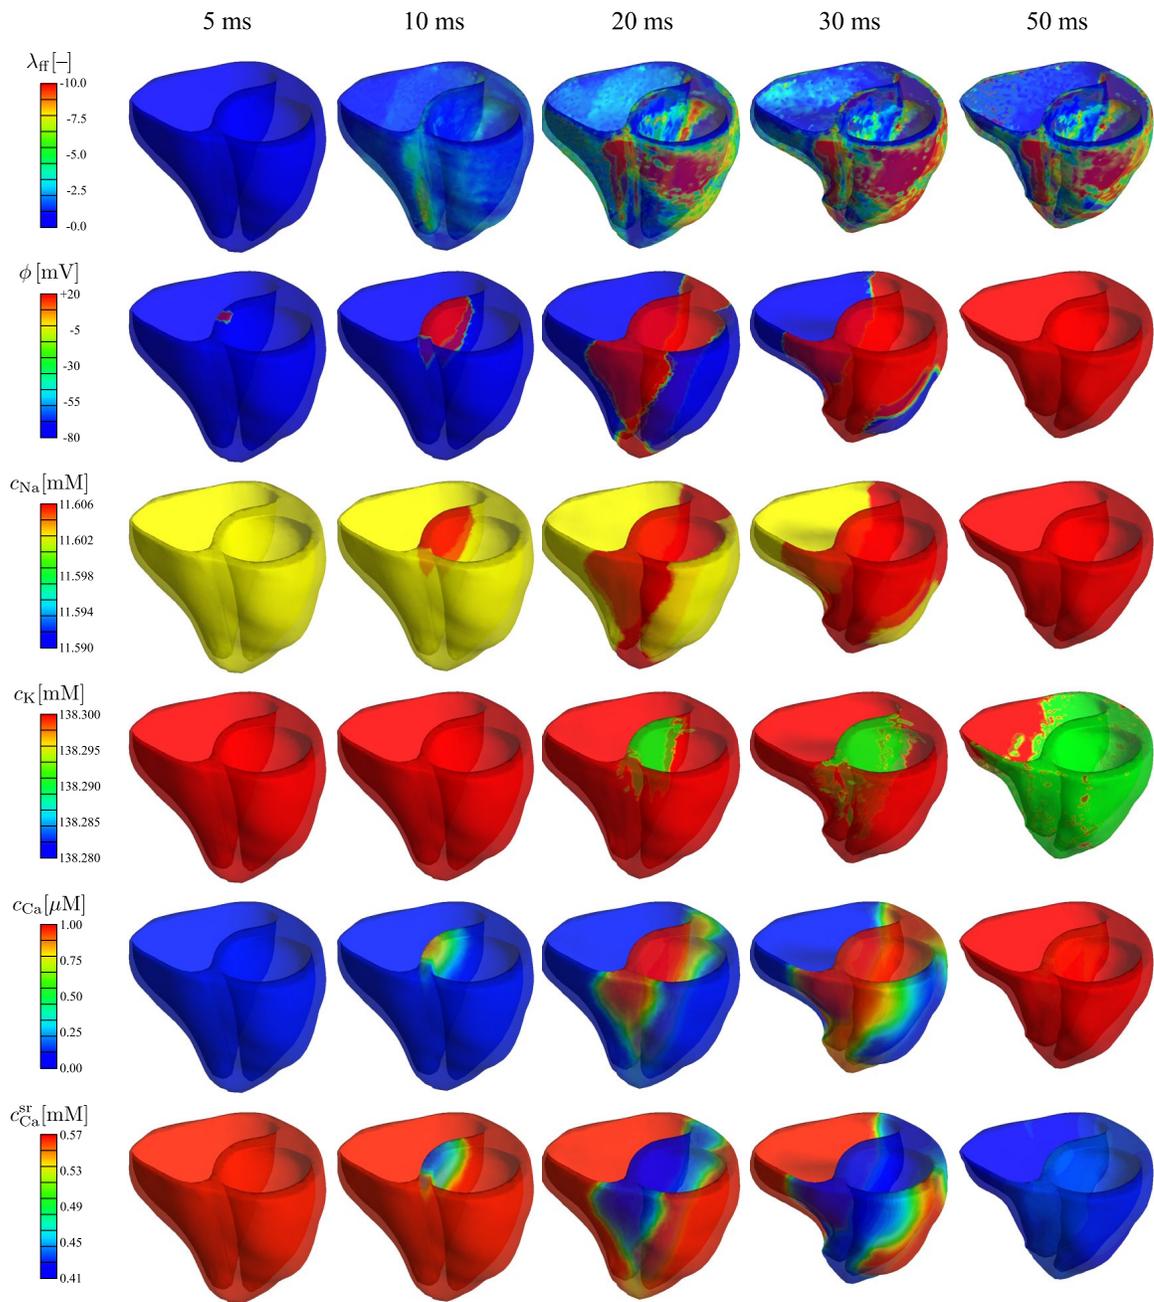


Figure 10.6: Chemo-electro-mechanical coupling in the human heart. Spatio-temporal evolution of the fiber contraction λ_{ff} , the transmembrane potential ϕ , the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_K , and c_{Ca} , and the calcium concentration in the sarcoplasmic reticulum c_{Ca}^{sr} during the rapid depolarization phase of the cardiac cycle.

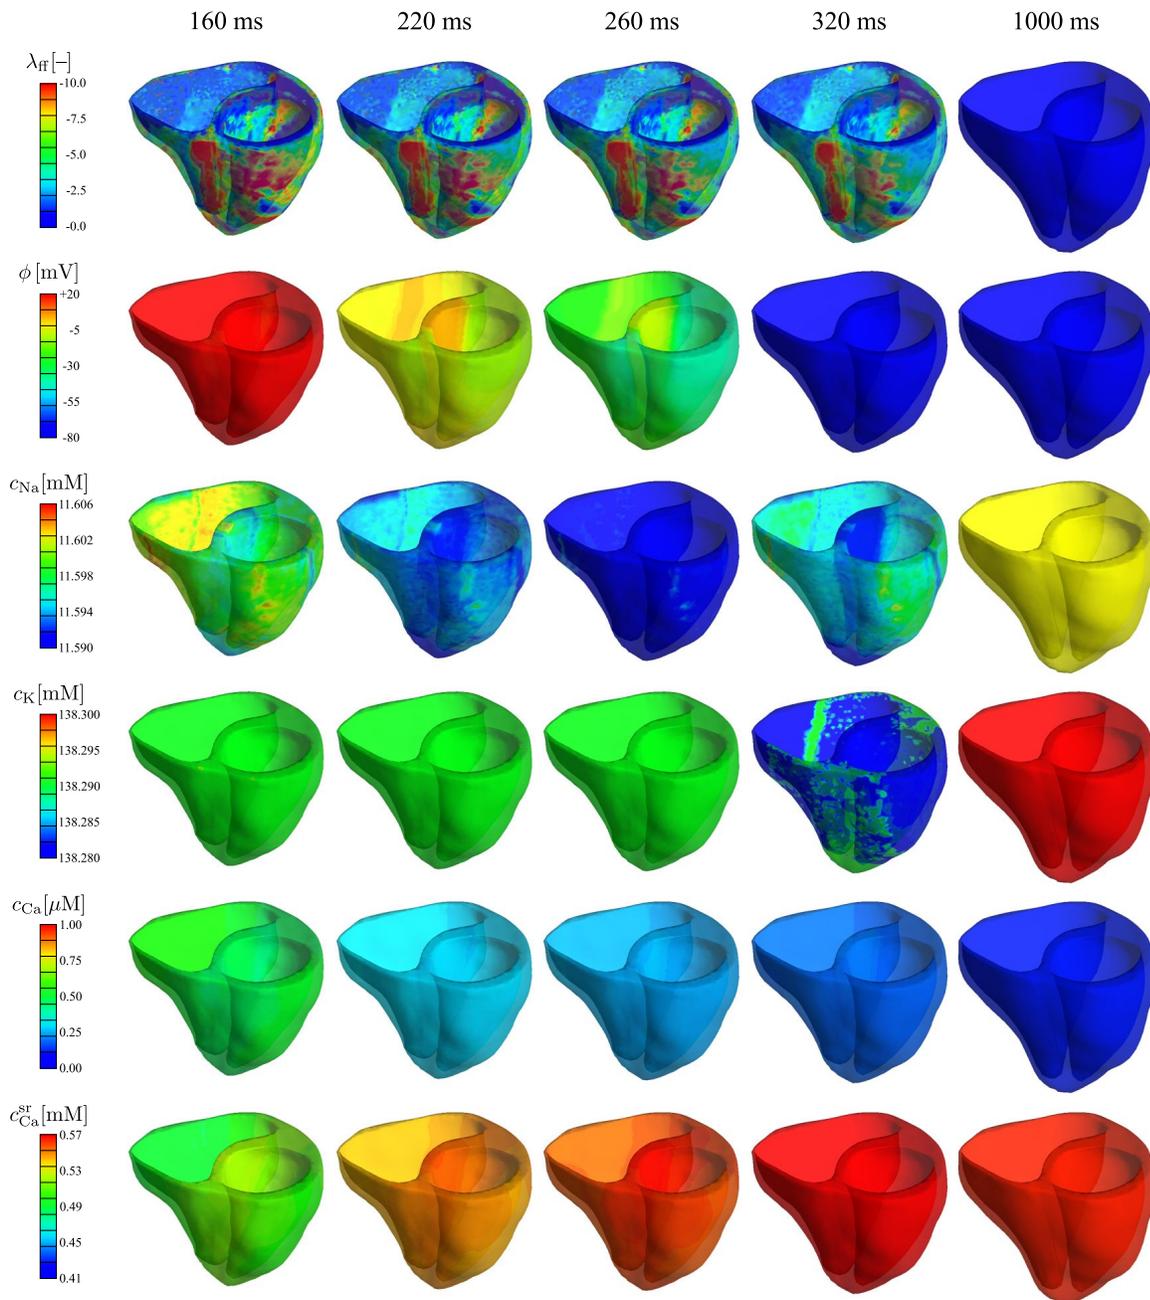


Figure 10.7: Chemo-electro-mechanical coupling in the human heart. Spatio-temporal evolution of the fiber contraction λ_{ff} , the transmembrane potential ϕ , the intracellular sodium, potassium, and calcium concentrations c_{Na} , c_K , and c_{Ca} , and the calcium concentration in the sarcoplasmic reticulum c_{Ca}^{sr} during the gradual repolarization phase of the cardiac cycle. Changes in the individual ion concentrations initiate a slow decrease in the transmembrane potential ϕ from +20 mV to -86 mV. A decrease in the intracellular calcium concentration c_{Ca} initiates mechanical relaxation with λ_{ff} returning gradually to 0%. During the filling phase, the apex moves away from the base and the heart undergoes a counterclockwise rotation back to its original position.

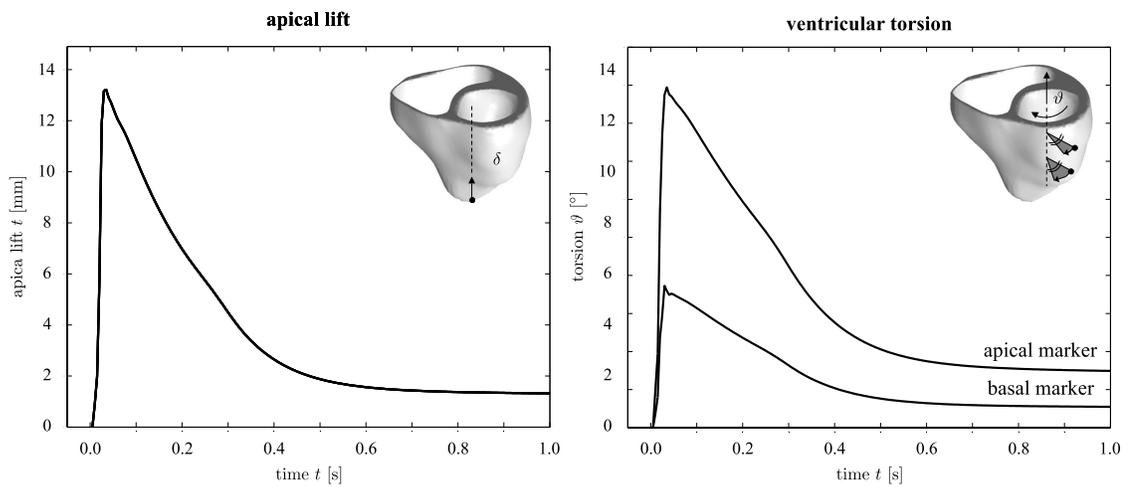


Figure 10.8: Mechanical contraction in a human heart. Temporal evolution of apical lift δ and of ventricular torsion ϑ . Shortly after the onset of excitation, the apex lifts rapidly towards the base moving upward by approximately 14mm, left. Simultaneously, the heart twists rapidly about its long axis rotating clockwise by approximately 8° and 17° , with the amount of torsion increasing from the fixed base to the free apex, right. Both apical lift and ventricular torsion then decrease gradually as the heart returns to its original position.

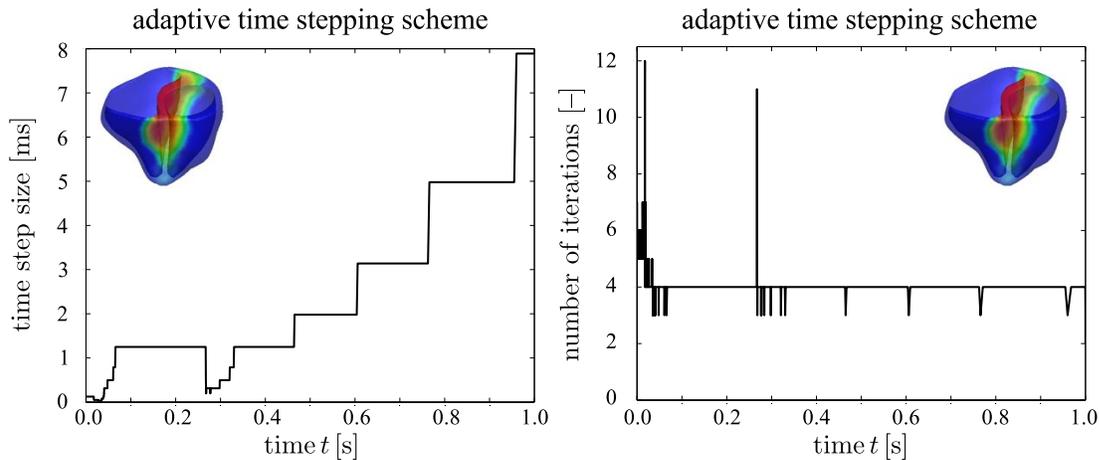


Figure 10.9: Algorithmic performance. Time step size and number of iterations for adaptive time stepping scheme. The algorithm typically converges within four Newton Raphson iterations. For more iterations, the adaptive algorithm automatically decreases the time step size, for example, during the rapid upstroke phase before $t = 0.05$ s and during the repolarization phase between $t = 0.25$ s and $t = 0.32$ s. For less iterations, the adaptive algorithm automatically increases the time step size, for example during the plateau phase, between $t = 0.05$ s and $t = 0.25$ s and during the resting phase after $t = 0.32$ s. The total number of time increments is 1,288, and the overall run time is 51.97 hours, calculated on a single core of an i7-950 3.06 GHz desktop with 12GB of memory.

Chapter 11

GPU applications

11.1 Introduction

The recent use of graphics processors (GPUs) for scientific applications has opened up possibilities in achieving near-real-time simulations, and has been shown to provide remarkable speed ups for many different applications. Scientific applications using GPUs range from computing the flow over hypersonic vehicles [51] to finite element simulations of virtual hearts [74] and of viscoelastic properties of soft tissue [178] to medical imaging applications for cancer treatment [148]. There is a developing body of literature concerning GPU implementation of finite elements. Most works have concentrated mainly on the implementation of GPU conjugate-gradient (CG) and multigrid solvers [25, 21, 64]. Recently, important investigations have also been made in optimizing GPU finite element assembly algorithms [21, 31]. While GPUs are already used for many inherently parallel operations, such as material point ODE solvers for finite elements [112, 190] and for finite differences [9], total performance gains for finite elements on GPUs can only be realized by efficiently combining the different aspects: element calculations, global assembly, and solver.

However, it is generally difficult to re-implement CPU algorithms for the GPU as architectural differences pose difficulties in developing GPU methods that outperform their CPU counterparts. Difference in GPU memory architecture, speed of individual GPU cores, and the communication between GPU cores, GPU memory,

and host (CPU) memory can greatly impact the performance of GPU algorithms. Because these differences, much effort has been spent in developing new algorithms designed for general purpose use. For example, scan primitives are useful in a large range of applications ranging from prefix-sum, sparse matrix vector multiplications (SPMV), and sorting [166]. However, the performance of many of these algorithms are dependent on the structure of the data. In particular, sparse matrices can be represented by a variety of formats and the computational performance may differ greatly depending on the format used[13].

In developing a GPU-based finite element code for near-real-time heart simulations, we developed an alternative GPU algorithm for SPMV that works well for finite element unstructured grids. While other algorithms have also targeted finite element sparse matrices by grouping rows with the same row structure [192], or by encoding sparse matrices in blocks [25], our algorithms mainly exploit row-length regularity for padding efficiency, data reordering, and localize computation within the intrinsic hardware-based Single-Instruction Multiple-Thread (SIMT) architecture. Surprisingly, the algorithm also performs well over finite element matrices as well as sparse matrices in general. Within the finite element method, SPMV routines can be leveraged in both the global assembly operation and the solver to gain additional computational performance.

In this work, we first describe the real-time heart simulation problem within the context of a traditional finite element framework. We then highlight costs of both the assembler and solver and highlight differences between CPU and GPU implementations. A survey of different GPU SPMV algorithms then follows. Our new set of SPMV algorithms are then introduced, and a variety of SPMV algorithms are benchmarked against our novel algorithms. Different aspects of optimizations used in the development of our algorithms are investigated. We discuss the utility of our spmv algorithms, and specifically, how they may be applied within a finite element context, and in a general computational setting. Lastly, we end with concluding remarks and possible future improvements to our algorithms and also possible improvements in leveraging GPUs for finite element simulations.

11.2 Description of physical problem and finite element breakdown

Our original aim was to utilize GPUs for achieving near-real-time electrophysiological simulations of the heart. The electrical propagation of voltage within the heart is described by the following two-variable phenomenological governing differential equations for nonlinear mono-domain excitable tissues.

$$\dot{\phi} = \operatorname{div} \mathbf{q}(\phi) + f^\phi(\phi, r) \quad (11.1)$$

$$\dot{r} = f^r(\phi, r) \quad (11.2)$$

The transmembrane voltage, ϕ , is the voltage difference across the cell membrane, while the phenomenological variable, r , describes the recovery behavior of cardiac tissue. To account for conduction throughout the tissue, a phenomenological potential flux term $\operatorname{div} \mathbf{q}$ is added.

$$\mathbf{q} = \mathbf{D} \cdot \nabla \phi \quad (11.3)$$

A phenomenological diffusion tensor $\mathbf{D} = d_{\text{iso}} \mathbf{I} + d_{\text{ani}} \mathbf{n} \otimes \mathbf{n}$ allows for cell-to-cell electrical coupling across cellular gap junctions. d_{iso} and d_{ani} are the respective isotropic and anisotropic conduction terms and \mathbf{n} is the preferred direction of anisotropic conduction.

In this paper, we will use the classical Aliev-Panfilov model [4] for convenience to evaluate the effectiveness of our GPU algorithms. The source terms for the Aliev-Panfilov model are

$$f^\phi = c \phi [\phi - \alpha][1 - \phi] - r \phi \quad (11.4)$$

$$f^r = \left[\gamma + \frac{\mu_1 r}{\mu_2 + \phi} \right] [-r - c \phi [\phi - b - 1]] \quad (11.5)$$

While the mono-domain equations (11.1) and (11.2) can be solved simultaneously, a

global-local internal variable splitting approach [72, 197] is taken because it yields a global symmetric tangent matrix for the global degree of freedom, ϕ , such that a fast iterative GPU solver can be effectively employed. With the global-local splitting approach, Equation (11.2) is solved locally at each integration point within a finite element in an iterative manner using Newton's method, while Equation (11.1) is solved globally using the finite element method over the domain. Both the local and global degrees of freedom are updated per global Newton-Raphson iteration to ensure the consistency of the solution. The scheme is summarized in Table 11.1. The

Table 11.1: Algorithmic treatment of transmembrane voltage in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ , is introduced globally on the node point level whereas the phenomenological recovery variable, r , is introduced locally on the integration point level. Element assembly is highlighted in green, and the solution update is highlighted in red.

initialize nodal degrees of freedom ϕ_J initialize internal variable r
global Newton iteration
loop over all elements
loop over all integration points
local Newton iteration
calculate local recovery variable residual, $R^r = [r - r^n]/\Delta t - f^r$ and local iteration matrix $[K^r] = d_r R^r$
update ion concentrations $r \leftarrow r - [K^r]^{-1} R^r$
calculate source term $f^\phi(\phi, r)$ and its linearization $d_\phi f^\phi$
calculate element residuals $R_I^{\phi e}$ and element matrices $K_{IJ}^{\phi e} = d_{\phi_J}^e R_I^{\phi e}$
calculate global residual R_I^ϕ and global iteration matrix $K_{IJ}^\phi = d_{\phi_J} R_I^\phi$
update membrane potential $\phi_J \leftarrow \phi_J - K_{IJ}^{\phi-1} R_I^\phi$

heart mesh is discretized using Lagrangian shape functions in space and tetrahedral

representations are used in this paper.

$$\delta\phi^h|_{\mathcal{B}^e} = \sum_{i=1}^{n_{\text{en}}} N^i \delta\phi_i \quad \phi^h|_{\mathcal{B}^e} = \sum_{j=1}^{n_{\text{en}}} N^j \phi_j \quad (11.6)$$

As the coupled nonlinear problem is stiff, backward Euler implicit time integration is used where $\dot{\phi} = \frac{1}{\Delta t}[\phi - \phi_n]$. Implicit time integration can be further leveraged in utilization of adaptive time-stepping schemes to further reduce computation time. Lastly reduced-order integration is used, whereby a linear tetrahedral element will only have one integration point. All implementations developed in this study are scientifically accurate and run at double precision. This summarizes are algorithmic framework for solving the physical problem.

11.2.1 Computational cost of the Finite Element method

While the biophysical problem may initially seem rather specific, the problem has been structured in the standard non-linear solid mechanics finite element framework. For example, such local-global splitting schemes are commonly used in the field of plasticity [168]. As such, a comparison between computation cost of our heart simulation for various refinements of the same mesh are shown in Figure 11.1 for a single CPU finite element method implementation of the electrophysiological problem using a standard scientific code, PETSC [7] and for our custom single GPU implementation. Our GPU implementation utilizes sparse matrix vector multiplication operations to assemble the global tangent and residual quantities, and SPMV is also employed in our standard jacobi preconditioned conjugate gradient solver. The use of SPMV within the context of CG solvers is prevalent [128, 192, 25, 21, 64]; however the benefit in using SPMV for global assembly on GPUs [31] is not particularly obvious and will be explained in more detail in the following sections.

