

Proceedings of the FEniCS Conference 2017

University of Luxembourg, 12-14 June 2017

Post-conference version. 10.6084/m9.figshare.5086369.v2

Licensed under a CC BY-SA 4.0 License.
<https://creativecommons.org/licenses/by-sa/4.0/>

Copyright remains with the authors of each article.

Jack S. Hale, Editor.

#	Authors	Title
1	Martin Sandve Alnæs, Vidar Fauske and Min Ragan-Kelley	3D visualization with FEniCS in Jupyter Notebooks
2	Martin Sandve Alnæs	FEniCS: Sustainable Software Development Practices
3	Henning Bonart, Christian Kahle and Jens-Uwe Repke	Simulating Moving Contact Line Problems Using the Cahn-Hilliard-Navier-Stokes Equations
4	Wietse Boon and Jan Nordbotten	Mixed-Dimensional Linear Elasticity with Relaxed Symmetry
5	Jan Blechta and Martin Řehoř	FENaPack - FEniCS Navier-Stokes preconditioning package
6	Davide Baroli, Lars A. A. Beex, Jack S. Hale and Stéphane P. A. Bordas	iRANitsche: reduced assembly with a Nitsche-based reduced basis method
7	Michele Cascio, Ioannis Derezitis, Giuseppe Falci, Salvatore Francesco Lombardo, Antonino Magliano, Silvia	MD-FEM Force Field Simulations of Particle Subjected to Dielectrophoresis Interactions
8	Manuel Cremades Buján, Jerónimo Rodríguez García and Pedro Fontán Muíños	Automated Solution of Parabolic-Elliptic Systems for Control Applications with Embedded Runge-Kutta Methods
9	Joshua W. Chen, Andrew Drach, Umberto Villa, Reza Avazmohammadi, David Li, Omar Ghattas and Michael S. Sacks	Identification of Mechanical Properties of 3D Myocardial Tissue: An Inverse Modeling and Optimal Experimental Design Problem
10	Diako Darian	Kinetic simulation of plasma-object interaction with FEniCS
11	Jørgen S. Dokken, Simon W. Funke, August Johansson and Stephan Schmidt	Shape Optimization with Multiple Meshes
12	Cécile Daversin-Catty and Marie E. Rognes	Towards coupled mixed dimensional finite elements in FEniCS
13	Patrick Farrell and Matteo Croci	Geometric Multigrid in FEniCS
14	Simon Funke and Sebastian Mitusch	A new algorithmic differentiation tool (not only) for FEniCS
15	Thomas Gibson, David Ham and Colin Cotter	Automating Hybridization via Composing Abstractions
16	Christophe Geuzaine and Jean-François Remacle	Gmsh 3.0 - Gmsh goes boolean!

17	Benjamin Goldsberry and Michael Haberman	Finite Element Modeling of Elastic Wave Dispersion in Pre-strained Negative Stiffness Honeycombs Using FEniCS
18	Martin Genet, Lik Chuan Lee and Sebastian Kozerke	Equilibrated Warping: Finite Element Image Correlation with Mechanical Regularization
19	Hugo Hadfield, Juliette Unwin, Nathan Sime and Garth Wells	podS: A Parallel Multilevel Monte Carlo Framework
20	Miklós Homolya, Fabio Luporini, David Ham and Robert C. Kirby	FinAT: Automating Optimal Code Generation for High-Order Finite Element Methods
21	Karl Erik Holter, Miroslav Kuchta and Kent-Andre Mardal	Tracer transport in microvasculature: A case study on coupling 1D-3D problems in FEniCS
22	Marc Herrmann, Roland Herzog, Heiko Kröner, José Vidal and Stephan Schmidt	Total Variation Image Reconstruction on Smooth Surfaces
23	Michal Habara, Ondřej Souček, Josef Málek and Jaroslav Hron	Discontinuous solution of the Laplace equation in a mixed finite element method and level-set setting
24	Paul Hauseux, Jack S. Hale and Stéphane P.A. Bordas	Uncertainty quantification for soft tissue biomechanics
25	Karl Erik Holter, Miroslav Kuchta and Kent-Andre Mardal	Trace Constrained Problems in FEniCS
26	Tormod Landet, Mikael Mortensen and Kent-Andre Mardal	Slope Limiting of Divergence Free Discontinuous Galerkin Vector Fields in the Context of Two-Phase Flows
27	Anders Logg, Carl Lundholm and Magne Nordaas	Solving Poisson's equation on the Microsoft HoloLens
28	Marco Morandini	Handling of finite rotations in Dolfin
29	Lawrence Mitchell and Robert Kirby	Across the great divide: composable block preconditioning from UFL
30	Nicolas Marsic, Herbert De Gersem and Christophe Geuzaine	Computationally Efficient Numerical Quadratures for the Finite Element Assembly of High-Order non-Lagrange Elements
31	Jakob Maljaars, Robert Jan Labeur and Matthias Möller	A High-Order Particle-Mesh Operator Splitting Approach for the Incompressible Navier-Stokes Equations
32	Antonino La Magna, Salvatore Francesco Lombardo, Ioannis Deretzis, Armand Verstraete, Bobby Lespinasse and Karim	LIAB: a FEniCS based computational tool for laser annealing simulation
33	Hernán Mella, Joaquín Mura and Esteban Sáez	Implementation of Mixed and Hybrid Formulation of the PML Method in Elastodynamics Using FEniCS

34	Johannes Neumann	Topology Optimization under Uncertainties in FEniCS
35	Anders Logg, Carl Lundholm and Magne Nordaas	A Space-Time Cut Finite Element Method for the Heat Equation in FEniCS
36	Liubov Nikitushkina, Simon Funke, Henrik Finsberg, Lik Chuan Lee and Samuel Wall	Sensitivity Analysis of Cardiac Growth Models
37	Mariya Ptashnyk and Brian Seguin	The impact of microscopic structure on mechanical properties of plant cell walls and tissues
38	Chris Richardson, Jack Hale and Garth Wells	Uncontainable enthusiasm: the pain and joy of Docker on HPC and cloud
39	Marie Rognes	Accurate Numerical Modeling of Small Collections of Cardiac Cells
40	Miloslav Feistauer, Filip Roskovec and Anna-Margarete Sändig	Optimality Of Apriori Error Estimates Of Discontinuous Galerkin Method For Elliptic Problems With Nonlinear Newton Boundary
41	Miguel Rodriguez, Christoph Augustin and Shawn Shadden	FEniCS Mechanics: A Package for Continuum Mechanics Simulations
42	Ajay Mandyam Rangarajan, Georg May and Vit Dolejsi	Analytic Metric-Based Adaptation Using a Continuous Mesh Model
43	Ulrich Römer, Sebastian Schöps and Herbert De Gersen	A Finite Element Approach to Computational Magnetism with Defect Correction
44	Tianjiao Sun and Fabio Luporini	Loop fusion for finite element assembly in PyOP2
45	Martin Řehoř and Jan Blechta	Implementation of Diffuse Interface Models in FEniCS
46	Stephan Schmidt	UFL and the Automatic Derivation of Weak and Strong Shape Derivatives
47	Vahid Shahrabi, Davide Baroli, Jack Hale, Stéphane Bordas, Paul Hauseux and Frank Hertel	From medical images to uncertainty quantification with FEniCS and 3D Slicer
48	Satyendra Tomar, Jack S. Hale and Stéphane P. A. Bordas	Functional-type two-sided a posteriori error estimates for high-order discretizations
49	Juliette Unwin, Garth Wells, Nathan Sime and Hugo Hadfield	Multi-Level Monte Carlo for Large-scale Vibrations Problems Using FEniCS and podS
50	Maria Vasilyeva	Multiscale model reduction using GMsFEM for applied problems in heterogeneous domains

51	Anders Logg	In memory of Hans Petter Langtangen
52	Florian Wechsung and Patrick Farrell	Shape Optimization with Geometric Constraints
53	Miklós Homolya, Andreas Klöckner and Robert C. Kirby	Generating matrix actions with TSFC and Loo.py
54	Yang Zhou, Paul Lacko and Martine Baelmans	Grid Independence and Manufacturability in Microchannel Heat Sink Topology Optimization
55	Will Zhang and Michael Sacks	Simulating the Response of Bioprosthetics Heart Valve Tissues to Cyclic Loading in FEniCS

FENICS'17: THE FENICS 2017 CONFERENCE

PROGRAM INDEXES

PROGRAM

Days: [Monday, June 12th](#) [Tuesday, June 13th](#) [Wednesday, June 14th](#)

Monday, June 12th

View this program: [with abstracts](#) [session overview](#) [talk overview](#)

12:00-13:00 Registration and welcome lunch

13:00-13:10 Session 1: Welcome session

CHAIR: [Jack S. Hale](#)

13:10-14:10 Session 2: Software and Algorithms I

CHAIR: [Stephan Schmidt](#)

LOCATION: Salle Tavenas Main Theatre

13:10 [Miklós Homolya](#), [Andreas Klöckner](#) and [Robert C. Kirby](#)
Generating matrix actions with TSFC and Loo.py ([abstract](#))

13:25 [Nicolas Marsic](#), [Herbert De Gersen](#) and [Christophe Geuzaine](#)
Computationally Efficient Numerical Quadratures for the Finite Element Assembly of High-Order non-Lagrange Elements ([abstract](#))

13:40 [Miklós Homolya](#), [Fabio Luporini](#), [David Ham](#) and [Robert C. Kirby](#)
FInAT: Automating Optimal Code Generation for High-Order Finite Element Methods ([abstract](#))

13:55 [Thomas Gibson](#), [David Ham](#) and [Colin Cotter](#)
Automating Hybridization via Composing Abstractions ([abstract](#))

14:10-14:40 Coffee Break

14:40-15:25 Session 3: Software and Algorithms II

CHAIR: [Martin Genet](#)

LOCATION: Salle Tavenas Main Theatre

14:40 [Christophe Geuzaine](#) and [Jean-François Remacle](#)
Gmsh 3.0 - Gmsh goes boolean! ([abstract](#))

14:55 [Patrick Farrell](#) and [Matteo Croci](#)
Geometric Multigrid in FEniCS ([abstract](#))

15:10 [Lawrence Mitchell](#) and [Robert Kirby](#)
Across the great divide: composable block preconditioning from UFL ([abstract](#))

15:25-15:40 Short break, poster preparation

15:40-16:10 Session 4: Magnetics

CHAIR: [Simon Funke](#)

LOCATION: Salle Tavenas Main Theatre

15:40 [Antonino La Magna](#), [Salvatore Francesco Lombardo](#), [Ioannis Deretzis](#), [Armand Verstraete](#), [Bobby Lespinasse](#) and [Karim Huet](#)
LIAB: a FEniCS based computational tool for laser annealing simulation ([abstract](#))

15:55 [Ulrich Römer](#), [Sebastian Schöps](#) and [Herbert De Gersen](#)
A Finite Element Approach to Computational Magnetics with Defect Correction ([abstract](#))

16:10-16:50 Session 5: In memory of Hans Petter Langtangen

CHAIR: [Patrick Farrell](#)

LOCATION: Salle Tavenas Main Theatre

16:10 [Anders Logg](#)
In memory of Hans Petter Langtangen ([abstract](#))

16:50-18:30 Session 6: Poster Session and Refreshments

CHAIR: [Jack S. Hale](#)

LOCATION: Salle Tavenas Rear Room

16:50 [Davide Baroli](#), [Lars A. A. Beex](#), [Jack S. Hale](#) and [Stéphane P. A. Bordas](#)
iRANitsche: reduced assembly with a Nitsche-based reduced basis method ([abstract](#))

16:50 [Michele Cascio](#), [Ioannis Deretzis](#), [Giuseppe Falci](#), [Salvatore Francesco Lombardo](#), [Antonino Magliano](#), [Silvia Scalese](#) and [Antonino La Magna](#)
MD-FEM Force Field Simulations of Particle Subjected to Dielectrophoresis Interactions ([abstract](#))

- 16:50 [Joshua W. Chen](#), [Andrew Drach](#), [Umberto Villa](#), [Reza Avazmohammadi](#), [David Li](#), [Omar Ghattas](#) and [Michael S. Sacks](#)
Identification of Mechanical Properties of 3D Myocardial Tissue: An Inverse Modeling and Optimal Experimental Design Problem ([abstract](#))
- 16:50 [Diako Darian](#)
Kinetic simulation of plasma-object interaction with FEniCS ([abstract](#))
- 16:50 [Miloslav Feistauer](#), [Filip Roskovec](#) and [Anna-Margarete Sändig](#)
Optimality Of Apriori Error Estimates Of Discontinuous Galerkin Method For Elliptic Problems With Nonlinear Newton Boundary Conditions ([abstract](#))
- 16:50 [Simon Funke](#) and [Sebastian Mitusch](#)
A new algorithmic differentiation tool (not only) for FEniCS ([abstract](#))
- 16:50 [Michal Habara](#), [Ondřej Souček](#), [Josef Málek](#) and [Jaroslav Hron](#)
Discontinuous solution of the Laplace equation in a mixed finite element method and level-set setting ([abstract](#))
- 16:50 [Paul Hauseux](#), [Jack S. Hale](#) and [Stéphane P.A. Bordas](#)
Uncertainty quantification for soft tissue biomechanics ([abstract](#))
- 16:50 [Liubov Nikitushkina](#), [Simon Funke](#), [Henrik Finsberg](#), [Lik Chuan Lee](#) and [Samuel Wall](#)
Sensitivity Analysis of Cardiac Growth Models ([abstract](#))
- 16:50 [Jakob Maljaars](#), [Robert Jan Labeur](#) and [Matthias Möller](#)
A High-Order Particle-Mesh Operator Splitting Approach for the Incompressible Navier-Stokes Equations ([abstract](#))
- 16:50 [Martin Řehoř](#) and [Jan Blechta](#)
Implementation of Diffuse Interface Models in FEniCS ([abstract](#))
- 16:50 [Martin Sandve Alnæs](#)
FEniCS: Sustainable Software Development Practices ([abstract](#))
- 16:50 [Yahid Shahrabi](#), [Davide Baroli](#), [Jack Hale](#), [Stephane Bordas](#), [Paul Hauseux](#) and [Frank Hertel](#)
From medical images to uncertainty quantification with FEniCS and 3D Slicer ([abstract](#))
- 16:50 [Tianjiao Sun](#) and [Fabio Luporini](#)
Loop fusion for finite element assembly in PyOP2 ([abstract](#))
- 16:50 [Satyendra Tomar](#), [Jack S. Hale](#) and [Stéphane P. A. Bordas](#)
Functional-type two-sided a posteriori error estimates for high-order discretizations ([abstract](#))
- 16:50 [Yang Zhou](#), [Paul Lacko](#) and [Martine Baelmans](#)
Grid Independence and Manufacturability in Microchannel Heat Sink Topology Optimization ([abstract](#))

Tuesday, June 13th

View this program: [with abstracts](#) [session overview](#) [talk overview](#)

09:15-10:15 Session 7: Optimisation I

CHAIR: [Miklós Homolya](#)

LOCATION: Salle Tavenas Main Theatre

09:15 [Johannes Neumann](#)

Topology Optimization under Uncertainties in FEniCS ([abstract](#))

09:30 [Martin Genet](#), [Lik Chuan Lee](#) and [Sebastian Kozerke](#)

Equilibrated Warping: Finite Element Image Correlation with Mechanical Regularization ([abstract](#))

09:45 [Marc Herrmann](#), [Roland Herzog](#), [Heiko Kröner](#), [José Vidal](#) and [Stephan Schmidt](#)

Total Variation Image Reconstruction on Smooth Surfaces ([abstract](#))

10:00 [Manuel Cremades Buján](#), [Jerónimo Rodríguez García](#) and [Pedro Fontán Muiños](#)

Automated Solution of Parabolic-Elliptic Systems for Control Applications with Embedded Runge-Kutta Methods ([abstract](#))

10:15-10:45 Coffee Break

10:45-11:30 Session 8: Optimisation II

CHAIR: [Nathan Sime](#)

LOCATION: Salle Tavenas Main Theatre

10:45 [Stephan Schmidt](#)

UFL and the Automatic Derivation of Weak and Strong Shape Derivatives ([abstract](#))

11:00 [Florian Wechsung](#) and [Patrick Farrell](#)

Shape Optimization with Geometric Constraints ([abstract](#))

11:15 [Jørgen S. Dokken](#), [Simon W. Funke](#), [August Johansson](#) and [Stephan Schmidt](#)

Shape Optimization with Multiple Meshes ([abstract](#))

11:30 [Ajay Mandyam Rangarajan](#), [Georg May](#) and [Vit Dolejsi](#)

Analytic Metric-Based Adaptation Using a Continuous Mesh Model ([abstract](#))

11:45-12:00 Short break

12:00-12:45 Session 9: Biomechanics

CHAIR: [Juliette Unwin](#)

12:00 [Marie Rognes](#)

Accurate Numerical Modeling of Small Collections of Cardiac Cells ([abstract](#))

12:15 [Mariya Ptashnyk](#) and [Brian Seguin](#)

The impact of microscopic structure on mechanical properties of plant cell walls and tissues ([abstract](#))

12:30 [Will Zhang](#) and [Michael Sacks](#)

Simulating the Response of Bioprosthetics Heart Valve Tissues to Cyclic Loading in FEniCS ([abstract](#))

12:45-13:45 Lunch Break

13:45-14:45 Session 10: Visualisation and Large Scale Computing

CHAIR: [Lawrence Mitchell](#)

LOCATION: Salle Tavenas Main Theatre

13:45 [Hugo Hadfield](#), [Juliette Unwin](#), [Nathan Sime](#) and [Garth Wells](#)

podS: A Parallel Multilevel Monte Carlo Framework ([abstract](#))

14:00 [Anders Logg](#), [Carl Lundholm](#) and [Magne Nordaas](#)

Solving Poisson's equation on the Microsoft HoloLens ([abstract](#))

14:15 [Martin Sandve Alnæs](#), [Vidar Fauske](#) and [Min Ragan-Kelley](#)

3D visualization with FEniCS in Jupyter Notebooks ([abstract](#))

14:30 [Chris Richardson](#), [Jack Hale](#) and [Garth Wells](#)

Uncontainable enthusiasm: the pain and joy of Docker on HPC and cloud ([abstract](#))

14:45-15:15 Coffee Break

15:15-16:30 Session 11: Coupled and interface problems

CHAIR: [Tormod Landet](#)

LOCATION: Salle Tavenas Main Theatre

15:15 [Wietse Boon](#) and [Jan Nordbotten](#)

Mixed-Dimensional Linear Elasticity with Relaxed Symmetry ([abstract](#))

15:30 [Cécile Daversin-Catty](#) and [Marie E. Rognes](#)

Towards coupled mixed dimensional finite elements in FEniCS ([abstract](#))

15:45 [Karl Erik Holter](#), [Miroslav Kuchta](#) and [Kent-Andre Mardal](#)

Trace Constrained Problems in FEniCS ([abstract](#))

16:00 [Karl Erik Holter](#), [Miroslav Kuchta](#) and [Kent-Andre Mardal](#)

Tracer transport in microvasculature: A case study on coupling 1D-3D problems in FEniCS ([abstract](#))

16:15 [Anders Logg](#), [Carl Lundholm](#) and [Magne Nordaas](#)

A Space-Time Cut Finite Element Method for the Heat Equation in FEniCS ([abstract](#))

16:30-16:40 Session 12: Group photo

CHAIR: [Jack S. Hale](#)

LOCATION: Salle Tavenas Rear Room

16:40-18:15 Session 13: Meet the developers and coding

CHAIR: [Jack S. Hale](#)

LOCATION: Salle Tavenas Rear Room

18:15-19:15 Session 14: Optional: Guided walk to dinner via UNESCO Heritage sites.

CHAIR: [Jack S. Hale](#)

LOCATION: Salle Tavenas Rear Room

19:15-23:59 Session : Conference Dinner

CHAIR: [Jack S. Hale](#)

LOCATION: La Table du Belvédère

Wednesday, June 14th

View this program: [with abstracts](#) [session overview](#) [talk overview](#)

09:15-10:00 Session 15: Solid Mechanics I

CHAIR: [Davide Baroli](#)

LOCATION: Salle Tavenas Main Theatre

09:15 [Miquel Rodriguez](#), [Christoph Augustin](#) and [Shawn Shadden](#)

FEniCS Mechanics: A Package for Continuum Mechanics Simulations ([abstract](#))

09:30 [Marco Morandini](#)

Handling of finite rotations in Dolfin ([abstract](#))

09:45 [Maria Vasilyeva](#)

Multiscale model reduction using GMsFEM for applied problems in heterogeneous domains ([abstract](#))

10:00-10:30 Coffee break

10:30-11:15 Session 16: Solid Mechanics II

CHAIR: [Cécile Daversin](#)

LOCATION: Salle Tavenas Main Theatre

10:30 [Juliette Unwin](#), [Garth Wells](#), [Nathan Sime](#) and [Hugo Hadfield](#)

Multi-Level Monte Carlo for Large-scale Vibrations Problems Using FEniCS and podS ([abstract](#))

10:45 [Hernán Mella](#), [Joaquín Mura](#) and [Esteban Sáez](#)

Implementation of Mixed and Hybrid Formulation of the PML Method in Elastodynamics Using FEniCS ([abstract](#))

11:00 [Benjamin Goldsberry](#) and [Michael Haberman](#)

Finite Element Modeling of Elastic Wave Dispersion in Pre-strained Negative Stiffness Honeycombs Using FEniCS ([abstract](#))

11:15-11:30 Short break

11:30-12:15 Session 17: Fluid Mechanics

CHAIR: [Carl Lundholm](#)

LOCATION: Salle Tavenas Main Theatre

11:30 [Jan Blechta](#) and [Martin Řehoř](#)

FENaPack - FEniCS Navier-Stokes preconditioning package ([abstract](#))

11:45 [Henning Bonart](#), [Christian Kahle](#) and [Jens-Uwe Repke](#)

Simulating Moving Contact Line Problems Using the Cahn-Hilliard-Navier-Stokes Equations ([abstract](#))

12:00 [Tormod Landet](#), [Mikael Mortensen](#) and [Kent-Andre Mardal](#)

Slope Limiting of Divergence Free Discontinuous Galerkin Vector Fields in the Context of Two-Phase Flows ([abstract](#))

12:15-12:30 Session 18: Closing session

CHAIR: [Stéphane P. A. Bordas](#)

LOCATION: Salle Tavenas Main Theatre

12:30-13:30 Closing lunch

3D visualization with FEniCS in Jupyter Notebooks

Martin S. Alnæs, Simula Research Laboratory, martinal@simula.no,
Vidar T. Fauske, Simula Research Laboratory, vidar@simula.no,
Min Ragan-Kelley, Simula Research Laboratory, benjaminrk@simula.no

Keywords: *FEniCS*, *Jupyter*, *Visualization*

As part of the OpenDreamKit project (www.opendreamkit.org), we are working on improving the state of 3D visualization in Jupyter Notebooks (www.jupyter.org).

Jupyter Notebooks run from a web browser based user interface and can display anything a browser can display, which includes interactive javascript widgets and GPU accelerated 3D graphics through WebGL. These fundamental technologies have now matured somewhat and reached widespread support, including mobile targets.

The last couple of years a number of visualization projects have started to explore this area, including established projects such as VTK, Paraview and MayaVi, as well as newer smaller projects such as ipyvolume, k3d-jupyter, and SciviJS, the latter two being initiatives from the OpenDreamKit project.

In this talk we will present our current work on this front, and invite you as users and developers of FEniCS to provide your opinions on which future directions of this project that would be useful to you.

FEniCS: Sustainable Software Development Practices

Martin S. Alnæs (Simula Research Laboratory, martinal@simula.no),
Jan Blechta, Jack S. Hale, Anders Logg, Chris Richardson,
Johannes Ring, Marie Rognes, and Garth N. Wells.

Keywords: *FEniCS*, *Software Productivity*, *Software Development*

The FEniCS project aims to provide a high productivity environment for development of finite element based simulation software. Techniques applied to achieve this goal include mixed language programming and code generation, which enables writing high performance programs in a high level language. End-user productivity is a high priority goal in our software designs. To sustain the productivity of the multinational team of part-time developers (mainly researchers and students) is paramount to the long term survival of the project. To minimize the developer workload while making the process open and accessible to new contributors and users, we regularly question which tools are the best available for our needs. On this poster we will present our current tool choices and work flows for developers and the wider FEniCS community. This list includes version control, build systems, testing, release management, team communication, documentation, and end user support. The most recent addition to our toolbox are developer curated Docker images. We are investigating their usefulness in testing infrastructure, end user deployment, HPC cluster deployment, and as reproducible software environments to accompany journal publications. We welcome discussion on alternatives that can simplify our lives.

This poster was previously presented at SIAM CSE17 PP108 Minisymposium: Software Productivity and Sustainability for CSE and Data Science.