For our GPU heart simulations the portion of time spent on SPMV during assembly is relatively small compared to total the total assembly time and increases towards 33% within the CG solver as the level of refinement increases (Figure 11.2). Together, total time spent in our GPU implementation shows that up to 10% of

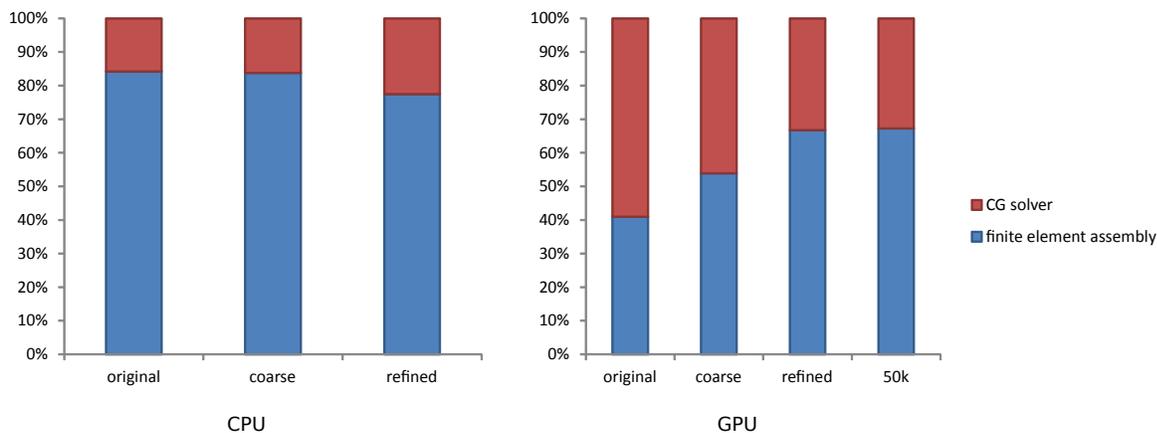


Figure 11.1: Percentage of computation time dedicated to finite element assembly (red) and the CG solver (blue) respectively on the CPU (left) and GPU (right).

computation time is spent purely on SPMV calculation. Thus, while general finite element problems may exhibit different solver-to-assembly ratios, improvements to SPMV operations will still be beneficial to the finite element frameworks. In the following sections, we will investigate various improvements and novel modifications to current SPMV algorithms and highlight their use within the finite element method and in other general SPMV applications.

11.3 GPU Architectural Organization

As GPUs are architecturally organized differently than traditional CPUs, a brief summary is given before we examine current state-of-the-art algorithms. These architectural differences are important in developing a substantially improved SPMV algorithm. Since we are using an Nvidia GPU, we will describe the general organization of an Nvidia GPU. The overall work-execution model on the GPU is composed of individual threads which each execute a specified kernel routine concurrently with respect to other threads (SIMT). Groups of 32 threads, called warps, are executed implicitly in synchronous. Furthermore, up to 16 warps can be grouped into a block. Threads from a block have access to shared memory within that block and may also

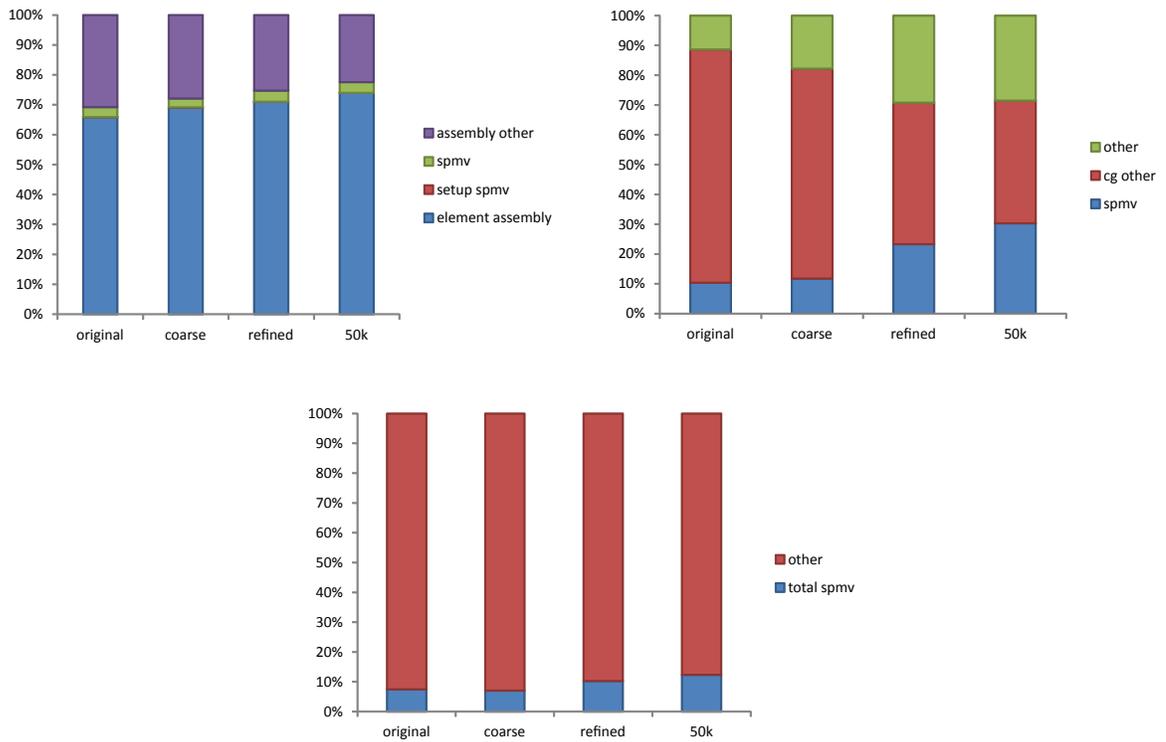


Figure 11.2: The proportion of time spent on different finite element assembly tasks within our GPU implementation for different refinements is shown (top left). Proportions of SPMV with relation to other tasks for the CG solver is shown (top right). Total proportional time spent on SPMV is shown (bottom).

be forced to synchronize with other threads within that block.

On the GPU, there are basically four types of memory available: global, shared, local, and register memory. In general, memory that is more accessible to more threads is slower than memory that is shared with a smaller group of threads. For example, global memory can be accessed by all threads on a GPU, but it is also the slowest form of memory available. Shared memory is shared within a block. It is faster than global memory when there are few memory bank conflicts and when threads are accessing the same portions of shared memory simultaneously. Local memory is global memory reserved for threads, where stores are kept in the L1 cache. Generally it is used in the background when a thread has fully utilized the maximum amount of memory registers available. Registers are the fastest form of memory, but register values cannot be shared with other threads and only 63 register are available for a given thread. Different memory caches are used to facilitate fetching memory accesses to and from the different types of memory to the individual threads. In this paper, the L1 and L2 caches are implicitly used and use of the texture cache, which is a special spatial locality-based global memory cache, is also investigated.

11.4 Description of current GPU methods

In order to achieve near-real-time finite element heart simulations, ideally both the finite element tangent matrix assembly and finite element solver will perform a majority of computations on the GPU to minimize CPU-to-GPU memory transfer overhead. Current GPU approaches and methods are summarized below.

For the finite element assembly of cardiac simulations, efforts have been made at leveraging GPUs to speed up the evaluation of the local ODE problem (11.2) [112, 190]. Since we have chosen a simple phenomenological model, these efforts are not directly applicable to this study. However, several different finite element assembly approaches have been recently investigated [21, 31] and show that assembly by non-zeros of the finite element tangent matrix is a better approach compared to assembly by elements, which is the traditional serial finite element approach. Furthermore, if we restructure the assembly such that shared memory is used instead of

global memory, further speed ups are possible. However, this improvement requires a substantial rewrite of the finite element method and is dependent on the particular type of element, thus the simpler, yet beneficial global memory approach was taken in this paper. Much work has also been done on finite element CG solvers for GPU, and even extensions to multi-grid solvers [25, 192, 64, 74]. However, both the finite element assembly on the GPU and conjugate gradient method utilize forms of sparse matrix vector product (SPMV) multiplications.

To avoid race conditions on the GPU, the finite element assembly procedure can be reorganized into two different SPMV operations to build the tangent matrix, \mathbf{K} , and residual vector, \mathbf{R} , for non-linear iterative finite element simulations that solve [31].

$$\mathbf{K}^{\text{elem}} \cdot \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \mathbf{K}$$

$$\mathbf{R}^{\text{elem}} \cdot \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \mathbf{R}$$

Likewise, the CG solver is then used to solve the canonical problem $\mathbf{K} \cdot \Delta \mathbf{x} = \mathbf{R}$, where $\Delta \mathbf{x}$ is the solution update for the Newton-Raphson method. At least one SPMV is performed per CG iteration step. Therefore, by improving SPMV algorithms on GPUs, we can improve different parts of the finite element method and also apply these algorithms to general SPMV linear algebra problems.

While there are many variations of sparse matrix vector product GPU algorithms [128, 25, 192], the majority can be summarized by the available implementations from several commonly available sparse matrix vector libraries: CUSPARSE [141], CUSP-library [14], and ModernGPU [10]. In [13], COO, CSR, ELL, and HYB formats

are examined and analyzed. A review of the different formats and algorithms is highlighted below. In this section, we will look at the classical SPMV problem $\mathbf{A} \cdot \mathbf{x} = \mathbf{y}$, where \mathbf{x} is given and the objective is to calculate the vector \mathbf{y} .

COO

The coordinate list (COO) sparse matrix format is composed of three array lists of row indices, column indices, and the corresponding list of matrix values. In the CUSP COO implementation, row indices are sorted in order. Each warp processes a section of the matrix and works over 32 non-zero values at a time. Each thread within the WARP performs a multiplication between the thread's value and the corresponding value from the vector and the product is store in shared memory. Intra-warp segmented reduction is performed in shared memory. The first thread within each warp determines whether to include results from a previous warp iteration in its rowsum, and if not, updates the solution vector. To handle rows that span across warps, intra-block segmented reduction is performed.

COO SPMV algorithms generally suffer from poor memory to computation ratios as row and column indices must be retrieved for each computation. Row indices are also required for row sum reduction, and additional explicit intra-block thread synchronization may be required depending on whether rows span multiple warps.

CSR

The compressed sparse row (CSR) matrix storage format is composed of three array lists of values, corresponding column indices, and corresponding row offsets that index the beginning of each row in the arrays of values. In the CUSP CSR vector implementation, a warp is assigned to each row in the matrix. A warp processes a continuous section of values and corresponding column indices in a coalesced manner. Intra-warp reduction is preformed and sums are accumulated in the solution vector.

The main issue with CSR algorithms is that while the storage is space efficient and the memory access pattern is contiguous, memory accesses are not aligned. The CUSP implementation will also suffer when the size of a row is less than the warp

size. However, unlike an alternative CSR algorithm where each thread is assigned to a row, the access pattern in CUSP is coalesced and contiguous.

ELL

In the ELLPACK format, the matrix is again described by two array lists of values, and column indices. The maximum size of longest row in the matrix is allocated for each row in the ELL format. Non-zero values are ordered contiguously, and the remaining values are typically zero-padded. Column indices are arranged and padded in a similar manner. Both values and columns are stored in column major order. The GPU kernel is fairly simple. Each thread is responsible for a row, and matrix vector products are summed in a coalesced manner.

The ELLPack format performs poorly when the row size varies from row-to-row resulting in less work being done. The format also suffers from issues with excess padding when the longest and shortest row are very different, and the average row length is significantly shorter than the longest row.

HYB

The Hybrid format is similar to the ELLPACK format. To reduce the amount of wasted padding due to row-to-row size differences, a certain number of values, determined empirically, is stored in ELLPACK, and the remaining entries are stored in COO format. The CUSP implementation is simply a combination of ELL and then the COO kernel.

Due to the more efficient storage scheme, the hybrid format requires launching of two kernels in order to perform the sum. While the values of the matrix are contiguous in each kernel, the access pattern for the vector will generally not be so.

ModernGPU

ModernGPU takes a different approach with its sparse matrix vector implementation called `mgpusparse` (MGPU). The algorithm is somewhat complicated but is essentially

constructed to address several issues with the algorithms above. It is a combination of a series of parallel and segmented scans. Data is partitioned into the list of values, list of corresponding column indices, and several lists are used to keep track of the ends of rows and row segments. Data is also arranged in a coalesced way similar to column-major storage, but for fixed-sized chunks of matrix values. Two kernels are utilized. In the first kernel, each thread within a warp is responsible for calculating `valuesPerThread` partial row sum entries. A serial scan is performed by each thread and the sums of different row segments are stored in shared memory. A number of flags is associated with each thread, such that each thread can determine where to store the partial sum for a given row segment in shared memory. An intra-warp parallel scan is then performed to reduce rows within each chunk quickly. Lastly an intra-block segmented reduction is performed to reduce rows the span multiple blocks.

MGPU also takes advantage of index representation compression and performs other compression tricks to increase efficiency. The algorithm generally performs well in comparison to the algorithms above. The main draw-back, however, is the complexity of the algorithm and the amount of book-keeping necessary. The main benefits are that the workload for threads within a block are very well distributed and very few threads will be idle compared to the rest within each block.

11.5 Development of novel GPU SPMV algorithms

There are several other variations of the algorithms above where most of them concentrate on padding efficiency, attaining less warp-divergence, and increasing the workload for a given block of threads [192, 25]. However, there are several key insights that have been made in ModernGPU, ELLPACK, and CSR vector implementations that may yield a better SPMV algorithm for finite elements. In ELL and MGPU algorithms, "column-major" coalesced data interleaving is leveraged to allow threads within a warp to work on different rows and row-segments in an efficient coalesced memory access pattern. In MGPU and in CSR vector, intra-warp reductions are utilized to avoid unnecessary thread synchronization which allow for warps and blocks to be executed without delays due to synchronization.

The variations in row length in ELL are acquiesced in the HYB matrix implementation, and are accounted for by segmented scan flags in `mgp sparse`. Both algorithms try to address the padding inefficiency in the simple ELL algorithm. In [13], it is noted that the impact of padding on well-behaved structured meshes should not be substantial as the row size should not vary excessively. In unstructured finite element meshes of solids, this appears to be partially true. Since the node connectivity of the mesh directly translates into the finite element tangent matrix, and element quality is usually controlled by meshing algorithms, the number of non-zeros per row is indirectly controlled by mesh quality algorithms. Therefore, the row lengths in a given unstructured mesh are usually constrained; however, there is no guarantee that there is row-to-row uniformity within the matrix for a given unstructured finite element mesh. In the case of assembly, the \mathbf{x} vector is simply a vector of ones, and one can simply avoid the multiplication of unity to the corresponding matrix value altogether. From Figure 11.3, it is clear that while row sizes do not vary greatly from

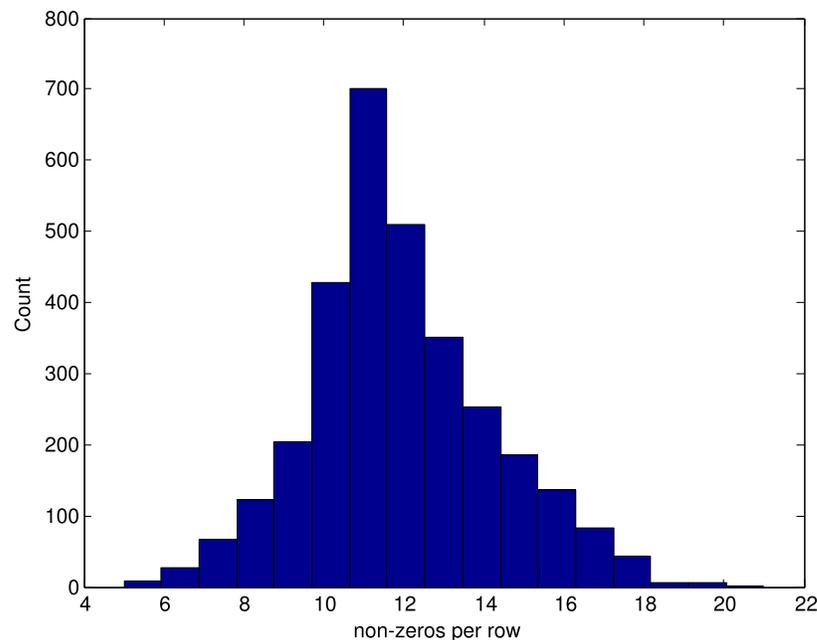


Figure 11.3: A histogram of the non-zeros per row in a small 3129 row stiffness matrix for a patient-specific simulation of the heart.

5 to 21 non-zeros per row, the actual distribution of non-zeros per row varies quite a bit from the median and mean row lengths overall.

Since MGPU out performs the simple kernel implementations of ELLPACK/HYB, this seems to indicate that there may be inefficiencies related to padding which are not an issue for the more complicated equal work distributed MGPU implementation. Therefore, we propose the following set of simple kernel algorithms that directly follow from the insights listed above, which will hopefully perform well compared to MGPU and HYB implementations.

11.5.1 WarpKernel1 (WPK1)

The proposed kernel algorithm is based on the ELL. To mitigate the observation of row lengths distribution irregularity, we first sort the rows by length from longest to shortest (Figure 11.4). A thread per row execution scheme is used, however,

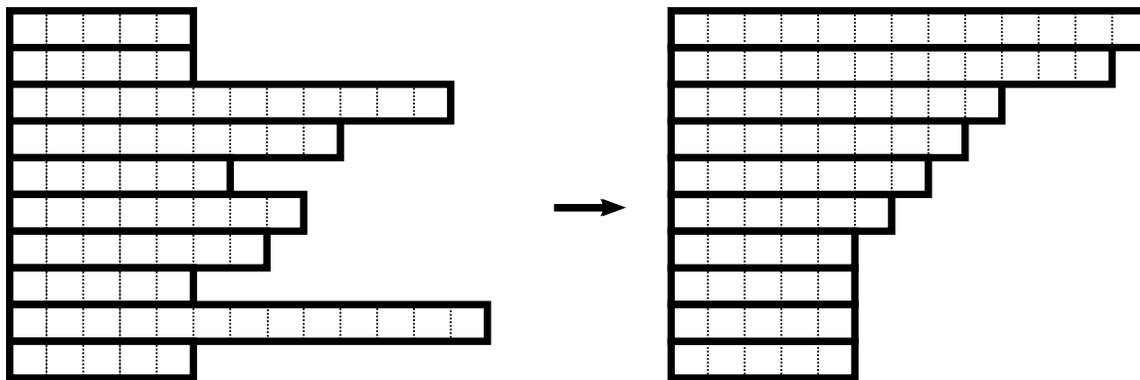


Figure 11.4: Sort rows of the matrix from longest to shortest.

instead of using the representation used in the ELL, where padding continues until the maximum matrix row length is reached, each warp is now padded up to the maximum row length of given warp (Figure 11.5). To ensure coalesced memory access, we permute the data within each warp in a "column-major" ordering similar to MGPU and ELL/HYB (Figure 11.6). Lastly, the results are mapped back to the original row-ordering when writing results to the solution vector y . This algorithm will be referred to as `warpkernel1` (WPK1). The code is shown in Listing 11.5.1.

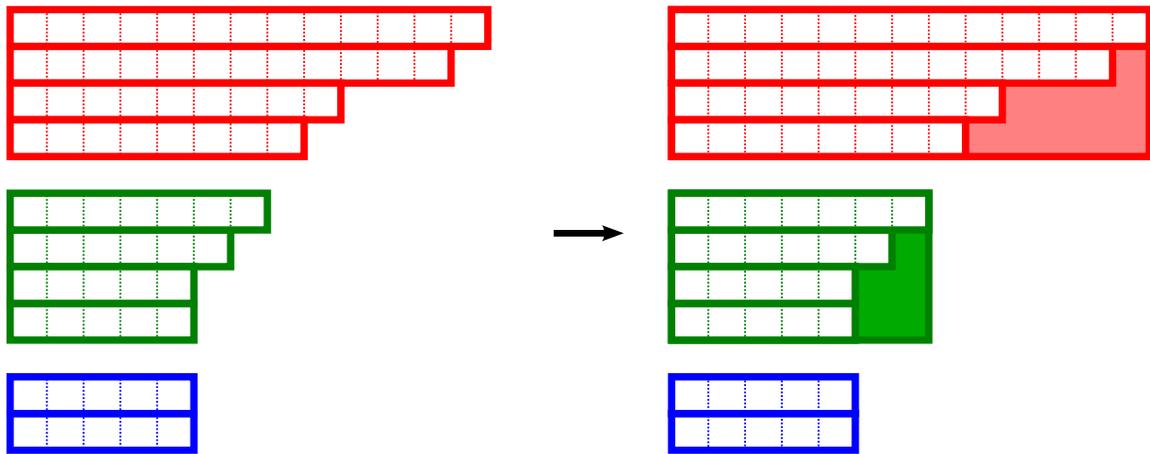


Figure 11.5: Arrange rows into groups of warp size and then pad accordingly. In this figure, we use a warp size = 4 only for purposes of illustration.

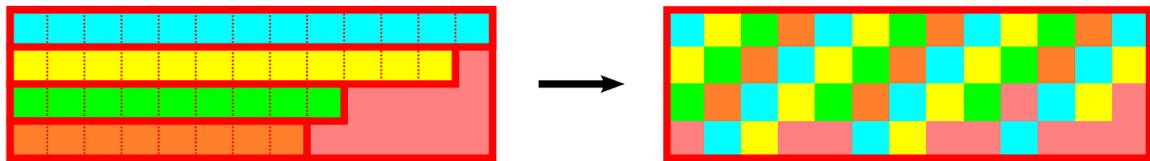


Figure 11.6: Reorder the first warp in Figure 11.5 in a coalesced column-major ordering.

Warpkernel1 requires a list of values and their column indices padded per warp, which are sorted by row length and arranged in "column-major" order. A list of the initial offsets for each warp is given, as well as a row mapping from sorted row index to the original row index.

The purpose of this arrangement is to exploit the relatively small variances in row lengths per warp to reduce unnecessary padding in ELL and also to increase the amount of threads performing meaningful operations. As values and columns indices are zero-padded, those values and column indices are cached and should not hurt the efficiency of this algorithm if the variances within each warp are in fact small. By having a fixed row length for each warp, one can reduce the amount of data needed to compute row sums in comparison to CSR and COO, while reducing the number of trivial instructions executed for padded entries in comparison to ELL. Thus, this kernel aims to achieve a similar equally distributed workload to MGPU in cases where

Listing 11.1: Code for warpkernel1.

```
1  template <uint WARP_SIZE, bool usecache>
2  __global__ void warpKernel1(ValueType* A,
3                               IndexType *colinds,
4                               IndexType *rowmap,
5                               uint* maxrows,
6                               ValueType* x,
7                               IndexType *warp_offset,
8                               ValueType* y,
9                               IndexType nrows) {
10
11  const uint tid = threadIdx.x;
12  const uint id = tid + blockIdx.x * blockDim.x;
13  const uint wid = tid & (WARP_SIZE-1);
14  const uint warpid = id / WARP_SIZE;
15
16  if (id < nrows) {
17
18      IndexType toffset = warp_offset[warpid] + wid;
19      uint maxnz = maxrows[warpid] * WARP_SIZE + toffset;
20
21      // Perform sequential sum
22      ValueType sum = A[toffset] * cache<usecache> (colinds[toffset],x);
23      for(toffset += WARP_SIZE; toffset < maxnz; toffset += WARP_SIZE) {
24          sum += A[toffset] * cache<usecache> (colinds[toffset],x);
25      }
26
27      // Store remapped result
28      y[rowmap[id]] = sum;
29  }
30 }
```

there is some degree of row length regularity within each warp. Unfortunately, the final non-coalesced memory write to the solution vector is a potential inefficiency for this kernel which is otherwise coalesced in terms of memory access to values and column indices.