Simulating Moving Contact Line Problems Using the Cahn-Hilliard-Navier-Stokes Equations

Henning Bonart, Technische Universität Berlin, henning.bonart@tu-berlin.de
Christian Kahle, Technische Universität München, christian.kahle@ma.tum.de
Jens-Uwe Repke, Technische Universität Berlin, jens-uwe.repke@tu-berlin.de

Keywords: *Cahn-Hilliard-Navier-Stokes, Moving contact lines, Supercomputing*

The wetting of liquid films on solid surfaces occurs in many industrial processes like coating or painting and apparatuses involving trickle films. At the intersection of the gas-liquid interface with the solid surface a moving contact line is formed. Applying the common no-slip boundary condition at the solid surface, a non-physical divergent stress at the contact line occurs. One possibility to circumvent this difficulty in the context of continuum mechanics is the coupling of the incompressible Navier-Stokes equations with the Cahn-Hilliard (CH) equation. The CH equation models the interface between the fluids with a diffuse interface of positive thickness and describes the distribution of the different fluids by a smooth indicator function. Especially, the CH equation allows the contact line to move naturally on the solid surface due to a diffusive flux across the interface, which is driven by the gradient of the chemical potential. The Cahn-Hilliard-Navier-Stokes (CHNS) model forms a tightly coupled and nonlinear system of four partial differential equations.

In this talk, we present a new FEniCS-based library for massive parallel simulations using CHNS models for problems involving moving contact lines. The library is programmed mainly in Python, whereas some time critical parts are in C++. Furthermore, Parallel HDF5 is applied for in- and output and PETSc's SNES with MUMPS is used for the solution of the nonlinear system. The library aims to be easily extendable and to serve as a platform for rapid benchmarking of different CHNS models and solution schemes. For efficient simulations of the thermodynamically consistent CHNS model by [1] for large density ratios, we employ a stable solution scheme as proposed in [2]. Besides the structure of the library, we discuss the compilation of the FEniCS toolchain and our own implementations on a Cray XC40/XC30 supercomputer system. Furthermore, we present some promising physical results on moving contact lines and discuss the good scalability on the supercomputer.

References

- [1] Helmut Abels, Harald Garcke, and Günther Grün. "Thermodynamically consistent, frame indifferent diffuse interface models for incompressible two-phase flows with different densities." *Mathematical Models and Methods in Applied Sciences* 22.03 (2012): 1150013.
- [2] Harald Garcke, Michael Hinze, and Christian Kahle. "A stable and linear time discretization for a thermodynamically consistent model for two-phase incompressible flow." *Applied Numerical Mathematics* 99 (2016): 151-171.

Mixed-Dimensional Linear Elasticity with Relaxed Symmetry

Wietse M. Boon, University of Bergen, wietse.boon@uib.no
Jan M. Nordbotten, University of Bergen, jan.nordbotten@uib.no

Keywords: *Mixed-dimensional, Mixed finite elements, Linear elasticity.*

Thin inclusions in materials are common in a variety of applications, ranging from steel plate reinforced concrete to composite materials and cemented fracture systems. To ease implementation issues such as mesh generation, mixed-dimensional representations of the material have become an attractive strategy. In such a representation, the thin inclusions are considered as lower-dimensional manifolds with significantly different material properties. The associated, governing equations are then fully coupled to the surroundings.

The explicit consideration of momentum conservation is especially important when considering stress states around thin inclusions. For example, in the context of fracture propagation, the region near fracture tips is of particular interest since the stress state contains a wealth of information. An accurate representation of the surrounding stress, possessing physical conservation properties, is therefore imperative.

In this work, we employ mixed finite elements to obtain a locally conservative discretization scheme. The symmetry of the stress tensor is then imposed in a weak sense, which allows for the use of familiar, conforming, finite elements with relatively few degrees of freedom. The mixed-dimensional representation is exploited to the fullest extent by considering subdomains of all dimensionalities. In particular, intersections between d -dimensional manifolds are considered as separate $(d - 1)$ -dimensional manifolds and all couplings are incorporated between domains with codimension one.

We present several theoretical results including well-posedness of the variational formulation and the mixed finite element discretization scheme. Moreover, numerical examples and their implementations using FEniCS are considered in two and three dimensions.

FENaPack - FEniCS Navier-Stokes preconditioning package

Jan Blechta, Charles University, blechta@karlin.mff.cuni.cz,
Martin Řehoř, Charles University, rehor@karlin.mff.cuni.cz

Keywords: *Navier-Stokes equations, PCD preconditioner, HPC*

In this contribution we present a novel theoretical analysis of PCD (pressure-convection-diffusion) preconditioner for approximation of Navier-Stokes equations, present some important implementational details, and provide an open-source implementation suitable for deployment on HPC systems.

We consider a linearization of Navier-Stokes equations written in a block form

$$\begin{pmatrix} -\nu\Delta + \mathbf{v} \cdot \nabla & \nabla \\ -\operatorname{div} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ 0 \end{pmatrix}. \quad (1)$$

Formally we can express a Schur complement of the system as

$$S = -\operatorname{div} (-\nu\Delta + \mathbf{v} \cdot \nabla)^{-1} \nabla. \quad (2)$$

Reasonable strategy for solving (1) is to employ GMRES with right preconditioner given by

$$\begin{pmatrix} -\nu\Delta + \mathbf{v} \cdot \nabla & \nabla \\ 0 & -S \end{pmatrix}^{-1}. \quad (3)$$

The problem is that an action of S^{-1} needed to express an action of (3) cannot be efficiently implemented because S given by (2) is a dense matrix when discretized. A possible remedy is to swap the order of operators in (2) and hope for a stability of this approximation. This gives rise to a pair of possible Schur complement approximations

$$S \approx -\Delta (-\nu\Delta + \mathbf{v} \cdot \nabla)^{-1} \equiv X \quad (4)$$

$$S \approx (-\nu\Delta + \mathbf{v} \cdot \nabla)^{-1} (-\Delta) \equiv Y. \quad (5)$$

The immediate question is what boundary conditions should one enforce in a Laplacian solve needed for action of X^{-1} or Y^{-1} . This has been a topic of long-standing discussions in the literature using different heuristic arguments. We provide a new rigorous analysis advocating boundary conditions which appeared in [4, Section 9.3.2] and [5, section 2.1] under certain reasonable assumptions which were not known before. Under such conditions we are able to prove that X and Y are continuous isomorphisms on $L^2(\Omega)$ which results in L^2 -norm equivalence with Schur complement S . This is a good heuristics giving mesh-independent convergence rates of GMRES.

We provide an open-source implementation [3] built on top of the FEniCS project [1] and PETSc [2]. Thanks to the abstraction the package is designed with, an implementation of a particular variant of PCD preconditioner can look like

Python code

```
1 class PCDPC_BRM2(BasePCDPC):
2     def apply(self, pc, x, y):
3         # Fetch work vector
4         z0, z1 = self.get_work_vecs(x, 2)
5
6         # Apply PCD
7         self.ksp_Mp.solve(x, y) # y = M_p^{-1} x
8         y.copy(result=z0) # z0 = y
9         self.mat_Kp.mult(z0, z1) # z1 = K_p z0
10        self.bcs_applier(z1) # apply bcs to z1
11        self.ksp_Ap.solve(z1, z0) # z0 = A_p^{-1} z1
12        y.axpy(1.0, z0) # y = y + z0
13        y.scale(-1.0) # y = -y
```

The base class `BasePCDPC` provides members to apply particular ingredients of PCD, e.g., `mat_Ap`, `bcs_applier`, etc., taking care of matrix-vector operations based on user-provided UFL input. This allows to easily modify the preconditioner for different needs, e.g., time-dependent problem, spatially variable viscosity, etc. Furthermore approximation of required solves can be setup with an ease using the PETSc parameter system


Python code

```

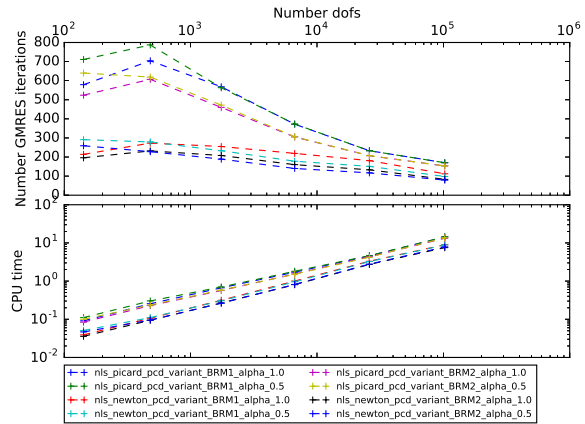
1 PETScOptions.set("fieldsplit_u_ksp_type", "richardson")
2 PETScOptions.set("fieldsplit_u_ksp_max_it", 1)
3 PETScOptions.set("fieldsplit_u_pc_type", "hypre")
4 PETScOptions.set("fieldsplit_u_pc_hypre_type", "boomeramg")
5 PETScOptions.set("fieldsplit_p_PCD_Ap_ksp_type", "richardson")
6 PETScOptions.set("fieldsplit_p_PCD_Ap_ksp_max_it", 2)
7 PETScOptions.set("fieldsplit_p_PCD_Ap_pc_type", "hypre")
8 PETScOptions.set("fieldsplit_p_PCD_Ap_pc_hypre_type", "boomeramg")
9 PETScOptions.set("fieldsplit_p_PCD_Mp_ksp_type", "chebyshev")
10 PETScOptions.set("fieldsplit_p_PCD_Mp_ksp_max_it", 5)
11 PETScOptions.set("fieldsplit_p_PCD_Mp_ksp_chebyshev_eigenvalues", "0.5, 2.0")
12 PETScOptions.set("fieldsplit_p_PCD_Mp_pc_type", "jacobi")

```

An example of obtained solution is in the Figure 1a. Figure 1b demonstrates mesh independence of the preconditioner.



(a) Flow over a backward facing step in 3D



(b) Scaling of different variants of PCD preconditioner to solve the nonlinear Navier-Stokes problem with prescribed tolerance. Top figure shows mesh independence of number of total GMRES iterations. This figure is produced automatically using a CI service.

References

- [1] M. ALNÆS, J. BLECHTA, J. HAKE, A. JOHANSSON, B. KEHLET, A. LOGG, C. RICHARDSON, J. RING, M. ROGNES, AND G. WELLS, *The FEniCS Project version 1.5*, Archive of Numerical Software, 3 (2015). URL: <http://journals.ub.uni-heidelberg.de/index.php/ans/article/view/20553>, doi:10.11588/ans.2015.100.20553.
- [2] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, V. EIJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, K. RUPP, B. F. SMITH, AND H. ZHANG, *PETSc Web page*, 2017. URL: <http://www.mcs.anl.gov/petsc>.
- [3] J. BLECHTA AND M. ŘEHOŘ, *FENaPack - FEniCS Navier-Stokes preconditioning package*. URL: <https://github.com/blechta/fenapack>.
- [4] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics*, Oxford University Press, 2nd ed., 2014.
- [5] M. A. OLSHANSKII AND Y. V. VASSILEVSKI, *Pressure Schur complement preconditioners for the discrete Oseen problem*, SIAM Journal on Scientific Computing, 29 (2007), pp. 2686–2704. doi:10.1137/070679776.

iRANitsche: reduced assembly with a Nitsche-based reduced basis method

Davide Baroli, University of Luxembourg, davide.baroli@uni.lu,
Lars A. A. Beex, University of Luxembourg, lars.beex@uni.lu,
Jack S. Hale, University of Luxembourg, jack.hale@uni.lu,
Stéphane P.A. Bordas, University of Luxembourg,
School of Engineering, Cardiff University, Wales, UK
Adjunct Professor, Intelligent Systems for Medicine Laboratory
School of Mechanical and Chemical Engineering, The University of Western Australia, Australia
stephane.bordas@uni.lu

Keywords: *model order reduction, Nitsche coupling, MultiMesh.*

Model order reduction methods are a family of well-established techniques for reducing the computational cost of parametrized PDEs. They provide a good trade-off between fast computation and high-fidelity approximation while preserving reliability.

One popular model order reduction technique is the reduced basis method (RB). This approach leads to a set of global RB basis functions that optimally represent the solution and also provides a reliability estimate of the gap between high fidelity solution and reduced one of the parametrized PDE, for a given range of possible sample parameters.

In this poster, we present a novel alchemy of model reduction, reduced assembly [Schenone et al.,2016], and domain decomposition method that we call iRANitsche [Baroli et al].

Rather than using globally reduced basis function computed offline, we design a spatially local reduced basis space and we apply the ‘reduced assembly’ (RA) technique [Schenone et al.,2016], which is used to efficiently reduce the assembly complexity cost of online Galerkin projection.




Figure 1: At left the picture shows a Gaussian function supported on the finest grid, at right the coarsening representation is obtained via reduced assembly.

Our approach instead deals with reducing the degree of freedom of domain preserving a certain

accuracy of the locally reduced basis function. However, it leads to non-matching mesh and in principle may lead to discontinuous basis function across the interface of each pair of sub-domains.

For this reason, we employ a Nitsche-based gluing formulation using the MultiMesh [Massing et al. 2013, Johansson et al. 2015] framework developed recently in DOLFIN. The robustness of this coupling scheme is determined by the penalty coefficients that are chosen using ghost penalty technique in [E.Burman 2015].

We will show how we develop our package at the top of DOLFIN [Logg et al. 2012], SLEPc4py and PETSC4py[Dalcin et al. 2011].

The numerical tests performed in 2D and 3D on an academic and the patient-specific problems (Figure 2) provided by Dr. Hertel of Hospital Center de Luxembourg shows the good performance of the method and reduction of computation cost of some order of magnitude with respect the high-fidelity approximation.




Figure 2: The figure shows the deformation obtain by iRBNitsche of post-augumented vertebra

Acknowledgement: We thank the financial support of the European Research Council Starting Independent Research Grant (ERC Stg grant agreement No. 279578) entitled Towards real time multiscale simulation of cutting in non-linear materials with applications to surgical simulation and computer guided surgery.. We thank the support of the Computational Sciences Priority of the University of Luxembourg. Jack S. Hale is supported by the National Research Fund, Luxembourg, and cofunded under the Marie Curie Actions of the European Commission (FP7-COFUND) Grant No. 6693582. The experiments presented in this paper were carried out using the HPC facilities of the University of Luxembourg.

References

- [Dalcin et al. 2011] Dalcin, L. D., Paz, R. R., Kler, P. A., and Cosimo, A. (2011). Parallel distributed computing using python. *Advances in Water Resources*, 34(9):1124–1139.
- [Baroli et al] D.Baroli, L.A.A. Lars, S. Reduced basis nitsche-based domain decomposition. in preparation.
- [Schenone et al.,2016] E. Schenone, L.A.A. Beex, J. H. S. B. Proper orthogonal decomposition with reduced integration method. application to nonlinear problems. manuscript.
- [E.Burman 2015] E.Burman, S. Claus, P. H. M. L. A. M. (2015). Cutfem: Discretizing geometry and partial differential equations. *International Journal for Numerical Methods in Engineering*, 104(7):472–501.
- [Logg et al. 2012] Logg, A., Mardal, K.-A., and Wells, G. N. (2012). *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media.
- [Massing et al. 2013] Massing, A., Larson, M. G. and Logg, A. (2013). *Efficient implementation of finite element methods on nonmatching and overlapping meshes in three dimensions*, *SIAM Journal on Scientific Computing*, 33(1):C23-C47.
- [Johansson et al. 2015] Johansson, A, Larson, M. G. and Logg, A. (2015). *High order cut finite element methods for the Stokes problem*, *Advanced Modeling and Simulation in Engineering Sciences*, 2(1):1–24, 2015.

MD-FEM Force Field Simulations of Particle Subjected to Dielectrophoresis Interactions

M. Cascio^{1,2,*}, G. Falci^{1,2}, S. F. Lombardo^{1,2}, A. Magliano¹,
I. Deretzis¹, S. Scalese¹ and A. La Magna¹,

¹CNR-IMM, Catania, Italy

²Department of Physics and Astronomy and INFN, University of Catania, Italy

*michele.cascio@ct.infn.it

Keywords: *MD-FEM, electromechanical particles, dielectrophoresis*

Particles with sizes that range from sub-micrometers to about 1 millimeter and with particular electrical and/or magnetic properties, experience mechanical forces and torques when they are subjected to electromagnetic fields (this type of particles are called “electromechanical particles”).

The theoretical study of this large class of complex systems is possible thanks to the development of real-system models and numerical simulations of the stable (multi-particle) configurations and their dynamics.

One of the phenomena that affect electromechanical particles is the dielectrophoresis (DEP)[1, 2].

A branch of emerging application relates to controlled manipulation of particles dispersed in colloidal solutions (i.e. biological particles such as cells or DNA), since the strong selectivity of the response depends on the particle volume, shape and composition [3]. Application fields of dielectrophoresis include cell partitioning/isolation for the capture/separations without the use of biomarkers [4].

This contribution focuses on the theoretical study of the dynamics of micro-sized spherical biological particles suspended in a colloidal solution, which are subjected to dielectrophoresis in the presence of non-homogeneous and non-uniform variable electric fields.

Most DEP models in the literature are based on particles in the diluted solution limit [5]; in this case the forces are calculated using an approximate method (standard DEP). The electric field is applied through the electrodes present in a microfluidic channel in which the colloidal solution flows. Particle manipulation and characterization using DEP is generally performed in a confined region near the electrodes, so that the interaction between the particle and the surrounding walls can be significant. In this work, we present numerical simulations of the movement of MDA-MB-231 tumor cells near electrode edges; we run a more detailed study, with a non-approximate calculation of the dielectrophoretic force: DEP forces are estimated by integrating the Maxwell tensor [6]; it leads to an overall DEP force independent of the complex dielectric permittivity of the particles and suspending medium and depends only on the type of boundaries (conductive or isolating) and on the ratio between the particle and electrode dimensions.

The dynamics is simulated by techniques borrowed from Molecular Dynamics (MD): our goal is to first evaluate the forces acting on electromechanical particles starting from the initial configuration of the system (position of the particles, geometry of the electrodes, electrical potentials applied), and then calculate the dynamics of the particles through the integration of the equations of motion using MD-like techniques. The Coupled MD-FEM study of particles’ kinetics consists of a loop with the following steps: initial positions of the particles; calculation of forces; calculation of acceleration; integration of the equations of motion; new positions. We use the numerical integration technique called Verlet Method. The coupled MD-FEM algorithm and its implementation in the FEniCS environment are presented. Realistic simulated cases will be discussed, showing also the difficulties of the methodic implementation in 3D domains.

We carry out simulations of the movement of MDA-MB-231 tumor cells near the electrode edges, based both on the standard DEP theory and on the non approximate theoretical model (MST). We find that, in the case of standard DEP, the cells experience an attractive force that traps them near the electrodes, while in the case of the MST-DEP force, the cells also form chains due to dipole-dipole interactions and some escape from the attraction of the electrodes.

Our work shows the potential of coupled MD-FEM study of electromechanical particles.

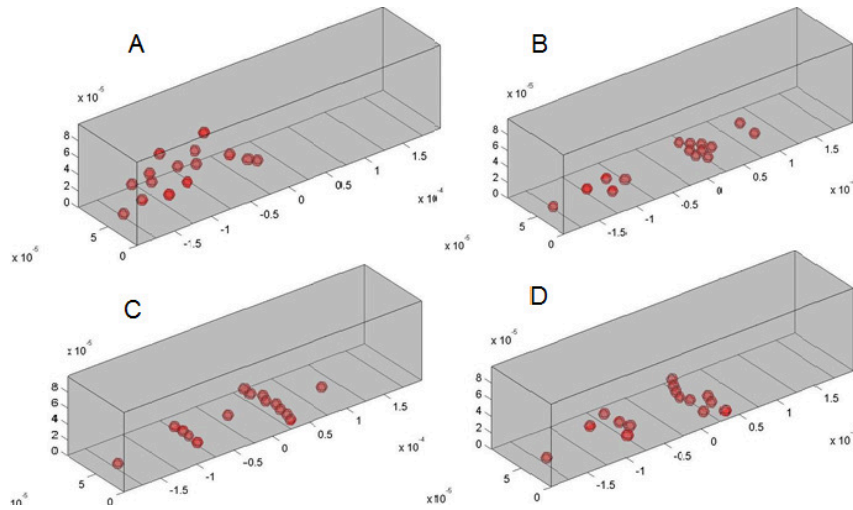


Figure 1: Standard DEP. A): 1 sec. B): 100 sec. C): 200 sec. D) 300 sec. In the presence of standard DEP force, the cells are attracted by the electrodes.

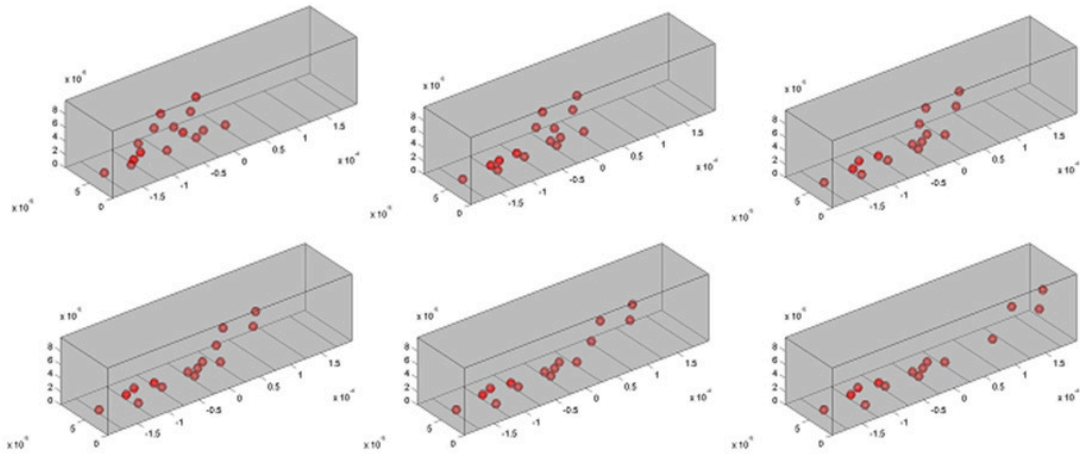


Figure 2: MST-DEP. A): 1 sec. B): 100 sec C): 200 sec. D) 300 sec. E): 400 sec. F): 450 sec. In the presence of the DEP force calculated by the Maxwell Stress Tensor, the cells are attracted by the electrodes and form chains due to dipole-dipole interactions; some cells escape the attraction of the electrodes.

References

- [1] P. R. C. Gascoyne, J. Noshari, T. J. Anderson, F. F. Becker, *Electrophoresis* **30**, 1388 (2009)
- [2] Dho-Hyoung Lee, Chengjie Yu, Papazoglou E., Bakhtier F., Hongseok M., *Electrophoresis* **32**, 2298-2306 (2011)
- [3] B. Cetin and D. Li, *Electrophoresis* **32**, 2410 (2011)
- [4] M. P. Huges, *Electrophoresis* **23**, 2569 (2002)
- [5] O. E. Nicotra, A. La Magna, S. Coffa, *Appl. Phys. Lett.* **95**, 073702 (2009)
- [6] M. Camarda, S. Scalese and A. La Magna, *Electrophoresis* **36**, 13961404 (2015)

Automated Solution of Parabolic-Elliptic Systems for Control Applications with Embedded Runge-Kutta Methods

Manuel Cremades, USC, ITMATI, manuel.cremades@usc.es,
 Pedro Fontán Muiños, USC, ITMATI, pedro.fontan@usc.es,
 Jerónimo Rodríguez, USC, ITMATI, jeronimo.rodriguez@usc.es

Keywords: *Partial Differential-Algebraic Equations (PDAEs), Diagonally Implicit Runge-Kutta (DIRK) methods, Rosenbrock-Wanner (ROW) methods, adjoint state.*

The objective of this work is to present a Python code developed for the solution of differential-algebraic systems of the form

$$\begin{cases} M(\mathbf{y}(t)) \frac{d\mathbf{y}}{dt}(t) = \mathbf{f}(t, \mathbf{y}(t)), \\ \mathbf{y}(0) = \mathbf{y}_0, \end{cases} \quad (1)$$

with a possibly singular matrix function $M : \mathbb{R}^d \rightarrow \mathcal{M}_{d \times d}(\mathbb{R})$ of constant rank in a neighbourhood of the solution. This kind of systems arise from the space discretization of transient partial differential equations (see [1] for a different approach where the time stepping is applied first) such as, for example, the following non-linear heat equation (in this case \mathbf{y} is associated to (u, \mathbf{v}))

$$\begin{cases} p(u) \frac{\partial u}{\partial t} - \operatorname{div}(\mathbf{v}) = f, & \text{in } \Omega \times (0, T], \\ \mathbf{v} - q(u) \nabla u = \mathbf{0}, & \text{in } \Omega \times (0, T], \end{cases} \quad \begin{cases} u = 0, & \text{in } \partial\Omega \times (0, T], \\ u(t=0) = u_0, & \text{in } \Omega, \end{cases} \quad (2)$$

by means of the finite element method. To this end, the finite element library FEniCS [2] is used. For the discretization in time we are interested in implicit stiffly-accurate one-step methods with a low computational cost. In particular, we consider Diagonally Implicit Runge-Kutta (DIRK) methods

$$\begin{aligned} M(\mathbf{y}_n + \sum_{j=1}^i a_{ij} \mathbf{k}_j) \mathbf{k}_i &= h_n \mathbf{f}(t_n + c_i h_n, \mathbf{y}_n + \sum_{j=1}^i a_{ij} \mathbf{k}_j), \quad i = 1, \dots, s, \\ \mathbf{y}_{n+1} &= \mathbf{y}_n + \sum_{j=1}^n b_j \mathbf{k}_j, \end{aligned} \quad (3)$$

sub-families of them (such as SDIRK and ESDIRK) and Rosenbrock-Wanner methods (ROW). It is worth noting that the jacobian matrices needed to solve the non-linear systems in (3) can be easily evaluated using the library FEniCS [2]. In both cases, an embedded pair of methods is used to compute an error estimator and adapt the time step with a marginal extra computational cost, as explained in [3].