11.5.2 Warpkernel1 - reordered (WPK1r)

To mitigate the necessity of performing a non-coalesced write within our kernel, the column indices of the matrix and the vector \mathbf{x} can be reordered, such that the numbering of corresponding entries in \mathbf{x} and the solution vector are consistent. This is commonly performed in finite element computations by renumbering the unknowns during the setup phase of most codes. Reordering the columns of the matrix and entries of \mathbf{x} does not change the result of the solution vector.

$$A_{i,j}x_j = y_i \quad A_{i,f(j)}x_{f(j)} = y_i \quad (11.7)$$

By performing this reordering, we can essentially ignore the original row mapping necessary in `warpkernel1` as the sorted solution vector is now numbered consistently with the newly arranged \mathbf{x}' vector. Unfortunately, reordering of the column indices, now produces coalesced non-ordered access from the \mathbf{x}' vector. While the effects may be mitigated by using a texture cache, this is somewhat undesirable.

To further illustrate the reordering scheme, we look a CSR ordered matrix for convenience. If x^{num} is the original ordering for \mathbf{x} , then x^{renum} is the new ordering such that the rows will be sorted accordingly from longest to shortest.

$$\begin{array}{c|c} x^{\text{num}} & \left[\begin{array}{ccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{array} \right] \\ \hline x^{\text{renum}} & \left[\begin{array}{ccccccc} 5 & 1 & 4 & 6 & 2 & 7 & 3 \end{array} \right] \end{array}$$

Given the following one-row CSR matrix \mathbf{A} , a dense vector of column indices \mathbf{c} , and a dense vector \mathbf{x}

$$\mathbf{A} = \left[\begin{array}{cccccc} 7 & 8 & 9 & 10 & 2 \end{array} \right]$$

$$\mathbf{c} = [1 \ 2 \ 4 \ 5 \ 6]$$

$$\mathbf{x} = [1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7]$$

the WPK1r scheme described above will re-arrange the data such that the following results.

$$\mathbf{c}' = [5 \ 1 \ 6 \ 2 \ 7]$$

$$\mathbf{x}' = [2 \ 5 \ 7 \ 3 \ 1 \ 4 \ 6]$$

This reordering scheme does not require that \mathbf{A} be ordered, but results in non-ordered column indices, \mathbf{c}' , which is possibly detrimental to maintaining ordered access from \mathbf{x} .

11.5.3 Warpkernel1 - reordered sorted (WPK1rs)

The following is a variation on WPK1r that attempts to fix the issue with non-ordered access from the \mathbf{x} vector. The solution is to reorder the values of the matrix in each row, such that the reordering of the column indices is sorted in order, after the reordering of the vector \mathbf{x} .

Again using the previous one-row CSR matrix as an example, we sort \mathbf{c}' such that it is ordered from smallest to largest indices within each row.

$$\mathbf{c}'' = [1 \ 2 \ 5 \ 6 \ 7]$$

$$\mathbf{A}'' = [8 \ 10 \ 7 \ 9 \ 2]$$

This \mathbf{x} remains the same; however, now it is necessary to sort the values of \mathbf{A} within each row to properly account for the changes in \mathbf{c}' . The result is an intra-row reordered CSR matrix, \mathbf{A}'' , and an ordered set of corresponding column-indices, \mathbf{c}'' .

11.5.4 Warpkernel2 (WPK2)

To accommodate the possibility of larger intra-warp row length variances, a multiple threads per row variation of warpkernel1 is proposed. For example, if there exist a small number of rows that are abnormally long compared to the rest of the rows in the matrix, it may be advantageous to process those rows with more threads, while allowing the rest of the rows in the matrix to be processed by WPK1.

A threshold value for the maximum number of entries a thread should be given is prescribed. If a row does not meet this criterion, the row is subdivided in half recursively until it meets this requirement. However, to avoid explicit thread synchronization within a block, if a row is longer than $32 \times \text{threshold}$, one warp will at the least process the entire row to maintain warp-to-warp row independence. As seen in Figure 11.7, this variation inevitably increases the amount of padding in comparison to WPK1 depending on the threshold value chosen. However, it allows unusually long rows to be processed more efficiently, while smaller rows retain the same padding as WPK1. The threshold parameter forces almost all warps to basically process the same number of non-zero entries. The kernel now includes a simple parallel reduction loop which reduces rows in a given warp efficiently using shared memory. This is seen in Figure 11.8 for the first row of the representative warp in our illustrative example.

As rows are sorted initially, it is relatively easy to reprocess the original warp-sized row groupings from WPK1 and subdivide the numbers of rows per warp as needed to meet the threshold criterion. Extra information per warp is necessary to dictate the row number per warp and also the number of reductions needed per warp. However, this extra indexing information is minimal and results in two extra integers retrieved from global memory per warp.

11.5.5 Warpkernel2 variations (WPK2r) and (WPK2rs)

Lastly, to avoid the final row mapping from sorted to non-sorted rows, the same variations to WPK1 can be simply applied directly to WPK2 as the main difference between WPK2 and WPK1 is how data is repackaged into blocks of data that are

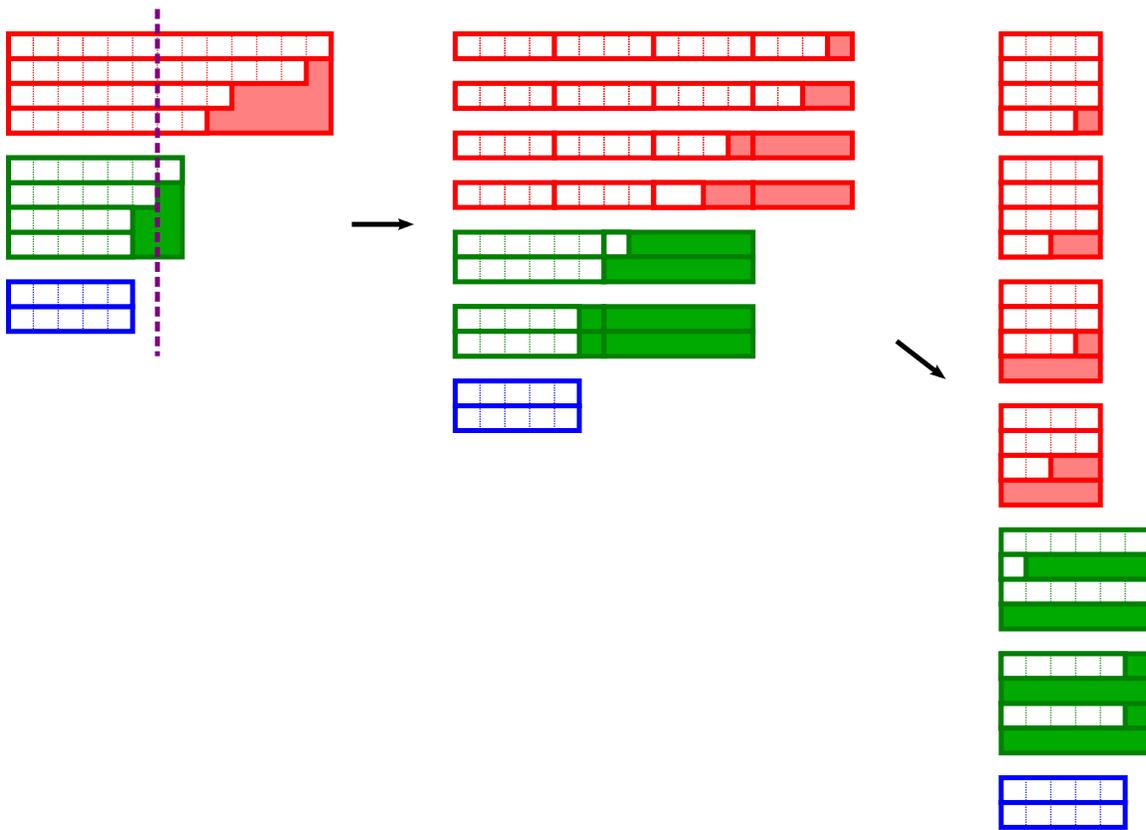


Figure 11.7: A threshold value (purple dashed line) is given and the original padding in WPK1 is padded and grouped into new warps.

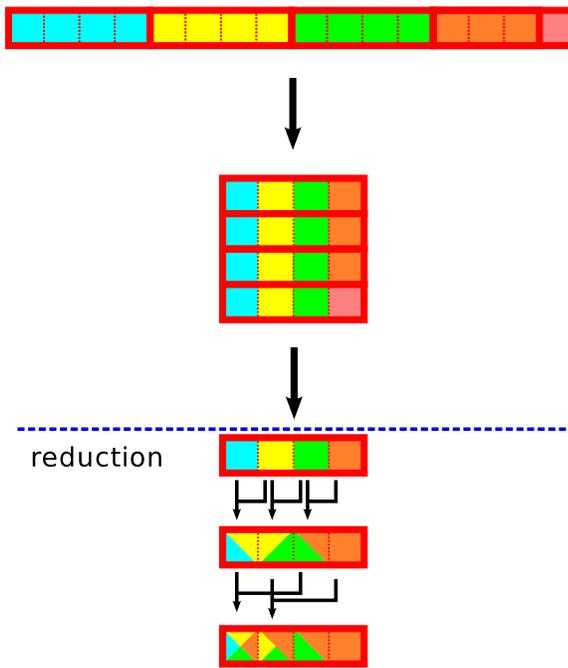


Figure 11.8: Coloring represents the data for a particular thread. In this instance, 4 threads are used to process the representative row. Data is arranged in column-major order, and then a parallel reduction step occurs resulting with the final sum in the first entry of the output row.

Listing 11.2: Code for warpkernel2.

```

1  template <uint WARP_SIZE, bool usecache>
2  __global__
3  void ellwarpKernel2(ValueType* A, IndexType *colinds,
4                    IndexType *rowmap, uint* maxrows,
5                    ValueType* x, IndexType *warp_offset,
6                    ValueType* y, IndexType nrows,
7                    uint* reduction, uint* rows_offset_warp,
8                    int nwarps) {
9
10     const uint tid = threadIdx.x;
11     const uint id = tid + blockIdx.x * blockDim.x;
12     const uint wid = tid & (WARP_SIZE-1);
13     const uint warpid = id / WARP_SIZE;
14
15     extern volatile __shared__ ValueType sumvalues[];
16
17     if (warpid >= nwarps) return;
18     const uint offsets = reduction[warpid];
19     const uint row_start = rows_offset_warp[warpid];
20     const uint rowid = row_start + wid/offsets;
21
22     if (rowid < nrows) {
23
24         IndexType toffset = warp_offset[warpid] + wid;
25         const uint maxnz = maxrows[warpid] * WARP_SIZE + toffset;
26         ValueType sum = A[toffset] * cache<usecache> (colinds[toffset],x);
27
28         for(toffset += WARP_SIZE; toffset<maxnz; toffset += WARP_SIZE) {
29             sum += A[toffset] * cache<usecache> (colinds[toffset],x);
30         }
31         sumvalues[tid] = sum;
32
33         for (int i = 1; i< offsets; i <<= 1) {
34             if (offsets > i) {
35                 sum += sumvalues[tid+i];
36                 sumvalues[tid] = sum;
37             }
38         }
39
40         if ((wid & (offsets-1)) == 0) {
41             y[rowmap[rowid]] = sum;
42         }
43     }
44 }

```

more regular across blocks rather than simply by rows of most similar length.

11.6 Benchmarking sparse-matrix vector multiplication methods

Our novel algorithms above are compared against several available standard SPMV algorithms using a set of sparse matrices commonly used in benchmarking SPMV methods [13]. Since the WPK1 and WPK2 kernels were developed for finite element simulations of the heart, we have also included 3 different refinements of a patient-specific heart mesh as part of the set of benchmark matrices. General information about each of the benchmark matrices is included in Table 11.2. Matrix structure

Table 11.2: Matrix Benchmark Information. The number of nonzeros (nz), the number of rows (nrows), bytes, minimum row length (minrow), maximum row length (maxrow), and average number of non-zero values per row (nz/nrows) is reported for each benchmark matrix.

MatrixName	nz	nrows	bytes	minrow	maxrow	nz/nrows
Circuit	958,936	170,998	19,178,720	1	353	6
Economics	1,273,389	206,500	25,467,780	1	44	7
Epidemiology	2,100,225	525,825	42,004,500	2	4	4
FEMAccelerator	2,624,331	121,192	52,486,620	8	81	22
FEMCantilever	4,007,383	62,451	80,147,660	1	78	65
FEMHarbor	2,374,001	46,835	47,480,020	4	145	51
FEMShip	7,813,404	140,874	156,268,080	24	102	56
FEMSpheres	6,010,480	83,334	120,209,600	1	81	73
Heart	37,035	3,129	740,700	5	21	12
HeartCoarse	52,715	4,563	1,054,300	6	22	12
HeartRefined	367,443	28,639	7,348,860	6	24	13
Protein	4,344,765	36,417	86,895,300	18	204	120
QCD	1,916,928	49,152	38,338,560	39	39	39
Webbase	3,105,536	1,000,005	62,110,720	1	4700	4
WindTunnel	11,634,424	217,918	232,688,480	2	180	54

properties vary substantially from matrix to matrix. It is apparent that certain matrices, such as Circuit and Webbase, have a very skewed non-zero row length distribution. On the other hand, some matrices such as Epidemiology and QCD, have very regular row length distributions.

The following kernels were then chosen for comparison: CUSP-CSR and CUSP-HYB [14], CUSPARSE-CSR [141], MGPU [10], WPK1/r/rs, and WPK2/r/rs. The published MGPU results report the kernel time over an average of 1200 iterations. Thus this averaging scheme is one of the averaging methods used for computing kernel times in our benchmarks. For completeness, we investigate whether averaged kernel times are greatly affected by the number of iterations for different kernels, and we will determine whether or not 1200 iterations is a good model for finite element SPMV operations.

For each matrix used in the benchmark, kernel times are averaged over a specified number of iteration runs. For the WPK1/r/rs kernels, different block sizes are varied from 1 to 8 blocks. Likewise the same is done for the WPK2/r/rs kernels except , in addition, different maximum non-zero thresholds are varied from the minimum row length of the matrix to the maximum row length of the matrix. Lastly for MGPU, the following number of values per thread are used : 4, 6, 8, 10, 12, and 16. After the parameters for WPK1, WPK2, and MGPU are acquired, the fastest averaged times are found and reported in Figures 11.9, 11.10, and 11.11.

All benchmarks were run on a single Asus ENGTX480 graphics card (CUDA Compute compatibility 2.0) and on a PC with an Intel I7 950 with 12 GB of memory. Kernels are compiled with optimizations enabled and kernels use the texture cache when possible. The results for 1, 50, and 1200 iterations are shown in Figures 11.9, 11.10, and 11.11 respectively.

Overall the results are fairly consistent. MGPU and WPK1/2 perform well over these benchmark sparse matrices and one of the two is the best for each benchmarked matrix. There are some slight differences between 1, 50, and 1200 iterations. Initially, MGPU is slightly slower than WPK2 for 1 iteration for the Circuit case, but is slightly faster when averaged over 50 and 1200 iterations. For the Epidemiology benchmark, MGPU becomes faster than CUSP-HYB at 1200 iterations. MGPU also

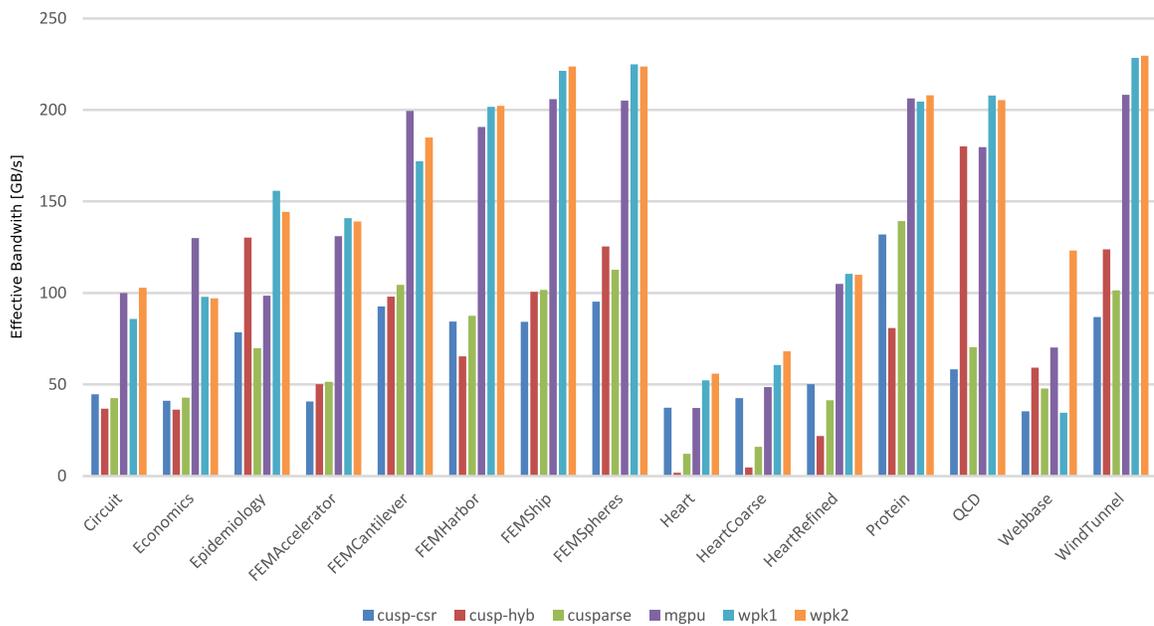


Figure 11.9: Benchmark results over matrices for 1 iteration.

slightly outperforms WPK1/2 only at 1200 iterations for FEMAccelerator. In the HeartCoarse benchmark, MGPU becomes slightly faster than WPK1 at 50 iterations, but never reaches the performance of WPK2 at 50 or 1200 iterations; however in HeartRefined, MGPU performs slightly better than WPK1/2 starting at 50 iterations.

In general, WPK1 and WPK2 perform the best for all 15 matrices for 1, 50, and 1200 iterations. For 1 iteration, a WPK1/2 kernel is fastest for 13 matrices, while at 50 iterations, a WPK1/2 kernel is only fastest for 10 matrices. At 1200 iterations, a MGPU kernel is fastest for a total of 6 matrices: Circuit, Economics, FEMAccelerator, FEMCantilever, HeartRefined, and Protein. On the otherhand, WPK1/2 kernel is fastest for a total of 9 matrices: Epidemiology, FEMHarbor, FEMShip, FEMSpheres, Heart, HeartCoarse, QCD, Webbase, and WindTunnel.

Next the parameters for MGPU, WPK1 and WPK2 are shown for each of the best performing kernels for each matrix in Tables 11.3, 11.4, and 11.5. In Table 11.4, WPK1 parameters rarely change between 1, 50, and 1200 iterations. WPK2 parameters are fairly constant, although a few matrices have threshold and block

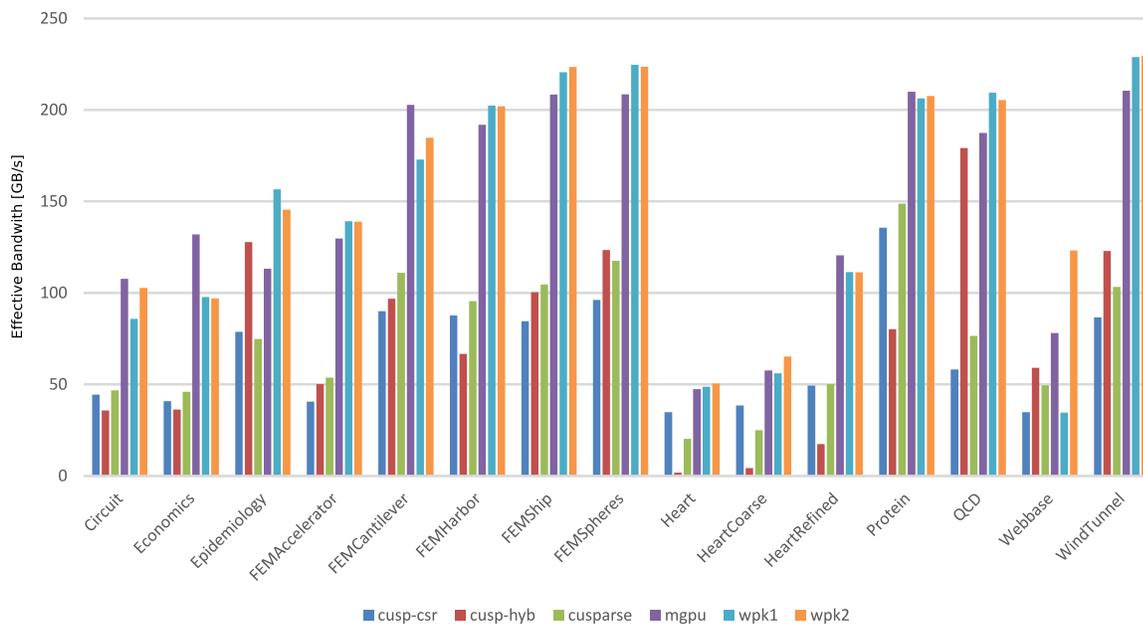


Figure 11.10: Benchmark results over matrices for 50 iteration.

parameters that vary between 1, 50, and 1200 iterations (Table 11.5). However, overall the results are fairly consistent across the range of iterations. The best performing MGPU kernels vary for several matrices also, but are constant for others.

Next, the affect on effective bandwidth for each best performing kernel type is examined over the range of 1, 50, and 1200 iterations. The benchmark results from Figures 11.9, 11.10, and 11.11 are reorganized by each kernel over different iterations. These results are shown in Figures 11.12 and 11.13. In general, CUSP-CSR, CUSP-HYB, WPK1, and WPK2 perform consistently over 1, 50, and 1200 iterations. Conversely, both CUSPARSE-CSR and MGPU are faster as the number of iterations is increased.

11.7 Performance Analysis

In this section, we evaluate different aspects of performance. We first investigate whether different kernel level optimizations and the variations to WPK1 and WPK2

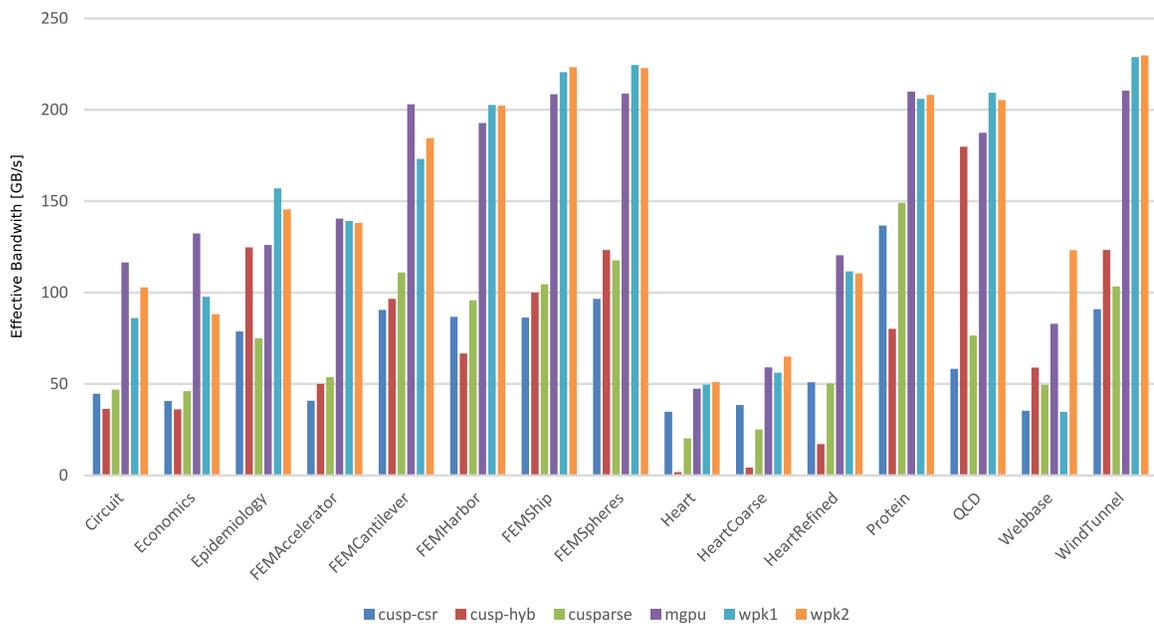


Figure 11.11: Benchmark results over matrices for 1200 iteration.

provide an improvement in performance. The cost-effectiveness of GPU data reordering is then compared and evaluated against CPU data reordering implementations. Lastly, the performance of our fully-GPU finite element and hybrid CPU assembly - GPU solver implementations are compared against a PETSc multi-core CPU implementation using 4 different mesh refinements for our nonlinear-coupled finite element electrical heart simulation of one complete cardiac cycle.

11.7.1 Effects of optimizations taken

A subset of the benchmark matrices were chosen to determine the effects of different kernel optimizations taken in the development of WPK1/2 and their variants. The matrices chosen were Circuit, Epidemiology, FEMHarbor, Heart, HeartCoarse, HeartRefined, and QCD. The following factors were measured: coalesced vs. non-coalesced memory access patterns, sorted vs. unsorted rows, kernels that involve remapping vs. those that renumber column indices and reorder the vector \mathbf{x} (WPK vs. WPKr), and lastly the differences between two possible reordering schemes for

Table 11.3: Benchmark best MGPU valuesPerThread parameter.

MatrixName	mgpu [1] best	mgpu [50] best	mgpu [1200] best
Circuit	6	6	4
Economics	6	6	6
Epidemiology	6	4	4
FEMAccelerator	6	6	4
FEMCantilever	10	10	10
FEMHarbor	10	12	16
FEMShip	10	10	10
FEMSpheres	10	12	12
Heart	12	10	10
HeartCoarse	10	8	8
HeartRefined	10	10	6
Protein	10	16	16
QCD	10	10	10
Webbase	6	4	4
WindTunnel	10	10	10

the vector \mathbf{x} (WPKr vs. WPKrs).

First, we investigate the importance of coalesced memory access in Figure 11.14. In this test, we compared our WPK1 and WPK2 kernel against the same kernel where data was left in the standard CSR "row-major" form instead of a "column-major" ordering that results in coalesced memory access within each warp. As expected, coalesced memory access patterns result in a dramatic reduction in computation time in comparison to non-coalesced memory access patterns. In fact, a coalesced memory access pattern accounts for 1.5 to 9 fold increase in speed when using texture memory access, and for 1.5 to 10 fold increase in speed without textures.

We also tested the effect of sorting matrix rows from longest to shortest (WPK1) and compared it to the equivalent algorithm where the ordering of rows of the original matrix is preserved. As the non-sorted implementation must result in more zero-padding, the table below (Table 11.6) reports the difference in padding, along with the percentage of time-difference with respect to a non-row ordered algorithm. The effect of sorting on performance is not definitively clear. For the majority of matrices,

Table 11.4: Benchmark best WPK1 blocksize parameter.

MatrixName	wpk1 [1] best	wpk1 [50] best	wpk1 [1200] best
Circuit	64	64	64
Economics	64	64	64
Epidemiology	128	128	128
FEMAccelerator	64	64	64
FEMCantilever	64	64	64
FEMHarbor	96	96	96
FEMShip	96	96	96
FEMSpheres	96	96	96
Heart	32	32	32
HeartCoarse	64	32	64
HeartRefined	64	64	64
Protein	64	96	96
QCD	224	224	224
Webbase	32	32	32
WindTunnel	96	96	96

sorting the matrices resulted in an increase in performance when using the texture memory cache; however sorting may have detrimental effects on other matrices as shown for Epidemiology even when using the texture cache. On the other hand, without texture cache access, non-ordered rows result in significantly faster kernel times. Except for the Epidemiology matrix, the texture cache helps significantly, and row-sorted texture cache kernel times are faster than non-sorted kernel times without texture cache. However, from Table 11.6, it is evident that Epidemiology and QCD have very similar row length regularity, and therefore the effects of sorting may actually ruin the better access pattern for \mathbf{x} . Texture access is therefore not as helpful in those cases. Overall, sorting the rows greatly reduces the zero-padding necessary, which is important in terms of applying SPMV operations in real situations.

Next we examine the benefits of avoiding the final solution reordering necessary in WPK1 and compare it to a pre-reordered \mathbf{x} and pre-renumbered \mathbf{c} . The results are shown in Figure 11.15 and 11.16. The first reordered \mathbf{x} kernel, WPK1r, outperforms WPK1 with texture cache access. Without the use of the texture cache, WPK1r out

Table 11.5: Benchmark WPK2 best threshold and blocksize parameters.

MatrixName	wpk2 [1]		wpk2 [50]		wpk2 [1200]	
	threshold	blocksize	threshold	blocksize	threshold	blocksize
Circuit	80	160	38	128	112	160
Economics	7	96	7	96	7	96
Epidemiology	4	128	4	128	4	128
FEMAccelerator	58	64	57	64	58	64
FEMCantilever	19	224	19	224	19	224
FEMHarbor	128	96	132	96	139	96
FEMShip	28	128	27	128	27	128
FEMSpheres	47	128	81	96	81	96
Heart	7	256	7	256	6	192
HeartCoarse	10	192	13	96	6	192
HeartRefined	21	64	17	64	17	64
Protein	172	96	167	96	171	96
QCD	39	96	39	224	39	128
Webbase	658	128	682	128	672	128
WindTunnel	38	96	37	96	38	96

Table 11.6: Benchmark data regarding the effects of sorting for WPK1. Padding difference designates the extra padding that results when rows are not sorted. The time difference percentage is defined as the WPK1 kernel time subtracted from the non-sorted kernel time divided by the non-sorted kernel time.

	Circuit	Epidemiology	FEMHarbor	Heart	HeartCoarse	HeartRefined	QCD
Padding Difference	1443456	3040	1206432	15712	27008	157824	0
Time difference (textured)	21.16%	-2.10%	0.00%	3.00%	22.67%	5.79%	0.47%
Time difference	-17.38%	-4.35%	0.00%	-28.22%	10.42%	-46.80%	0.49%

performws WPK1, except in the case of the HeartRefined mesh. Likewise, with the texture cache enabled, WPK2r out performs WPK2, but is sometimes slower without texture cache access. The effect of reordering \mathbf{x} is generally beneficial, but is similar to the effect of sorting by row length where texture cache access for the well ordered matrices such as Epidemiology and QCD result in only slight differences between texture and non-texture cache results.

To determine the cost of the final solution remapping, the WPK1r kernel was modified to perform the final solution remapping that results in a solution \mathbf{y} that matches the original numbering for \mathbf{x} and the results are reported in Figure 11.17.

In all cases, the cost of remapping is more than WPK1r. For WPK1r with use of the texture cache, final solution remapping costs between 1% to 11.5% of the computation time of WPK1r, whereas without the texture cache, remapping costs between 0.8% and 8% of the computation time of WPK1r. Similar results were found for WPK2r.

Next, we examine the improvement of WPK1rs over WPK1r and of WPK2rs over WPK2r in Figures 11.18 and 11.19. Other than the QCD and Epidemiology matrix cases which are already very well ordered, the values and row sorted versions of both kernels are marginally faster. WPK1rs is faster than WPK1r by 1.75% to 7.25% with the texture cache, and by 4% to 23.5% without the texture cache for matrices other than QCD and Epidemiology. WPK2rs is faster than WPK1r by 1% to 7.5% with the texture cache and 0.25% to 21% without the texture cache for matrices other than QCD and Epidemiology. WPK1rs and WPK2rs are slower than their WPKr counterparts for the Epidemiology case by less than 0.25% with and without the texture cache. For the QCD case, the row-sorted reordered kernels are slower by 0.25% to 3.25%.

Lastly, we look at the difference between using a register to buffer the shared memory reduction in WPK2/r/rs kernels in Figure 11.20. The difference between use of the register and without is generally by about 1% of the computation time with use of the register. However, in certain cases, use of the register is slower by a factor of 3.6% with texture cache access, and is faster by a factor of 4.7% without texture cache access. Overall, the differences in parallel reduction using a register as a buffer is small.

Cost of reordering matrix values

To determine whether it is beneficial overall to reorder the data in a "column-major" coalesced pattern in SPMV applications for finite elements, we consider the following. If WPK1 and WPK2 are used to calculate sparse matrix vector multiplications in the conjugate gradient method, only one initial transpose of values is necessary at the beginning of each Newton-Raphson iteration with the assumption that the assembler passes a CSR formatted matrix to the solver. Column indices do not need to be

reordered, as we assume the connectivity of the Langrangian mesh does not change during the simulation; therefore, only the values of the tangent matrix change, while the matrix structure remains constant. We can then determine the number of CG iterations necessary, such that WPK1 and WPK2 will out perform CUSPARSE by the following

$$t_{\text{reorder}} + \alpha t_{\text{wpk}} \leq \alpha t_{\text{cusparse}} \quad (11.8)$$

where α is the number of iterations necessary such that a reordering of data from CSR to "column-major" provides a benefit for the CG solver. The following data is used to compare the differences between reordering on the GPU and on the CPU. As our algorithm produces a reordering mapping for all non-zero entries during the initial scan, scatter operations are simply performed on the CPU using a simple for-loop, and by using `thrust::scatter()` on the GPU [82]. From Tables 11.7 and 11.8,

Table 11.7: Comparison of reordering times and WPK1 and cusparse kernel SPMV times.

MatrixName	GPU [s]	CPU [s]	WPK1 [s]	cusparse [s]	α_{GPU}	α_{CPU}
Circuit	4.60E-04	2.60E-03	2.08E-04	3.82E-04	3	15
Economics	7.43E-04	3.27E-03	2.43E-04	5.15E-04	3	12
Epidemiology	4.43E-04	5.30E-03	2.50E-04	5.23E-04	2	20
FEMAccelerator	1.62E-03	6.54E-03	3.51E-04	9.11E-04	3	12
FEMCantilever	2.82E-03	1.18E-02	4.32E-04	6.73E-04	12	50
FEMHarbor	1.60E-03	6.55E-03	2.18E-04	4.63E-04	7	27
FEMShip	5.58E-03	1.93E-02	6.60E-04	1.39E-03	8	27
FEMSpheres	4.18E-03	1.52E-02	4.98E-04	9.53E-04	10	34
Heart	1.31E-05	7.00E-05	1.42E-05	3.40E-05	1	4
HeartCoarse	1.68E-05	9.96E-05	1.75E-05	3.92E-05	1	5
HeartRefined	1.66E-04	7.40E-04	6.14E-05	1.37E-04	3	10
Protein	3.16E-03	1.45E-02	3.92E-04	5.44E-04	21	96
QCD	1.01E-03	4.78E-03	1.70E-04	4.67E-04	4	17
Webbase	1.22E-03	8.71E-03	1.67E-03	1.17E-03	∞	∞
WindTunnel	8.49E-03	2.93E-02	9.47E-04	2.10E-03	8	26

it is fairly clear that CPU reordering times are substantially slower than reordering directly on the GPU even for small matrices (Heart, HeartCoarse, etc...). However,

Table 11.8: Comparison of reordering times and WPK2 and cusparse kernel SPMV times.

MatrixName	GPU [s]	CPU [s]	WPK2 [s]	cusparse [s]	α_{GPU}	α_{CPU}
Circuit	4.57E-04	2.44E-03	1.74E-04	3.82E-04	3	12
Economics	5.93E-04	3.41E-03	2.44E-04	5.15E-04	3	13
Epidemiology	4.52E-04	5.32E-03	2.69E-04	5.23E-04	2	21
FEMAccelerator	1.67E-03	6.54E-03	3.52E-04	9.11E-04	3	12
FEMCantilever	2.36E-03	1.05E-02	4.04E-04	6.73E-04	9	39
FEMHarbor	1.62E-03	6.65E-03	2.19E-04	4.63E-04	7	28
FEMShip	5.25E-03	1.78E-02	6.51E-04	1.39E-03	8	24
FEMSpheres	4.11E-03	1.53E-02	5.01E-04	9.53E-04	10	34
Heart	1.10E-05	6.85E-05	1.37E-05	3.40E-05	1	4
HeartCoarse	1.57E-05	9.80E-05	1.51E-05	3.92E-05	1	5
HeartRefined	1.69E-04	7.30E-04	6.15E-05	1.37E-04	3	10
Protein	3.17E-03	1.39E-02	3.90E-04	5.44E-04	21	91
QCD	1.05E-03	4.89E-03	1.74E-04	4.67E-04	4	17
Webbase	1.21E-03	8.36E-03	4.69E-04	1.17E-03	2	12
WindTunnel	8.41E-03	2.68E-02	9.44E-04	2.10E-03	8	24

even when reordering on the CPU at every Newton-Raphson iteration for the finite element benchmark matrices, α_{CPU} is within the general number of iterations for most CG problems. On the other hand, GPU reordering is very fast, and benefits should be noticeable within 10 iterations for both WPK1 and WPK2 on average. Thus, for the remaining number of iterations, the benefit of WPK1 and WPK2 over other kernels will be evident.

From Tables 11.7 and 11.8, one initially may discount the benefit of using any reordered kernels for global finite element tangent assembly. However, the ordering for finite element assembly can be arranged, such that the resulting ordering is already "column-major" ordered. Likewise the result of the assembly of the tangent matrix can also be arranged such that results are already coalesced and ordered properly; thus, bypassing the need for reordering the matrix \mathbf{A} altogether.

11.7.2 Finite Element Comparison between GPU and multi-core PETSc

In the final set of results, we compare the performance of our finite element GPU finite element framework against a multi-core PETSc implementation of the same code. For the GPU implementation, we use the best parameters found for WPK1 and WPK2 from Tables 11.4 and 11.5 for 50 iterations. For MGPU, we simply have chosen `valuesPerThread=6` and 10 for reference. We look at 4 different successive refinements of the Heart mesh. The benchmark problem results in the same parameters for the 50k mesh as found for HeartRefined for WPK1 and WPK2 SPMV kernels. For the multi-core PETSc implementation, we used a jacobi preconditioned CG solver when possible, as there were convergence issues with the simple jacobi preconditioner for the 50k mesh. However, for the purposes of comparison, we also show results of our PETSc implementation using the block jacobi preconditioner for all 4 meshes. In these results, we use a CPU based reordering for MGPU, WPK1, and WPK2. On the CPU, the following results are reported for the PETSc jacobi and block preconditioners at 1, 2, 3, and 4 cores.

Several general observations can be made. First, the multi-core CPU PETSc implementation does not scale linearly. The block Jacobi preconditioner CG solver is slightly faster than the Jacobi preconditioner counterpart, but at each multi-core setting, they are roughly the same. The GPU CG solver with CPU finite element assembly implementation is marginally slower than the triple-core PETSc implementation. The GPU finite element implementations are at least twice as fast compared to the quad-core PETSc implementations.

The resulting speedup for the different tests with a single-core PETSc block-jacobi solver with standard finite element assembly as reference is shown in Figure 11.22. Overall, the GPU CG solver implementations with finite element assembly start with from a two-fold speedup at the lowest 3129-node mesh and increase to a factor of 2.5 for the 50,000 node heart mesh. CUSP-CSR and CUSPARSE-CSR seem to perform better in comparison to WPK1 and WPK2 as the number of nodes increases. MGPU also performs well, but is the poorest performing kernel in all

cases. On the other hand, the fully-GPU finite element implementation for WPK1/2 ranges from a speedup factor of 5 to 12 as the mesh is subsequently refined. Again, CUSPARSE-CSR and CUSP perform better relative to MGPU as the mesh size increases; however, WPK1 and WPK2 kernels are initially slower than CUSP-CSR, but at larger mesh sizes 30k and 50k, WPK1 and WPK2 out perform CUSP-CSR, even when using CPU data reordering.

However, the ASUS ENGTX480 GPU has a thermal design power (TDP) of 295W, while the I7 950 has a TDP of 130W. We then normalize the speed up factor results to account for the difference in power consumption in Figure 11.23 to determine the computing effectiveness of our GPU algorithms as compared to computations on the CPU. For the GPU CG solver, we assume that the total CPU assembly and GPU CG solver implementation consume the maximum TDP of our CPU or GPU at worst. After power normalization, WPK1 and WPK2 kernels are computationally power normalized equivalent to the CPU starting at HeartCoarse and onwards. The CUSP-CSR and CUSPARSE GPU CG solvers are power-effective for the HeartRefined and Heart50k simulation. On the other hand, MGPU kernels only become power effective for the Heart50k simulation. For the fully-GPU implementations, all implementations are power-effective starting with a power normalized speedup factor of at least 1.5 and reach a power normalized speedup factor over 5 at the highest mesh refinement.

Lastly, we evaluate the effect of using CPU or GPU reordering on a finite element heart simulations in Figure 11.24. Again we see the same trends as before, with slight increases in performance for both WPK1 and WPK2 implementations. WPK1 and WPK2 implementations are more competitive in comparison to the surprising performance of CUSP-CSR over a larger range of mesh refinements in light of the initial benchmarks in Figures 11.9, 11.10, and 11.11.

11.8 Discussion

The results reaffirm several common GPU conventions, while also highlighting several interesting findings. The synthetic benchmarks using the Nvidia set of sparse matrices were used to find the "optimal" test parameters throughout this study for MGPU,

WPK1, and WPK2 kernels. The effect of different optimizations and variations to WPK1 and WPK2 kernels were performed to examine the increase in performance of the naive WPK1 and WPK2 implementations, and also suggest possible alternative SPMV implementations that may not be obvious in the traditional SPMV sense. The cost of reordering is also studied to determine the utility of the developed kernels. Lastly, fully-GPU finite element and hybrid CPU-GPU solver implementations were compared against multi-core implementations to study the utility of GPU finite element implementations as a whole, to evaluate the utility of the synthetic SPMV benchmarks in comparison to actual SPMV application, and to highlight possible FEA implementation improvements.

11.8.1 Synthetic Benchmarks

From the synthetic SPMV benchmarks, it was found that CUSP-CSR, CUSP-HYB, WPK1 and WPK2 had fairly consistent results regardless of the number of iterations. This of course makes these SPMV algorithms attractive as they perform consistently. However, CUSPARSE and MGPU algorithms tend to level off between 50 and 1200 iterations; thus having a high number of repetitive SPMV operations using the same matrix \mathbf{A} should allow those SPMV algorithms to also function more predictably.

Overall, the performance of our novel kernels is favorable with respect to MGPU. The benchmarks show that WPK1 and WPK2 perform very well over different numbers of iterations and over an assortment of sparse matrix structures, and in most cases is substantially faster than MGPU. In many ways, it is quite surprising that such a simple kernel algorithm (22 lines) has comparable performance to a sophisticated segmented scan algorithm in the case of MGPU (145 lines). Unlike other GPU SPMV algorithms, WPK1 and WPK2 are monolithic kernels and extra kernel invocations are unnecessary, in comparison to MGPU. As the synthetic benchmarks only account for the computation of the SPMV operation and do not account for data reordering, the optimal results shown in the synthetic benchmarks can only realistically be obtained when finite element degrees of freedom are renumbered in a fashion that results in

longest to shortest row length distributions.

Since WPK1 and WPK2 are partially warp-based variations of the ELL SPMV implementation and thus related to the HYB format, the effect of warp-level organization vs global data restructuring in HYB can be inferred. The results from the two variations of ELL show drastically different performance results for the majority of the benchmarked sparse matrices. WPK1/2 reduce the amount of padding with respect to ELL and simultaneously reduce the amount of memory transactions needed for computing the SPMV operation in comparison to HYB; thereby providing a dramatic increase in performance. WPK2 is used when there is a large difference in row length regularity and effectively handles "outlier" rows in a hybrid ELL-CSR manner, instead of using the CSR format directly as is done in HYB. Together, WPK1 and WPK2 are substantially better than ELL and HYB for sparse matrices.

In gathering the results for the finite element heart simulations over a range of different mesh refinements, it was found that the results from the 50 iteration synthetic benchmark seem to produce relatively fast computation times for the full GPU finite element implementation, but not necessarily in the case of the hybrid GPU Solver-CPU finite element implementation. Slightly larger blocksizes compared to the parameters found from the synthetic benchmarks (Table 11.4 and 11.5) have been found to provide a substantial increase in speed that outperforms CUSP-CSR and CUSPARSE at every mesh refinement. However, this is found by running a separate parameter search on the actual finite element simulation at each level of mesh refinement. The difference in parameters may be due to the way the synthetic benchmark test was performed. Several matrix formats were loaded onto the GPU and remain in global memory while other kernels run. This may be similar to the memory utilization of the fully-GPU finite element implementation. In the case of the GPU solver and CPU finite element assembly implementation, much less memory is required. These differences in memory utilization and memory transactions may affect the performance of the different caches on the GPU and result in the necessity of using different parameters.

11.8.2 Optimizations and Variations

In terms of optimizations to WPK1/2, our results show that column-major coalesced memory access is very important. This corroborates key insights made in developing the ELL, HYB, and MGPU SPMV algorithms [13, 10]. In general, it seems that sorting rows by row length and then subsequently reordering the column indices and renumbering the solution vector \mathbf{x} is beneficial when the texture cache is used and improves performance slightly (WPKr). In addition further reordering of the matrix \mathbf{A} may provide additional improvements (WPKrs). The cost of the final solution reordering in WPK1 can be expensive depending on the particular matrix, up to an additional 11.5% of computation time. Unfortunately, embedding alternate renumbering schemes adds some additional complexity to the finite element implementation. Therefore, if every bit of performance must be obtained, one would ideally use a pre-ordered kernel variation (WPKr/rs). While the speed up gains are compounded between WPK vs. WPKr and WPKr vs. WPKrs, proper modifications should yield substantial improvements in SPMV computation speed. However, as a general drop-in SPMV replacement, WPK1 seems to perform quite well for many applications.