Figure 1: Structure of the `solver` class

These families of methods have been implemented in an object-oriented Python code that, providing the matrix function $M(\cdot)$ and the field $f(\cdot, \cdot)$ in (1) and the Butcher tableau defining the embedded methods,

allows the user to automatically integrate the equation in time with an adaptive time step adjusted to a prescribed tolerance (see Figure 1 for the structure of the `solver` class).

Let us give further details on how $M(\cdot)$ and $f(\cdot, \cdot)$ can be generated for the model problem in (2) which admits the following weak formulation: Find $(u, \mathbf{v}) \in L^2((0, T]; H_0^1(\Omega)) \times L^2((0, T]; (L^2(\Omega))^d)$ such that

$$\left| \begin{aligned} \int_{\Omega} p(u) \frac{\partial u}{\partial t} \tilde{u} \, d\Omega - \int_{\Omega} \mathbf{v} \cdot \nabla \tilde{u} \, d\Omega - \int_{\Omega} f \tilde{u} \, d\Omega &= 0, \quad \forall \tilde{u} \in H_0^1(\Omega), \\ \int_{\Omega} \mathbf{v} \cdot \tilde{\mathbf{v}} \, d\Omega - \int_{\Omega} q(u) \nabla u \cdot \tilde{\mathbf{v}} \, d\Omega &= 0, \quad \forall \tilde{\mathbf{v}} \in (L^2(\Omega))^d. \end{aligned} \right. \quad (4)$$

This weak formulation must be written in UFL language [4] as non-linear with the time derivative declared as a trial function in order to be able to extract the mass matrix and right hand side:

$$\left. \begin{aligned} \int_{\Omega} p(u) \frac{\partial u}{\partial t} \tilde{u} \, d\Omega + \int_{\Omega} \mathbf{0} \cdot \tilde{\mathbf{v}} \, d\Omega &\xrightarrow{\text{assemble()}} M(\mathbf{y}(t)), \\ \int_{\Omega} \mathbf{v} \cdot \nabla \tilde{u} \, d\Omega + \int_{\Omega} f \tilde{u} \, d\Omega - \int_{\Omega} \mathbf{v} \cdot \tilde{\mathbf{v}} \, d\Omega + \int_{\Omega} q(u) \nabla u \cdot \tilde{\mathbf{v}} \, d\Omega &\xrightarrow{\text{assemble()}} \mathbf{f}(t, \mathbf{y}(t)), \end{aligned} \right. \quad (5)$$

where the mass matrix will clearly be singular with constant rank.

This procedure can be used to solve problems of the form

$$\left| \begin{aligned} \text{minimize } \Psi(\mathbf{m}) &= \Phi(\mathbf{y}(T), \mathbf{m}) + \int_0^T g(\mathbf{y}(t), \mathbf{m}) \, dt \\ \text{s.t. } \mathbf{m} &\in \mathcal{A} \\ M(\mathbf{y}(t)) \frac{d\mathbf{y}}{dt}(t) &= \mathbf{f}(t, \mathbf{y}(t), \mathbf{m}) \\ \mathbf{y}(0) &= h(\mathbf{m}) \end{aligned} \right. \quad (6)$$

arising from the space discretization of optimal control problems governed by transient partial differential equations. The gradient of the discrete cost function with respect to the degrees of freedom of the control \mathbf{m} (needed by most of the iterative optimization algorithms) is computed by means of the discrete adjoint state, deduced in [5] for Runge-Kutta methods and a particular case of this kind of problems.

During the talk we will show for different problems such as the non-linear heat equation described before or the transient Stokes equations

$$\left| \begin{aligned} \frac{\partial \mathbf{u}}{\partial t} - \nu \Delta \mathbf{u} + \nabla p &= \mathbf{f} \quad \text{in } \Omega \times (0, T], \\ \text{div}(\mathbf{u}) &= 0 \quad \text{in } \Omega \times (0, T], \\ \mathbf{u} &= \mathbf{g} \quad \text{in } \Gamma \times (0, T], \end{aligned} \right. \quad (7)$$

that the methods presented here perform well converging with the expected order of accuracy, being the Rosenbrock-Wanner methods specially attractive for the solution of high-scale transient partial differential equations due to their low computational cost.

References

- [1] K. E. Skare, *Gryphon - a Module for Time Integration of Partial Differential Equations in FEniCS*, Norwegian University of Science and Technology, 2012.
- [2] M. S. Alnaes, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, G. N. Wells: The FEniCS Project Version 1.5, *Archive of Numerical Software*, 3 (2015).
- [3] E. Hairer, G. Wanner, *Solving Ordinary Differential Equations II, Stiff and Differential-Algebraic Problems*, Springer, 1996.
- [4] M. S. Alnaes, A. Logg, K. B. Igaard, M. E. Rognes and G. N. Wells: Unified Form Language: A domain-specific language for weak formulations of partial differential equations, *ACM Transactions on Mathematical Software*, 40 (2014).
- [5] J. M. Sanz-Serna: Symplectic Runge-Kutta Schemes for Adjoint Equations, Automatic Differentiation, Optimal Control and More, *SIAM Review*, 58 (2016), 3-33.

Identification of Mechanical Properties of 3D Myocardial Tissue: An Inverse Modeling and Optimal Experimental Design Problem

Joshua W. Chen, joshua@ices.utexas.edu¹, Dr. Andrew Drach, andrewdrach@utexas.edu¹, Dr. Umberto Villa, uvilla@ices.utexas.edu¹, Dr. Reza Avazmohammadi, rezaavaz@ices.utexas.edu¹, David Li, davidqli@utexas.edu², Dr. Omar Ghattas, omar@ices.utexas.edu¹, and Dr. Michael S. Sacks, msacks@ices.utexas.edu¹

¹*Institute for Computational Engineering and Sciences*

²*Department of Biomedical Engineering, The University of Texas at Austin*

Keywords: Optimal Experimental Design, Myocardium Strain Energy Function, Material Model Parameter Identification, Bayesian Inverse Problem, FEniCS, hIPPYlib

In this work, we present novel computational tools for bayesian parameter inversion and optimal experimental design for determining the 3D strain energy function parameters for myocardial tissue. In previous work[1], we presented a framework for parameter inversion and optimal design based on an experimental setup for 1 cm³ myocardial tissue samples that applies tri-axial Dirichlet boundary conditions incrementally in quasi-static equilibrium and measures resultant surface traction forces (Figure 1 a).

In the current work, we extend the framework to a Bayesian inversion context for parameter estimation and we consider more advanced techniques for optimal experimental design. We use the inverse problems Python library hIPPYlib[2] a toolbox for FEniCS that enables easy development of scalable algorithms for PDE-based deterministic and Bayesian inverse problems. In what follows, we introduce further the functionality of hIPPYlib, we then present the forward problem, the Bayesian inverse problem, and some thoughts on the optimal design of experiments.

hIPPYlib. Using hIPPYlib, the user expresses the forward problem PDE and the likelihood in weak form using the friendly, compact, near-mathematical notation of FEniCS, which will then automatically generate efficient code for the discretization. Given a prior distribution, determined by our prior knowledge of the model parameters, we take advantage of hIPPYlib's robust implementation of the inexact Newton-conjugate gradient algorithm to compute the maximum a posterior (MAP) point. Invoking the Laplace approximation to the posterior, we approximate the posterior distribution as a Gaussian distribution with mean equal to the map point and covariance operator given by the inverse of the Hessian evaluated at the MAP point. Such distribution can then be used directly as a *surrogate* for the true posterior or as a proposal distribution for Markov Chain Monte Carlo (MCMC). The reduced gradient and Hessian actions are automatically computed via their weak form specification in FEniCS by constraining the state and adjoint variables to satisfy the forward and adjoint problem.

Forward Problem. An experimental setup presented in previous work is briefly described. We built a tri-axial experimental setup for 1 cm³ soft tissue samples, capable of incrementally applying a combination of multi-axial tension, compression, and shear Dirichlet boundary conditions in a quasi-static equilibrium context and measuring tissue surface traction forces. Assuming an incompressible orthotropic Fung strain energy function $\Psi_{Fung}(u(c_k), \xi)$, where c_k are parameters which represent shear and elastic moduli: $c_k, k = 0, 1, \dots, 6 \in \mathbf{C}$, the parameter space, and ξ are the local material orientations, we can model **forward problem** as: find the displacement field u that minimizes the energy functional Ψ_{Fung} s.t. the incompressibility constraint and the Dirichlet boundary conditions are satisfied. Specifically, it reads:

Find u in $[H^1(\Omega)]^3$ verifying

$$\boxed{\begin{aligned} \min_u \int_{\Omega} \Psi_{Fung}(u(c_k), \xi) dV \\ \text{such that} \\ \det(\mathbf{F}) = 1 \text{ in } \Omega \\ u = g \text{ on } \Gamma \subset \partial\Omega, \end{aligned}} \tag{1}$$

where

$$\mathbf{F} = \mathbf{I} + \nabla u.$$

Inverse Problem. In our experiments we impose boundary conditions g_i and we observe forces f_i . We define the pairs (g_i, f_i) as our collection of experiments $i \in \mathbf{E}$ (the experiments space). Given ξ , the *inverse problem* for global material parameters $c_k, k = 0, 1, \dots, 6$ is:

Find $c \in \mathbf{C}$ verifying

$$\min_c J(c) := \frac{1}{2} \sum_i^{|E|} \left(\int_{\Gamma} \mathbf{T}(u_i) \cdot n_i ds - f_i \right)^2 \quad (2)$$

subject to the forward problem,

where \mathbf{T} is Cauchy stress. Notice from Figure 1 b), that since the amount of experimental data is large, i.e. $j = 0, 1, \dots [150 - 300]$, the inverse problem is constrained by over 1,300 PDE's, arriving from data from 7 experiments, and, intuitively, the most important data for the inverse problem lie toward the right end of the loading curves, where the stress/strain relationship is *most nonlinear*.




Figure 1: **a)** A photograph of the experimental setup, with Dirichlet boundary conditions applied to two sets of 9 circular pins on opposing faces. **b)** Experimental data. There are >1300 data points from 7 experiments, each plotted in a different color **c)** FEniCS forward problem with Dirichlet boundary conditions at the 9 circular pins on two opposite faces of a cube of tissue.

Optimal Experimental Design. Since excised tissue has a short lifespan, only a few experiments are feasible (6 or 7). Therefore it is crucial to find a set of loading conditions that best allow for inversion of material parameters. In this work, we solve an *optimal experimental design* problem to find such a near-optimal set. In particular, we seek to maximize the expected information gain from the experimental data, that is to maximize the Kullback-Leibler divergence between the prior and posterior distributions.

References

- [1] R. Avazmohammadi, D. Li, T. Leahy, E. Shih, J. Soares, G. Gorman, and M. Sacks. An Integrated Inverse Model-Experimental Approach to Determine Soft Tissue Three-Dimensional Strain Energy Density Function Parameters: Application to Post-Infarcted Myocardium. *Biomechanics and Modeling in Mechanobiology*, 2017. Accepted for publication.
- [2] U. Villa, N. Petra, and O. Ghattas. hIPPYlib: an Extensible Software Framework for Large-scale Deterministic and Linearized Bayesian Inversion. <http://hippylib.github.io>, 2016.

Shape Optimization with Multiple Meshes

Jørgen S. Dokken, Simula Research Laboratory, dokken@simula.no

Simon Funke, Simula Research Laboratory, simon@simula.no

August Johansson, Simula Research Laboratory, august@simula.no

Stephan Schmidt, University of Würzburg, stephan.schmidt@mathematik.uni-wuerzburg.de

Keywords: *multi-mesh, FEniCS, Optimization, Dynamic domains*

For shape optimization problems, the computational domain is the design variable. Changing the shape of an airfoil in a channel to minimize drag is such a problem. The evolving domains complicate the numerical solution of shape optimization problems, and typically require large mesh deformations with quality checks and a re-meshing software as a fallback. We propose an approach for solving shape optimization problems on multiple overlapping meshes. In this approach, each mesh can be moved freely and hence the multi-mesh approach allows larger deformation of the domain than standard single-mesh approaches. The approach has been implemented in FEniCS and dolfin-adjoint, by employing the already tested environment for multi-mesh [1]. We give examples of derivation of the shape-optimization problem for a Stokes flow and present implementation of this in FEniCS.

Consider a general PDE constrained shape optimization problem

$$\min_{\Omega, u} J(u(\Omega), \Omega) \quad \text{subject to } F(u(\Omega), \Omega) = 0, \quad (1)$$

where J is the goal functional, $F(u, \Omega)$ is the state equations, u is the solution of the state equation, Ω is the domain of state equations.

We choose to divide the domain Ω into two non-overlapping domains by creating an artificial interface Γ , s.t. $\Omega = \Omega_0 \cup \Omega_1$ and $\Gamma = \Omega_0 \cap \Omega_1$, as depicted in Figure 1. Extension to an arbitrary number of overlapping domains is possible.

The weak formulation of the state equations are then formulated and the continuity over the artificial boundary is enforced by using Nitsches method.

For minimization, we choose a gradient based scheme, and find the gradient by using the adjoint method, $\mathcal{L} = J + \lambda^* F$,

$$dJ(x)[V] = \partial \mathcal{L} / \partial \Omega, \quad (2)$$

$$(\partial F / \partial u)^* \lambda = -(\partial J / \partial u)^*. \quad (3)$$

Here Equation 3 is called the adjoint equation and λ the adjoint variable. By employing the Hadamard formulas for Volume and Surface objective functions[2] one can achieve the functional sensitivities as a function of the moving boundary and not the domain.

A concrete example of this approach is the shape optimization of an obstacle in Stokes-flow in the domain specified in Figure 2, namely

$$\min_{u, \Omega} J(u, \Omega) = \int_{\Omega} \sum_{i,j=1}^d \left(\frac{\partial u_i}{\partial x_j} \right) d\Omega, \quad (4)$$

subject to

$$\begin{aligned} -\Delta u + \nabla p &= 0 \text{ in } \Omega, & \nabla \cdot u &= 0 \text{ in } \Omega, \\ u &= 0 \text{ on } \Gamma_1 \cup \Gamma_2, & u &= u_0 \text{ on } \Gamma_3, & p &= 0 \text{ on } \Gamma_4, & \text{Vol} &= \text{Vol}_0, \end{aligned} \quad (5)$$

where Vol_0 is the volume of the obstacle.

For deformation of the mesh, we have used two different deformation equations, a Laplacian smoothing and a set of Eikonal convection equations[3]. For the multi-mesh problem, deformation is only done one the front mesh, while the background mesh is stationary. Figure 3 shows that with the Laplacian deformation the mesh degenerates in both the single-mesh and multi-mesh-case. Figure 4 shows that the Eikonal convection equations preserves the mesh-quality in the multi-mesh-case, but not in the single-mesh case, where the mesh degenerates at the boundary. We conclude that with a multi-mesh-approach, the meshes are preserved better than with a single-mesh approach.




Figure 1: Left: Example of how to divide the domain into two pieces. Center: The specification of two new domains, that will be used for meshing. Right: Example of how to mesh the domains $\hat{\Omega}_0$ and $\hat{\Omega}_1$. Note that the boundaries of $\hat{\mathcal{K}}_{h,1}$ does not align with the edges of $\hat{\mathcal{K}}_{h,0}$.




Figure 2: Illustration of initial domain for the shape optimization problem in Stokes-flow.




Figure 3: Left: Optimal shape with Laplacian deformation for the Stokes-optimization problem on a single mesh. Right: multi-mesh solution of similar problem. Both cases degeneration of the mesh near the front of the geometry. The quality is measured in as the radius ratio of each triangle.




Figure 4: Left: Optimal shape with Eikonal convection deformation for the Stokes-optimization problem on a single mesh. Right: multi-mesh solution of similar problem. In the multi-mesh-case the mesh does not degenerate. The quality is measured in as the radius ratio of each triangle.

References

- [1] A. Johansson, M.G. Larson, A. Logg, *High order cut finite element methods for the Stokes problem*, Advanced Modeling and Simulation in Engineering Sciences, 2 (2015).
- [2] S. Schmidt, *Efficient large scale aerodynamic design based on shape calculus*, 2010.
- [3] S. Schmidt, *A Two Stage CVT/Eikonal Convection Mesh Deformation Approach for Large Nodal Deformations*, arXiv preprint arXiv:1411.7663, 2014.

Towards coupled mixed dimensional finite elements in FEniCS

Cécile Daversin-Catty, Simula Research Laboratory, cecile@simula.no,
Marie E. Rognes, Simula Research Laboratory, meg@simula.no

Keywords: *FEniCS, mixed dimensional, mixed finite elements, non-local operators*

In the brain, the vasculature twist and turn through the brain tissue as topologically one-dimensional structures embedded in three dimensions. In the Earth's crust, topologically two-dimensional faults and one-dimensional fractures cut through the three-dimensional rock. In general, there is an abundance of examples that call for mathematical and numerical models coupling $n < m$ and m dimensional spaces.

As one example, consider the model presented by [1]: consider a topologically one-dimensional subdomain Λ embedded in a three-dimensional domain $\Omega \subseteq \mathbb{R}^3$ and consider real function spaces V defined over Ω and K defined over Λ . The problem reads as find $(u, l) \in V \times K$ such that

$$\int_{\Omega} (\nabla u \cdot \nabla v + uv) \, dx + \int_{\Lambda} \beta (\Pi u - l) v \, ds = F(v) \quad \forall v \in V, \quad (1)$$

$$\int_{\Lambda} \left(\frac{\partial}{\partial s} l \frac{\partial}{\partial s} k + lk \right) \, ds - \int_{\Lambda} \beta (\Pi u - l) k \, ds = G(k) \quad \forall k \in K. \quad (2)$$

In (1) and (2), $\Pi : V \rightarrow L^2(\Lambda)$ e.g. represents a non-local averaging operator of the form:

$$\Pi u(s) = \frac{1}{2\pi} \int_0^{2\pi} u(x(s, R, \theta)) \, d\theta, \quad (3)$$

where the radius R and the angles θ define circles normal to the curve Λ .

This talk is intended to continue the discussion on how best to support such finite element spaces and associated finite element formulations in the FEniCS framework. Emphasis will be put on key motivating examples, key abstractions and structures and key challenges.

References

- [1] D'Angelo, C. and Quarteroni, A. (2008). On the coupling of 1D and 3D diffusion-reaction equations: application to tissue perfusion problems. *Mathematical Models and Methods in Applied Sciences*, 18(08), 1481-1504.

This research is supported by European Research Council grant #714892 (Waterscales).

Simulations of plasma-object interaction using particle-in-cell method with FEniCS

Diako Darian, University of Oslo, diakod@math.uio.no
Sigvald Marholm, University of Oslo, sigvald.marholm@fys.uio.no
Wojciech J. Miloch, University of Oslo, w.j.miloch@fys.uio.no
Mikael Mortensen, University of Oslo, mikaem@math.uio.no

Keywords: *Plasma-object interaction, Charging of objects in plasmas, Finite element method*

The understanding of plasma-object interaction is of importance for men and space technology in harsh and inhospitable space environments. The flowing plasma in the upper-atmosphere contains energetic megaelectronvolt electrons and ions that may cause damage to spacecrafts and satellites, and represents serious health hazard to astronauts. Due to the complexity of the plasma-object interaction, in order to study realistic space weather conditions, it is advisable to use first-principle numerical simulations, such as with particle-in-cell (PIC) method, where the dynamics of plasma particles can be studied in self-consistent force fields. The use of finite element method to study plasma-object interaction has the advantage of giving considerable flexibility to have an accurate representation of the complex spacecraft or satellite geometries [1].

This poster presents a particle-in-cell code written in FEniCS environment to study plasma-object interaction.

References

- [1] R. Marchand, *PTetra, a tool to simulate low orbit satellite-plasma interaction*, IEEE Transactions on Plasma Science, **40** 2, (2012) 217-229.

Geometric Multigrid in FEniCS

Patrick Farrell, University of Oxford, patrick.farrell@maths.ox.ac.uk
Matteo Croci, University of Oxford, matteo.croci@maths.ox.ac.uk

Keywords: *geometric, multigrid*

In this talk we describe a recent addition to FEniCS: native support for geometric multigrid via PETSc.

Rather than building an additional software project on top of FEniCS, we have included support for GMG inside FEniCS itself. This is achieved by teaching PETSc how to construct the necessary prolongation and restriction operators. That is, we enable the use of `pc_type "mg"` in PETSc.

The main technical difficulty is building a sparse matrix that represents the prolongation operation between two function spaces. Our code for the construction of the prolongation operator works in parallel, for any Lagrange element. The code works even when the domain decompositions of the meshes do not overlap, at the cost of some communication.

The code relies on the user supplying the hierarchy of meshes. The code does *not* rely on the hierarchy being nested or constructed via uniform refinement. In fact, the code works even when the meshes describe different domains (under certain well-understood conditions). This is frequently necessary for complex domains arising in industrial practice.

Its current restrictions are that it does not work for non-Lagrange elements or for mixed finite element problems. Its extension to mixed finite elements is relatively straightforward (and has been achieved in `defcon`, the author's package for bifurcation analysis), but relies on some modifications to `dolfin` that are not yet ready for merging. This may or may not be resolved by the time of the talk.

Contributions to support other element types and multigrid for $H(\text{div})$ and $H(\text{curl})$ problems would be extremely welcome.

A new algorithmic differentiation tool (not only) for FEniCS

Simon W. Funke, Simula Research Laboratory, simon@simula.no,
Sebastian Mitusch, Simula Research Laboratory, sebastkm@student.matnat.uio.no

Keywords: *adjoint model, algorithmic differentiation, higher-order derivatives, dolfin-adjoint, FEniCS*

The derivation and implementation of adjoint models for time-dependent, non-linear PDEs is a challenging task. A common strategy is to apply an algorithmic differentiation tool (AD) which (semi-)automatically derives the adjoint model from the forward model. Specifically for finite-element models, [1] proposed a high-level AD approach, which derives the adjoint by analysing and exploiting the high-level mathematical structure inherent in finite element methods. This idea has shown to provide major benefits compared to traditional low-level AD tools, including near to theoretically optimal performance and natural support of parallelism. However, the high-level AD tool for FEniCS, *dolfin-adjoint*, lacks important features such as differentiation with respect to Dirichlet boundary conditions and higher-order derivatives.

To overcome these limitations, we propose a new algorithmic differentiation tool for FEniCS. The core of this tool is formed by *pyadjoint*, a generic operator overloading AD tool written in Python. *pyadjoint* considers the model as a sequence of arbitrary operations with inputs and outputs. This abstraction can be seen as a generalisation of low and high-level AD tools: operations can be individual floating point operations (as for traditional AD tools), entire systems of differential equations (as for high-level AD tools), or a mix of both. The adjoint developer must overload each relevant model function according to the *pyadjoint* API, and in particular provide implementations for their derivatives. With this information *pyadjoint* records a tape of model operations at runtime and automatically derives and executes the associated adjoint model. Specifically, the support for adjoint FEniCS model is achieved by overloading of the FEniCS API, in particular the creating of new objects such as *Functions*, *Constants*, and overloading operators such as *assemble*, *project* and *solve*, see figure 1.

Python code

```
1 from fenics import *
2 from fenics_adjoint import *
3
4 mesh = UnitSquareMesh(10, 10, 0.1, 0.1)
5 V = FunctionSpace(mesh, "Lagrange", 1)
6
7 u = TrialFunction(V)
8 v = TestFunction(V)
9
10 u_ = Function(V, name="PDE Solution")
11 f = Constant(1.0, name="Control")
12
13 a = inner(grad(u), grad(v))*dx
14 L = f*v*dx
15 bc = DirichletBC(V, 1, "on_boundary")
16
17 solve(a == L, u_, bc)
18
19 J = assemble(u_**2*dx)
20
21 tape = get_working_tape()
22 J.set_initial_adj_input(1.0)
23 tape.evaluate()
24 dJdf = f.get_adj_output()
```

(a)

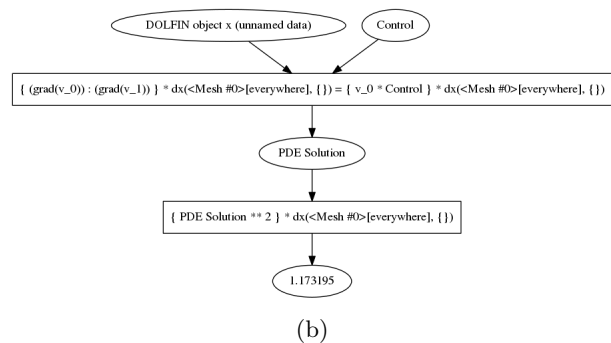


Figure 1: (a) example FEniCS model with adjoint annotations, and (b) the recorded tape visualised as a graph

This poster will present the initial implementation of this new AD tool and demonstrate its highlights:

- Automatic derivation of adjoints of FEniCS models with parallel support;

- Derivatives with respect to nearly any model input, including Dirichlet boundary conditions;
- Support for providing a user-defined Riesz-map for representing the gradient;
- Support for higher-order derivatives in development;
- Extensions to other models (such as Firedrake) possible;
- Small code base

The proposed AD tool is available as open-source under the LGPL 3 license and we warmly welcome contributions.