11.8.3 Finite Element Simulations

From the finite element heart simulation results, even with a naive solution remapping kernel (WPK1 and WPK2), there is already a clear benefit in using these novel algorithms within the GPU CG solver and in fully-GPU finite element implementations. Even when factoring in the cost of reordering the matrix \mathbf{A} , there is a reduction in computational benefit after a reasonable number of CG iterations, especially when using a GPU data reordering method.

A single GTX480 GPU finite element implementation can out perform a comparable 4 core PETSc implementation on a i7 950 CPU by a factor of at least factor of 3 for the most refined mesh (Heart50k). Even for very small matrices (Heart, Heart-Coarse), there is a perceived benefit in using the GPU. The scaling will continue to favor GPUs with increasing matrix size. When power usage is considered, fully-GPU finite element implementations are very power-effective, while hybrid CG solver

implementations are at worst power-equivalent to a group of CPUs. In our power-normalized analysis, we made a worst-case assumption for the hybrid GPU/CPU finite element implementation. As the heart simulation problem leads to a computationally expensive finite element assembly operation, in general, the power consumption of the GPU only applies to a small portion of the actual computation time, and therefore the power-normalization for the hybrid implementation is a worst-case scenario. Overall, this further affirms the use of GPUs in finite element simulations.

While WPK1 performs equivalently to WPK2 for finite element simulations, WPK2 is still a very useful variation as seen in the synthetic benchmarks. WPK2 may not seem as useful for the benchmarked finite element simulations because of the small row-length regularity of the mesh matrices. However, from the benchmark results, WPK2 is a generally useful algorithm that degrades smoothly to the WPK1 kernel, when a threshold parameter is chosen to be the size of the maximum row length. WPK2 may be more useful on different finite element situations where there is a small number of highly connected nodes within the finite element mesh.

11.8.4 Improvements

Reordering of values is still expensive compared to the actual SPMV computation time. Luckily it is possible to hide the reordering bottleneck within the assembly routine. This can be done by first reordering the location where each element stores its elemental stiffness and residual entries such that the kernel can directly use the stored matrix to generate a resulting global stiffness and residual matrix that is already reordered for the CG solver.

Since this particular physical problem is assembly heavy, improvements and optimizations to finite element assembly will result in very substantial speed improvements in addition to improvements to SPMV operations. The shared-memory non-zero assembly operation [31] can be used to reduce global memory use and increase computational density.

Lastly, the slight differences between optimal GPU CG solvers and optimal GPU finite element implementations, highlight potential avenues of improvement for our

novel algorithms. Since the structure and access pattern of each mesh and matrix can be analyzed beforehand, it would be extremely beneficial to develop a metric for determining good partitioning parameters for WPK1 and WPK2 algorithms a priori. Given the relative simplicity of this algorithm, such a study, should be possible, and would further increase the utility of these warpkernel algorithms.

11.9 Conclusion

In conclusion, we have shown how key insights in the ELL, HYB, and MGPU SPMV algorithms have led to the development of new, simple, yet efficient SPMV algorithms that perform well over a large range of sparse matrices. The effects of different optimizations have been examined and ultimately lead to faster SPMV computation times. While there are computational costs associated with data-reordering, GPU coalesced data reordering can mitigate some of these costs. Ideally, it should be possible, albeit a little more complicated, to simply avoid coalesced data reordering altogether by embedding the reordering within the finite element stiffness and residual assembly operations. Lastly, this study highlights the potential use of GPUs for general finite element simulations.

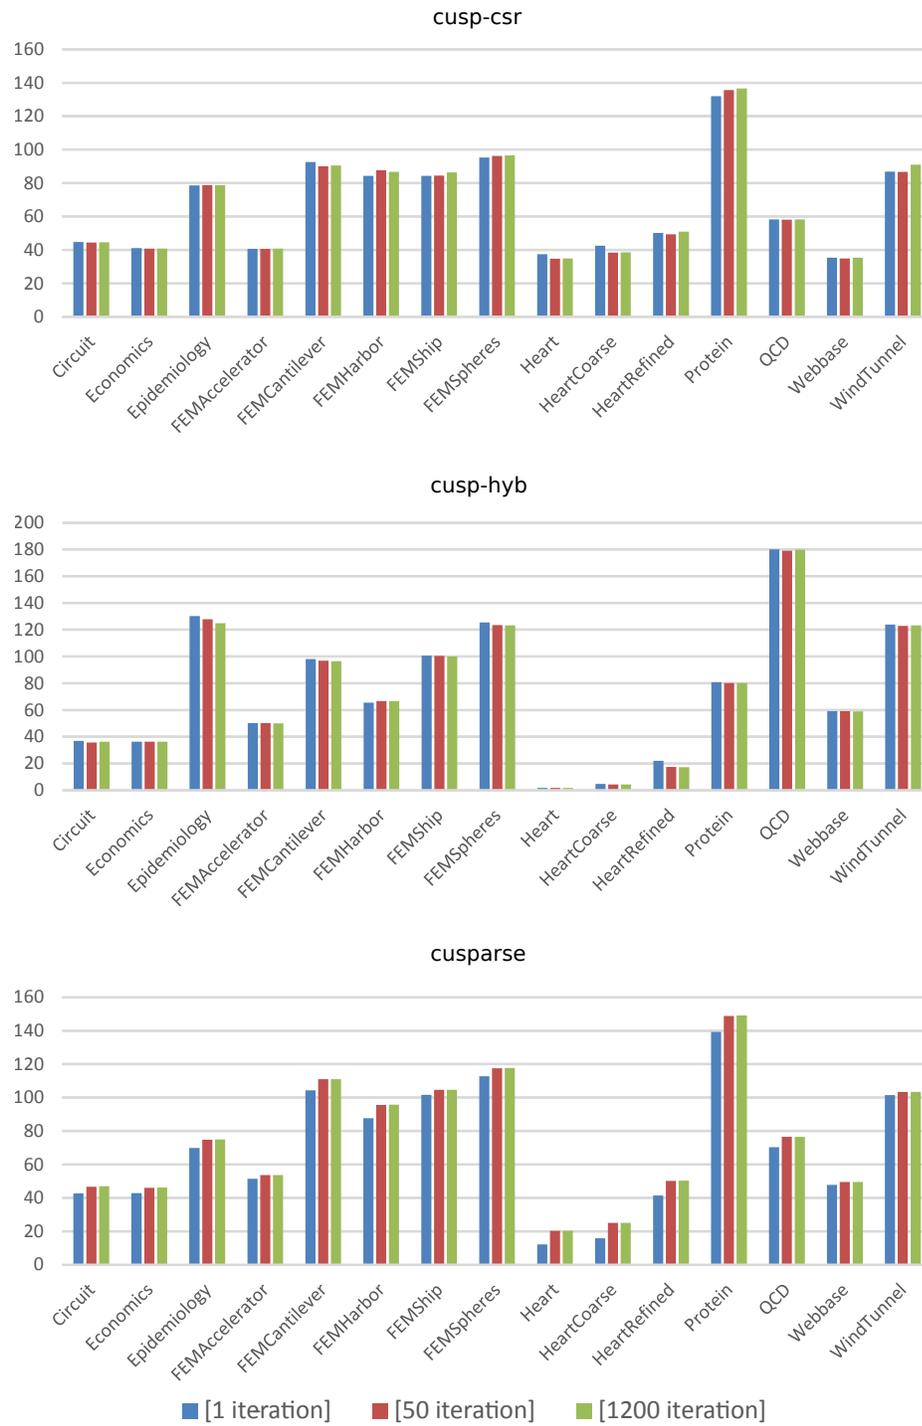


Figure 11.12: Effective bandwidth [GB/s] comparison over 1, 50, and 1200 iterations for cusp-csr, cusp-hyb, and cusparse implementations.

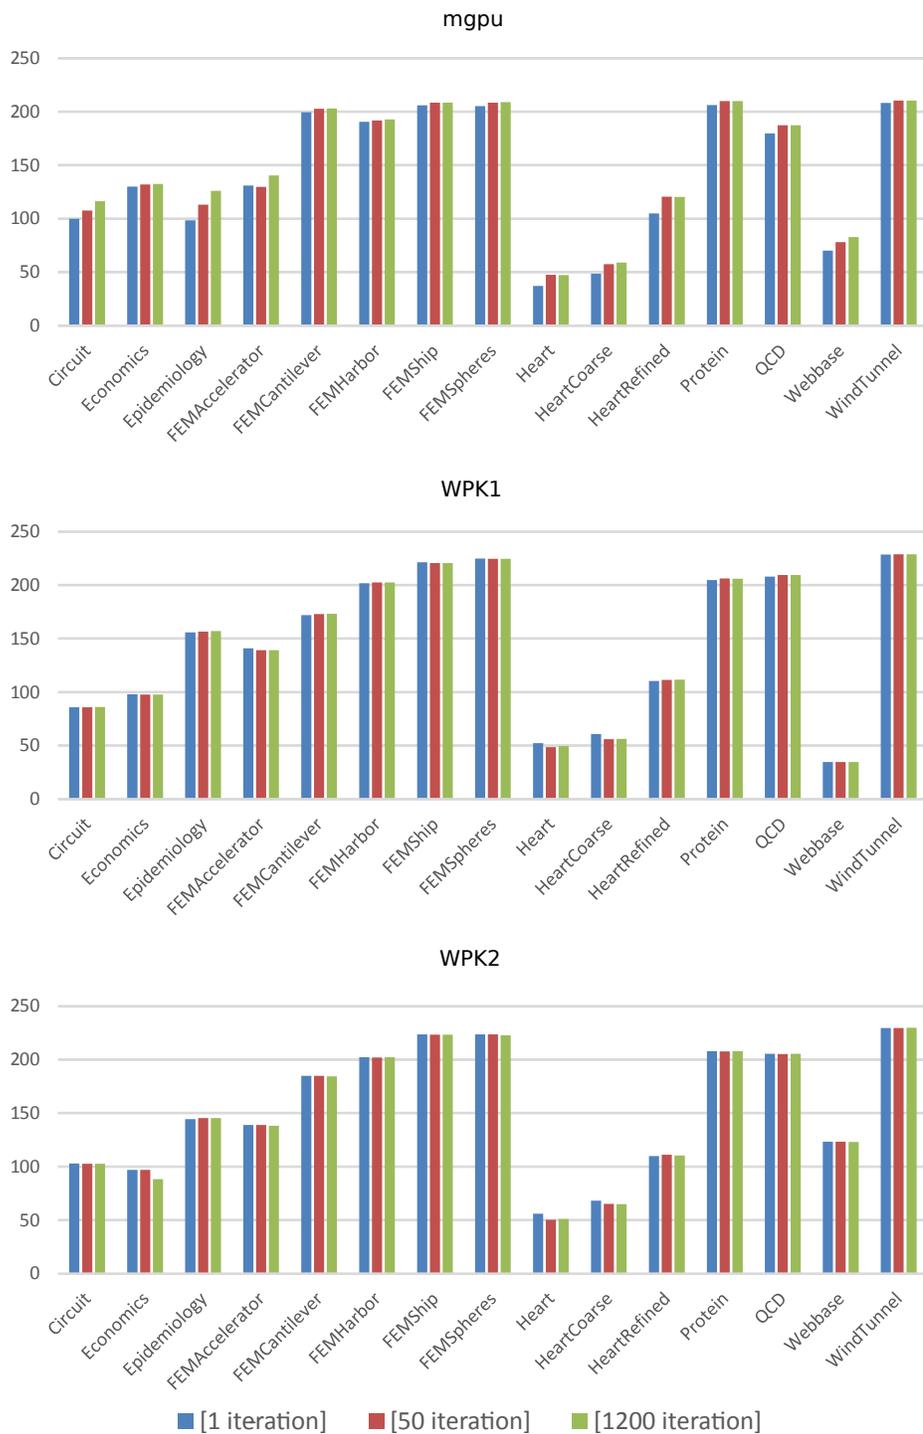


Figure 11.13: Effective bandwidth comparison over 1, 50, and 1200 iterations for mgpu, WPK1, and WPK2 implementations.

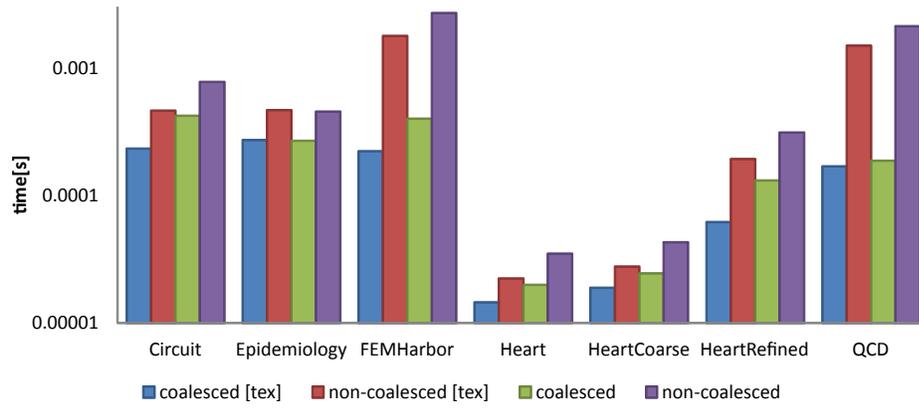


Figure 11.14: Comparison of coalesced "column-major" data vs. non-coalesced data ordering kernel times.

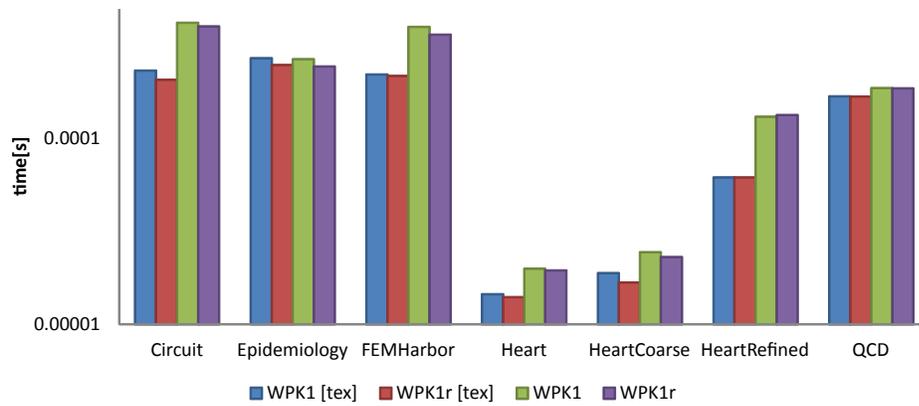


Figure 11.15: Comparison of WPK1 vs WPK1r kernel times.

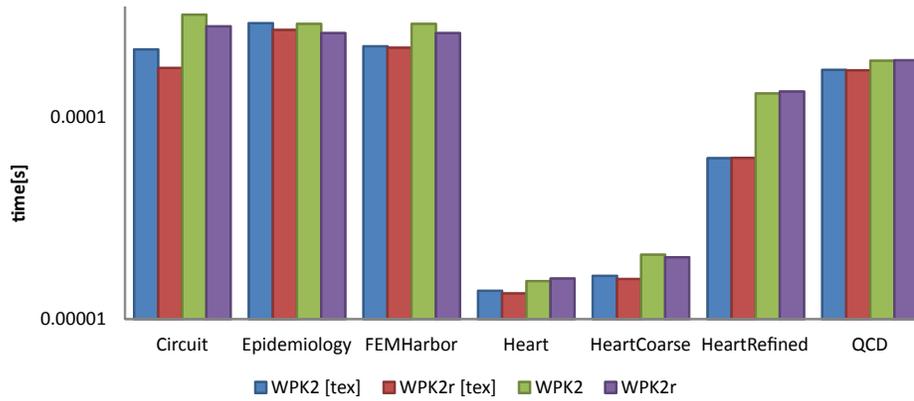


Figure 11.16: Comparison of WPK2 vs WPK2r kernel times.

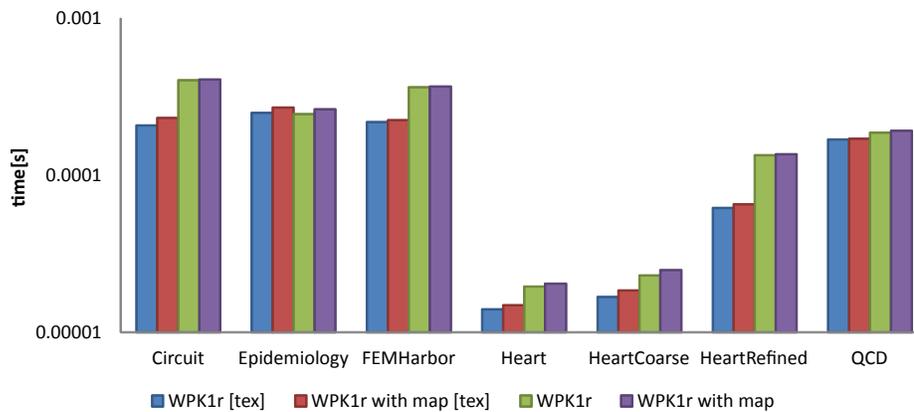


Figure 11.17: Comparison of kernel times to investigate the cost of the final remapping to the original reordering vs WPK1r.

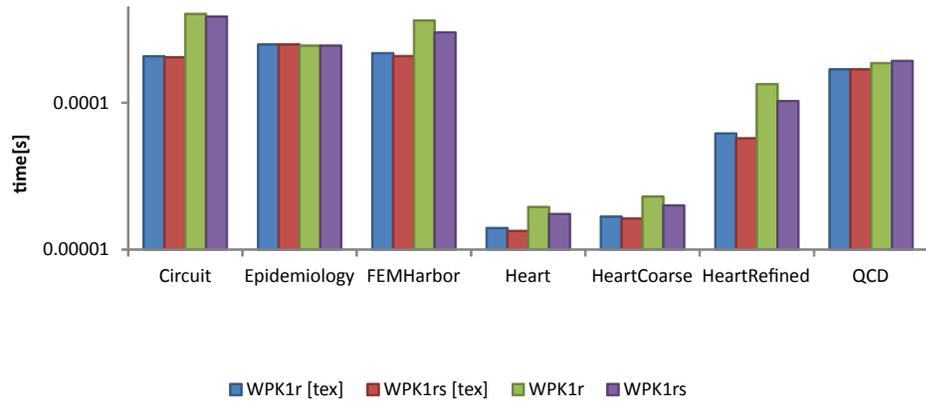


Figure 11.18: Comparison of WPK1r vs WPK1rs kernel times.

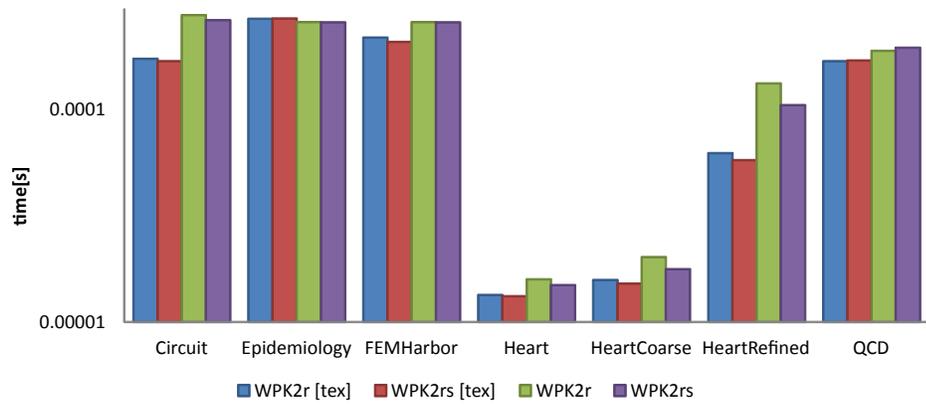


Figure 11.19: Comparison of WPK2r vs WPK2rs kernel times.

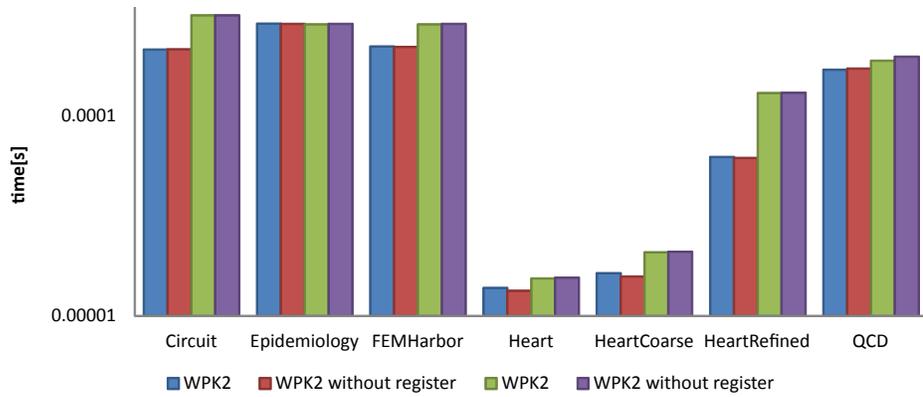


Figure 11.20: Comparison of WPK2 vs WPK2 without shared memory buffering register.

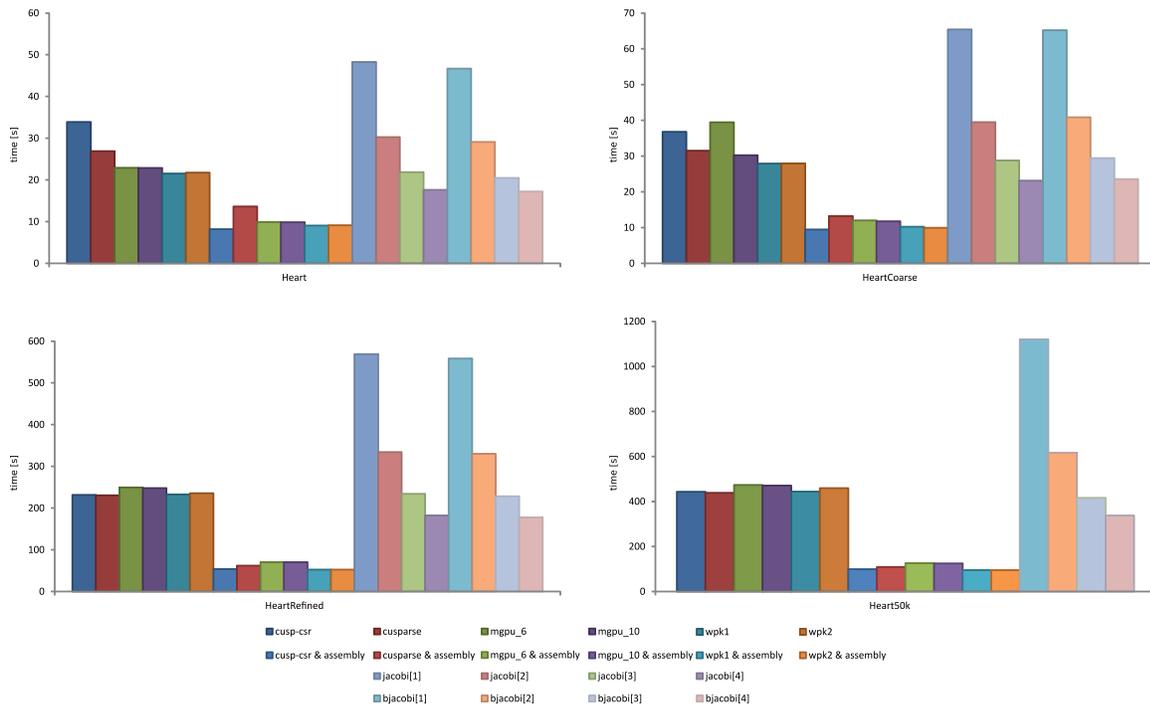


Figure 11.21: Running time in [s] of the a single cardiac cycle (1 second) for different finite element tested configurations: GPU CG solver, GPU finite element implementation, and multi-core PETSc implementation.