References

- [1] Patrick E. Farrell, David A. Ham, Simon W. Funke and Marie E. Rognes, *Automated derivation of the adjoint of high-level transient finite element programs*, SIAM Journal on Scientific Computing 35.4, pp. C369-C393, 2013.

Automating Hybridization via Composing Abstractions

Thomas H. Gibson, Imperial College London, t.gibson15@imperial.ac.uk,
David A. Ham, Imperial College London, david.ham@imperial.ac.uk,
Colin J. Cotter, Imperial College London, colin.cotter@imperial.ac.uk

Keywords: *Hybridization, mixed finite elements, elliptic problems.*

The first hybridization of a finite element method was proposed in 1977 for a numerical method for solving the equations of linear elasticity [1]. As a sub-set of static condensation methods, its primary advantage allows one to reduce the number of globally-coupled degrees of freedom. In practice, this requires algebraically manipulating the local discretized systems during the equation assembly process to produce the reduced global problem. Using conventional model development techniques, hybridization of complex discretizations requires manual intervention in intricate numerical code, and this intervention must be repeated every time the model is modified, extended, or debugged.

In contrast, the Firedrake [2] and FEniCS [3] projects take the discretized equations in symbolic form as input, and automatically generates high performance parallel code from this mathematical specification. In this talk, we present a robust abstraction framework within Firedrake for specifying local operations on finite element tensors. By introducing these symbolic operations, and generating code from them, we successfully extend hybridized finite elements for automated simulation.

References

- [1] B. M. Fraeijis De Veubeke. Displacement and equilibrium models in the finite element method, *Stress Analysis*, Wiley, New York, 145-197, 1977.
- [2] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. Mcrae, G. Bercea, G. R. Markall, and P. H. J. Kelly. Firedrake: automating the finite element method by composing abstractions. *ACM Trans. Math. Softw.*, 43(3):24:124:27, 2016.
- [3] A. Logg, G. N. Wells, and J. Hake. Automated Solution of Differential Equations by the Finite Element Method. *Springer*, 2012.

Gmsh 3.0 - Gmsh goes boolean!

Christophe Geuzaine, Université de Liège, cgeuzaine@ulg.ac.be
Jean-François Remacle, Université catholique de Louvain, jean-francois.remacle@uclouvain.be

Keywords: *Gmsh, mesh generation, constructive solid geometry*

In this talk I will give a brief overview of the Gmsh project, and present the new constructive solid geometry features introduced in Gmsh 3.0.

Finite Element Modeling of Elastic Wave Dispersion in Pre-strained Negative Stiffness Honeycombs Using FEniCS

Benjamin M. Goldsberry, The University of Texas at Austin, goldsbe@arlut.utexas.edu,
Michael R. Haberman, The University of Texas at Austin, haberman@utexas.edu

Keywords: *acoustic metamaterials, nonlinear solid mechanics, acoustics, Bloch waves*

An acoustic metamaterial (AMM) is an engineered composite material, often constructed using a periodic lattice, that behaves as an effective material with unconventional dynamic effective properties. AMMs can manipulate and control acoustic and/or elastic waves in ways that are not possible with conventional materials, and are therefore of interest for applications such as acoustic lenses that defy the diffraction limit, material slabs that prevent certain frequencies from propagating, and acoustic cloaks that redirect acoustic signals around an object [1]. Reconfigurable AMMs allow even greater control over the direction of energy propagation by altering the structure through an external stimulus, such as mechanical deformation, which fundamentally changes the effective dynamic properties [2]. In this work, a periodic arrangement of pre-curved beams in a honeycomb configuration, commonly referred to as negative stiffness honeycombs (NSH), is studied as a candidate reconfigurable AMM [3]. This work seeks to explore the degree to which exotic dynamic properties and extreme changes in the modal dispersion can be elicited by uniaxial mechanical deformation of the NSH.

The representative unit cell under consideration is shown in Figure 1(a). The nonlinear mechanical deformation of the unit cell is modeled in FEniCS using the principle of stationary potential energy. In this work, the St. Venant-Kirchhoff strain energy density is used:

$$W(\mathbf{E}) = \frac{\lambda}{2} \text{tr}(\mathbf{E})^2 + \mu \text{tr}(\mathbf{E}^2), \quad (1)$$

where \mathbf{E} is the Lagrangian strain tensor, and λ and μ are the first and second Lamé parameters, respectively. The interior boundaries are stress free, while displacement (Dirichlet) boundary conditions are enforced on the top and bottom boundaries, simulating a uniaxial compression of the NSH. Newton's method is used to solve the nonlinear problem, yielding an equilibrium displacement solution. An example of the displacement field solution for a pre-strain of 3% is shown in Figure 1(c). The time-harmonic acoustic problem is then considered by linearizing the equations of motion about the pre-strain configuration, yielding the vector Helmholtz equation with non-constant coefficients. The resulting expression is written in index notation as [4]:

$$(L_{ijkl}u_{k,l})_{,j} = \rho\omega^2u_i, \quad (2)$$

$$L_{ijkl} = \frac{\partial^2 W}{\partial F_{ij} \partial F_{kl}} \Big|_{u=u_{\text{eq}}}, \quad (3)$$

where \mathbf{L} is the fourth-order mixed elasticity tensor, $\mathbf{F} = \nabla \vec{u} + \mathbf{I}$ is the deformation gradient, \vec{u} is the acoustic displacement vector, \vec{u}_{eq} is the equilibrium displacement vector, ρ is the density, and ω is the angular frequency. The elasticity tensor can be easily constructed in FEniCS using the *diff* function:

Python code

```

1 I = Identity(2)
2 F = grad(u_eq) + I #Equilibrium deformation gradient
3 F = variable(F) #redefine as a variable
4 C = F.T*F #Right Cauchy-Green deformation tensor
5 E = 0.5*(C-I) #Lagrangian strain tensor
6 W = (lmbda/2*(tr(E)**2 + mu*tr(E*E)) #St. Venant-Kirchhoff strain energy density
7 S = diff(W,F) #First Piola-Kirchhoff stress tensor
8 L = diff(S,F) #Mixed elasticity tensor

```

Due to the periodicity of the geometry, the Bloch wave theorem is used to calculate the vibrational modes of the NSH. The assumed form of the acoustic displacement is:

$$\vec{u}(\vec{X}) = \vec{U}(\vec{X})e^{i\vec{K}\cdot\vec{X}}, \quad (4)$$

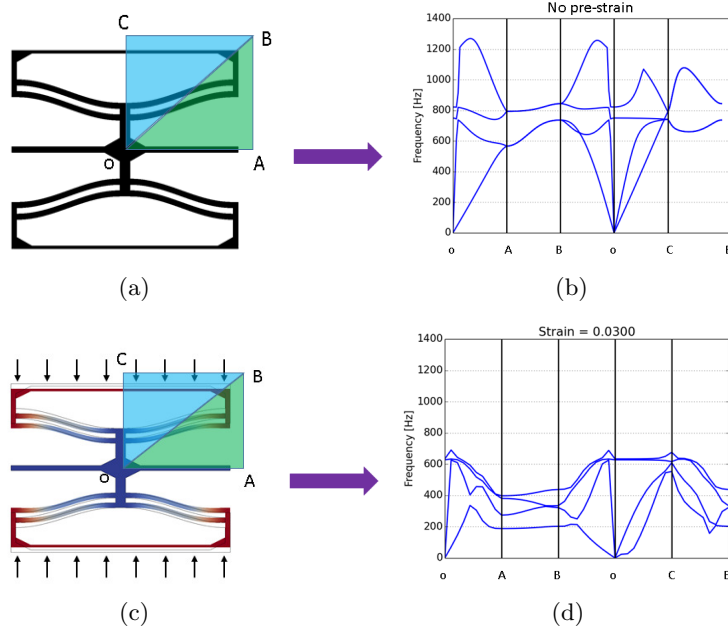


Figure 1: (a) Unit cell with irreducible Brillouin zone overlaid. (b) Dispersion plot displaying the first four eigenfrequencies for the undeformed NSH. (c) Magnitude of the displacement for the deformed unit cell. (d) Dispersion plot displaying the first four eigenfrequencies of the deformed NSH.

where $\vec{U}(\vec{X})$ is a periodic function of the unit cell, and $e^{i\vec{K}\cdot\vec{X}}$ accounts for the phase change across each unit cell. Due to the simplicity of defining periodic boundary conditions in FEniCS, the periodic function $\vec{U}(\vec{X})$ is sought rather than solving for the acoustic displacement field \vec{u} and imposing Floquet boundary conditions. The weak form of Eq. (2) is derived by the Hermitian inner product of both sides with a test vector v . Application of Green's first identity yields the following integral equation in index notation:

$$\int_{\Omega} [L_{ijkl} (U_{k,l}V_{i,j} + K_l K_j U_k V_i)] d\Omega + i \int_{\Omega} [L_{ijkl} (U_k K_l V_{i,j} - U_{k,l} K_j V_i)] d\Omega = \rho \omega^2 \int_{\Omega} U_i V_i d\Omega. \quad (5)$$

Each term of the above equation is discretized and assembled into a compressed sparse column matrix to yield the following general eigenvalue problem:

$$AU = \omega^2 BU \quad (6)$$

$$A = A_r + iA_i, \quad (7)$$

where A_r and A_i are the matrix discretization of the first and second term of Eq. (5), respectively. The complete band diagram can be characterized by solving Eq. (6) for Bloch wave vectors that span the irreducible Brillouin zone, overlaid on the unit cell in Figure 1(a), which is determined by the reciprocal lattice [5]. The effect of pre-strain on the dispersion properties of the NSH structure is shown in Figure 1(d). One clearly observes that the band structure is significantly altered due to the application of a small uniaxial pre-strain.

This talk will highlight the use of FEniCS to compute the dispersion properties of NSH, as well as describe how this model can fit into a design framework for finding optimal NSH configurations for particular applications.

References

- [1] S. A. Cummer, J. Christensen, A. Alú: Controlling sound with acoustic metamaterials, *Nature Reviews Materials*, 1 (2016), 1-13.
- [2] M. Haberman, A. Norris: Acoustic Metamaterials, *Acoustics Today* 12, no. 3 (2016), 31-39.
- [3] D. M. Correa, T. Klatt, S. Cortes, M. Haberman, D. Kovar, C. Seepersad: Negative stiffness honeycombs for recoverable shock isolation, *Rapid Prototyping Journal* 21, no. 2, 193-200.
- [4] E. B. Tadmor, R. E. Miller, R. S. Elliot, *Continuum Mechanics and Thermodynamics: From Fundamental Concepts to Governing Equations*, United Kingdom: Cambridge University Press, 2012.
- [5] L. Brillouin, *Wave Propagation in Periodic Structures*, McGraw-Hill, 1953.

Equilibrated Warping: Finite Element Image Correlation with Mechanical Regularization

Martin Genet, *École Polytechnique*, martin.genet@polytechnique.edu,
Lik Chuan Lee, *Michigan State University*, lcee@egr.msu.edu,
Sebastian Kozerke, *ETH Zurich*, kozerke@biomed.ee.ethz.ch

Keywords: *Finite elements, Image Correlation, Mechanical Regularization, Hyperelasticity*

Context. Image correlation/registration is playing an increasing role in many domains such as biomedical engineering [1]. Despite significant progress in the past decades, robustness, efficiency and precision of existing methods and tools must still be improved to translate them into medical and engineering applications. This abstract describes a finite element-based image correlation method with, as regularization, a novel continuum large deformation formulation of the equilibrium gap principle (introduced in [2] at the discrete level for linearized elasticity).

Methods. Let us denote I_0 & I_t as the intensity fields of two images representing an object occupying the domains Ω_0 & Ω_t in the reference and deformed states, respectively. The problem is to find the mapping $\underline{\varphi}$ between Ω_0 & Ω_t , or equivalently the displacement field \underline{U} ($\underline{\varphi}(\underline{X}) = \underline{X} + \underline{U}(\underline{X})$):

$$\text{find } \underline{U} = \operatorname{argmin}_{\{\underline{U}^*\}} \left\{ J^2(\underline{U}^*) = \frac{1-\beta}{2} \int_{\Omega_0} (I_t \circ \underline{\varphi}^* - I_0)^2 d\Omega_0 + \beta \psi^{\text{reg}}(\underline{U}^*) \right\}, \quad (1)$$

where ψ^{reg} is required to regularize the otherwise ill-posed problem, and β is the regularization strength. Many regularizers have been proposed. One common approach, called hyperelastic warping [3], is to use the strain energy potential directly, i.e., $\psi^{\text{reg,hyper}} = \rho_0 \psi$, thus penalizing strain. Here we propose an alternate regularizer, which essentially penalizes any deviation from the solution of an hyperelastic body in equilibrium with arbitrary external loads: $\psi^{\text{reg,equil}} = \frac{1}{2} \|\underline{\operatorname{div}}(\underline{P})\|_{L^2(\Omega_0)}^2$, where $\underline{P} = \frac{\partial \rho_0 \psi}{\partial \underline{F}}$ is the first Piola-Kirchhoff stress tensor. However, we discretize Problem (1) using standard Lagrange elements, so that \underline{P} belongs to $L^2(\Omega_0)$ but not $H(\operatorname{div}; \Omega_0)$. Thus, the following equivalent norm is used instead:

$$\psi^{\text{reg,equil}} = \sum_K \frac{1}{2} \|\underline{\operatorname{div}}(\underline{P})\|_{L^2(K)}^2 + \sum_F \frac{1}{2h} [[\underline{P} \cdot \underline{N}]]_{L^2(K)}^2, \quad (2)$$

where K denotes the set of finite elements, F the interior faces, h a characteristic length of the mesh.

Results on synthetic data. Here we consider the simple problem of a uniformly compressed square domain, and study the influence of the regularization strength β on the computed strain. Figure 1 shows the initial and final images superimposed with the undeformed and deformed mesh obtained.

In case of hyperelastic warping, if the regularization strength is close to 1, the mesh does not deform. And when regularization strength decreases, measured strain converges toward the exact value. For noise-free images, it does converge exactly. For noisy images, there is an optimum where the mean strain is close to the exact solution and standard deviation is still limited.

Conversely, with equilibrated warping, the registration is almost perfect over a wide range of regularization strengths, even on noisy images.

Results on *in vivo* images. Here we consider 3D CSPAMM cardiac magnetic resonance images of a healthy human subject. Figure 2 shows the resulting strains computed by both methods. Main features of left ventricular deformation are well captured. Equilibrated warping produces larger absolute strain values than hyperelastic warping, closer to expected values [1].

Conclusion. Equilibrated warping is a powerful method for non-rigid registration of images involving large deformation. Penalizing the equilibrium gap regularizes the image correlation problem, even in presence of noise, and without affecting strain measurement. The method has been implemented based on FEniCS and VTK, providing an efficient tool for 2D & 3D images registration.

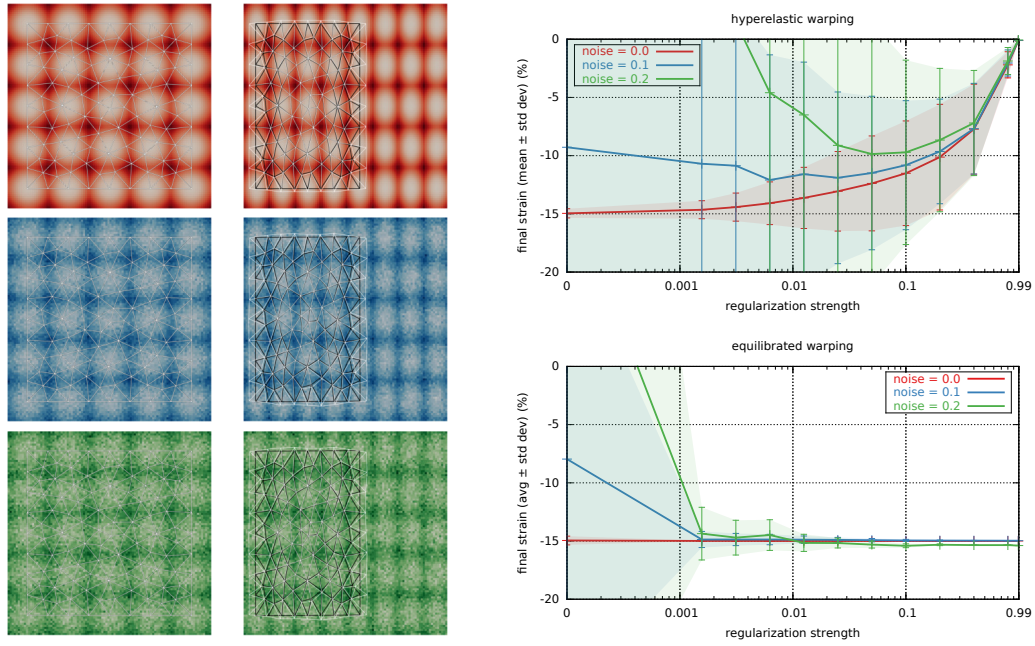


Figure 1: Results on synthetic data. Left: hyperelastic (white mesh) *vs.* equilibrated (black mesh) warping, for a regularization strength of 0.1. Right: influence of regularization strength on hyperelastic (top) and equilibrated (bottom) warping strains. Ground truth is -15% homogeneous strain.

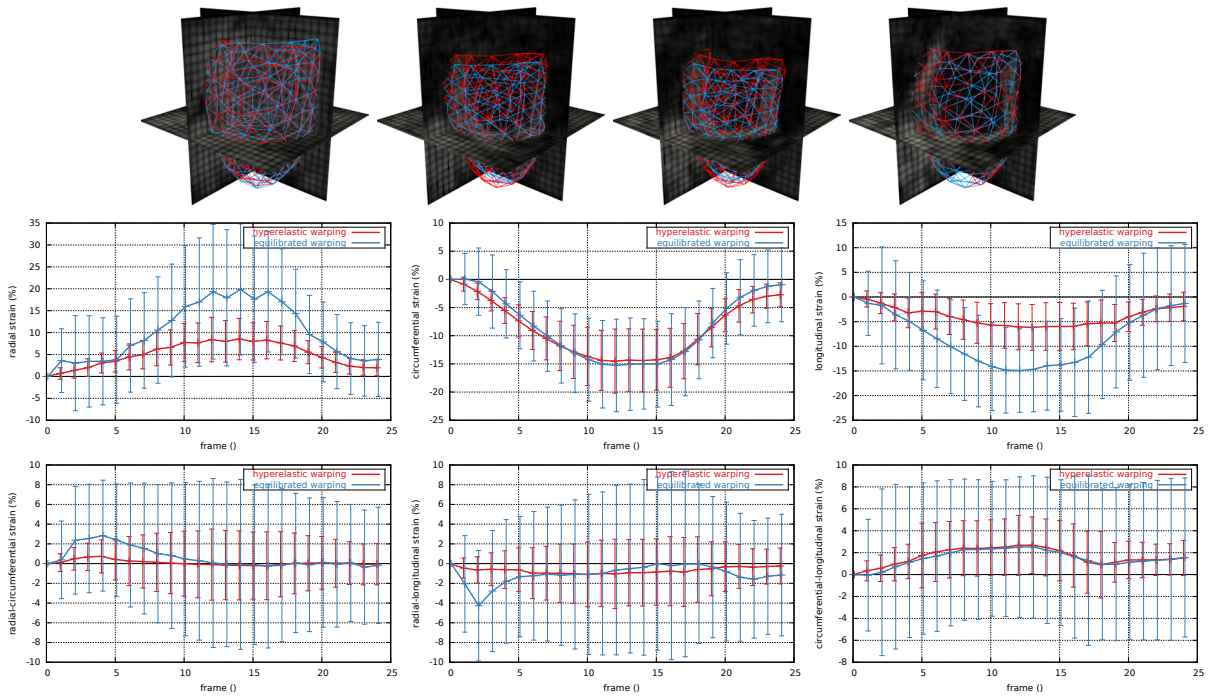


Figure 2: Results on *in vivo* data, hyperelastic (red) *vs.* equilibrated (blue) warping. Top: Sequence of 3D CSPAMM images with superimposed mesh. Bottom: Sequence of local strain components.

References

1. Tobon-Gomez, C. *et al.* Benchmarking Framework for Myocardial Tracking and Deformation Algorithms: An Open Access Database. *Medical Image Analysis* **17**. 00057, 632–648 (Aug. 2013).
2. Claire, D. *et al.* A Finite Element Formulation to Identify Damage Fields: The Equilibrium Gap Method. *International Journal for Numerical Methods in Engineering* **61**. 00164, 189–208 (2004).
3. Veress, A. I. *et al.* Measurement of Strain in the Left Ventricle during Diastole with Cine-MRI and Deformable Image Registration. *Journal of Biomechanical Engineering* **127**, 1195 (July 2005).

podS: A Parallel Multilevel Monte Carlo Framework

Hugo Hadfield, University of Cambridge, hh409@cam.ac.uk
H. Juliette T. Unwin, University of Cambridge, hjtu2@cam.ac.uk,
Nathan Sime, University of Cambridge, njcs4@cam.ac.uk,
Garth N. Wells, University of Cambridge, gnw20@cam.ac.uk,

Keywords: *uncertainty quantification, multilevel Monte Carlo, MPI*

Multilevel Monte Carlo (MLMC) is an improvement upon traditional Monte Carlo methods. It is designed to dramatically reduce the computational cost of solving certain stochastic partial differential equations (SPDEs) and other problems involving uncertainty. MLMC is an embarrassingly parallel technique. In theory any number of problem realisations could be solved in parallel. Reality forces researchers to use computer systems with finite memory and processing cores. The problem we face is: given the architecture of the computer system on which we wish to compute a stochastic approximation using the MLMC method, how do we best assign system resources? To this end, we have developed the C++ library podS. podS is designed to automatically and optimally allocate resources for an MLMC problem over a variety of computer architectures, from laptop to supercomputer. The goal is to allow researchers to focus on the problems they want to solve and not the intricacies of parallelising code on different hardware.

We introduce podS and show examples of its use with FEniCS and Docker to compute approximations of SPDEs. We also discuss the technical challenges that we faced designing the library.

FInAT: Automating Optimal Code Generation for High-Order Finite Element Methods

Miklós Homolya, Imperial College London, m.homolya14@imperial.ac.uk,
Fabio Luporini, Imperial College London, f.luporini12@imperial.ac.uk,
David Ham, Imperial College London, david.ham@imperial.ac.uk,
Robert C. Kirby, Baylor University, Robert_Kirby@baylor.edu

Keywords: *tensor product elements, high order elements, sum factorisation, code generation*

FEniCS and Firedrake [1] have demonstrated that code generation is a key technology in enabling the productive exploitation of advanced numerical methods for complex systems of equations. However, no existing high-level code generation systems facilitate the generation of optimal implementations of high-order finite element methods.

The optimal algorithms for high-order finite elements rely on the exploitation of structure within the finite elements. For example, the tabulation matrix Φ of a tensor product element can be written as the outer product of the tabulation matrices $\Phi^{(1)}$ and $\Phi^{(2)}$ of the factor elements, that is

$$\Phi = \Phi^{(1)} \otimes \Phi^{(2)}, \quad \Phi_{(p_1, p_2), (i_1, i_2)} = \Phi_{p_1, i_1}^{(1)} \Phi_{p_2, i_2}^{(2)}.$$

This structure enables the reordering of assembly loop nests to avoid redundant computations, a technique known as *sum factorisation*. Considering its implementation in either FEniCS or Firedrake, a particular issue is that FIAT [2], the finite element library, is unable to express such structure within the elements.

Here we present FInAT [3], a smarter library of finite elements. While FIAT evaluates the element basis as a table of numerical values, FInAT provides a symbolic expression for this evaluation. Thus the tensor product structure is preserved, so that the form compiler or a suitable optimisation tool can sum factorise the assembly kernel.

Through its integration with TSFC, FInAT is incorporated in the Firedrake system. TSFC [4] is the form compiler in Firedrake, which together with COFFEE [5] can restructure the assembly loop nests to produce sum factored algorithms with optimal complexity.

References

- [1] Rathgeber et al.: Firedrake: automating the finite element method by composing abstractions, *ACM TOMS*, 43 (2016), 24:1–24:27.
- [2] Robert C. Kirby: FIAT, A New Paradigm for Computing Finite Element Basis Functions, *ACM TOMS*, 30 (2004), 502–516.
- [3] FInAT Is not A Tabulator (FInAT), <https://github.com/FInAT/FInAT>.
- [4] Homolya et al.: TSFC, a structure-preserving form compiler. *In preparation*.
- [5] Luporini et al.: An algorithm for the optimization of finite element integration loops, *ACM TOMS*, 44 (2017), 3:1–3:26.