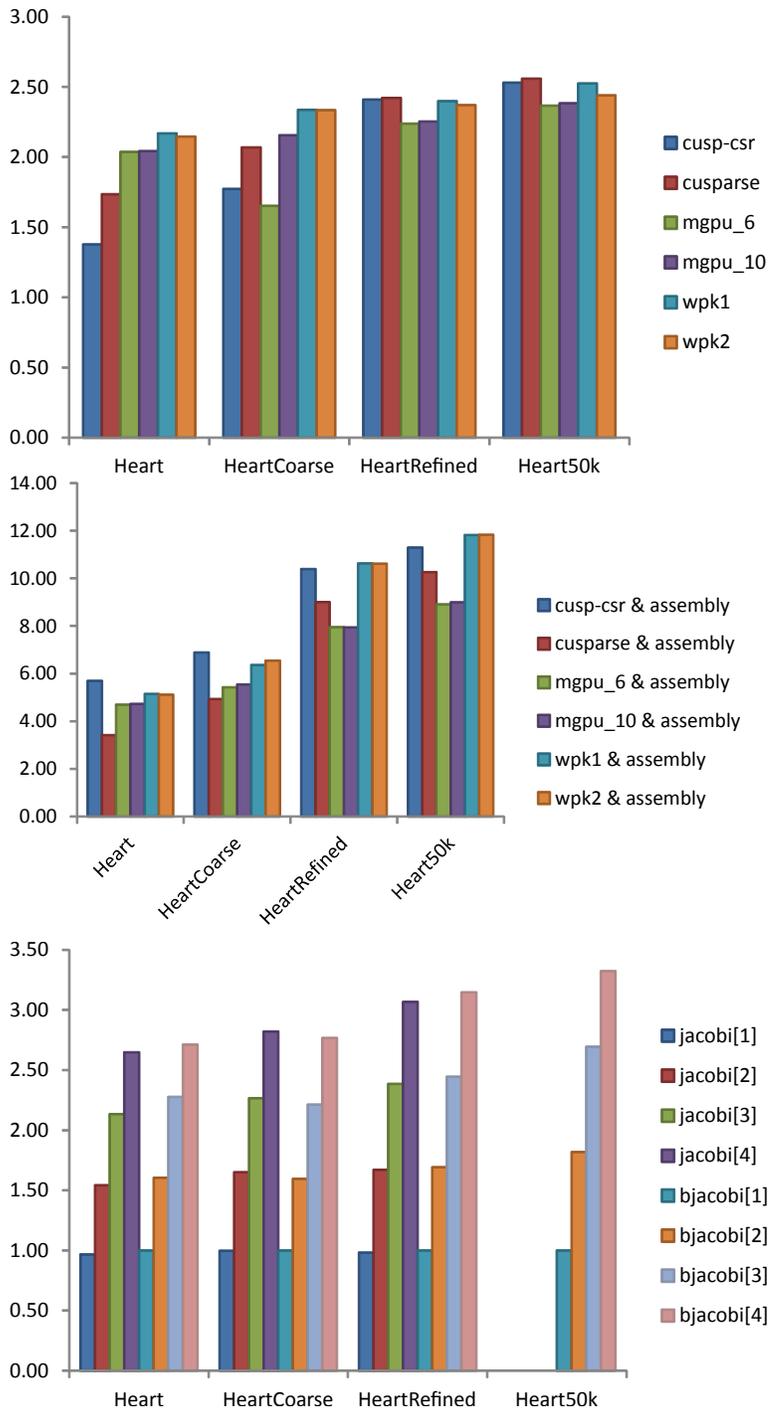


Figure 11.22: The resulting factor of increase in speed is shown for the GPU CG Solver and GPU finite element method with a single core PETSc block-jacobi CG solver as reference. PETSc implementation speedups are also shown.

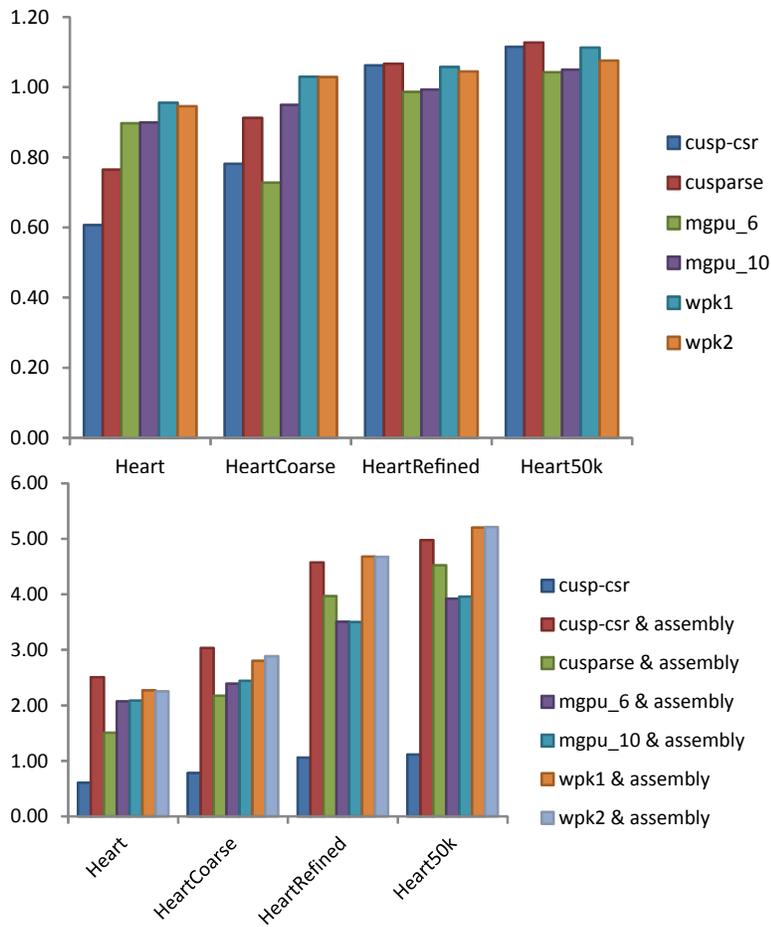


Figure 11.23: TDP normalized resulting factor of increase in speed is shown for the GPU CG Solver and GPU finite element method with a single core PETSc block-jacobi CG solver as reference.

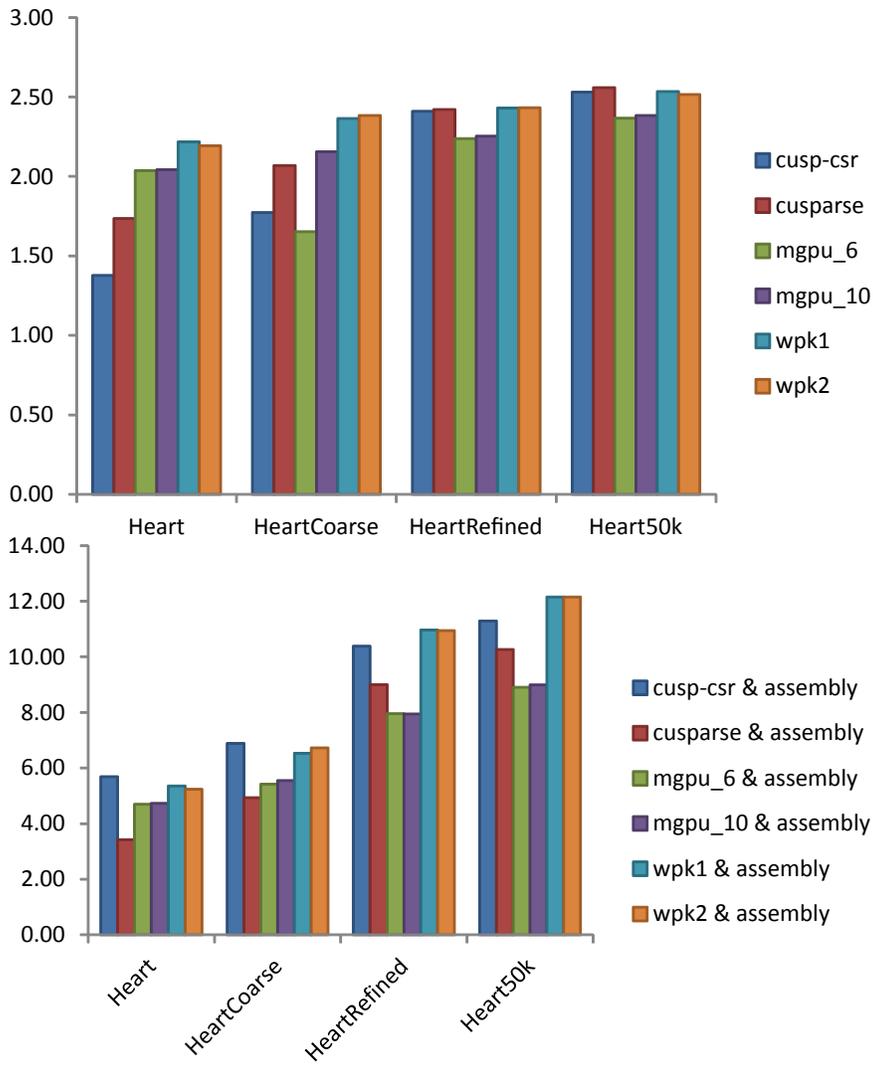


Figure 11.24: Speed up factor for GPU reordered WPK1 and WPK2 CG and finite element implementations.

Chapter 12

Concluding remarks and future work

In conclusion, several techniques were proposed and developed to address issues relating to the efficiency, modularity, flexibility, stability, and accessibility of cardiac electromechanical models and were applied to various applications. The monolithic modular algorithmic framework yields very good results and has been shown to, in general, be very robust for different electromechanical simulations with different material models. The scheme is modular and can be easily extended to encompass different types of biological phenomena for other electromechanically excitable tissues such as skeletal muscle. The framework also lends itself to multi-scale modeling approaches and has been shown through the optogenetic examples to help bridge the mesoscopic scale gap assuming relatively continuum behavior of material properties. While the scheme may initially seem computationally expensive, the reduction in global degrees of freedom seems to be cost-effective and may yield a symmetric tangent matrix which allows iterative solvers to more easily be applied. The algorithmic splitting also reduces the computational cost of ionic models by partitioning computation of internal variables based on their dependence on different ionic properties. Variables that are not coupled to other ionic properties can be computed cheaply, while the resulting set of nonlinear coupled ODEs can be solved within a banded tangent matrix structure. Thus additional complexity in terms of local ODE solution will scale linearly instead

of quadratically.

The choice in using implicit time-stepping integration schemes has proven to be extremely beneficial. The numerical stability due to implicit schemes should not be understated and has led to the development of robust strongly-coupled electromechanical cardiac simulations with physiologically accurate macroscopic torsional properties. The anisotropic structural properties of the heart, large deformations present during pumping, and the fast nonlinear behavior of the electrical wavefront have caused a multitude numerical problems in integrating multiphysics models; however, implicit methods provide the stable framework necessary for developing new models. Lastly, the stability of implicit time-stepping allows for trivial application and use of adaptive time-stepping methods based on energy norm convergence. This has proved especially useful not only for speeding up calculations, but also in selecting proper initial time-step sizes for the stiff nonlinear cardiac differential equations where boundary conditions vary wildly, i.e. the initiation of the excitation wave within the heart, through use of adaptive time-stepping methods.

A general fiber orientation distribution method based on Poisson interpolation using finite elements has been developed. This method is available and accessible to those with access to thermal finite elements, which are prevalent in finite element codes. The method also allows researchers and clinicians to perform approximate simulations on patient specific geometries even when only surface imaging has been used to examine the heart. This fiber orientation distribution can also be applied to a variety of materials as it is geometrically-based on the unstructured mesh, unlike other material designating algorithms that must perform cumbersome mapping transformations. This method has been applied in other recent simulations in our lab, and has found to produce good results. Lastly, the general method can also use experimentally obtained fiber measurements and directly interpolate those. Therefore, it is a general use interpolating method developed within the convenience of existing finite element codes.

Lastly, a set of GPU kernels have been proposed to increase the performance of sparse matrix vector multiplication operations on GPUs. While the algorithms were developed for use on finite element unstructured meshes, they performs very well on

a large range of sparse matrices. With these improvements in sparse matrix vector multiplication, and an initial conversion of CPU assembler and solver routines to the GPU, dramatic speed-ups were obtained both in the case of hybrid CPU-GPU and fully-GPU finite element implementations. Our implementation of the CG solver and assembly routines on the GPU have been shown to be power-effective compared to CPUs. This work has highlighted potential of GPUs in scientific computing, finite elements of solids, and near-real-time simulations of complex multi-field biological phenomena. Most importantly the particular GPU used in this dissertation is a commonly available consumer desktop gaming graphics card. This means that everyone now has access to computational resources that can be used for efficient scientific modeling of the heart.

While these results are decent and have been shown to provide many benefits in improving the difficult nature of electromechanical heart modeling, there is a plethora of future research topics that arise. While a monolithic modular framework has been developed, the algorithmic framework does not specify what models to use for particular simulations. There is a large need, especially with the pervasiveness of stem cell and IPS stem cell culture research, in developing methods to create useful ionic cell models for basic electromechanical studies of bioengineered cell type. Since bioengineered cell types exhibit a wide range of behaviors and have unique properties, it is important to develop accurate models quickly and to evaluate their functional utility especially in the case of tissue grafts of the heart or other important organs. Along such lines, as most ionic models contain a myriad of material parameters that will generally not be measured in experiments, hybrid methods should be developed to balance computational cost, experimental parameter identification correlation, and accuracy in characterizing complex indirectly measured phenomena such as electrical cell restitution. One particular approach may be to combine ionic and phenomenological cell models, such that there is a smooth transition from simple to advanced models based on the amount of experimental data available.

Another future area of interest may be in the development of electromechanically-based growth models of the heart. Through histological and explanted tissue studies on diseased patients, perhaps common fiber distribution patterns can be obtained for

the heart. Alternatively different electromechanical disease progression theories can be tested within the developed electromechanical framework. With the improvements in GPU performance, inverse studies may now be feasible and accessible to clinicians and experimentalists to perform parameter identification or to possibly even evaluate different surgical procedures in silico. While the developed models may not always yield quantitative matching results in cases where parameters are not calibrated to the experimental conditions, they may still provide qualitative insight and serve as a metric of principle validation for the scientific and medical community.

In the end, the focus of this thesis work has been to develop accessible, efficient, and robust tools for clinicians and experimentalists for the benefit of society. It is hoped that the material in this dissertation, and especially the near-real-time arrhythmia generating Kinect demonstration during my dissertation defense, will pique further interest in developing accessible computational biological models that are tightly integrated within experimental and clinical environments. While computational models are not an end-all solution to the worlds problems, they are a powerful class of tools that are relatively cheap to develop and use. They are simply the manifestation of current scientific knowledge condensed into more rigid mathematical rules. Thus, computational models of the heart should be made to be accessible, efficient and robust for experimentalists, clinicians, and scientists alike.

Bibliography

- [1] Oscar J. Abilez, Jonathan Wong, Rohit Prakash, Karl Deisseroth, Christopher K. Zarins, and Ellen Kuhl. Multiscale computational models for optogenetic control of cardiac function. *Biophysical Journal*, 101(6):1326–1334, September 2011.
- [2] David I. Abramson and Sydney Margolin. A purkinje conduction network in the myocardium of the mammalian ventricles. *Journal of Anatomy*, 70(Pt 2):250–259, January 1936. PMID: 17104583 PMCID: 1249121.
- [3] Fadi G. Akar, Bradley J. Roth, and David S. Rosenbaum. Optical measurement of cell-to-cell coupling in intact heart using subthreshold electrical stimulation. *American Journal of Physiology - Heart and Circulatory Physiology*, 281(2):H533–H542, August 2001.
- [4] Rubin R. Aliev and Alexander V. Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7(3):293–301, March 1996.
- [5] Aristides B. Arrenberg, Didier Y. R. Stainier, Herwig Baier, and Jan Huisken. Optogenetic control of cardiac function. *Science*, 330(6006):971–974, November 2010.
- [6] T. Arts, K. D. Costa, J. W. Covell, and A. D. McCulloch. Relating myocardial laminar architecture to shear strain and muscle fiber orientation. *American Journal of Physiology - Heart and Circulatory Physiology*, 280(5):H2222–H2229, May 2001.

- [7] Satish Balay, Jed Brown, , Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.3, Argonne National Laboratory, 2012.
- [8] Christian Bamann, Taryn Kirsch, Georg Nagel, and Ernst Bamberg. Spectral characteristics of the photocycle of channelrhodopsin-2 and its implication for channel function. *Journal of Molecular Biology*, 375(3):686–694, January 2008.
- [9] E. Bartocci, E. M. Cherry, J. Glimm, R. Grosu, S. A. Smolka, and F. H. Fenton. Toward real-time simulation of cardiac dynamics. In *Proceedings of the 9th International Conference on Computational Methods in Systems Biology*, CMSB '11, page 103112, New York, NY, USA, 2011. ACM.
- [10] Sean Baxter. Modern gpu libraries, 2011.
- [11] P V Bayly, B H KenKnight, J M Rogers, R E Hillsley, R E Ideker, and W M Smith. Estimation of conduction velocity vector fields from epicardial mapping data. *IEEE transactions on bio-medical engineering*, 45(5):563–571, May 1998. PMID: 9581054.
- [12] G. W. Beeler and H. Reuter. Reconstruction of the action potential of ventricular myocardial fibres. *The Journal of Physiology*, 268(1):177–210, June 1977.
- [13] Nathan Bell and Michael Garland. Implementing sparse matrix-vector multiplication on throughput-oriented processors. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC '09, page 18:118:11, New York, NY, USA, 2009. ACM.
- [14] Nathan Bell and Michael Garland. Cusp: Generic parallel algorithms for sparse matrix and graph computations, 2012. Version 0.3.0.
- [15] Andr Berndt, Matthias Prigge, Dietrich Gradmann, and Peter Hegemann. Two open states with progressive proton selectivities in the branched

- channelrhodopsin-2 photocycle. *Biophysical Journal*, 98(5):753–761, March 2010.
- [16] Andr Berndt, Philipp Schoenenberger, Joanna Mattis, Kay M. Tye, Karl Deisseroth, Peter Hegemann, and Thomas G. Oertner. High-efficiency channelrhodopsins for fast neuronal stimulation at low light levels. *Proceedings of the National Academy of Sciences*, 108(18):7595–7600, May 2011.
- [17] O. Bernus, R. Wilders, C. W. Zemlin, H. Vershelde, and A. V. Panfilov. A computationally efficient electrophysiological model of human ventricular cells. *American Journal of Physiology - Heart and Circulatory Physiology*, 282(6):H2296–H2308, June 2002.
- [18] Donald M. Bers. *Excitation-Contraction Coupling and Cardiac Contractile Force*. Springer, August 2001.
- [19] Markus Böl and Stefanie Reese. Micromechanical modelling of skeletal muscles based on the finite element method. *Computer methods in biomechanics and biomedical engineering*, 11(5):489–504, October 2008. PMID: 19230146.
- [20] Markus Böl, Maike Sturmat, Christine Weichert, and Cornelia Kober. A new approach for the validation of skeletal muscle modelling using MRI data. *Comput. Mech.*, 47(5):591601, May 2011.
- [21] Jeff Bolz, Ian Farmer, Eitan Grinspun, and Peter Schroder. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Trans. Graph.*, 22(3):917924, July 2003.
- [22] Muriel Boulakia, Serge Cazeau, Miguel Fernandez, Jean-Frdric Gerbeau, and Nejib Zemzemi. Mathematical modeling of electrocardiograms: A numerical study. *Annals of Biomedical Engineering*, 38(3):1071–1097, 2010.
- [23] Edward S. Boyden, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9):1263–1268, 2005.

- [24] Marco Brera, Joseph W. Jerome, Yoichiro Mori, and Riccardo Sacco. A conservative and monotone mixed-hybridized finite element approximation of transport problems in heterogeneous domains. *Computer Methods in Applied Mechanics and Engineering*, 199(4144):2709–2720, October 2010.
- [25] Luc Buatois, Guillaume Caumon, and Bruno Levy. Concurrent number cruncher: a GPU implementation of a general sparse linear solver. *Int. J. Parallel Emerg. Distrib. Syst.*, 24(3):205223, June 2009.
- [26] N Bursac, K K Parker, S Iravanian, and L Tung. Cardiomyocyte cultures with controlled macroscopic anisotropy: a model for functional electrophysiological studies of cardiac muscle. *Circulation research*, 91(12):e45–54, December 2002. PMID: 12480825.
- [27] Markus Bl, Stefanie Reese, Kevin Kit Parker, and Ellen Kuhl. Computational modeling of muscular thin films for cardiac repair. *Computational Mechanics*, 43(4):535–544, September 2008.
- [28] Sherry L Caffrey, Paula J Willoughby, Paul E Pepe, and Lance B Becker. Public use of automated external defibrillators. *The New England journal of medicine*, 347(16):1242–1247, October 2002. PMID: 12393821.
- [29] Patrizia Camelliti, Andrew D. McCulloch, and Peter Kohl. Microstructured cocultures of cardiac myocytes and fibroblasts: A Two-Dimensional in vitro model of cardiac tissue. *Microscopy and Microanalysis*, 11(03):249–259, 2005.
- [30] Feng Cao, Shuan Lin, Xiaoyan Xie, Pritha Ray, Manishkumar Patel, Xianzhong Zhang, Micha Drukker, Scott J. Dylla, Andrew J. Connolly, Xiaoyuan Chen, Irving L. Weissman, Sanjiv S. Gambhir, and Joseph C. Wu. In vivo visualization of embryonic stem cell survival, proliferation, and migration after cardiac delivery. *Circulation*, 113(7):1005–1014, February 2006.
- [31] Cris Cecka, Adrian J. Lew, and E. Darve. Assembly of finite element methods on graphics processors. *International Journal for Numerical Methods in Engineering*, 85(5):640669, 2011.