Tracer transport in microvasculature: A case study on coupling 1D-3D problems in FEniCS

Karl Erik Holter, University of Oslo, karleh@math.uio.no,
Miroslav Kuchta, University of Oslo, miroslav.kuchta@gmail.com,
Kent-Andre Mardal, University of Oslo, kent-and@math.uio.no

Keywords: *blood perfusion, codimension 2 subdomains, microvasculature, fenicsii*

Medical imaging in which a tracer is injected in the blood stream in order to assess how the tracer diffuses into the microstructure of the surrounding tissue is popular in diagnostics. In mathematical models of this process, the viscous flow in the vasculature is coupled to the surrounding porous tissue. These models are in principle well established, but a main challenge is that the vasculature forms a dense network of almost space-filling curves.

In this work, we reduce the order of such a model by replacing the 3D vessels with 1D line segments. As the blood flow in the vessels we study is approximately Poiseuille and the vessels are already only a few microns wide, this is a good approximation. This reduces the computational load, but introduces some difficulties with coupling a 1D problem (fluid flow in the vessels) with a 3D problem (fluid flow in the surrounding tissue).

Our approach follows that of [2] and [3] in which fluid flow is Poiseuille in the vessels and Darcy in the surrounding tissue, and tracer distribution is driven by the convection-diffusion equation. In this work we compare the modeling with well-known results from Tofts type of modelling in DCE-MR imaging used for diagnostics of cancer. As FEniCS has limited support for coupling 1D and 3D problems, we use the `fenicsii` package [1] developed by M. Kuchta.

In our talk, we give an overview of the underlying biological motivation, the mathematical formulation of the coupled 1D-3D problem and the challenges involved in solving it in FEniCS. We close with a discussion of our results.

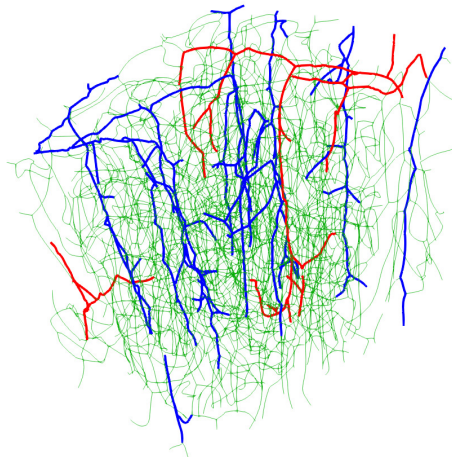


Figure 1: The microvasculature in a 0.22 mm^3 section of rodent cortex.

References

- [1] M. Kuchta, `fenicsii`. https://github.com/MiroK/fenics_ii
- [2] C. D'Angelo, A. Quarteroni: On the coupling of 1D and 3D diffusion-reaction equations. Applications to tissue perfusion problems *MOXReport*, 28 (2008).
- [3] L. Cattaneo, P. Zunino: A computational model of drug delivery through microcirculation to compare different tumor treatments. Applications to tissue perfusion problems *Int. J. Numer. Meth. Biomed. Engng.*, 30 (2014), 1347-1371.

Total Variation Image Reconstruction on Smooth Surfaces

Marc Herrmann, University of Wuerzburg, marc.herrmann@uni-wuerzburg.de,
Roland Herzog, University of Chemnitz, roland.herzog@mathematik.tu-chemnitz.de,
Heiko Kröner, University of Hamburg, heiko.kroener@uni-hamburg.de,
Stephan Schmidt, University of Wuerzburg, stephan.schmidt@uni-wuerzburg.de,
José Vidal, University of Chemnitz, jose.vidal-nunez@mathematik.tu-chemnitz.de

Keywords: *total bounded variation, Fenchel predual problem, interior-point methods, image reconstruction, image denoising, image inpainting, surfaces*

An analog of the total variation image reconstruction approach [7] for images f , defined on smooth surfaces, is introduced, which can be used as an application for 3D scanner data of objects and their textures. A field, where the texture is of great importance, is neuroimaging, where random noise enters due to eddy-current distortions or physical motion during the magnetic resonance imaging (MRI) process [2]. For this purpose, we consider the image reconstruction problem

$$\begin{cases} \text{Minimize} & \frac{1}{2} \int_S |Ku - f|^2 ds + \frac{\alpha}{2} \int_S |u|^2 ds + \beta \int_S |\nabla u| \\ \text{over} & u \in BV(S), \end{cases} \quad (1)$$

where $S \subset \mathbb{R}^3$ is a smooth, compact, orientable and connected surface without boundary. The observed data $f \in L^2(S)$, parameters $\beta > 0$, $\alpha \geq 0$ and the observation operator $K \in \mathcal{L}(L^2(S))$ are given. Furthermore, $BV(S)$ denotes the space of functions of bounded variation on the surface S . A function $u \in L^1(S)$ belongs to $BV(S)$ if the TV-seminorm defined by

$$\int_S |\nabla u| = \sup \left\{ \int_S u \operatorname{div} \boldsymbol{\eta} ds : \boldsymbol{\eta} \in \mathbf{V} \right\} \quad (2)$$

is finite, where $\mathbf{V} = \{\boldsymbol{\eta} \in \mathbf{C}_c^\infty(\operatorname{int} S, TS) : \boldsymbol{\eta} \text{ is a vector field, } |\boldsymbol{\eta}(p)|_2 \leq 1 \text{ for all } p \in S\}$ with TS as the *tangent bundle* of S . Note that $BV(S) \hookrightarrow L^2(S)$ and hence, the integrals in (1) are well defined.

The non-smoothness of TV-seminorm is dealt with a duality approach, see [1, 3, 4]. This leads to the predual problem of (1), which is a quadratic optimization problem for the vector field $\boldsymbol{p} \in \mathbf{H}(\operatorname{div}; S) := \{\boldsymbol{v} \in \mathbf{L}^2(S; TS) \mid \operatorname{div} \boldsymbol{v} \in L^2(S)\}$ with pointwise inequality constraints on the surface.

For the application, we concentrate on the classical denoising problem, where $K = \operatorname{id}$ holds, see Figure 1. The numerical studies are based on geometries obtained by scanning real objects with the Artec Eva 3D scanner. The scanner provides Wavefront .obj files, which contain a description of the geometry via vertices and triangles. The surface texture is provided as a 2D flat bitmap file and a mapping of each physical surface triangle into said bitmap. Thus, originally the textured object is described by a varying number of pixels glued onto each surface triangle. Due to the impossibility of continuously mapping a closed surface onto the flat bitmap, there are discontinuities in the bitmap. Essentially, two adjacent triangles on the surface can be part of discontinuous regions in the texture file. In order to apply our novel denoising scheme, the above mentioned Wavefront object including the texture needs to be made available to the finite element library used to discretize the predual problem. One way of achieving this is to provide the texture data f at each quadrature point. However, for ease of implementation and processing within the finite element framework FEniCS, we converted the textured object into the finite element setting. To account for both natural discontinuities in the texture as well as the discontinuity of the surface-to-texture mapping, we chose a discontinuous Lagrange finite element representation of the texture data f . Thus, u and f are elements of the \mathcal{DG}_r finite element space on the surface.

To carry out the texture preprocessing, we compute the spatial location for each degree of freedom of the surface DG function f within the texture bitmap and use the respective gray value. For color textures, this is realized via a vector valued DG function on the surfaces with values in the RGB color space. In the original Wavefront object, each surface triangle will usually obtain data from multiple texture pixels. Thus, in order to maintain a similar quality of the texture in the DG setting, higher order finite element spaces are needed, depending on the quality of the scan. Although in the original Wavefront object, the number of pixels per triangle may vary significantly, whereas we use a constant finite element order, typically $r = 2$ or $r = 3$. Because we actually solve the predual problem via an interior-point method,



Figure 1: Test-case: noise free, noisy and denoised for $\beta = 0.2$ and $\beta = 0.5$. The scan-data are provided by the Artec Group inc. (www.artec3d.com) under the Creative Commons Attribution 3.0 Unported License

we need to find $\mathbf{p} \in \mathbf{H}(\text{div}; S)$ before recovering the image u . We employ a conforming discretization by surface Raviart–Thomas finite elements. The Raviart–Thomas element space \mathcal{RT}_{r+1} is designed to be the smallest polynomial space with $\mathcal{RT}_{r+1}|_K \subset \mathcal{P}_{r+1}$ for every triangle K such that the divergence maps onto \mathcal{P}_r [5]. For denoising, we have the relation $u = \text{id}^{-1}(\text{div } \mathbf{p} + f)$. Except of the divergence, there is no differentiation. Therefore we choose to discretize $u \in \mathcal{DG}_r$, $\mathbf{p} \in \mathcal{RT}_{r+1}$ and $f \in \mathcal{DG}_r$.

References

- [1] T.F. Chan, G.H. Golub, P. Mulet: A nonlinear primal-dual method for total variation-based image restoration, *SIAM Journal on Scientific Computing*, 6 (1999), 1964-1977.
- [2] Y.N. Chang and H.H. Chang: Automatic brain MR image denoising based on texture feature-based artificial neural networks, *Bio-Medical Material and Engineering*, 26 (2015), 1275-1282.
- [3] M. Hintermüller, K. Kunisch: Total bounded variation regularization as a bilaterally constrained optimization problem, *SIAM Journal on Applied Mathematics*, 64 (2004), 1311-1333.
- [4] R. Lai and T. F. Chan: A framework for intrinsic image processing on surfaces, *Computer Vision and Image Understanding*, 115 (2011), 1647-1661.
- [5] A. Logg, K.A. Mardal, G. Wells, eds., *Automated Solution of Differential Equations by the Finite Element Method*, vol. 84 of Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2012.
- [6] M.E. Rognes, D.A. Ham, C.J. Cotter, and A.T.T. McRae: Automating the solution of PDEs on the sphere and other manifolds in FEniCS 1.2, *Geoscientific Model Development*, 6 (2013), 2099-2119.
- [7] L.I. Rudin, S. Osher, and E. Fatemi: Nonlinear total variation based noise removal algorithms, *Physica D*, 60 (1992), 259-268.
- [8] A. Schiela: Barrier methods for optimal control problems with state constraints, *SIAM Journal on Optimization*, 20 (2009), 1002-1031.

Discontinuous solution of the Laplace equation in a mixed finite element method and level-set setting

Michal HABERA, Charles University, habera@karlin.mff.cuni.cz,
Ondřej SOUČEK, Charles University, soucek@karel.troja.mff.cuni.cz,
Josef MÁLEK, Charles University, malek@karlin.mff.cuni.cz,
Jaroslav HRON, Charles University, hron@karlin.mff.cuni.cz

Keywords: *Nonlinear elliptic interface, level-set, electrochemistry, moving discontinuity*

Electrochemical sciences (batteries, chemical cells, corrosion, metal oxidation) are mostly based on Butler-Volmer kinetic relation [2], which couples electric current, \mathbf{j} , passing through some interface with an electrostatic potential jump, $[[\phi]]$, on the interface, i.e.

$$\mathbf{j} \cdot \mathbf{n} \propto \sinh([[\phi]]).$$

Additionally, in the real world problems, the interface of discontinuity evolves in time and is implicitly given. This type of issue could be characterised as a nonlinear elliptic interface problem.

In this contribution we show very natural interpretation of this problem in terms of mixed Poisson problem (saddle-point structure), where couple $(\mathbf{j}, \phi) \in \mathbf{H}(\text{div}) \times L^2$ is sought for, such that

$$\begin{aligned} a_1(\mathbf{j}, \boldsymbol{\tau}) + a_2(\mathbf{j}, \boldsymbol{\tau}) + b(\boldsymbol{\tau}, \phi) &= F(\boldsymbol{\tau}) \quad \forall \boldsymbol{\tau} \in \mathbf{H}, & (1) \\ b(\mathbf{j}, v) &= 0 \quad \forall v \in Q & (2) \end{aligned}$$

with a_2 being our special nonlinear functional representing the jump condition.

We prove the existence and uniqueness for infinite dimensional problem for Butler-Volmer kinetic relation (Lipschitz continuous and monotone).

From numerical point of view, stable finite element discretization of (1) and (2) requires the use of elements satisfying discrete inf-sup condition. Here FEniCS [1] comes in handy, where $\mathbf{H}(\text{div})$ conforming families $\mathcal{BDM}_k, \mathcal{BDFM}_k, \mathcal{RT}_k$ are available.

The implicit interface description is handled with characteristic level-set method (meshless technique). This method belongs to group of diffuse-interface methods while thickness of the interface is proportional to minimal element diameter.

The level-set function is evolved according to the normal component of a found solution, $\mathbf{j} \cdot \mathbf{n}$, and reinitialization [3] maintains its characteristic level-set property.

In summary, we are proposing a general scheme for a nonlinear electrochemical elliptic interface problems and we show its well-posedness. With the aid of FEniCS, orders of convergence for steady case and simple numerical examples are computed.




Figure 1: Example of a one-dimensional numerical solution. Several mesh densities are plotted, where h denotes a minimal element diameter.

References

- [1] M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. Rognes, and G. Wells, The FEniCS Project version 1.5, *Archive of Numerical Software*, 3 (2015). URL: <http://journals.ub.uni-heidelberg.de/index.php/ans/article/view/20553>, doi:10.11588/ans.2015.100.20553.
- [2] John S. Newman and Karen E. Thomas-Alyea. *Electrochemical systems*. 3rd ed. Hoboken, N.J.: J. Wiley, c2004. ISBN 0471477567.
- [3] E. Olsson and G. Kreiss. A conservative level-set method for two phase flow. *Journal of Computational Physics*[online]. 2005, 210(1): 225-246. DOI: 10.1016/j.jcp.2005.04.007.

Uncertainty quantification for soft tissue biomechanics

Paul Hauseux, University of Luxembourg, paul.hauseux@uni.lu.

Jack. S. Hale, University of Luxembourg, jack.hale@uni.lu.

Stéphane P. A. Bordas, University of Luxembourg,

University of Western Australia, Cardiff University, stephane.bordas@uni.lu.

Keywords: *Monte Carlo methods, sensitivity derivatives, global sensitivity analysis, parallel computing, hyperelastic soft tissue models.*

Monte Carlo methods are used to assess the effects of uncertainty in material parameters in soft tissue models. We present a sensitivity derivative Monte Carlo method [1, 2] to provide effectively statistical results. This technique reduces the error in the Monte Carlo estimator and therefore the workload compared to the standard Monte Carlo approach. A global sensitivity analysis is also performed to study how the uncertainty of the inputs influence the outputs and to determine the uncertain parameters which have the most influence on the variance of a given quantity of interest. We implement our forward and tangent linear model solvers using DOLFIN [3] and we use the Python toolbox chaospy [4] to generate stochastic objects. Numerical results of a FE stochastic analysis of brain deformation are presented and discussed. We implemented two hyperelastic soft tissue models: a classical Mooney-Rivlin material and an anisotropic Holzapfel and Ogden model [4].

References

- [1] Y. Cao, M. Y. Hussaini and T. A. Zang, Exploitation of Sensitivity Derivatives for Improving Sampling Methods, *AIAA Journal*, Vol 42, 4, 2004
- [2] P. Hauseux, J.S. Hale and S.P.A Bordas, Accelerating Monte Carlo estimation with derivatives of high-level finite element models, *Computer Methods in Applied Mechanics and Engineering*, Vol 318, 2017
- [3] A. Logg, K.A. Mardal and G. N. Wells: Automated Solution of Differential Equations by the Finite Element Method, *Springer*, Vol 84, 2012
- [4] J. Feinberg, H. P. Langtangen, Chaospy: An open source tool for designing methods of uncertainty quantification, *Journal of Computational Science*, Vol 11, 2015
- [5] G.A. Holzapfel and R.W. Ogden, Constitutive modelling of passive myocardium: a structurally based framework for material characterization, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol 367, (1902), 2009

Trace Constrained Problems in FEniCS

Karl Erik Holter, University of Oslo & Simula Research Laboratory, karleh@ifi.uio.no,
Miroslav Kuchta, University of Oslo, mirok@math.uio.no,
Kent-Andre Mardal, University of Oslo & Simula Research Laboratory,
kent-and@simula.no

Keywords: *trace operator, preconditioning, saddle point systems*

Mixed function spaces with subspaces defined over domains with different topological dimensions arise in numerous applications. In [1] the Lagrange multiplier defined on domain boundary is used to enforce Dirichlet boundary conditions on the solution of the Poisson problem. More recently, [4, 5] seek electric potential of the extracellular/cellular domains in spaces of functions over \mathbb{R}^d while the potential of the membrane separating the two media is represented as function defined over $d - 1$ dimensional embedded manifold. In reduced order modeling of interstitial flows [2] model tissue as a three dimensional structure while the vasculature is modeled as a one dimensional curve.

Due to the dimensional gap between the domains variational forms stemming from these problems are currently not supported in FEniCS and as such the platform cannot be used for the applications unless various hacks are employed. In this talk we present our package (FEniCS)_{ii} which offers set of primitives, in particular the trace operator and assembler, allowing for proper discretization of the trace constrained problems in FEniCS. The package is built on top of cbc.block [3] and therefore the assembled systems are to be solved iteratively. For the above listed applications we shall discuss suitable preconditioners.

References

- [1] I. Babuška, *The finite element method with Lagrangian multipliers*, Numerische Mathematik, 20 (1973), pp. 179–192.
- [2] C. D’Angelo and A. Quarteroni, *On the coupling of 1D and 3D diffusion-reaction equations: Application to tissue perfusion problems*, Mathematical Models and Methods in Applied Sciences, 18 (2008), pp. 1481–1504.
- [3] K.-A. Mardal and J. B. Hoga, *Block preconditioning of systems of PDEs*, in Automated Solution of Differential Equations by the Finite Element Method, G. N. Wells et al. A. Logg, K.-A. Mardal, ed., Springer, 2012.
- [4] A. Agudelo-Toro and A. Neef, *Computationally efficient simulation of electrical activity at cell membranes interacting with self-generated and externally imposed electric fields*, Journal of Neural Engineering, 10 (2013), p. 026019.
- [5] A. Tveito, K. Horgmo Jæger, M. Kuchta, K.-A. Mardal and M. Rognes, *Accurate numerical modeling of small collections of cardiac cells*, in preparation (2017).

Generating matrix actions with TSFC and Loo.py

Miklós Homolya, Imperial College, London, m.homolya14@imperial.edu,
Andreas Klöckner, University of Illinois at Urbana-Champaign, andreask@illinois.edu,
Robert C. Kirby, Baylor University, robert_kirby@baylor.edu

Keywords: *code generation, form compiler, platform-independence*

Generating efficient finite element code for multiple architectural targets requires great flexibility. In this work, we describe a chain of transformations that starts with UFL [1] and goes through TSFC [2] to produce low-level code for computing the matrix-free action of bilinear forms. We convert TSFC’s intermediate representation of an element-local kernel into an equivalent kernel in Loo.py [], which is a transformation-based tool that allows tedious and error-prone conversion of high-level loop descriptions into tight, low-level code in OpenCL, CUDA, or other device languages.

The element kernel requires several kinds of transformations before its compilation and use, all of which demonstrate important features of Loo.py. For one, we “batch” the element kernel to iterate over many cells. Then, we generate and fuse in additional kernels that scatter and gather between element-level and assembled processor-global storage for each field. Beyond these, we may apply transformations to produce intra-node parallelism, loop-unrolling, or other performance optimizations. The resulting kernel can be invoked on each MPI rank as a kind of Firedrake “implicit matrix” and hence used inside a Krylov method.

References

- [1] Martin S. Alnæs, Anders Logg, Kristian B. Ølgaard, Marie E. Rognes, and Garth N. Wells: Unified Form Language: A domain-specific language for weak formulations of partial differential equations, ACM TOMS, 40:2 (2014), 9:1-9:37.
- [2] Homolya et al: TSFC, a structure-preserving form compiler. *In preparation*.
- [3] Andreas Klöckner, Loo.py: transformation-based code generation for GPUs and CPUs, Proceedings of ARRAY’14: ACM SIGPLAN Workshop on Libraries, Languages, and Compilers for Array Programming (2014).

Solving Poisson's equation on the Microsoft HoloLens

Anders Logg, Chalmers University of Technology, logg@chalmers.se

Carl Lundholm, Chalmers University of Technology, carlun@chalmers.se

Magne Nordaas, Chalmers University of Technology, magne@chalmers.se

Keywords: HoloLens, augmented reality, Poisson's equation, finite element method

We develop an app for solving Poisson's equation with the finite element method using Microsoft's augmented reality glasses HoloLens, see e.g. [1]. The idea with the HoloLens app is to set up and solve a Poisson problem in a real world room where the HoloLens user is located and then visualize the computed solution in the room. The app works by first letting the HoloLens create a surface mesh of the surroundings through a spatial scan, see Figure 1. The surface mesh is then used to construct a geometry that defines the solution domain by a polyhedral representation of the room. A tetrahedral mesh is then generated from the geometry. The user may provide problem data by placing sources in the room and setting boundary conditions on the walls, floor, and ceiling. The finite element method is then used to assemble the Poisson system before it is solved. Finally the computed solution may be visualized in the room.

The development environment for HoloLens apps consists of the game engine Unity and the IDE Microsoft Visual Studio. Projects are initially started in Unity and then exported to Visual Studio where the coding takes place. The programming language used is C#. For the finite element assembly of the system we use the FEniCS form compiler (FFC) to compute the element stiffness matrix. Not only does this automatically take care of steps in going from the variational formulation to the linear system, but it also makes it easier to generalize the app to other types of differential equations. Say for example that we would like to know how a dangerous substance spreads in a room after a leak has sprung. The heat equation could be used as a simplistic model for describing this situation, potentially making apps like this useful in building planning and safety engineering.



Figure 1: *Left:* Microsoft HoloLens. *Right:* Conceptual demonstration of a spatial scan.

References

- [1] <https://www.microsoft.com/microsoft-hololens/en-us>

Slope Limiting of Divergence Free Discontinuous Galerkin Vector Fields in the Context of Two-Phase Flows

Tormod Landet, University of Oslo, tormodla@math.uio.no,
Mikael Mortensen, University of Oslo, mikaem@math.uio.no,
Kent-Andre Mardal, University of Oslo & Simula, kent-and@simula.no

Keywords: *two phase flow, slope limiter, mass conservation, Navier-Stokes*

We investigate an exactly mass conserving higher order interior penalty discontinuous Galerkin method for the incompressible and variable density Navier-Stokes equations. Our intended application is low viscosity flows with density fields containing large and sharp jumps such as simulation of multi phase flows with a water/air free surface embedded in the computational domain. The equation system to be solved is

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} \right) = \nabla \cdot \mu \nabla \mathbf{u} - \nabla p + \rho \mathbf{g}, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2)$$

$$\frac{\partial \rho}{\partial t} + \mathbf{u} \cdot \nabla \rho = 0. \quad (3)$$

Sharp jumps in the solution gives rise to convective instabilities due to Gibbs phenomena. For our application we see these instabilities clearly in the momentum equation when used with a discontinuous density field. In most cases the simulations will almost immediately break down due to exponential growth of the velocity in localised regions near the discontinuity due to the factor 1000 jump in density.

In the context of DG methods there are existing methods known as *slope limiters* for eliminating these non-linear convective instabilities in scalar transport equations. We extend the existing work with an approximate slope limiter for divergence free, solenoidal, vector fields and compare it with an existing hierarchical Taylor based scalar limiter applied to each of the velocity components.

We investigate the combination of the solenoidal approximate limiter with the stable hierarchical Taylor limiter and establish a higher order method for two phase flows that is mass conserving and has no unphysical smearing of the fluid properties near the free surface as well as no Gibbs instabilities due to the jump in density.

Across the great divide: composable block preconditioning from UFL

Lawrence Mitchell, Imperial College London, lawrence.mitchell@imperial.ac.uk,
Robert Kirby, Baylor University, robert_kirby@baylor.edu

Keywords: *Block preconditioning, multiphysics, UFL*

UFL and the FEniCS problem solving language present a beautiful set of abstractions for succinctly writing variational problems. This removes a large challenge in the development of complex models. Preconditioning the resulting systems of equations is often still problematic. Where monolithic algebraic approaches are available, life is simple: just hand off the Jacobian to your favourite solver library. Problems with coupled systems of variables do not usually fit this paradigm. The state of the art here is often block preconditioning, relying on block factorisations and approximations of the block inverses. If we are lucky, it is possible construct such block preconditioners algebraically (by manipulation of the assembled Jacobian). But usually, we are not lucky.

In this talk, I discuss how we overcome this problem in Firedrake. Leveraging PETSc's flexibility in defining operators and preconditioners, and using UFL to simplify the assembly of any auxiliary operators our preconditioner might need.

This approach allows the development of matrix-free iterative solvers with assembly of only the necessary preconditioning blocks. Moreover, our approach is easily extensible, allowing development and composition of complex preconditioners for subblocks as we add new couplings to our model. I will illustrate the usage with examples from fluid convection and phase separation.

Computationally Efficient Numerical Quadratures for the Finite Element Assembly of High-Order non-Lagrange Elements

Nicolas Marsic, Technische Universität Darmstadt, marsic@temf.tu-darmstadt.de,
Herbert De Gerssem, Technische Universität Darmstadt, degersem@temf.tu-darmstadt.de,
Christophe Geuzaine, Université de Liège, cgeuzaine@ulg.ac.be

Keywords: *high-order finite element method, high performance computing, non-Lagrange elements*

1 Introduction

There is a growing consensus in the finite element community that state-of-the-art low-order finite element (FE) technology requires, and will continue to require, too extensive computational resources to provide the necessary resolution for complex simulations, even at the rate of computational power increase [1]. The requirement for precise resolution naturally leads us to consider methods with a higher order of grid convergence than the classical second-order provided by most industrial grade codes. In particular, for high-frequency time-harmonic simulations, high-order schemes allow to efficiently resolve the rapid small-scale spatial oscillations of the solution and allow to alleviate the pollution effect [2]. However, with this approach, the computational cost of solving the linear system of equations rapidly becomes overshadowed by the cost of assembling the finite element matrix itself, as the order of the basis functions increases [3]. The aim of this abstract is to present a reformulation of the assembly procedure in a computationally more efficient way.

2 Classical finite element assembly

By applying the classical FE scheme, the solution is computed using elementary integrals $\mathcal{T}_{i,j}^e$, where each $\mathcal{T}_{i,j}^e$ is giving the contribution of the degrees of freedom (DOF) i and j of the mesh element e . The classical FE assembly algorithm computes on-the-fly the $\mathcal{T}_{i,j}^e$ terms for every pair of DOF i and j on every element e of the mesh. It is worth noticing that increasing the FE basis order will have two impacts on the computation time: each element will have more DOFs and the numerical quadrature will require more points. Both phenomena will substantially increase the assembly time, as shown in Figure 1.

Furthermore, when handling non-Lagrange basis functions, the orientation of the different functions needs to be taken into account. This additional complexity is also classically treated on-the-fly during the assembly, by analysing the orientation an edge (or face) on which a given basis function is defined. Let us note that non-Lagrange bases are commonly encountered when constructing discrete subspace of $H(\mathbf{curl})$ in electromagnetic applications for instance.

3 Efficient assembly

The key idea of a fast assembly procedure is to compute all the $\mathcal{T}_{i,j}^e$ terms using computationally efficient BLAS3 operations, as proposed by [3, 4] for standard nodal Lagrange finite elements. The $\mathcal{T}_{i,j}^e$ terms can be computed by the product of two matrices. The first matrix will be composed of the Jacobian matrices and non-linear terms evaluated at each quadrature point. The second matrix will be composed of only the FE basis functions defined over the reference element and also evaluated at each quadrature point.

Unfortunately, in this approach, non-Lagrange basis functions cannot be treated because of the orientation problem. To circumvent this limitation, the previous solution has been adapted to handle more than one reference space, leading thus to more than one matrix-matrix product: one per possible orientation. However, with this new solution, the orientations of the edges and faces cannot be determined on-the-fly, and need to be considered during a pre-processing step. This can be efficiently achieved by exploiting a newly designed orientation dictionary structure and additional geometrical transformations [5].

4 Numerical results

Figure 1 presents the assembly times of the on-the-fly and matrix assembly procedures for an increasing FE basis order. The FE matrix is assembled for a high-frequency electromagnetic cavity problem. The tests were done on an Intel Xeon E5645 and by using the OpenBLAS implementation of the matrix-matrix product with 6 threads. It is worth mentioning that the classical implementation also uses 6 threads for the assembly. It can be seen from Figure 1 that the matrix procedure is much faster than the classical one for high-order solutions. For instance, the speedup for a sixth-order problem, with more than 900.000 unknowns, is around 20.

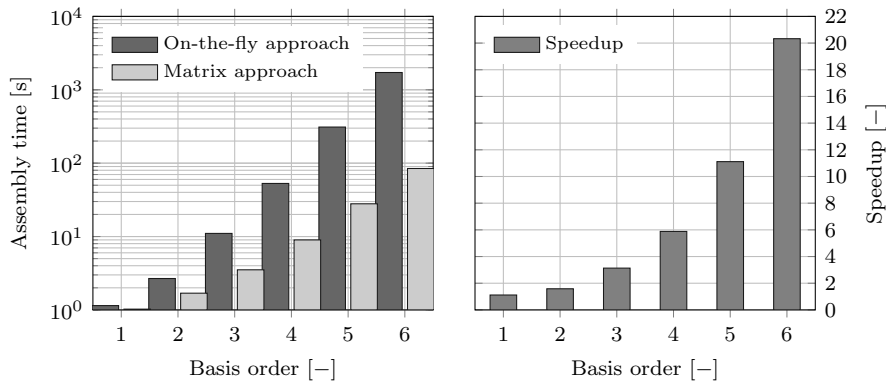


Figure 1: Assembly time and speedup for the classical and fast procedures

Acknowledgements

This work was supported in part by the Belgian Science Policy Office under grant IAP P7/02 and the Walloon Region under grant WIST3 DOMHEX. N. Marsic was a fellowship beneficiary with the Belgian Research Training Fund for Industry and Agriculture (FRIA).

References

- [1] P. E. Vincent and A. Jameson, “Facilitating the adoption of unstructured high-order methods amongst a wider community of fluid dynamicists,” *Mathematical Modelling of Natural Phenomena*, vol. 6, no. 3, pp. 97–140, 2011.
- [2] F. Ihlenburg and I. Babuška, “Finite element solution of the Helmholtz equation with high wave number part II: The h-p version of the fem,” *SIAM Journal on Numerical Analysis*, vol. 34, no. 1, pp. 315–358, 1997.
- [3] K. Hillewaert, “Exploiting data locality in the DGM discretisation for optimal efficiency,” in *ADIGMA*, vol. 113 of *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, pp. 11–23, Springer Berlin Heidelberg, 2010.
- [4] J. Lambrechts, *Finite element methods for coastal flows: application to the great barrier reef*. PhD thesis, Université catholique de Louvain, Belgique, 2011.
- [5] N. Marsic and C. Geuzaine, “Efficient finite element assembly of high order Whitney forms,” *IET Science, Measurement and Technology*, vol. 9, no. 2, pp. 204–210, 2015.

A High-Order Particle-Mesh Operator Splitting Approach for the Incompressible Navier-Stokes Equations

Jakob Maljaars, Delft University of Technology, j.m.maljaars@tudelft.nl,
Robert Jan Labeur, Delft University of Technology, r.j.labeur@tudelft.nl
Matthias Möller, Delft University of Technology, m.moller@tudelft.nl

Keywords: *incompressible Navier-Stokes equations, operator-splitting, hybrid finite elements, Lagrangian-Eulerian, particle-in-cell*

The poster will present a particle-mesh method for simulating incompressible fluid flows. Building upon particle-in-cell concepts, the method is formulated as an operator splitting strategy in which Lagrangian particles are used to discretize an advection operator, and an Eulerian mesh-based method is employed for the constitutive modeling in order to account for the inter-particle interactions. An immediate advantage of this hybrid Lagrangian-Eulerian approach is that no additional stabilization techniques are required in the advective limit.

Key to the presented methodology is the generic variational framework for the coupling between the scattered particles and the mesh in terms of L^2 -minimization problems. These projections remain local and can be efficiently combined with the hybridized Discontinuous Galerkin method for solving the constitutive equations [1]. An efficient solution strategy for the HDG equations is presented by using a static condensation procedure. The model is implemented using tools from the FEniCS-project [4].

As shown by means of various numerical examples the developed methodology overcomes various issues which were to date persistent to particle-mesh methods, such as: the presented particle-mesh coupling results in optimal spatial accuracy and the particle distribution remains uniform over time, even for complicated problems such as the flow around a circular cylinder, Figure 1. Heuristic particle shifting algorithms as often applied in existing particle(-mesh) methods (see e.g. [2, 3]) are thus not required.




Figure 1: Flow past a circular cylinder: particle velocity field at $t = 8$ for $Re = 100$; a detailed view of the particle distribution near the circular cylinder is shown at the right.

References

- [1] R.J. Labeur, G.N. Wells, Energy stable and momentum conserving hybrid finite element method for the incompressible Navier-Stokes equations, *SIAM Journal on Scientific Computing* 34 (2) (2012) 889-913.
- [2] Q. Chen, J. Zang, A. S. Dimakopoulos, D. M. Kelly, C. J. Williams, A Cartesian cut cell based two-way strong fluid-solid coupling algorithm for 2D floating bodies, *Journal of Fluids and Structures* 62 (2016) 252271.
- [3] R. Xu, P. Stansby, D. Laurence, Accuracy and stability in incompressible SPH (ISPH) based on the projection method and a new approach, *Journal of Computational Physics* 228 (18) (2009) 67036725.
- [4] A. Logg, K.-A. Mardal, G. N. Wells, *Automated Solution of Differential Equations by the Finite Element Method*, Vol. 84, Springer Science & Business Media, 2012.

LIAB: a FEniCS based computational tool for laser annealing simulation

A. La Magna¹, S. F. Lombardo^{1,2}, I. Deretzis¹,
A. Verstraete³, B. Lespinasse³ and K. Huet³,

¹CNR-IMM, Catania, Italy

²Department of Physics and Astronomy and INFN, University of Catania, Italy

³Screen-Lasse, France

*antonino.lamagna@imm.cnr.it

Keywords: *Laser Annealing, Self Consistent solutions, Multiple phases and dopants*

Laser annealing (LA), with space uniform beams (over $\sim\text{cm}^2$ areas) and nanosecond range pulses, is the reference annealing technique in micro- and nano- electronics when strongly confined heating is needed in the process flows. Indeed, due to its low in-depth thermal diffusion, laser anneal is nowadays widely applied as a post-fabrication annealing step to activate isolated doped regions (e.g. back junctions) with a null or strongly reduced heating of the other zones of the devices. The application of the process in future electronic device generation is hindered by the difficulties in the process control. Indeed, LA process is highly influenced by the interaction between electromagnetic (e.m.) field and complex device structures. Models of laser annealing process have been developed by our team for particular limited applications and implemented in academic or commercial package [1, 2, 3, 4]. However, several limitations remain (see e.g. discussion in Ref. [5]) in the previous modeling approach for the general application in future devices, characterized by complex structures with nm wide elements made of different materials/phases.

In the paper we present a FEniCS based tool (named LIAB: LASSE Innovation Application Booster), in the development stage, for the simulation of LA process. This is a complex self-consistent problem, where the heating is evaluated by mean of the time harmonic solution of the Maxwell equations and the source term in the heat equation is

$$S_{laser}(t, \mathbf{r}) = \frac{\varepsilon''}{2\rho} |\mathbf{E}_{t-h}|^2 \quad (1)$$

Where ε'' is the imaginary part of the complex dielectric constant $\varepsilon'' = \varepsilon' + j\varepsilon''$ of the material and \mathbf{E}_{t-h} is the time harmonic electric field $\mathbf{E} = \mathbf{E}_{t-h} \times \exp(-j\omega t + \phi)$. The self consistency derives from the dependence of the optical constant on the temperature field (varying in the range 300-2000K), the phases and the alloy fraction.

The main features of the LIAB package include the following:

- Versatile Graphical User Interface for the structure design, the material assignment and the simulation analysis;
- Interface with the FEniCS solver for the automatic generation of the mesh and the runtime control;
- Many materials calibration (optical and thermal properties and mass transport) as a function of temperature and phases;
- Efficient coupling with Electromagnetic Simulation for the self-consistent source estimate (i.e. power dissipation) in nano-structured topographies;
- Experimental validation in nanostructured samples;
- Multiple-dopant models simulating dopant atoms redistribution including diffusion solubility and segregation;
- Alloy model e.g. SiGe (where melting point depends on the alloy fraction);
- Multiple phases (e.g. amorphous, liquid, crystal).

Some application cases will be discussed in order to demonstrate the potentiality of the package (see attached figures). Possible future extension will be outlined in the framework of the development of the FEniCS's project.

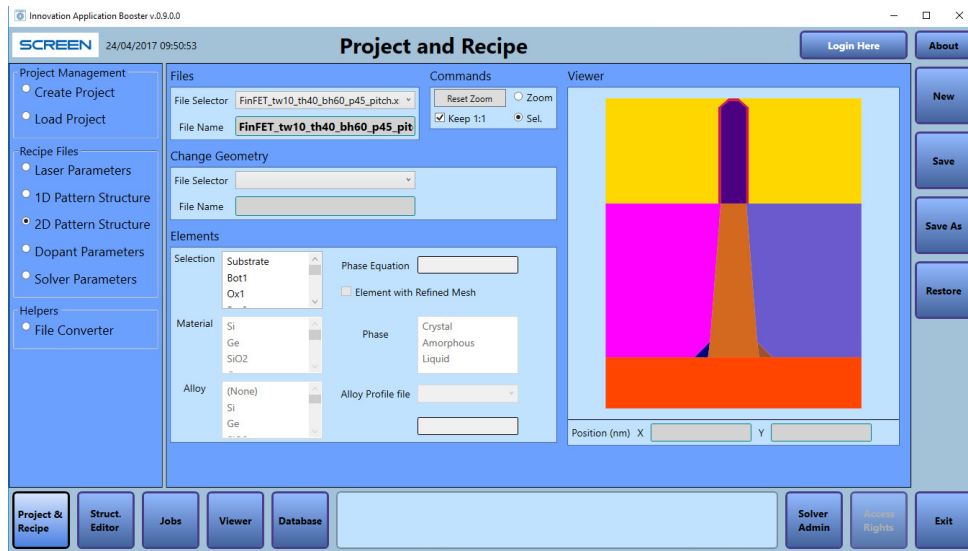


Figure 1: Screenshot of the LIAB Graphical User Interface, a FINEFET device is shown and the different colors indicate the domains, which can be initialize with different materials properties.




Figure 2: Simulation example of the Laser Annealing process in a FINEFET device structure, from left to right the heat sources, the phase (liquid phase in blue, solid one in red) and temperature. Is also shown the mesh used. The correlation between phase and heat source in the figure demonstrates the role of self-consistency in the simulation.

References

- [1] A. La Magna, P. Alippi, V. Privitera, G. Fortunato, M. Camalleri, B. Svensson *Journal of Applied Physics* **95**, 4806 (2004)
- [2] A. La Magna, P. Alippi, V. Privitera, G. Fortunato, (16) *Applied Physics Letters* **86**, 4738 (2005)
- [3] K. Huet, G. Fiscaro, J. Venturini, H. Besaucle, A. La Magna, *Applied Physics Letters* **95**, 231901 (2009)
- [4] G. Fiscaro, K. Huet, R. Negru, M. Hackenberg, P. Pichler, N. Taleb, A. La Magna, *Phys. Rev. Lett.* **110**, 117801 (2013)
- [5] S.F. Lombardo, F. Cristiano, G. Fiscaro, G. Fortunato, M.G. Grimaldi, G. Impellizzeri, M. Italia, A. Marino, R. Milazzo, E. Napolitani, V. Privitera, A. La Magna *Mat. Sci. Semi. Proc* **62**, 80 (2017)

Implementation of Mixed and Hybrid Formulation of the PML Method in Elastodynamics Using FEniCS

Hernán Mella, Pontificia Universidad Católica de Chile, hmella@uc.cl.
 Joaquín Mura, Pontificia Universidad Católica de Chile, jamura@uc.cl.
 Esteban Sáez, Pontificia Universidad Católica de Chile, esaez@ing.puc.cl.
Hernán Mella, Pontificia Universidad Católica de Chile, hmella@uc.cl

Keywords: *Perfectly matched layers, elastodynamics, FEniCS, mixed formulation.*

Introduction

The simulation of propagation of elastic waves in unbounded domain is still an open problem. In the past years different kinds of absorbing boundary conditions (ABC) has been developed, such as Gaussian taper method, high-order ABC and paraxial approximations. One of the most successful is the perfectly matched layers (PML) method, which has been widely used in elastodynamics [3], acoustic [2], poroelasticity [5] and other research fields. This method was first introduced by Berenger in the context of electromagnetism, simulating the propagation of linear electromagnetic waves on unbounded domains [1].

A PML is an absorbing layer model for linear wave equations that absorbs, quite efficiently, propagating waves of all non-tangential angles of incidence and of all non-zero frequencies.

Theoretical framework

Following the work done by Kucukcoban & Kallivokas in [3], the modified mixed formulation of the elastodynamics equations using the PML method, defined on the domain $\Omega = \Omega_f \cup \Omega_p$ that is composed by the absorbing boundary layer Ω_p and the truncated physical domain Ω_f (see Figure 1a), are:

$$\nabla \cdot [\dot{\mathbf{S}}^T \Lambda_e + \mathbf{S}^T \Lambda_p] = \rho(a\ddot{\mathbf{u}} + b\dot{\mathbf{u}} + c\mathbf{u}) \quad \text{on } \Omega \quad (1a)$$

$$\mathbf{D}[(a\ddot{\mathbf{S}} + b\dot{\mathbf{S}} + c\mathbf{S})] = \frac{1}{2}[(\nabla \mathbf{u})\Lambda_p + \Lambda_p(\nabla \mathbf{u})^T + (\nabla \dot{\mathbf{u}})\Lambda_e + \Lambda_e(\nabla \dot{\mathbf{u}})^T] \quad \text{on } \Omega, \quad (1b)$$

with initial and boundary conditions


$$\mathbf{u} = \mathbf{0} \quad \text{on } \Gamma_d \times [t_0, T[\quad (2a)$$

$$(\dot{\mathbf{S}}^T \Lambda_e + \mathbf{S}^T \Lambda_p)\mathbf{n} = \mathbf{g} \quad \text{on } \Gamma_n \times [t_0, T[\quad (2b)$$


$$\mathbf{u} = \dot{\mathbf{u}} = \mathbf{0} \quad \text{on } \Omega \times \{t_0\} \quad (2c)$$

$$\mathbf{S} = \dot{\mathbf{S}} = \mathbf{0} \quad \text{on } \Omega \times \{t_0\} \quad (2d)$$

where \mathbf{u} and \mathbf{S} are the displacements field and the stress history, the scalar functions a , b , c and the tensors Λ_p and Λ_e are stretching functions defined on Ω_p and \mathbf{D} is a compliance operator. For further details see [3].



(a) PML-truncated semi infinite domain



(b) Homogeneous domain with explosive volumetric source

Considering that the functions a , b , c , Λ_p and Λ_e only modifies the problem on Ω_p , an hybrid formulation can be stated using a Dirichlet-Neumann coupling.

$$\begin{aligned} \nabla \cdot [\mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) + \lambda \mathbf{I} \nabla \cdot \mathbf{u}] &= \rho \ddot{\mathbf{u}} & \text{on } \Omega_f & \quad (3) \\ \text{Eq. 1} & & \text{on } \Omega_p & \quad (4) \end{aligned}$$

subject to boundary, initial and interface conditions

$$\mathbf{u} = \mathbf{0} \quad \text{on } \Gamma_d \times [t_0, T[\quad (5a)$$

$$[\mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) + \lambda \mathbf{I} \nabla \cdot \mathbf{u}] \mathbf{n} = \mathbf{g} \quad \text{on } \Gamma_n \times [t_0, T[\quad (5b)$$

$$\mathbf{u} = \dot{\mathbf{u}} = \mathbf{0} \quad \text{on } \Omega \times \{t_0\} \quad (5c)$$

$$\mathbf{S} = \dot{\mathbf{S}} = \mathbf{0} \quad \text{on } \Omega_p \times \{t_0\} \quad (5d)$$

$$\mathbf{u}^+ = \mathbf{u}^- \quad \text{on } \Gamma_i \times [0, T[\quad (5e)$$

$$[\mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) + \lambda \mathbf{I} \nabla \cdot \mathbf{u}] \mathbf{n} = -(\dot{\mathbf{S}}^T \Lambda_e + \mathbf{S}^T \Lambda_p) \mathbf{n} \quad \text{on } \Gamma_i \times [0, T[\quad (5f)$$

Implementation

This problem can be easily implemented using the functionalities of the FEniCS libraries. In order to illustrate this, in Figure 2 it is presented a snippet code of the UFL implementation of the weak form of 4.

Python code

```

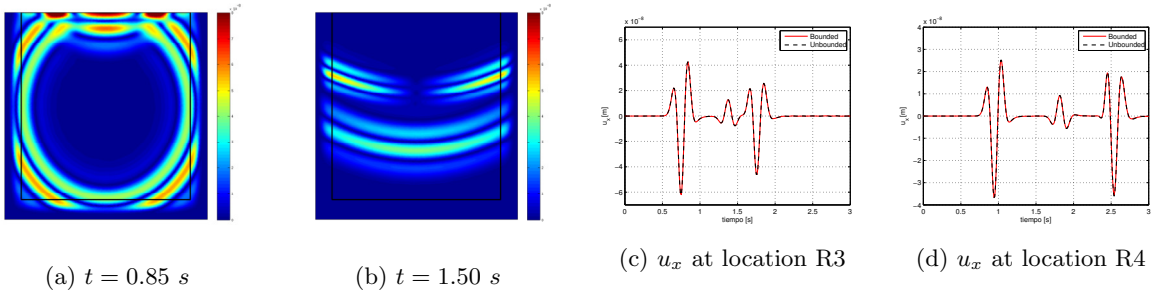
1 # Truncated domain problem
2 F_f = rho*inner(a*N_ddot(u) + b*N_dot(u) + c*u, w)*dx \
3 + inner(N_dot(S).T*Lambda_e + S.T*Lambda_p, grad(w))*dx \
4 - inner(excitation, w)*ds(1)
5
6 # PML domain problem
7 F_p = inner(compliance(a*N_ddot(S) + b*N_dot(S) + c*S), T)*dx \
8 - 0.5*inner(grad(u)*Lambda_p + Lambda_p*grad(u).T + grad(N_dot(u))*Lambda_e \
9 + Lambda_e*grad(N_dot(u)).T, T)*dx
10
11 af, Lf = lhs(F_f), rhs(F_f)
12 ap, Lp = lhs(F_p), rhs(F_p)

```

Figure 2: UFL implementation of the weak forms

Results

One 2D numerical experiment was performed (Figure 1b), in which a volumetric explosive Ricker source at the center of Ω_f and homogeneous domain was considered.



In Figures 3a and 3b two snapshots of the displacements magnitude at different times are presented, whereas in 3c and 3d the obtained results and a simulation with an enlarged domain are compared. In 3b the reflection of waves with the traction-free boundary and the attenuation of outgoing waves can be appreciated.

Conclusions and perspectives

A widely used ABC was implemented in FEniCS using a mixed and hybrid formulation. This method allow us to simulate semi-infinite mediums using truncated domains with a high level of accuracy and low computational cost.

References

- [1] Bérenger J. P. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, 1994.

- [2] Qi Q. and Geers T. L. Evaluation of the perfectly matched layer for computational acoustics. *Journal of Computational Physics*, 139(1):166–183, 1998.
- [3] Kucukcoban S. and Kallivokas L. F. Mixed perfectly-matched-layers for direct transient analysis in 2D elastic heterogeneous media. *Computer Methods in Applied Mechanics and Engineering*, 50:57–76, 2011.
- [4] Kucukcoban S. and Kallivokas L. F. A symmetric hybrid formulation for transient wave simulations in PML-truncated heterogeneous media. *Wave Motion*, 50:57–79, 2013.
- [5] Zeng Y. Q., He J. Q., and Liu Q. H. The application of the perfectly matched layer in numerical modeling of wave propagation in poroelastic media. *Geophysics*, 66(4):1258–1266, 2001.

Handling of finite rotations in Dofin

Marco Morandini, Politecnico di Milano, marco.morandini@polimi.it

Keywords: *finite rotations, beams*

The configuration of a so-called intrinsic beam model with shear deformation is defined by the position vector \mathbf{x} and the associated orientation tensor $\boldsymbol{\alpha} \in \text{SO}(3)$. Let identify the position and orientation in the deformed configuration with an appended prime, \mathbf{x}' and $\boldsymbol{\alpha}'$, with \mathbf{x} and $\boldsymbol{\alpha}$ the corresponding position and orientation in the reference configuration. The corresponding linear form reads

$$\int_L (\delta \boldsymbol{\varepsilon} \mathbf{T} + \delta \boldsymbol{\beta} \mathbf{M}) ds = \int_L (\delta \mathbf{x}' \mathbf{t} + \boldsymbol{\varphi}_\delta \mathbf{m}) ds + \sum_i \delta \mathbf{x}'_i \mathbf{F}_i + \boldsymbol{\varphi}_{\delta i} \mathbf{C}_i \quad (1)$$

where s is the arc length and

$$\begin{aligned} \boldsymbol{\varepsilon} &= \boldsymbol{\alpha}'^T \mathbf{x}'_{/s} - \boldsymbol{\alpha}^T \mathbf{x}_{/s} \\ \boldsymbol{\beta} &= \boldsymbol{\alpha}'^T \text{ax}(\boldsymbol{\alpha}'_{/s} \boldsymbol{\alpha}'^T) - \boldsymbol{\alpha}^T \text{ax}(\boldsymbol{\alpha}_{/s} \boldsymbol{\alpha}^T) \\ &= \boldsymbol{\alpha}^T \boldsymbol{\Phi}^T \text{ax}(\boldsymbol{\Phi}_{/s} \boldsymbol{\Phi}^T) \end{aligned} \quad (2)$$

are the generalized deformation measures energetically conjugated to the internal forces and moments \mathbf{T} and \mathbf{M} ; the axial operator $\mathbf{a} = \text{ax}(\mathbf{A})$ is the operator extracting the vector \mathbf{a} that characterize the skew symmetric part of tensor \mathbf{A} , i.e. $\mathbf{a} \times = 1/2(\mathbf{A}^T + \mathbf{A})$; $\boldsymbol{\Phi} = \boldsymbol{\alpha}' \boldsymbol{\alpha}^T$ is the rotation tensor bringing $\boldsymbol{\alpha}$ into $\boldsymbol{\alpha}'$. Being $\boldsymbol{\alpha}$ and $\boldsymbol{\Phi}$ orthogonal, tensors $\boldsymbol{\alpha}'_{/s} \boldsymbol{\alpha}'^T$, $\boldsymbol{\alpha}_{/s} \boldsymbol{\alpha}^T$ and $\boldsymbol{\Phi}_{/s} \boldsymbol{\Phi}^T$ are skew-symmetric. Vectors \mathbf{t} and \mathbf{m} are the the external forces and moments per unit length, respectively; vectors \mathbf{F}_i and \mathbf{C}_i are the concentrated external forces and moments. Finally, the virtual rotation vector $\boldsymbol{\varphi}_\delta$ characterizes the virtual rotation of the orientation tensor in the deformed configuration, $\boldsymbol{\varphi}_\delta = \text{ax}(\delta \boldsymbol{\alpha}' \boldsymbol{\alpha}'^T) = \text{ax}(\delta \boldsymbol{\Phi} \boldsymbol{\Phi}^T)$. The internal forces and moments, \mathbf{T} and \mathbf{M} , are functions of the generalized strains $\boldsymbol{\varepsilon}$ and $\boldsymbol{\beta}$. A linear constitutive law can often be assumed.