- [32] Marvin G. Chang, Leslie Tung, Rajesh B. Sekar, Connie Y. Chang, Josh Cysyk, Peihong Dong, Eduardo Marbn, and M. Roselle Abraham. Proarrhythmic potential of mesenchymal stem cell transplantation revealed in an in vitro coculture model. *Circulation*, 113(15):1832–1841, April 2006.
- [33] D. Chapelle, J.-F. Gerbeau, J. Sainte-Marie, and I. Vignon-Clementel. A poroelastic model valid in large strains with applications to perfusion in cardiac modeling. *Computational Mechanics*, 46(1):91–101, 2010.
- [34] Dominique Chapelle, Miguel A. Fernandez, Jean-Frdric Gerbeau, Philippe Moireau, Jacques Sainte-Marie, and Nejib Zemzemi. Numerical simulation of the electromechanical activity of the heart. In *Proceedings of the 5th International Conference on Functional Imaging and Modeling of the Heart, FIMH '09*, page 357365, Berlin, Heidelberg, 2009. Springer-Verlag.
- [35] F.F.-T. Ch'en, R.D. Vaughan-Jones, K. Clarke, and D. Noble. Modelling myocardial ischaemia and reperfusion. *Progress in Biophysics and Molecular Biology*, 69(23):515–538, March 1998.
- [36] Michael Q. Chen, R. Hollis Whittington, Peter W. Day, Brian K. Kobilka, Laurent Giovangrandi, and Gregory T. A. Kovacs. A device for separated and reversible co-culture of cardiomyocytes. *Biotechnology Progress*, 26(4):11641171, 2010.
- [37] Michael Q Chen, Jonathan Wong, Ellen Kuhl, Laurent Giovangrandi, and Gregory T A Kovacs. Characterisation of electrophysiological conduction in cardiomyocyte co-cultures using co-occurrence analysis. *Computer methods in biomechanics and biomedical engineering*, October 2011. PMID: 21970595.
- [38] Allen Cheng, Frank Langer, Filiberto Rodriguez, John C Criscione, George T Daughters, D Craig Miller, and Jr Ingels, Neil B. Transmural sheet strains in the lateral wall of the ovine left ventricle. *American journal of physiology. Heart and circulatory physiology*, 289(3):H1234–1241, September 2005. PMID: 15879489.

- [39] William C. Claycomb, Nicholas A. Lanson, Beverly S. Stallworth, Daniel B. Egeland, Joseph B. Delcarpio, Anthony Bahinski, and Nicholas J. Izzo. HL-1 cells: A cardiac muscle cell line that contracts and retains phenotypic characteristics of the adult cardiomyocyte. *Proceedings of the National Academy of Sciences of the United States of America*, 95(6):2979–2984, March 1998. PMID: 9501201 PMCID: PMC19680.
- [40] R.H. Clayton and A.V. Panfilov. A guide to modelling cardiac electrical activity in anatomically detailed ventricles. *Progress in Biophysics and Molecular Biology*, 96(13):19–43, January 2008.
- [41] Sonia Cortassa, Miguel A Aon, Eduardo Marbn, Raimond L Winslow, and Brian O'Rourke. An integrated model of cardiac mitochondrial energy metabolism and calcium dynamics. *Biophysical journal*, 84(4):2734–2755, April 2003. PMID: 12668482.
- [42] K. D. Costa, P. J. Hunter, J. S. Wayne, L. K. Waldman, J. M. Guccione, and A. D. McCulloch. A Three-Dimensional finite element method for large elastic deformations of ventricular myocardium: II—Prolate spheroidal coordinates. *Journal of Biomechanical Engineering*, 118(4):464–472, November 1996.
- [43] Marc Courtemanche, Rafael J. Ramirez, and Stanley Nattel. Ionic mechanisms underlying human atrial action potential properties: Insights from a mathematical model. *American Journal of Physiology - Heart and Circulatory Physiology*, 275(1):H301–H321, July 1998.
- [44] Hüsnu Dal, Serdar Göktepe, Michael Kaliske, and Ellen Kuhl. A fully implicit finite element method for bidomain models of cardiac electrophysiology. *Computer methods in biomechanics and biomedical engineering*, 15(6):645–656, June 2012. PMID: 21491253.
- [45] Larry S. Davis, Steven A. Johns, and J. K. Aggarwal. Texture analysis using generalized Co-Occurrence matrices. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-1(3):251–259, July 1979.

- [46] Karl Deisseroth. Optogenetics. *Nature Methods*, 8(1):26–29, 2011.
- [47] D DiFrancesco. Pacemaker mechanisms in cardiac tissue. *Annual Review of Physiology*, 55(1):455–472, 1993.
- [48] D. DiFrancesco and D. Noble. A model of cardiac electrical activity incorporating ionic pumps and concentration changes. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 307(1133):353–398, January 1985.
- [49] Socrates Dokos, Branko Celler, and Nigel Lovell. Ion currents underlying sinoatrial node pacemaker activity: A new single cell mathematical model. *Journal of Theoretical Biology*, 181(3):245–272, August 1996.
- [50] Socrates Dokos, Bruce H. Smaill, Alistair A. Young, and Ian J. LeGrice. Shear properties of passive ventricular myocardium. *American Journal of Physiology - Heart and Circulatory Physiology*, 283(6):H2650–H2659, December 2002.
- [51] Erich Elsen, Patrick LeGresley, and Eric Darve. Large calculation of the flow over a hypersonic vehicle using a GPU. *Journal of Computational Physics*, 227(24):10148–10161, December 2008.
- [52] Daniel B. Ennis, Tom C. Nguyen, Jonathan C. Riboh, Lars Wigstrm, Katherine B. Harrington, George T. Daughters, Neil B. Ingels, and D. Craig Miller. Myofiber angle distributions in the ovine left ventricle do not conform to computationally optimized predictions. *Journal of Biomechanics*, 41(15):3219–3224, November 2008.
- [53] Marc Ethier and Yves Bourgault. Semi-Implicit Time-Discretization schemes for the bidomain model. *SIAM Journal on Numerical Analysis*, 46(5):2443–2468, January 2008.
- [54] Vadim V. Fedorov, Ilya T. Lozinsky, Eugene A. Sosunov, Evgeniy P. Anyukhovskiy, Michael R. Rosen, C. William Balke, and Igor R. Efimov. Application of blebbistatin as an excitation-contraction uncoupler for electrophysiologic study of rat and rabbit hearts. *Heart Rhythm*, 4(5):619–626, May 2007.

- [55] Matthew Fisher, Peter Schröder, Mathieu Desbrun, and Hugues Hoppe. Design of tangent vector fields. In *ACM SIGGRAPH 2007 papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.
- [56] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, July 1961.
- [57] P. J. Flory. Thermodynamic relations for high elastic materials. *Transactions of the Faraday Society*, 57(0):829–838, January 1961.
- [58] M. R. Franz, K. Bargheer, W. Rafflenbeul, A. Haverich, and P. R. Lichtlen. Monophasic action potential mapping in human subjects with normal electrocardiograms: direct evidence for the genesis of the t wave. *Circulation*, 75(2):379–386, February 1987.
- [59] Y. C. Fung. *Biomechanics: Mechanical Properties of Living Tissues*. Springer, 2nd edition, June 1993.
- [60] Giedrius Gaudesius, Michele Miragoli, Stuart P. Thomas, and Stephan Rohr. Coupling of cardiac electrical activity over extended distances by fibroblasts of cardiac origin. *Circulation Research*, 93(5):421–428, September 2003.
- [61] L. Geerts, P. Bovendeerd, K. Nicolay, and T. Arts. Characterization of the normal cardiac myofiber field in goat measured with MR-diffusion tensor imaging. *American Journal of Physiology - Heart and Circulatory Physiology*, 283(1):H139–H145, July 2002.
- [62] Kristin H. Gilchrist, Valerie N. Barker, Lauren E. Fletcher, B.Derek DeBusschere, Pejman Ghanouni, Laurent Giovannardi, and Gregory T.A. Kovacs. General purpose, field-portable cell-based biosensor platform. *Biosensors and Bioelectronics*, 16(78):557–564, September 2001.
- [63] Steven D. Girouard, Joseph M. Pastore, Kenneth R. Laurita, Kenton W. Gregory, and David S. Rosenbaum. Optical mapping in a new guinea pig model of

- ventricular tachycardia reveals mechanisms for multiple wavelengths in a single reentrant circuit. *Circulation*, 93(3):603–613, February 1996.
- [64] D. Goddeke, S.H.M. Buijssen, H. Wobker, and S. Turek. GPU acceleration of an unmodified parallel finite element navier-stokes solver. In *International Conference on High Performance Computing Simulation, 2009. HPCS '09*, pages 12–21, June 2009.
- [65] S. Göktepe, S. N. S Acharya, J. Wong, and E. Kuhl. Computational modeling of passive myocardium. *International Journal for Numerical Methods in Biomedical Engineering*, 27(1):1–12, January 2011.
- [66] Serdar Göktepe and Ellen Kuhl. Electromechanics of the heart: a unified approach to the strongly coupled excitation-contraction problem. *Computational Mechanics*, 45(2-3):227–243, November 2009.
- [67] Viviana Gradinaru, Feng Zhang, Charu Ramakrishnan, Joanna Mattis, Rohit Prakash, Ilka Diester, Inbal Goshen, Kimberly R. Thompson, and Karl Deisseroth. Molecular and cellular approaches for diversifying and extending optogenetics. *Cell*, 141(1):154–165, April 2010.
- [68] Joseph L Greenstein and Raimond L Winslow. An integrative model of the cardiac ventricular myocyte incorporating local control of ca^{2+} release. *Biophysical Journal*, 83(6):2918–2945, December 2002. PMID: 12496068 PMCID: PMC1201479.
- [69] J M Guccione and A D McCulloch. Mechanics of active contraction in cardiac muscle: Part i—constitutive relations for fiber stress that describe deactivation. *Journal of biomechanical engineering*, 115(1):72–81, February 1993. PMID: 8445901.
- [70] Lisa A. Gunaydin, Ofer Yizhar, Andr Berndt, Vikaas S. Sohal, Karl Deisseroth, and Peter Hegemann. Ultrafast optogenetic control. *Nature Neuroscience*, 13(3):387–392, 2010.

- [71] Viatcheslav Gurev, Ted Lee, Jason Constantino, Hermenegild Arevalo, and Natalia A. Trayanova. Models of cardiac electromechanics based on individual hearts imaging data. *Biomechanics and modeling in mechanobiology*, 10(3):295–306, June 2011. PMID: 20589408 PMCID: PMC3166526.
- [72] S. Gktepe and E. Kuhl. Computational modeling of cardiac electrophysiology: A novel finite element approach. *International Journal for Numerical Methods in Engineering*, 79(2):156–178, February 2009.
- [73] Serdar Gktepe, Jonathan Wong, and Ellen Kuhl. Atrial and ventricular fibrillation: computational simulation of spiral waves in cardiac tissue. *Archive of Applied Mechanics*, 80(5):569–580, 2010.
- [74] Gundolf Haase, Manfred Liebmann, Craig Douglas, and Gernot Plank. A parallel algebraic multigrid solver on graphics processing units. In Wu Zhang, Zhangxin Chen, Craig Douglas, and Weiqin Tong, editors, *High Performance Computing and Applications*, volume 5938 of *Lecture Notes in Computer Science*, pages 38–47. Springer Berlin / Heidelberg, 2010.
- [75] O. P. Hamill, A. Marty, E. Neher, B. Sakmann, and F. J. Sigworth. Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches. *Pflügers Archiv European Journal of Physiology*, 391(2):85–100, 1981.
- [76] R.M. Haralick. Statistical and structural approaches to texture. *Proceedings of the IEEE*, 67(5):786 – 804, May 1979.
- [77] Robert M. Haralick, K. Shanmugam, and Its’Hak Dinstein. Textural features for image classification. *Systems, Man and Cybernetics, IEEE Transactions on*, SMC-3(6):610 –621, November 1973.
- [78] P. Hegemann, W. Grtner, and R. Uhl. All-trans retinal constitutes the functional chromophore in chlamydomonas rhodopsin. *Biophysical Journal*, 60(6):1477–1489, December 1991. PMID: 19431816 PMCID: PMC1260206.

- [79] Peter Hegemann, Sabine Ehlenbeck, and Dietrich Gradmann. Multiple photocycles of channelrhodopsin. *Biophysical Journal*, 89(6):3911–3918, December 2005. PMID: 16169986 PMCID: PMC1366958.
- [80] Peter Hegemann and Andreas Mglych. Channelrhodopsin engineering and exploration of new optogenetic tools. *Nature Methods*, 8(1):39–42, 2011.
- [81] Melonie Heron, Donna L Hoyert, Sherry L Murphy, Jiaquan Xu, Kenneth D Kochanek, and Betzaida Tejada-Vera. Deaths: final data for 2006. *National vital statistics reports: from the Centers for Disease Control and Prevention, National Center for Health Statistics, National Vital Statistics System*, 57(14):1–134, April 2009. PMID: 19788058.
- [82] Jared Hoberock and Nathan Bell. Thrust: A parallel template library, 2010. Version 1.3.0.
- [83] A. L Hodgkin and A. F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, August 1952.
- [84] Gerhard A. Holzapfel and Ray W. Ogden. Constitutive modelling of passive myocardium: a structurally based framework for material characterization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1902):3445–3475, 2009.
- [85] Thomas J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Publications, August 2000.
- [86] T J Hund, J P Kucera, N F Otani, and Y Rudy. Ionic charge conservation and long-term steady state in the Luo-Rudy dynamic cell model. *Biophysical Journal*, 81(6):3324–3331, December 2001. PMID: 11720995 PMCID: 1301789.
- [87] Sharon Ann Hunt, William T. Abraham, Marshall H. Chin, Arthur M. Feldman, Gary S. Francis, Theodore G. Ganiats, Mariell Jessup, Marvin A. Konstam, Donna M. Mancini, Keith Michl, John A. Oates, Peter S. Rahko, Marc A.

- Silver, Lynne Warner Stevenson, Clyde W. Yancy, Elliott M. Antman, Sidney C. Smith, Cynthia D. Adams, Jeffrey L. Anderson, David P. Faxon, Valentin Fuster, Jonathan L. Halperin, Loren F. Hiratzka, Sharon Ann Hunt, Alice K. Jacobs, Rick Nishimura, Joseph P. Ornato, Richard L. Page, and Barbara Riegel. ACC/AHA 2005 guideline update for the diagnosis and management of chronic heart failure in the adult a report of the american college of Cardiology/American heart association task force on practice guidelines (Writing committee to update the 2001 guidelines for the evaluation and management of heart failure): Developed in collaboration with the american college of chest physicians and the international society for heart and lung transplantation: Endorsed by the heart rhythm society. *Circulation*, 112(12):e154–e235, September 2005.
- [88] Peter J. Hunter, Andrew J. Pullan, and Bruce H. Smaill. Modeling total heart function. *Annual Review of Biomedical Engineering*, 5(1):147–177, 2003.
- [89] Toru Ishizuka, Masaaki Kakuda, Rikita Araki, and Hiromu Yawo. Kinetic evaluation of photosensitivity in genetically engineered neurons expressing green algae light-gated channels. *Neuroscience Research*, 54(2):85–94, February 2006.
- [90] Akinobu Itoh, Gaurav Krishnamurthy, Julia C. Swanson, Daniel B. Ennis, Wolfgang Bothe, Ellen Kuhl, Matts Karlsson, Lauren R. Davis, D. Craig Miller, and Neil B. Ingels. Active stiffening of mitral valve leaflets in the beating heart. *American Journal of Physiology - Heart and Circulatory Physiology*, 296(6):H1766–H1773, June 2009.
- [91] M. J. Janse and F. J. van Capelle. Electrotonic interactions across an inexcitable region as a cause of ectopic activity in acute regional myocardial ischemia. a study in intact porcine and canine hearts and computer models. *Circulation Research*, 50(4):527–537, April 1982.
- [92] Michiel J. Janse, Tobias Opthof, and Andr G. Klber. Animal models of cardiac arrhythmias. *Cardiovascular Research*, 39(1):165–177, July 1998.

- [93] Barbara M. Johnston, Peter R. Johnston, and David Kilpatrick. A solution method for the determination of cardiac potential distributions with an alternating current source. This work was supported by the Australian Research Council. *Computer Methods in Biomechanics and Biomedical Engineering*, 11(3):223–233, 2008.
- [94] Peter R Johnston. A finite volume method solution for the bidomain equations and their application to modelling cardiac ischaemia. *Computer methods in biomechanics and biomedical engineering*, 13(2):157–170, 2010. PMID: 19639486.
- [95] B. Julesz. Visual pattern discrimination. *Information Theory, IRE Transactions on*, 8(2):84–92, February 1962.
- [96] James Keener and James Sneyd. *Mathematical Physiology (Interdisciplinary Applied Mathematics) 2 Vol Set*. Springer, 2nd edition, December 2008.
- [97] Izhak Kehat, Leonid Khimovich, Oren Caspi, Amira Gepstein, Rona Shofti, Gil Arbel, Irit Huber, Jonathan Satin, Joseph Itskovitz-Eldor, and Lior Gepstein. Electromechanical integration of cardiomyocytes derived from human embryonic stem cells. *Nature Biotechnology*, 22(10):1282–1289, 2004.
- [98] R. Keldermann, M. Nash, and A. Panfilov. Pacemakers in a Reaction-Diffusion mechanics system. *Journal of Statistical Physics*, 128(1):375–392, 2007.
- [99] Mohammed N Khan, George Joseph, Yaariv Khaykin, Khaled M Ziada, and Bruce L Wilkoff. Delayed lead perforation: a disturbing trend. *Pacing and clinical electrophysiology: PACE*, 28(3):251–253, March 2005. PMID: 15733190.
- [100] Andr G. Klber and Yoram Rudy. Basic mechanisms of cardiac impulse propagation and associated arrhythmias. *Physiological Reviews*, 84(2):431–488, April 2004.
- [101] Peter Kohl, Patrizia Camelliti, Francis L. Burton, and Godfrey L. Smith. Electrical coupling of fibroblasts and myocytes: relevance for cardiac propagation. *Journal of Electrocardiology*, 38(4, Supplement):45–50, October 2005.

- [102] M. Kotikanyadanam, S. Gktepe, and E. Kuhl. Computational modeling of electrocardiograms: A finite element approach toward cardiac excitation. *International Journal for Numerical Methods in Biomedical Engineering*, 26(5):524–533, June 2009.
- [103] M. Kotikanyadanam, S. Gktepe, and E. Kuhl. Computational modeling of electrocardiograms: A finite element approach toward cardiac excitation. *International Journal for Numerical Methods in Biomedical Engineering*, 26(5):524–533, June 2009.
- [104] V I Krinsky. Mathematical models of cardiac arrhythmias (spiral waves). *Pharmacology & therapeutics. Part B: General & systematic pharmacology*, 3(4):539–555, 1978. PMID: 366630.
- [105] Michael A. Laflamme, Kent Y. Chen, Anna V. Naumova, Veronica Muskheli, James A. Fugate, Sarah K. Dupras, Hans Reinecke, Chunhui Xu, Mohammad Hassanipour, Shailaja Police, Chris O’Sullivan, Lila Collins, Yinhong Chen, Elina Minami, Edward A. Gill, Shuichi Ueno, Chun Yuan, Joseph Gold, and Charles E. Murry. Cardiomyocytes derived from human embryonic stem cells in pro-survival factors enhance function of infarcted rat hearts. *Nature Biotechnology*, 25(9):1015–1024, 2007.
- [106] Pierre Lafortune, Ruth Ars, Mariano Vzquez, and Guillaume Houzeaux. Coupled electromechanical model of the heart: Parallel finite element formulation. *International Journal for Numerical Methods in Biomedical Engineering*, 28(1):7286, 2012.
- [107] W. J. Lammers, M. J. Schaliij, C. J. Kirchhof, and M. A. Allesie. Quantification of spatial inhomogeneity in conduction and initiation of reentrant atrial arrhythmias. *American Journal of Physiology - Heart and Circulatory Physiology*, 259(4):H1254–H1263, October 1990.

- [108] M.A. Lawson, D.N. Zacks, F. Derguini, K. Nakanishi, and J.L. Spudich. Retinal analog restoration of photophobic responses in a blind *chlamydomonas reinhardtii* mutant. evidence for an archaebacterial like chromophore in a eukaryotic rhodopsin. *Biophysical Journal*, 60(6):1490–1498, December 1991.
- [109] Jin Hyung Lee, Remy Durand, Viviana Gradinaru, Feng Zhang, Inbal Goshen, Dae-Shik Kim, Lief E. Fenno, Charu Ramakrishnan, and Karl Deisseroth. Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature*, 465(7299):788–792, June 2010.
- [110] I J LeGrice, B H Smaill, L Z Chai, S G Edgar, J B Gavin, and P J Hunter. Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *The American journal of physiology*, 269(2 Pt 2):H571–582, August 1995. PMID: 7653621.
- [111] John Y. Lin, Michael Z. Lin, Paul Steinbach, and Roger Y. Tsien. Characterization of engineered channelrhodopsin variants with improved properties and kinetics. *Biophysical Journal*, 96(5):1803–1814, March 2009.
- [112] Fred V. Lionetti, Andrew D. McCulloch, and Scott B. Baden. Source-to-source optimization of CUDA c for GPU accelerated cardiac cell modeling. In *Proceedings of the 16th international Euro-Par conference on Parallel processing: Part I*, EuroPar’10, page 3849, Berlin, Heidelberg, 2010. Springer-Verlag.
- [113] Donald Lloyd-Jones, Robert J. Adams, Todd M. Brown, Mercedes Carnethon, Shifan Dai, Giovanni De Simone, T. Bruce Ferguson, Earl Ford, Karen Furie, Cathleen Gillespie, Alan Go, Kurt Greenlund, Nancy Haase, Susan Hailpern, P. Michael Ho, Virginia Howard, Brett Kissela, Steven Kittner, Daniel Lackland, Lynda Lisabeth, Ariane Marelli, Mary M. McDermott, James Meigs, Dariush Mozaffarian, Michael Mussolino, Graham Nichol, Veronique Roger, Wayne Rosamond, Ralph Sacco, Paul Sorlie, Randell Stafford, Thomas Thom, Sylvia Wasserthiel-Smoller, Nathan D. Wong, Judith Wylie-Rosett, on behalf of the

- American Heart Association Statistics Committee, and Stroke Statistics Subcommittee. Heart disease and stroke statistics 2010 update. a report from the american heart association. *Circulation*, December 2009.
- [114] Y. T. Lu, H. X. Zhu, S. Richmond, and J. Middleton. Modelling skeletal muscle fibre orientation arrangement. *Computer Methods in Biomechanics and Biomedical Engineering*, 14(12):1079–1088, 2011.
- [115] Tenneille E. Ludwig, Veit Bergendahl, Mark E. Levenstein, Junying Yu, Mitchell D. Probasco, and James A. Thomson. Feeder-independent culture of human embryonic stem cells. *Nature Methods*, 3(8):637–646, August 2006.
- [116] C. H. Luo and Y. Rudy. A model of the ventricular cardiac action potential. depolarization, repolarization, and their interaction. *Circulation Research*, 68(6):1501–1526, June 1991.
- [117] CH Luo and Y Rudy. A dynamic model of the cardiac ventricular action potential. i. simulations of ionic currents and concentration changes. *Circ Res*, 74(6):1071–1096, June 1994.
- [118] Jerrold E. Marsden and Thomas J. R. Hughes. *Mathematical Foundations of Elasticity*. Courier Dover Publications, February 1994.
- [119] A Mathur and JF Martin. Stem cells and repair of the heart. *The Lancet*, 364(9429):183–192, October 2004.
- [120] Akemi Matsuno-Yagi and Yasuo Mukohata. Two possible roles of bacteriorhodopsin; a comparative study of strains of halobacterium halobium differing in pigmentation. *Biochemical and Biophysical Research Communications*, 78(1):237–243, September 1977.
- [121] Nelson Max. Weights for computing vertex normals from facet normals. *J. Graph. Tools*, 4(2):16, March 1999.