The rotation tensor $\boldsymbol{\Phi}$ belongs to the Special Orthogonal Group $\text{SO}(3)$. Tensors fields belonging to the intrinsically nonlinear group $\text{SO}(3)$ needs to be treated differently than vectors/tensors that belongs to linear spaces, such as the position or displacement fields. Thus, the solution of linear forms involving finite rotations can be challenging when using finite element libraries that assume – from the ground up – that all unknown fields do belong to a linear space. In a nutshell, a parametrization of $\text{SO}(3)$ is often chosen, and the parameters are interpolated using standard finite elements, that are linear with respect to the unknown field. The rotation vector $\boldsymbol{\varphi}$ is one of the most often chosen parametrization. The corresponding rotation tensor $\boldsymbol{\alpha}$ can be computed using the exponential map,

$$\boldsymbol{\Phi} = \exp(\boldsymbol{\varphi} \times) = \sum_{n=0}^{\infty} \frac{1}{n!} (\boldsymbol{\varphi} \times)^n = \mathbf{I} + \frac{\sin \varphi}{\varphi} \boldsymbol{\varphi} \times + \frac{1 - \cos \varphi}{\varphi^2} \boldsymbol{\varphi} \times \boldsymbol{\varphi} \times \quad (3)$$

with φ the norm of $\boldsymbol{\varphi}$; the virtual rotation vector $\boldsymbol{\varphi}_\delta$ can be computed from the variation of the rotation parameter $\delta \boldsymbol{\varphi}$ as $\boldsymbol{\varphi}_\delta = \boldsymbol{\Gamma}(\boldsymbol{\varphi}) \delta \boldsymbol{\varphi}$, with tensor $\boldsymbol{\Gamma}$ defined by

$$\boldsymbol{\Gamma}(\boldsymbol{\varphi}) = \text{dexp}(\boldsymbol{\varphi} \times) = \sum_{n=0}^{\infty} \frac{1}{(n+1)!} (\boldsymbol{\varphi} \times)^n = \mathbf{I} + \frac{1 - \cos \varphi}{\varphi^2} \boldsymbol{\varphi} \times + \frac{\varphi - \sin(\varphi)}{\varphi^3} \boldsymbol{\varphi} \times \boldsymbol{\varphi} \times; \quad (4)$$

Equation 1 and its linearization need to be written as a function of $\boldsymbol{\varphi}$, with test and trial functions $\delta \boldsymbol{\varphi}$ and $\partial \boldsymbol{\varphi}$, respectively; tensor $\boldsymbol{\Gamma}$ enters the formulation because $\boldsymbol{\varphi}_\delta = \boldsymbol{\Gamma} \delta \boldsymbol{\varphi}$, $\boldsymbol{\varphi}_\partial = \boldsymbol{\Gamma} \partial \boldsymbol{\varphi}$ and $\text{ax}(\boldsymbol{\Phi}_{/s} \boldsymbol{\Phi}^T) = \boldsymbol{\Gamma} \boldsymbol{\varphi}_{/s}$. The parametrization is singular for 2π : given an orientation tensor $\boldsymbol{\alpha}$ the corresponding rotation vector is defined up to a coaxial rotation of 2π ; furthermore, and more importantly, tensor $\boldsymbol{\Gamma}$ is singular for $\varphi = 2\pi$. Note also that the coefficients $\sin(\varphi)/\varphi$, $(1 - \cos(\varphi))/\varphi^2$ and $(\varphi - \sin(\varphi))/\varphi^3$ need to be evaluated near $\varphi = 0$ using their Maclaurin series; the expansion must be truncated to an order that is high enough to guarantee the sought precision not only for the coefficients themselves, but also for their second derivative, that is required to compute the linearization $\partial \delta \boldsymbol{\beta}$ of $\delta \boldsymbol{\beta} = \boldsymbol{\alpha}^T \delta(\boldsymbol{\Phi}^T \text{ax}(\boldsymbol{\Phi}_{/s} \boldsymbol{\Phi}^T)) = (\boldsymbol{\alpha}'^T \boldsymbol{\Gamma} \boldsymbol{\varphi}_{/s}) \times \boldsymbol{\alpha}'^T \boldsymbol{\Gamma} \delta \boldsymbol{\varphi} + \boldsymbol{\alpha}'^T \otimes \boldsymbol{\varphi}_{/s} : \delta \boldsymbol{\Gamma} + \boldsymbol{\alpha}'^T \boldsymbol{\Gamma} \delta \boldsymbol{\varphi}_{/s}$.

Summing up, and after writing a python library that allows to deal with finite rotations and the above sketched parametrization, the above formulation can be implemented in FEniCS provided that




Figure 1: Wrench beam testcase, first load steps.

- the expressions of Φ and Γ are computed correctly, with a truncated Maclaurin series for $\varphi \approx 0$;
- the magnitude of the rotation vector is kept lower than 2π , with ideally $\varphi \in [-\pi, \pi]$.

The first requirement can be accomplished by resorting to Uflacs' excellent handling of conditional expressions. For example, coefficient $\sin(\varphi)/\varphi$ can be computed as

Python code

```
1 def a(phi2): return conditional(le(phi2, ath), aexp(phi2), sin(sqrt(phi2))/sqrt(phi2))
```

where `phi2` = φ^2 , `ath` is the threshold for switching from the Maclaurin series to the trigonometric expansion, and `aexp(phi2)` computes the Maclaurin truncated series. Using these conditionals with the library methods written to compute Φ and Γ allows to easily write the linear form Eq. 1 and have the corresponding bilinear form automatically generated.

Dealing with the second requirement is a bit more tricky. The chosen approach is suitable if the rotation field is approximated using Lagrange elements. In other words, it is doomed to fail for e.g. Nedelec elements. First of all, a C++ `Expression` computes either the rotation vector φ or, if $\varphi > \pi$, its complementary vector $\varphi_c = \varphi(1 - 2\pi/\varphi)$. This `Expression` is used to replace the values of φ with its complementary rotation angle when $\varphi > \pi$. This well known trick [1] allows to keep the rotation vector magnitude under control, and works perfectly well if all the rotation vector unknowns of an element get changed into their complementary. Unfortunately one must take care of the elements for which only some of the nodal rotation vector unknowns have been changed into their complementary, see e.g. [2]. This would result in an incorrectly interpolated field. To fix this problem, a custom `Assembler` is required, whose `assemble` method checks, for each element, whether only some of the rotation vector unknowns have been changed. If this is the case then it changes back, in place, the data that gets passed to the element's `tabulate_tensor`. Then, after calling `tabulate_tensor`, it modifies the assembled vector/tensor in order to account that the locally changed rotation vector unknowns are actually a function of the complementary vector. This is accomplished by accounting for the derivative of the complementary vector

$$\frac{\partial \varphi_C}{\partial \varphi} = \left(1 + \frac{2\pi}{\varphi}\right) \mathbf{I} - \frac{2\pi}{\varphi^3} \varphi \otimes \varphi,$$

either on the left for the test functions $\delta\varphi$ or on the right for the trial functions $\partial\varphi$.

This rotation handling strategy is not free from problems, see e.g. [3] and references therein. Nonetheless, the resulting finite elements, whose complex Jacobian matrix is computed automatically, are able to withstand finite rotations of any magnitude and to correctly solve complex nonlinear problems. As an example, the very first deformed configurations of a classic benchmark test are plotted in Fig. 1.

The library is general enough to be suitable for the implementation of intrinsic shell models with the so-called drilling degree of freedom. In this case, however, special care should be given to alleviate the shear locking problem, that is rather easily dealt with for beam finite elements.

References

- [1] M. Ritto-Corrêa, D. Camotim: On the differentiation of the Rodrigues formula and its significance for the vector-like parameterization of Reissner-Simo beam theory, *IJNME*, 55 (2002), 1005-1032.
- [2] J. Mäkinen: Total Lagrangian Reissner's geometrically exact beam element without singularities, *IJNME*, 70 (2007), 1009-1048.
- [3] T. Merlini, M. Morandini: On successive differentiations of the rotation tensor: An application to nonlinear beam elements, *JOMMS*, 8 (2013), 305-340.

Topology Optimization under Uncertainties in FEniCS

Johannes Neumann, WIAS Berlin

Keywords: *topology optimization, uncertainty quantification, phase field, Allen-Cahn, mesh adaptivity, Monte Carlo*

The application of phase field models to shape optimization allows for the implicit treatment of the topology of the work piece. Gradient descent approaches lead to an Allen-Cahn gradient flow type problem. We treat additional constraints of the model either with a projection step in each iteration or with Lagrange multipliers in an augmented linear system. Adaptive methods for the resulting time-stepping and the underlying finite element mesh considerably speed up the computation. Furthermore, additional adaptive choices of key model parameters allow for robust convergence in practice.

The modeling of material and load uncertainties leads to a stochastic cost functional for the optimization which requires a proper risk measure in turn. The conditional value at risk is here a practical choice which allows to modify the robustness of the resulting shape by some risk parameter. The resulting stochastic gradient descent solves an augmented optimization problem with Monte Carlo sampling or tensor reconstruction where we employ parallel computations.

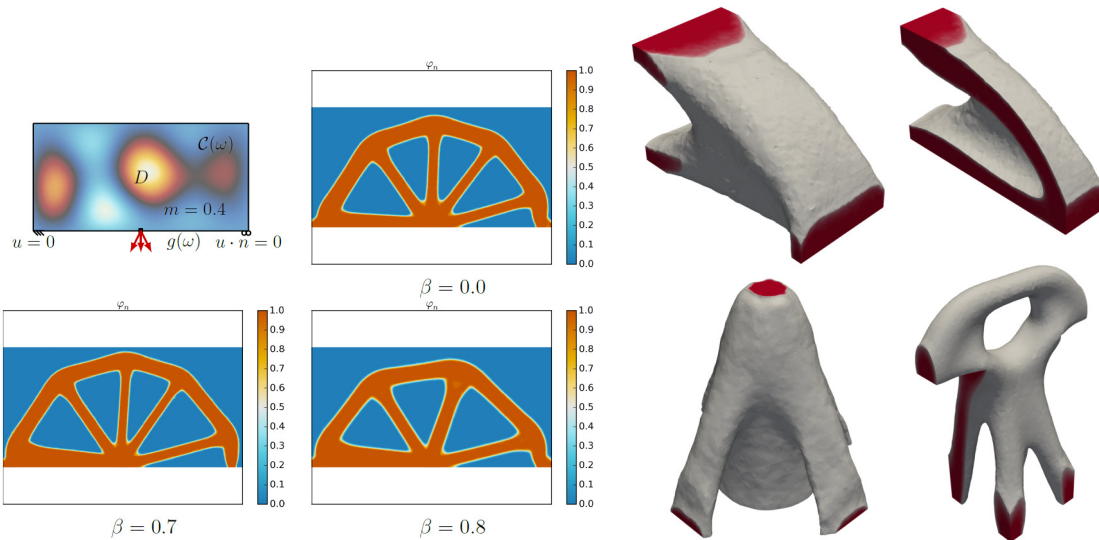


Figure 1: Bridge example and solutions for various degrees of robustness (left) and solutions in three dimensions for various examples (right).

References

- [1] M. Eigel, J. Neumann, R. Schneider, S. Wolf: Stochastic topology optimisation with hierarchical tensor reconstruction, *WIAS Preprint* No. 2362, 2016
- [2] F. Anker, Ch. Bayer, M. Eigel, J. Neumann, J. Schoenmakers: Adaptive SDE based interpolation for random PDEs, *WIAS Preprint* No. 2200, 2015 (accepted)
- [3] M. Eigel, Ch. Merdon, J. Neumann: An adaptive multi level Monte-Carlo method with stochastic bounds for quantities of interest in groundwater flow with uncertain data, *SIAM ASA J. Uncertain. Quantif.* 4 (2016) pp. 1219–1245.

A Space-Time Cut Finite Element Method for the Heat Equation in FEniCS

Anders Logg, Chalmers University of Technology, logg@chalmers.se
 Carl Lundholm, Chalmers University of Technology, carlun@chalmers.se
Magne Nordaas, Chalmers University of Technology, magne@chalmers.se

Keywords: *space-time finite elements, cut-FEM, overlapping meshes, moving meshes, FEniCS multi-mesh, parabolic PDEs*

We present a space-time cut finite element method for parabolic PDEs on overlapping meshes, implemented using the FEniCS multimesh functionality. This work extends the space-time method developed in [1] to two and three spatial dimensions and to more than two meshes.

Here the concept of overlapping meshes means that the computational domain consists of two or more overlapping meshes, and each mesh has a prescribed time-dependent movement. An example where such use of overlapping meshes could be advantageous over traditional methods, is when the domain contains a moving object, see Figure 1. The idea is first to generate one mesh around the object and background mesh for the empty solution domain. The mesh-enclosed object is then brought back into the computational domain by placing it “on top” of the background mesh. Finally the mesh enclosed object can be moved around in the domain as seen in Figure 2. By using overlapping meshes, mesh generation only has to be done initially as opposed to traditional methods where a new mesh has to be generated when the old mesh becomes too deformed.




Figure 1: An application where traditional methods render a mesh unsuitable for computations.




Figure 2: Alternative approach to the same application with overlapping meshes.

Finite element methods may then be derived for the overlapping mesh environment by using Nitsche-based techniques to account for the interface between two meshes. Using $cG(p)$ elements in space and $dG(q)$ elements in time, we let V_h be a corresponding finite element space that permits a discontinuity on the interface between two meshes.

The finite element variational formulation used here reads: Find $u_h \in V_h$ such that for all $v \in V_h$,

$$\begin{aligned} & \sum_{i=1}^2 \sum_{n=1}^N \left(([u_h]_{n-1}, v_{n-1}^+)_{\Omega_i(t_{n-1})} + \int_{t_{n-1}}^{t_n} (\dot{u}_h, v)_{\Omega_i(t)} + (\nabla u_h, \nabla v)_{\Omega_i(t)} dt \right) \\ & + \sum_{n=1}^N \left(\int_{\Gamma_n} -\bar{n}^t [u_h] v_\sigma - \langle \partial_{\bar{n}^x} u_h \rangle [v] - \langle \partial_{\bar{n}^x} v \rangle [u_h] + |\bar{n}^x| \gamma h_K^{-1} [u_h] [v] d\bar{s} \right) \\ & = \sum_{i=1}^2 \sum_{n=1}^N \int_{t_{n-1}}^{t_n} (f, v)_{\Omega_i(t)} dt. \end{aligned} \quad (1)$$

Here, Ω_1 denotes the non-overlapped part of the background mesh, Ω_2 denotes the moving and overlapping mesh, Γ_n denotes the space-time interface between the meshes on the time interval $(t_{n-1}, t_n]$, and $\bar{n} = (\bar{n}^t, \bar{n}^x)$ denotes the unit vector normal to the space-time surface Γ_n .

The method presented above has been implemented in FEniCS using the multimesh functionality. The first step is to replace the integral in time with a suitable quadrature rule. Afterwards, the individual parts of the variational form are expressed with Python code, mimicking the mathematical syntax and making use of the multimesh integral types `dX` and `dI`. For example,

$$\sum_{i=1}^2 \int_{\Omega_i} \nabla u_h \cdot \nabla v dx \quad \text{inner(grad(u_h), grad(v_h)) * dX}$$

$$\int_{\Gamma} \langle \partial_{\bar{n}^x} u_h \rangle [v] ds \quad \text{avg(inner(nbar_x, grad(u_h))) * jump(v) * dI}$$

With this implementation in FEniCS, we further investigate the aspects of using space-time cut-FEM for time-dependent problems. An interesting matter is for example how the speed of the moving mesh influences the error convergence. In [1], a drop in the order of convergence with respect to the time-step was observed when the mesh speed became sufficiently high, as shown in figure 3. Previously, numerical experiments for this method have been limited to the case with one spatial dimension, cG(1) elements in space, and dG(0) and dG(1) elements in time. In the present work, we study the numerical error convergence in two or three spatial dimensions, and we also consider higher order elements in time.




Figure 3: *Left:* A cG(1)–dG(1) space-time finite element solution u_h for a heat equation problem in one spatial dimension and with one overlapping mesh. *Right:* The order of convergence of the error with respect to the time-step versus the speed of the overlapping mesh using dG(1) in time for a problem in one spatial dimension.

References

- [1] C. Lundholm, *A space-time cut finite element method for a time-dependent parabolic model problem*, Master’s thesis, Chalmers University of Technology, 2015.

Sensitivity Analysis of Cardiac Growth Models

Liubov Nikitushkina, University of Oslo, liubovn@ifi.uio.no,
Simon W. Funke, Simula Research Laboratory, simon@simula.no,
Henrik Finsberg, Simula Research Laboratory, henriknf@simula.no,
Lik Chuan Lee, Michigan State University, llee@egr.msu.edu,
Espen W. Remme, Oslo University Hospital, espen.remme@medisin.uio.no,
Samuel T. Wall, Simula Research Laboratory, samwall@simula.no

Keywords: *Cardiac Mechanics, Finite Element Method, Sensitivity Analysis, FEniCS, dolfin-adjoint*

1 Introduction

The heart is a dynamic organ capable of changing its shape in response to the body's demands. For example, the human heart continuously adapts in size and geometry to meet greater blood flow needs of the growing body during normal development. In this case, a gradually imposed volume overload leads to progressive chamber enlargement. Another example of a normal physiological growth can be found in the athlete's heart, where a sustained elevated chamber pressure results in the chamber wall thickening and an overall increase in cardiac mass. Growth processes, however, can also be maladaptive as found in many cardiovascular diseases where structural changes in the heart progressively decompensate cardiac function.

In order to better understand this balance between adaptive and maladaptive cardiac growth, we examine the effect of known growth stimuli using a mechanical model of the heart. We perform a sensitivity analysis of existing growth models in order to assess the relative importance of model parameters and respective mechanisms. This work can eventually lead to simplifications in the model systems for prediction of growth, or help in localizing shortcomings that need to be addressed in the existing modeling frameworks.

2 Methods

In order to simulate the motion of the heart throughout the cardiac cycle, we use a nonlinear finite element (FE) model of a realistic left ventricle (LV) coupled to a lumped-parameter model of the systemic circulation. Under the quasi-static assumption, this problem is reduced to finding the displacement \mathbf{u} , hydrostatic pressure p and LV pressure p_{LV} that minimize the incompressible strain energy functional Π parameterised by the LV volume V_{LV} [1]:

$$\Pi(\mathbf{u}, p, p_{LV}; V_{LV}) = \int_{\Omega} \Psi(\mathbf{F}_e) + p(\det(\mathbf{F}_e) - 1) \, dx + p_{LV}(V(\mathbf{u}) - V_{LV}). \quad (1)$$

The muscular tissue of the heart is modeled as a transversely isotropic hyperelastic material via the strain energy density Ψ [2]:

$$\Psi(\mathbf{C}_e, \mathbf{m}) = \frac{a}{2b}(e^{b(I_1-3)} - 1) + \frac{a_f}{2b_f}(e^{b_f(I_{4,f_0})^2} - 1), \quad (2)$$

where $\mathbf{m} = (a, b, a_f, b_f)$ is a set of passive material parameters, \mathbf{C}_e is the elastic right Cauchy-Green tensor, $I_1 = \text{tr} \mathbf{C}_e$ and $I_{4,f_0} = \mathbf{f}_0(\mathbf{C}_e \mathbf{f}_0)$. By introducing an activation parameter γ representing the active shortening in the fiber direction \mathbf{f}_0 at zero-load, the model incorporates muscle contraction using the active strain approach, which is based on a multiplicative decomposition of the deformation gradient $\mathbf{F} = \mathbf{I} + \text{Grad}(\mathbf{u})$ into an elastic and an active parts $\mathbf{F} = \mathbf{F}_e \mathbf{F}_a$ with \mathbf{F}_a defined as:

$$\mathbf{F}_a = (1 - \gamma) \mathbf{f}_0 \otimes \mathbf{f}_0 + \frac{1}{\sqrt{1 - \gamma}} (\mathbf{I} - \mathbf{f}_0 \otimes \mathbf{f}_0). \quad (3)$$

The dynamic changes in the ventricular blood pressure and volume over the entire cardiac cycle are modeled by a three-element Windkessel model described by a system of ordinary differential equations [3]. At each time step, the coupling between the FE model and the circulatory model is achieved through

an additional Lagrange multiplier p_{LV} which represents the LV cavity pressure. The problem (1) is solved such that the simulated LV cavity volume $V(\mathbf{u})$ matches the target volume value V_{LV} obtained from the circulatory model.

Growth process in the heart wall is modeled by deforming the reference unloaded geometry to a new grown configuration, again through a multiplicative decomposition of \mathbf{F} into an elastic and, this time, a growth part, where $\mathbf{F} = \mathbf{F}_e \mathbf{F}_g$. The constitutive laws for finite growth can be expressed using a generic format for the growth tensor $\mathbf{F}_g = \theta_f \mathbf{f}_0 \otimes \mathbf{f}_0 + \theta_s \mathbf{s}_0 \otimes \mathbf{s}_0 + \theta_n \mathbf{n}_0 \otimes \mathbf{n}_0$. The evolution of the local tissue growth parameter $\boldsymbol{\theta}_g = [\theta_f, \theta_s, \theta_n]^T$ can be defined in terms of a growth activation function $\phi(\mathbf{F}_e)$ and a growth rate function $\mathbf{k}(\boldsymbol{\theta}_g)$ which specifies the speed and nonlinearity of the growth process [4].

Using the above described model it is possible to simulate various physiological conditions together with the associated structural adaptation of the heart walls in response to change in loadings. In each case, a growth model can be chosen depending on the nature of the considered physiology. The sensitivity of the system's grown state to the prescribed growth model can then be estimated. For this, we define an objective functional $J(\mathbf{u})$, the model output of interest, which is to serve as a qualitative and/or quantitative measure of how well the growth model reproduces the expected behavior. More specifically, if we are to compare the model response to a real measurement, then the objective functional can be defined as the mismatch between the simulated \mathbf{u} and the measured \mathbf{u}_{exp} grown states at a reference time t_{ref} :

$$J(\mathbf{u}) = \int_{\Omega} |\mathbf{u} - \mathbf{u}_{\text{exp}}|^2(t_{ref}) dx. \quad (4)$$

Finally, the sensitivities of J to $\boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is a set of growth parameters specific to a given model, are defined as $\frac{dJ(\mathbf{u})}{d\boldsymbol{\theta}}$.

The solver has been developed within the FEniCS [5] framework and the functional gradients are computed by solving an automatically derived adjoint equations [6].

3 Results

We first focus on implementing and testing a strain-driven growth law to simulate a volume overload state of the left ventricle. To achieve this, the initial physiological equilibrium of the heart is altered by increasing a diastolic filling pressure. This initiates cardiac growth that continues until a new equilibrium state is reached.

The model is able to reproduce qualitatively experimental observations reported in the literature, such as LV cavity dilation due to fiber over-stretching and a gradual increase in myocardium volume. At the next step, we perform a sensitivity analysis of the model with respect to the growth model parameters and evaluate its performance in reproducing the expected growth behavior.

References

- [1] Gjerard S, Hake J, Pezzuto S, Sundnes J, Wall ST. Patient-specific parameter estimation for a transversely isotropic active strain model of left ventricular mechanics, *Statistical Atlases and Computational Models of the Heart-Imaging and Modelling Challenges*, 2009 Sep 13;367(1902):3445-75.
- [2] Holzapfel GA, Ogden RW. Constitutive modelling of passive myocardium: a structurally based framework for material characterization, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, pp. 93104. Springer (2015).
- [3] Nierop BJ. Coupling a finite element model of left ventricular mechanics in a closed loop lumped parameter model of circulatory hemodynamics, Eindhoven University of Technology, 2007.
- [4] Göktepe S, Abilez OJ, Parker KK, Kuhl E. A multiscale model for eccentric and concentric cardiac growth through sarcomerogenesis, *J Theor Biol.*, 2010 Aug 7;265(3):433-42.
- [5] Logg A, Mardal KA, Wells G, editors. Automated solution of differential equations by the finite element method: The FEniCS book, *Springer Science & Business Media*, 2012 Feb 24.
- [6] Farrell PE, Ham DA, Funke SW, Rognes ME. Automated derivation of the adjoint of high-level transient finite element programs, *SIAM Journal on Scientific Computing*, 2013 Jul 11;35(4): C369-93.