- [122] R E McAllister, D Noble, and R W Tsien. Reconstruction of the electrical activity of cardiac purkinje fibres. *The Journal of Physiology*, 251(1):1–59, September 1975. PMID: 1185607 PMCID: PMC1348375.
- [123] Andrew D. McCulloch and Jeffrey H. Omens. Myocyte shearing, myocardial sheets, and microtubules. *Circulation Research*, 98(1):1–3, January 2006.
- [124] Philippe Menasch, Albert A Hagge, Jean-Thomas Vilquin, Michel Desnos, Eric Abergel, Bruno Pouzet, Alain Bel, Sorin Sarateanu, Marcio Scorsin, Ketty Schwartz, Patrick Bruneval, Marc Benbunan, Jean-Pierre Marolleau, and Denis Duboc. Autologous skeletal myoblast transplantation for severe postinfarction left ventricular dysfunction. *Journal of the American College of Cardiology*, 41(7):1078–1083, April 2003.
- [125] C. Miehe. Aspects of the formulation and finite element implementation of large strain isotropic elasticity. *International Journal for Numerical Methods in Engineering*, 37(12):1981–2004, 1994.
- [126] L. Mirabella, F. Nobile, and A. Veneziani. An a posteriori error estimator for model adaptivity in electrocardiology. *Computer Methods in Applied Mechanics and Engineering*, 200(3740):2727–2737, September 2011.
- [127] Michele Miragoli, Giedrius Gaudesius, and Stephan Rohr. Electrotonic modulation of cardiac impulse conduction by myofibroblasts. *Circulation Research*, 98(6):801–810, March 2006.
- [128] Alexander Monakov and Arutyun Avetisyan. Implementing blocked sparse matrix-vector multiplication on NVIDIA GPUs. In Koen Bertels, Nikitas Dimopoulos, Cristina Silvano, and Stephan Wong, editors, *Embedded Computer Systems: Architectures, Modeling, and Simulation*, volume 5657 of *Lecture Notes in Computer Science*, pages 289–297. Springer Berlin / Heidelberg, 2009.
- [129] Christine Mummery, Dorien Ward-van Oostwaard, Pieter Doevendans, Rene Spijker, Stieneke van den Brink, Rutger Hassink, Marcel van der Heyden, Tobias

- Opthof, Martin Pera, Aart Brutel de la Riviere, Robert Passier, and Leon Tertoolen. Differentiation of human embryonic stem cells to cardiomyocytes role of coculture with visceral Endoderm-Like cells. *Circulation*, 107(21):2733–2740, June 2003.
- [130] Marilena Munteanu and Luca F. Pavarino. DECOUPLED SCHWARZ ALGORITHMS FOR IMPLICIT DISCRETIZATIONS OF NONLINEAR MONODOMAIN AND BIDOMAIN SYSTEMS. *Mathematical Models and Methods in Applied Sciences*, 19(07):1065, 2009.
- [131] Maria Murillo and Xiao-Chuan Cai. A fully implicit parallel algorithm for simulating the nonlinear electrical activity of the heart. *Numerical Linear Algebra with Applications*, 11(23):261–277, March 2004.
- [132] G. Nagel. Channelrhodopsin-1: A Light-Gated proton channel in green algae. *Science*, 296(5577):2395–2398, June 2002.
- [133] Georg Nagel, Tanjef Szellas, Wolfram Huhn, Suneel Kateriya, Nona Adeishvili, Peter Berthold, Doris Ollig, Peter Hegemann, and Ernst Bamberg. Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proceedings of the National Academy of Sciences of the United States of America*, 100(24):13940–13945, November 2003. PMID: 14615590 PMCID: 283525.
- [134] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, October 1962.
- [135] Martyn P. Nash and Alexander V. Panfilov. Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias. *Progress in Biophysics and Molecular Biology*, 85(2-3):501–522, June 2004.
- [136] Steven A. Niederer and Nicolas P. Smith. An improved numerical method for strong coupling of excitation and contraction models in the heart. *Progress in Biophysics and Molecular Biology*, 96(13):90–111, January 2008.

- [137] P. M. Nielsen, I. J. Le Grice, B. H. Smaill, and P. J. Hunter. Mathematical model of geometry and fibrous structure of the heart. *American Journal of Physiology - Heart and Circulatory Physiology*, 260(4):H1365–H1378, April 1991.
- [138] K. Nikolic, P. Degenaar, and C. Toumazou. Modeling and engineering aspects of ChannelRhodopsin2 system for neural photostimulation. In *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, pages 1626–1629, 2006.
- [139] Konstantin Nikolic, Nir Grossman, Matthew S Grubb, Juan Burrone, Chris Toumazou, and Patrick Degenaar. Photocycles of channelrhodopsin2. *Photochemistry and Photobiology*, 85(1):400–411, January 2009.
- [140] D. Noble. A modification of the HodgkinHuxley equations applicable to purkinje fibre action and pacemaker potentials. *The Journal of Physiology*, 160(2):317–352, February 1962. PMID: 14480151 PMCID: PMC1359535.
- [141] NVIDIA. Cuda toolkit 4.2 cusparse library, 2012.
- [142] Dieter;stoeckenius Oesterhelt. Rhodopsin-like protein from the purple membrane of halobacterium halobium. *Nature*, 233(39):149–152, September 1971.
- [143] Lionel H. Opie. *The Heart: Physiology, from Cell to Circulation*. Lippincott Williams & Wilkins, 3 sub edition, April 2003.
- [144] Amit N. Patel, Luis Geffner, Roberto F. Vina, Jorge Saslavsky, Harold C. Urschel Jr, Robert Kormos, and Federico Benetti. Surgical treatment for congestive heart failure with autologous adult stem cell transplantation: A prospective randomized study. *The Journal of Thoracic and Cardiovascular Surgery*, 130(6):1631–1638.e2, December 2005.
- [145] Pras Pathmanathan, Gary R. Mirams, James Southern, and Jonathan P. Whiteley. The significant effect of the choice of ionic current integration method in cardiac electro-physiological simulations. *International Journal for Numerical Methods in Biomedical Engineering*, 27(11):17511770, 2011.

- [146] Gernot Plank, Rebecca A. B. Burton, Patrick Hales, Martin Bishop, Tahir Mansoori, Miguel O. Bernabeu, Alan Garny, Anton J. Prassl, Christian Bollensdorff, Fleur Mason, Fahd Mahmood, Blanca Rodriguez, Vicente Grau, Jrgen E. Schneider, David Gavaghan, and Peter Kohl. Generation of Histo-Anatomically representative models of the individual heart: Tools and application. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1896):2257–2292, June 2009.
- [147] Robert Plonsey and Roger C. Barr. *Bioelectricity: A Quantitative Approach*. Springer, 2000.
- [148] Guillem Pratx and Lei Xing. GPU computing in medical physics: A review. *Medical Physics*, 38(5):2685–2697, 2011.
- [149] Leo Priebe and Dirk J. Beuckelmann. Simulation study of cellular electric properties in heart failure. *Circulation Research*, 82(11):1206–1223, June 1998.
- [150] S. Puwal and B.J. Roth. Optimization of feedback pacing for defibrillation. *Biomedical Engineering, IEEE Transactions on*, 56(2):532–534, February 2009.
- [151] Zhilin Qu and A. Garfinkel. An advanced algorithm for solving partial differential equation in cardiac conduction. *Biomedical Engineering, IEEE Transactions on*, 46(9):1166–1168, September 1999.
- [152] Charulatha Ramanathan, Ping Jia, Raja Ghanem, Kyungmoo Ryu, and Yoram Rudy. Activation and repolarization of the normal human heart under complete physiological conditions. *Proceedings of the National Academy of Sciences of the United States of America*, 103(16):6309–6314, April 2006. PMID: 16606830 PMCID: PMC1458874.
- [153] John Jeremy Rice, Fei Wang, Donald M. Bers, and Pieter P. de Tombe. Approximate model of cooperative activation and crossbridge cycling in cardiac muscle using ordinary differential equations. *Biophysical Journal*, 95(5):2368–2390, September 2008.

- [154] J M Rogers and A D McCulloch. A collocation–Galerkin finite element model of cardiac action potential propagation. *IEEE transactions on bio-medical engineering*, 41(8):743–757, August 1994. PMID: 7927397.
- [155] Jack M. Rogers. Wave front fragmentation due to ventricular geometry in a model of the rabbit heart. *Chaos (Woodbury, N.Y.)*, 12(3):779–787, September 2002. PMID: 12779606.
- [156] S. Rohr, D. M. Schly, and A. G. Klber. Patterned growth of neonatal rat heart cells in culture. morphological and electrophysiological characterization. *Circulation Research*, 68(1):114–130, January 1991.
- [157] Luca Romero, Esther Pueyo, Martin Fink, and Blanca Rodriguez. Impact of ionic current variability on human ventricular cellular electrophysiology. *American Journal of Physiology - Heart and Circulatory Physiology*, 297(4):H1436–H1445, October 2009.
- [158] M. B. Rook, A. C. van Ginneken, B. de Jonge, A. el Aoumari, D. Gros, and H. J. Jongsma. Differences in gap junction channels between cardiac myocytes, fibroblasts, and heterologous pairs. *American Journal of Physiology - Cell Physiology*, 263(5):C959–C977, November 1992.
- [159] Douglas A Rubinson, Christopher P Dillon, Adam V Kwiatkowski, Claudia Sievers, Lili Yang, Johnny Kopinja, Dina L Rooney, Mingdi Zhang, Melanie M Ihrig, Michael T McManus, Frank B Gertler, Martin L Scott, and Luk Van Parijs. A lentivirus-based system to functionally silence genes in primary mammalian cells, stem cells and transgenic mice by RNA interference. *Nature genetics*, 33(3):401–406, March 2003. PMID: 12590264.
- [160] S Rush and H Larsen. A practical algorithm for solving dynamic membrane equations. *IEEE transactions on bio-medical engineering*, 25(4):389–392, July 1978. PMID: 689699.

- [161] J. Sainte-Marie, D. Chapelle, R. Cimrman, and M. Sorine. Modeling and estimation of the cardiac electromechanical activity. *Computers & Structures*, 84(28):1743–1759, November 2006.
- [162] S. Scacchi. A multilevel hybrid NewtonKrylovSchwarz method for the bidomain model of electrocardiology. *Computer Methods in Applied Mechanics and Engineering*, 200(58):717–725, January 2011.
- [163] Simone Scacchi. A hybrid multilevel schwarz method for the bidomain model. *Computer Methods in Applied Mechanics and Engineering*, 197(4548):4051–4061, August 2008.
- [164] Boris Schmitt, Katsiaryna Fedarava, Jan Falkenberg, Kai Rothaus, Narendra K Bodhey, Carolin Reischauer, Sebastian Kozerke, Bernhard Schnackenburg, Dirk Westermann, Paul P Lunkenheimer, Robert H Anderson, Felix Berger, and Titus Kuehne. Three-dimensional alignment of the aggregated myocytes in the normal and hypertrophic murine heart. *Journal of applied physiology (Bethesda, Md.: 1985)*, 107(3):921–927, September 2009. PMID: 19628727.
- [165] Miguel Sena-Esteves, Jessica C. Tebbets, Sabine Steffens, Timothy Crombleholme, and Alan W. Flake. Optimized large-scale production of high titer lentivirus vector pseudotypes. *Journal of Virological Methods*, 122(2):131–139, December 2004.
- [166] Shubhabrata Sengupta, Mark Harris, Yao Zhang, and John D. Owens. Scan primitives for GPU computing. In *Graphics Hardware 2007*, pages 97–106, August 2007.
- [167] M. Sermesant, K. Rhode, G.I. Sanchez-Ortiz, O. Camara, R. Andriantsimiavona, S. Hegde, D. Rueckert, P. Lambiase, C. Bucknall, E. Rosenthal, H. Delingette, D.L.G. Hill, N. Ayache, and R. Razavi. Simulation of cardiac pathologies using an electromechanical biventricular model and XMR interventional imaging. *Medical Image Analysis*, 9(5):467–480, October 2005.

- [168] J.C. Simo. A framework for finite strain elastoplasticity based on maximum plastic dissipation and the multiplicative decomposition. part II: computational aspects. *Computer Methods in Applied Mechanics and Engineering*, 68(1):1–31, May 1988.
- [169] Juan C. Simo and Robert L. Taylor. Quasi-incompressible finite elasticity in principal stretches. continuum basis and numerical algorithms. *Computer Methods in Applied Mechanics and Engineering*, 85(3):273–310, February 1991.
- [170] J. H. Smith, C. R. Green, N. S. Peters, S. Rothery, and N. J. Severs. Altered patterns of gap junction distribution in ischemic heart disease. an immunohistochemical study of human myocardium using laser scanning confocal microscopy. *The American Journal of Pathology*, 139(4):801–821, October 1991. PMID: 1656760 PMCID: PMC1886321.
- [171] David E. Sosnovik, Ruopeng Wang, Guangping Dai, Teresa Wang, Elena Aikawa, Mikhail Novikov, Anthony Rosenzweig, Richard J. Gilbert, and Van J. Wedeen. Diffusion spectrum MRI tractography reveals the presence of a complex network of residual myofibers in infarcted myocardium. *Circulation: Cardiovascular Imaging*, March 2009.
- [172] Ms Spach, Wt Miller, Db Geselowitz, Rc Barr, Jm Kootsey, and Ea Johnson. The discontinuous nature of propagation in normal canine cardiac muscle. evidence for recurrent discontinuities of intracellular resistance that affect the membrane currents. *Circulation Research*, 48(1):39–54, January 1981.
- [173] Daniel D. Streeter and William T. Hanna. Engineering mechanics for successive states in canine left ventricular myocardium: II. fiber angle and sarcomere length. *Circulation Research*, 33(6):656–664, December 1973.
- [174] J. Sundnes, G.T. Lines, K.A. Mardal, and A. Tveito. Multigrid block preconditioning for a coupled system of partial differential equations modeling the electrical activity in the heart. *Computer Methods in Biomechanics and Biomedical Engineering*, 5(6):397–409, 2002.

- [175] Joakim Sundnes, Glenn Terje Lines, and Aslak Tveito. An operator splitting method for solving the bidomain equations coupled to a volume conductor model for the torso. *Mathematical Biosciences*, 194(2):233–248, April 2005.
- [176] Stephanie Szobota and Ehud Y Isacoff. Optical control of neuronal activity. *Annual review of biophysics*, 39:329–348, 2010. PMID: 20192766.
- [177] Kenshi Takayama, Takashi Ashihara, Takashi Ijiri, Takeo Igarashi, Ryo Haraguchi, and Kazuo Nakazawa. A sketch-based interface for modeling myocardial fiber orientation that considers the layered structure of the ventricles. *The journal of physiological sciences: JPS*, 58(7):487–492, December 2008. PMID: 18926007.
- [178] Z.A. Taylor, O. Comas, M. Cheng, J. Passenger, D.J. Hawkes, D. Atkinson, and S. Ourselin. On modelling of anisotropic viscoelasticity for soft tissue simulation: Numerical solution and GPU execution. *Medical Image Analysis*, 13(2):234–244, April 2009.
- [179] E.K. Theofilogiannakos, G.K. Theofilogiannakos, P.G. Danias, T.V. Yioultis, A. Anogeianaki, V. Stergiou-Michailidou, K. Kallaras, T. Xenos, and G. Anogianakis. A semi-automated approach towards generating three-dimensional mesh of the heart using a hybrid MRI/histology database. *Computer Methods in Biomechanics and Biomedical Engineering*, 14(4):349–358, 2011.
- [180] C.A. Thomas Jr., P.A. Springer, G.E. Loeb, Y. Berwald-Netter, and L.M. Okun. A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Experimental Cell Research*, 74(1):61–66, September 1972.
- [181] Nicolas Toussaint, Maxime Sermesant, Christian T. Stoeck, Sebastian Kozerke, and Philip G. Batchelor. In vivo human 3D cardiac fibre architecture: Reconstruction using curvilinear interpolation of diffusion tensor images. In Tianzi

- Jiang, Nassir Navab, Josien P. W. Pluim, and Max A. Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention MICCAI 2010*, volume 6361, pages 418–425. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [182] Nicolas Toussaint, Christian T. Stoeck, Maxime Sermesant, Sebastian Kozerke, and Philip G. Batchelor. Three-dimensional prolate spheroidal extrapolation for sparse dti of the in-vivo heart. In *Book of Abstracts, International Society in Magnetic Resonance in Medicine (ISMRM'10)*, Stockholm, 2-7 June 2010.
- [183] Natalia A. Trayanova, Jason Constantino, and Viatcheslav Gurev. Electromechanical models of the ventricles. *American Journal of Physiology - Heart and Circulatory Physiology*, 301(2):H279–H286, August 2011.
- [184] Chin-Feng Tsai, Ching-Tai Tai, Ming-Hsiung Hsieh, Wei-Shiang Lin, Wen-Chung Yu, Kwo-Chang Ueng, Yu-An Ding, Mau-Song Chang, and Shih-Ann Chen. Initiation of atrial fibrillation by ectopic beats originating from the superior vena cava : Electrophysiological characteristics and results of radiofrequency ablation. *Circulation*, 102(1):67–74, July 2000.
- [185] Alkiviadis Tsamis, Wolfgang Bothe, John-Peder Escobar Kvitting, Julia C. Swanson, D. Craig Miller, and Ellen Kuhl. Active contraction of cardiac muscle: In vivo characterization of mechanical activation sequences in the beating heart. *Journal of the Mechanical Behavior of Biomedical Materials*, 4(7):1167–1176, October 2011.
- [186] Alkiviadis Tsamis, Wolfgang Bothe, John-Peder Escobar Kvitting, Julia C Swanson, D Craig Miller, and Ellen Kuhl. Active contraction of cardiac muscle: in vivo characterization of mechanical activation sequences in the beating heart. *Journal of the mechanical behavior of biomedical materials*, 4(7):1167–1176, October 2011. PMID: 21783125.
- [187] Satoshi P. Tsunoda, David Ewers, Sabrina Gazzarrini, Anna Moroni, Dietrich Gradmann, and Peter Hegemann. H⁺-Pumping rhodopsin from the marine alga acetabularia. *Biophysical Journal*, 91(4):1471–1479, August 2006.

- [188] K. H. W. J. ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov. A model for human ventricular tissue. *American Journal of Physiology - Heart and Circulatory Physiology*, 286(4):H1573–H1589, April 2004.
- [189] K.H.W.J. Ten Tusscher and A.V. Panfilov. Modelling of the ventricular conduction system. *Progress in Biophysics and Molecular Biology*, 96(13):152–170, April 2008.
- [190] Edward J Vigmond, Patrick M Boyle, L Leon, and Gernot Plank. Near-real-time simulations of bioelectric activity in small mammalian hearts using graphical processing units. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, 2009:3290–3293, 2009. PMID: 19964295.
- [191] Stephen C Vlay. Complications of active-fixation electrodes. *Pacing and clinical electrophysiology: PACE*, 25(8):1153–1154, August 2002. PMID: 12358162.
- [192] Richard W. Vuduc and Hyun-Jin Moon. Fast sparse matrix-vector multiplication by exploiting variable block structure. In *Proceedings of the First international conference on High Performance Computing and Communications, HPC'05*, page 807816, Berlin, Heidelberg, 2005. Springer-Verlag.
- [193] Jonathan F. Wenk, Parastou Eslami, Zhihong Zhang, Chun Xu, Ellen Kuhl, Joseph H. Gorman III, J. Daniel Robb, Mark B. Ratcliffe, Robert C. Gorman, and Julius M. Guccione. A novel method for quantifying the In-Vivo mechanical effect of material injected into a myocardial infarction. *The Annals of Thoracic Surgery*, 92(3):935–941, September 2011.
- [194] R.H. Whittington, M.Q. Chen, L. Giovannardi, and G.T.A. Kovacs. Temporal resolution of stimulation threshold: A tool for electrophysiologic analysis. In *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, pages 3891–3894, September 2006.

- [195] R.H. Whittington, L. Giovangrandi, and G.T.A. Kovacs. A closed-loop electrical stimulation system for cardiac cell cultures. *Biomedical Engineering, IEEE Transactions on*, 52(7):1261–1270, July 2005.
- [196] R.H. Whittington and G.T.A. Kovacs. A Discrete-Time control algorithm applied to Closed-Loop pacing of HL-1 cardiomyocytes. *Biomedical Engineering, IEEE Transactions on*, 55(1):21–30, January 2008.
- [197] Jonathan Wong, Serdar Göktepe, and Ellen Kuhl. Computational modeling of electrochemical coupling: A novel finite element approach towards ionic models for cardiac electrophysiology. *Computer Methods in Applied Mechanics and Engineering*, 200(4546):3139–3158, October 2011.
- [198] Yuanfang Xie, Alan Garfinkel, Patrizia Camelliti, Peter Kohl, James N. Weiss, and Zhilin Qu. Effects of fibroblast-myocyte coupling on cardiac conduction and vulnerability to reentry: A computational study. *Heart Rhythm*, 6(11):1641–1649, November 2009.
- [199] Feng Yang, Yue-Min Zhu, Isabelle E. Magnin, Jian-Hua Luo, Pierre Croisille, and Peter B. Kingsley. Feature-based interpolation of diffusion tensor fields and application to human cardiac DT-MRI. *Medical Image Analysis*, 16(2):459–481, February 2012.
- [200] Lei Yang, Mark H. Soonpaa, Eric D. Adler, Torsten K. Roepke, Steven J. Kattman, Marion Kennedy, Els Henckaerts, Kristina Bonham, Geoffrey W. Abbott, R. Michael Linden, Loren J. Field, and Gordon M. Keller. Human cardiovascular progenitor cells develop from a KDR+ embryonic-stem-cell-derived population. *Nature*, 453(7194):524–528, April 2008.
- [201] F C Yin, C C Chan, and R M Judd. Compressibility of perfused passive myocardium. *The American journal of physiology*, 271(5 Pt 2):H1864–1870, November 1996. PMID: 8945902.

- [202] Feng Zhang, Li-Ping Wang, Edward S. Boyden, and Karl Deisseroth. Channelrhodopsin-2 and optical control of excitable cells. *Nature Methods*, 3(10):785–792, October 2006.
- [203] Feng Zhang, Li-Ping Wang, Martin Brauner, Jana F. Liewald, Kenneth Kay, Natalie Watzke, Phillip G. Wood, Ernst Bamberg, Georg Nagel, Alexander Gottschalk, and Karl Deisseroth. Multimodal fast optical interrogation of neural circuitry. *Nature*, 446(7136):633–639, April 2007.
- [204] Sharon Zlochiver, Viviana Muoz, Karen L. Vikstrom, Steven M. Taffet, Omer Berenfeld, and Jos Jalife. Electrotonic Myofibroblast-to-Myocyte coupling increases propensity to reentrant arrhythmias in Two-Dimensional cardiac monolayers. *Biophysical Journal*, 95(9):4469–4480, November 2008.
- [205] Steven W. Zucker and Demetri Terzopoulos. Finding structure in co-occurrence matrices for texture analysis. *Computer Graphics and Image Processing*, 12(3):286–308, March 1980.

Jonathan Joi-Mun Wong, Ph.D. Candidate

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ellen Kuhl) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Peter Pinsky)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Eric Darve)

Approved for the University Committee on Graduate Studies