The impact of microscopic structure on mechanical properties of plant cell walls and tissues

Mariya Ptashnyk, Brian Seguin, University of Dundee, mptashnyk@maths.dundee.ac.uk

Keywords: *plant biomechanics, linear elasticity, nonlinear PDEs, multiscale analysis, FEniCS*

To better understand plant growth and development it is important to analyse how the microscopic structure of plant tissues impacts their mechanical properties. The elastic properties of plant tissues are strongly determined by the mechanical properties of the cell walls surrounding plant cells and by the cross-linked pectin network of the middle lamella which joins individual cells together. Primary cell walls of plant cells consist mainly of oriented cellulose microfibrils, pectin, hemicellulose, proteins, and water. Since the turgor pressure acts isotropically, it is the microstructure of the cell walls, e.g. the orientation of the cellulose microfibrils, which determines the anisotropic deformation and expansion of plant cells. Also it is supposed that calcium-pectin cross-linking chemistry is one of the main regulators of plant cell wall elasticity and extension [1]. Pectin is deposited into the cell walls in a methylesterified form. In the cell walls and middle lamella, pectin can be modified by the enzyme pectin methylesterase, which removes methyl groups by breaking ester bonds. The de-esterified pectin is able to form calcium-pectin cross-links, and thus stiffen the cell wall and reduce its expansion.

Here we investigate the impact of the orientation of cellulose microfibrils on the elastic deformations of plant cell walls and tissues using multiscale modelling and numerical simulations. Equations of linear elasticity coupled with reaction-diffusion equations for chemical processes are used to model mechanical properties of plant cell walls:

$$\begin{aligned} \operatorname{div} \left((\mathbb{E}_M(b)\chi_{\Omega_M} + \mathbb{E}_F\chi_{\Omega_F}) \mathbf{e}(\mathbf{u}) \right) &= 0 && \text{in } \Omega, \\ (\mathbb{E}_M(b)\chi_{\Omega_M} + \mathbb{E}_F\chi_{\Omega_F}) \mathbf{e}(\mathbf{u}) \boldsymbol{\nu} &= \mathbf{F} && \text{on } \partial\Omega, \\ \partial_t b - \operatorname{div} (D_b \nabla b) &= \mathcal{R}_b(b, \mathbf{e}(\mathbf{u})) && \text{in } \Omega, \\ D_b \nabla b \cdot \boldsymbol{\nu} &= \mathcal{P}_b && \text{on } \partial\Omega, \end{aligned} \tag{1}$$

where Ω denotes a part of a plant cell wall or a domain composed of parts of eight cells connected by the middle lamella. In the microscopic problem we distinguish between mechanical properties of cell wall microfibrils, with elasticity tensor denoted by \mathbb{E}_F , and cell wall matrix and middle lamella, with elasticity tensor denoted by $\mathbb{E}_M(b)$. We also assume that the mechanical properties of the cell wall matrix and middle lamella depend on the density of calcium-pectin cross-links b . The dynamics of the density of calcium-pectin cross-links is described by the reaction-diffusion equation with reaction terms dependent on the elastic deformations of the plant cell walls. Here the parts of the domain Ω occupied by microfibrils and cell wall matrix and middle lamella are denoted by Ω_F and Ω_M , respectively.

In the microscopic model of plant cell wall biomechanics three different scenarios for the orientation of the microfibrils in the cell walls are considered, see Fig. 1. We also consider the situation where the microfibrils are rotated through the thickness of the cell wall. Using homogenization theory, the




Figure 1: Schematic of the Representative Volume Element (RVE) Y for three configurations of microfibrils in the plant cell walls, considered here.




Figure 2: A depiction of the displacements in the x_2 -direction with two different configurations of microfibrils in the upper and lower parts of the cell walls [3].

macroscopic model for plant cell wall and tissue biomechanics is derived from the microscopic description of the elastic properties of the cell wall microfibrils and wall matrix [2]. To determine macroscopic elasticity tensor for the plant cell walls the so-called ‘unit-cell’ problems

$$\begin{cases} \operatorname{div}_y [(\mathbb{E}_M(b)\chi_{Y_M} + \mathbb{E}_F\chi_{Y_F})(\mathbf{e}_y(\mathbf{w}^{kl}) + \mathbf{b}^{kl})] = \mathbf{0} & \text{in } Y, \\ \int_Y \mathbf{w}^{kl} dy = \mathbf{0}, & \mathbf{w}^{kl} \text{ is } Y\text{-periodic,} \end{cases} \quad (2)$$

with $\mathbf{b}^{kl} = \frac{1}{2}(\mathbf{b}^k \otimes \mathbf{b}^l + \mathbf{b}^l \otimes \mathbf{b}^k)$ and $k, l = 1, 2, 3$, where $(\mathbf{b}^1, \mathbf{b}^2, \mathbf{b}^3)$ is the standard basis in \mathbb{R}^3 , are solved numerically. Here Y denotes the Representative Volume Element (RVE) that corresponds to the configurations of the cell wall microfibrils considered here, Fig. 1.

Using the numerical results for the effective elasticity tensor, we consider numerical simulations for the macroscopic model for mechanical properties of the plant cell walls and tissues for different microfibril orientations in different parts of the cell walls, Fig. 2. Using numerical simulations for the macroscopic equations derived from the microscopic model (1) the impact of the orientation of microfibrils on the dynamics and spatial patterns in the distribution of calcium-pectin cross-links is analysed.

The numerical simulations were performed using FEniCS [4]. This involved discretizing the domain using a nonuniform mesh and applying the continuous Galerkin method to solve the equations of linear elasticity and reaction-diffusion equations. Semi-implicit time discretization scheme is applied to the coupled system of linear elasticity and reaction-diffusion equations. The resulting linear system was solved using the iterative Krylov solver, i.e. the general minimal residual method (GMRES), with an algebraic multigrid preconditioner.

References

- [1] S. Wolf, K. Hématy, H. Höfte, Growth control and cell wall signaling in plants, *Annu. Rev. Plant Biol.*, 63 (2012), 381-407.
- [2] M. Ptashnyk, B. Seguin, Homogenization of a system of elastic and reaction-diffusion equations modelling plant cell wall biomechanics, *ESAIM: M2AN*, 50 (2016), 593-631.
- [3] M. Ptashnyk, B. Seguin, The impact of microfibril orientations on the biomechanics of plant cell walls and tissues, *Bulletin of Mathematical Biology*, 78 (2016), 2135-2164.
- [4] A. Logg, K.A. Mardal, G.N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012.

FEniCS Mechanics: A Package for Continuum Mechanics Simulations

Miguel A. Rodriguez, University of California, Berkeley,
miguelr@berkeley.edu,
Christoph M. Augustin, University of California, Berkeley,
christoph.augustin@berkeley.edu,
Shawn C. Shadden, University of California, Berkeley,
shadden@berkeley.edu

Keywords: *solid mechanics, fluid mechanics, python*

This work presents a python package built on top of FEniCS, FEniCS Mechanics, developed to carry out computational mechanics simulations. The goal of this package is to facilitate defining and numerically solving mechanics problems through creating a rapid prototyping environment for users of all programming levels and mechanics knowledge. Problems in continuum mechanics are generally described by the weak form for the balance of linear momentum,

$$\int_{\mathcal{R}} \boldsymbol{\xi} \cdot \rho \mathbf{a} \, dv + \int_{\mathcal{R}} \frac{\partial \boldsymbol{\xi}}{\partial \mathbf{x}} \cdot \mathbf{T} \, dv = \int_{\mathcal{R}} \boldsymbol{\xi} \cdot \rho \mathbf{b} \, dv + \int_{\Gamma_q} \boldsymbol{\xi} \cdot \bar{\mathbf{t}} \, dv, \quad (1)$$

where \mathcal{R} denotes the current configuration of the body, ρ the density of the material, \mathbf{a} the acceleration, \mathbf{T} the Cauchy stress tensor, \mathbf{b} the body force applied, $\bar{\mathbf{t}}$ a prescribed Cauchy traction on $\Gamma_q \subset \partial\mathcal{R}$, and $\boldsymbol{\xi}$ an arbitrary, kinematically admissible, vector field. Note that Equation (1) is applicable to any continua since a constitutive relation is yet to be specified. This facilitates the implementation of various materials within the same framework. The user can also turn incompressibility on, triggering a mixed function space formulation that results in a block system. Furthermore, inverse elastostatics has been implemented as described by Govindjee and Mihalic [1]. This method is used to determine the unloaded geometry when the loaded configuration under a given load is known. In addition to supporting fluid and solid mechanics problems, the simulations can be specified as quasi-static or time-dependent, where a generalized- α method is used for time integration.

The input provided by the user is a python dictionary defining the problem. This dictionary contains the subdictionaries `material`, `mesh`, and `formulation`. The constitutive equation and relevant material parameters are given in `material`. Mesh information, as well as function space discretization is given in `mesh`. Lastly, boundary conditions, time integrating parameters, domain formulation, and initial conditions are given in `formulation`. The user then creates a problem object by providing this dictionary. During instantiation of this object, the values in the configuration dictionary, and their combination, are checked to ensure they are consistent with physical processes. Then, a solver object can be created and used to compute the solution to Equation (1).

To illustrate a simple use of this package, consider the quasi-static loading of an incompressible unit cube modeled as a neo-Hookean material. The material parameters are $\kappa = 10$ GPa, and $\mu = 1.5$ MPa, where κ and μ are the bulk and shear moduli of natural rubber, respectively. The cube is clipped at $x = 0$ m, while a traction of $\bar{\mathbf{t}} = (1 \text{ MPa})\mathbf{e}_1$ is applied at $x = 1$ m. The code for this problem is given below:

Python code


```
1 import dolfin as dlf
2 import fenicsmechanics as fm
3
4 config = {'material': {
5     'const_eqn': 'neo_hookean', # constitutive equation
6     'type': 'elastic', # choose between elastic and viscous
7     'incompressible': True, # toggle incompressibility
8     'density': 1.0,
9     'kappa': 10e9, # bulk modulus [Pa]
10    'mu': 1.5e6}, # shear modulus [Pa]
11    'mesh': {
12        'mesh_file': 'mesh.h5', # mesh file in xml or hdf5 format
13        'mesh_function': 'mesh_function.h5',
```

```

14         'element': 'p2-p1'}, # finite element choices,
15                                     # p1, p2, p2-p1, etc.
16     'formulation': {
17         'time': {'unsteady': False}, # quasistatic or unsteady using gen. alpha
18         'domain': 'eulerian', # lagrangian or eulerian
19         'inverse': False, # toggle inverse finite elastostatics
20         'bcs': {
21             'dirichlet': {
22                 'displacement': [dlf.Constant([0.,0.,0.])],
23                 'regions': [1]}, # surface tags
24             'neumann': {
25                 'regions': [2], # surface tags
26                 'types': ['cauchy'], # choose between piola, cauchy,
27                                     # and pressure
28                 'values': [dlf.Constant([1e6,0.,0.])] # traction vector [Pa]
29             }}}
30
31 problem = fm.SolidMechanicsProblem(config) # set up solid mechanics problem
32 solver = fm.SolidMechanicsSolver(problem) # set up solid mechanics solver
33 solver.full_solve(fname_disp='results/displacement.pvd') # Solve PDE and save

```

Running the above script in parallel yielded the results shown in Figure 1a, where the maximum displacement is 0.263 m.



(a) Elongation of a unit cube ($1\text{m} \times 1\text{m} \times 1\text{m}$) with the unloaded configuration shown in transparent gray.

(b) The idealized left ventricle in its unloaded (left) and loaded (right) states.

The code has been verified using several test cases and comparing simulation results against in-house computational tools, as well as against demos within FEniCS repositories. Furthermore, inflation of an idealized left ventricle was compared with results obtained from other groups [2], and is shown in Figure 1b. The simulation of the left ventricle was accomplished with minimal changes to the configuration dictionary shown above.

FEniCS Mechanics will be available through git repositories once further quantitative verification is performed for quality assurance. This code makes modeling of custom mechanics problems even more accessible to those with a limited programming background or knowledge of mechanics.

References

- [1] Govindjee, Sanjay, and Paul A. Mihalic. "Computational Methods for Inverse Finite Elastostatics." *Computer Methods in Applied Mechanics and Engineering* 136.1-2 (1996): 47-57.
- [2] Land S et al. "Verification of Cardiac Mechanics Software: Benchmark Problems and Solutions for Testing Active and Passive Material Behaviour." *Proceedings of the Royal Society of London A* 471 (2015): 20150641.

Uncontainable enthusiasm: the pain and joy of Docker on HPC and cloud

Chris Richardson, University of Cambridge, UK, chris@bpi.cam.ac.uk

Jack S. Hale, University of Luxembourg, jack.hale@uni.lu

Garth N. Wells, University of Cambridge, UK, gnw20@cam.ac.uk

Keywords: *docker, HPC, cloud computing, shifter, containers*

The FEniCS project has embraced the use of containers for Continuous Integration, and for distribution of pre-configured images for users. However, there is another area where docker containers hold great promise: High Performance Computing (HPC) [1].

Because of privilege restrictions on HPC, it is necessary to use a different runtime, such as Shifter or Singularity, rather than Docker. Performance can be very good, even for MPI jobs, provided that the right MPI libraries are used.

Finally, Microsoft Azure now provides Virtual Machines (VMs) with Infiniband interconnect. Here, we have full control, and can use Docker without restriction. Additionally, there is a useful project called *Batch Shipyard* [2] which simplifies the deployment of Docker on Azure. The performance is very good for medium sized problems.

References

- [1] J.S. Hale, L. Li, C.N. Richardson, G.N. Wells, *Containers for portable, productive and performant scientific computing* <https://arxiv.org/abs/1608.07573>
- [2] Park, A. *Batch Shipyard* <https://github.com/Azure/batch-shipyard>

Analytic Metric-Based Adaptation Using a Continuous Mesh Model

Ajay Rangarajan, RWTH Aachen, rangarajan@aices.rwth-aachen.de,
Georg May, RWTH Aachen, may@aices.rwth-aachen.de,
Vit Dolejsi, Charles University, dolejsi@karlin.mff.cuni.cz

Keywords: *Adaptation, Anisotropy, Metric, Discontinuous Galerkin*

Anisotropic adaptive meshing is widely recognized as an important tool in the numerical solution of convection-diffusion problems. In most flow problems we are not interested in the flow properties in the entire domain but instead we would like to know more about the flow in a smaller region. Using the anisotropy and size information from the underlying solution we can adapt towards an optimal mesh.

Moving from solution to an optimized mesh happens through error estimates based on the higher order derivatives. Here we present a generalized mesh adaptation and mesh optimization method for discontinuous Galerkin Schemes using a metric-based continuous mesh model [1]. The rationale behind the new error model is to incorporate as much as possible analytic optimization techniques acting on the continuous mesh, as opposed to numerical optimization acting on the discrete mesh. The metric based approach facilitates changing and manipulating the mesh in a general non-isotropic way. The metric is a symmetric positive definite matrix that encodes the information required to construct a mesh element as shown in Fig. 1




Figure 1: Relating metric and triangle

We also discuss target-based adaptation in this context. This is important in a variety of applications, where one is interested in computing certain solution-dependent functionals rather accurately, as opposed to minimizing the global error norm. To incorporate target-based adaptation we extend previous continuous mesh models to a weighted-norm optimization, where the weight comes from an adjoint solution providing sensitivities with respect to some relevant target functional [2].

Additionally, high-order consistent numerical schemes using piecewise polynomial approximation call for hp -adaptation to maximize efficiency. We present a two step algorithm using the continuous error estimate to produce a close-to-optimal hp mesh [3].

We present the formulation of the methods as well as numerical experiments in the context of convection-diffusion systems of the form:

$$\begin{aligned} \mathbf{q} &= \nabla u \\ \nabla \cdot (\mathbf{F}_c(u) - \mathbf{F}_v(u, \mathbf{q})) &= s(u, \mathbf{q}) \end{aligned}$$

with the appropriate boundary conditions. An example mesh produced using the analytic optimization is shown in Fig. 2

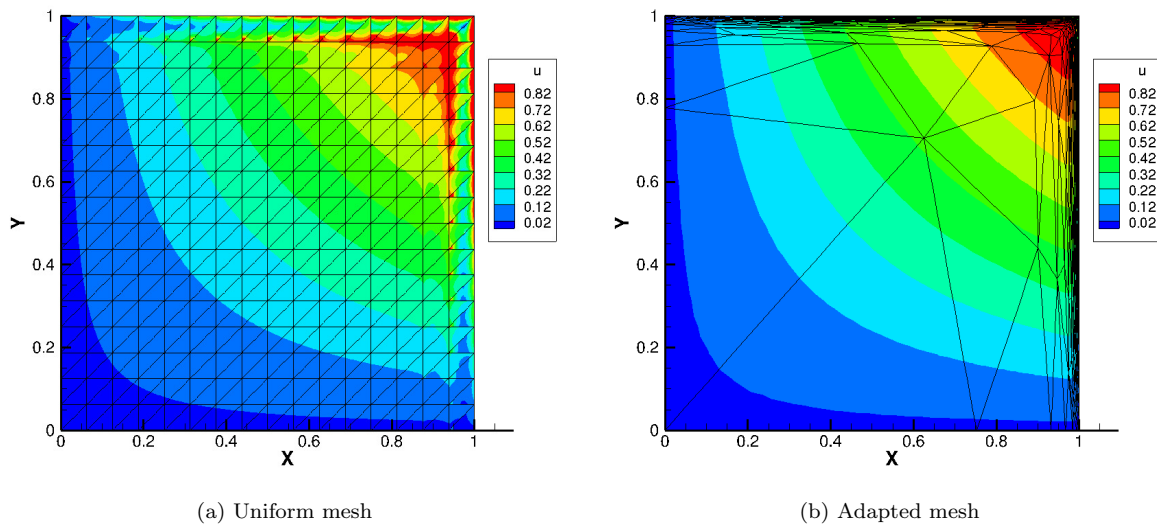


Figure 2: Mesh produced using analytic optimization

References

- [1] Ajay Mandyam Rangarajan, Aravind Balan, and Georg May, *Mesh Adaptation and Optimization for Discontinuous Galerkin Methods Using a Continuous Mesh Model*, AIAA Modeling and Simulation Technologies Conference, AIAA SciTech Forum, (AIAA 2016-2142)
- [2] Ajay Mandyam Rangarajan and Georg May, *Adjoint-based anisotropic mesh adaptation for Discontinuous Galerkin Methods Using a Continuous Mesh Model*, AIAA Aviation 2017 (Accepted)
- [3] Vit Dolejsi, Georg May and Ajay Mandyam Rangarajan, *A Continuous hp-mesh model for adaptive discontinuous Galerkin Schemes*, In preparation

Accurate Numerical Modeling of Small Collections of Cardiac Cells

Aslak Tveito, Simula Research Laboratory, aslak@simula.no,
Karoline Horgmo Jæger, Simula Research Laboratory, karolihj@simula.no
Miroslav Kuchta, University of Oslo, mirok@math.uio.no,
Kent-Andre Mardal, University of Oslo and Simula Research Laboratory, kent-and@simula.no,
Marie E. Rognes, Simula Research Laboratory, meg@simula.no

Keywords: *mixed finite elements, cardiac electrophysiology, interior interface, $H(\text{div})$*

Cardiac tissue is commonly modeled using the classical monodomain or bidomain models. These models have provided valuable insight in cardiac electrophysiology; in particular, the models have been useful tools for understanding the nature of excitation waves and how these waves are affected by mutation or by illness, and also how drugs can be applied to alleviate disease. However, the classical models of cardiac tissue fundamentally assume that the discrete nature of cardiac tissue can be represented by homogenized equations where the extracellular space, the cell membrane and the intracellular space are continuous and exist everywhere. This is a bold assumption.

In this talk, we will discuss a more accurate model of cardiac tissue based on a geometrically explicit representation of the extracellular space, the cell membrane and the intracellular domain. This *EMI model* of cardiac tissue combines elliptic equations in the extracellular and in the intracellular domains with a system of nonlinear equations modelling the electrochemical processes across the cell membrane. We will present the EMI equations, propose new families of mixed finite element methods for these equations, analyze these families of methods in terms of inf-sup stability and *a priori* convergence, and show some numerical results using FEniCS of course.




Figure 1: Membrane potential and surrounding extracellular potential for a simulation of two connected cylinder-shaped cells

References

- [1] A. Tveito, K. H. Jæger, M. Kuchta, K.-A. Mardal, and M. E. Rognes, Accurate Numerical Modeling of Small Collections of Cardiac Cells *in preparation*, 2017.

A Finite Element Approach to Computational Magnetics with Defect Correction

Ulrich Römer, Technische Universitaet Darmstadt, roemer@temf.tu-darmstadt.de,
 Sebastian Schöps, Technische Universitaet Darmstadt, schoeps@gsc.tu-darmstadt.de,
 Herbert De Gerssem, Technische Universitaet Darmstadt, degersem@temf.tu-darmstadt.de

Keywords: *Computational magnetics, defect correction, adjoint equation*

Computational magnetics is concerned with the numerical solution of the magnetostatic or magneto-quasistatic approximation of Maxwell's equations. Popular application examples are electrical machines, magnets and transformers. In this work, we focus on the stationary nonlinear problem

$$\nabla \times \left(\nu(|\nabla \times \vec{A}|) \nabla \times \vec{A} \right) = \vec{J}, \quad \text{in } \Omega \quad (1)$$

$$\vec{A} \times \vec{n} = 0, \quad \text{on } \partial\Omega, \quad (2)$$

on a domain Ω , with outer unit normal \vec{n} , where \vec{A} and \vec{J} refer to the magnetic vector potential and the source current density, respectively. In many applications it is important to take into account the nonlinearity, i.e., the dependency of the magnetic reluctivity ν on the solution, originating in the ferromagnetic behavior of iron or steel parts.

The solution is sought in the space $H(\text{curl}, \Omega)$ and the primary numerical method for this class of problems is the finite element method with Nédélec's elements. A numerical challenge, in addition to the nonlinearity, is the following: often the behavior of the device is fully characterized by the homogeneity of the solution in a small part of the computational domain Ω . Furthermore, the numerical accuracy in this region has to be very high, as for instance in particle accelerator magnets, very small inhomogeneities may have significant effects on the beam quality. Fig 1 depicts the computational model of a Stern-Gerlach deflection magnet, where Ω_0 is given by the beam area. In view of the high (local) accuracy requirements an h-p adaptive approach is promising, but remains difficult to realize. Additionally, several quantities of interest demand for a higher differentiability of the solution. For instance the field quality of deflection type magnets can be quantified with

$$Q(\vec{A}) = \frac{1}{|\Omega_0|} \int_{\Omega_0} \partial_{x_1} |\nabla \times \vec{A}| \, dx, \quad (3)$$

where x_1 represents the horizontal component in Figure 1. The expression for Q is not well-defined, even for higher order finite elements, and further post-processing is required.




Figure 1: Picture from [3]. Computational model of a Stern-Gerlach magnet with coil, air, iron and the beam area Ω_0 .

In this work we use defect correction to improve both the accuracy and differentiability of the finite element solution [1, 2]. The defect correction approach is summarized as follows: in a first step, the finite element solution is projected to a space of higher regularity. Typically, splines are employed to this end. However, spline reconstructions are limited to structured grids and we use radial basis functions

instead [3]. More precisely, in the two-dimensional case with finite element solution u_h , we determine the coefficients α of

$$\pi_h u_h(\vec{x}) = \sum_{i=1}^N \alpha_i \Phi(|\vec{x} - \vec{x}_i|) + p(\vec{x}), \quad (4)$$

with reconstruction operator π_h , by enforcing interpolation at the finite element nodes \vec{x}_i . The radial basis functions Φ considered in this work are polyharmonic splines, such as thin-plate spline $\Phi(r) = r^2 \log(r)$, cubics $\Phi(r) = r^3$ and quintics $\Phi(r) = r^5$. The polynomial p above is needed to ensure the existence of a unique solution of the interpolation problem. This reconstruction is then used in an extrapolation step to obtain a higher accuracy of the solution. The remaining error, after extrapolation, can be estimated using the adjoint equation, see, e.g., [4] and the references therein. After defect correction, the accuracy is determined by the approximation properties of the reconstruction operator π_h . This will be discussed theoretically and illustrated by numerical examples. In a last step, we localize the defect correction approach, to account for the local field homogeneity requirements of computational magnetics applications, outlined above. In this way, there is no need to implement a local p-refinement. The post-processing is very general as we only need to access the degrees of freedom of the method and assemble new right-hand-sides. Its main drawback lies in the computational cost for the reconstruction. On a structured grid, this cost is often negligible. On an unstructured grid, the cost can become high when the domain of reconstruction is large. This is due to the fact, that a dense system of equations needs to be solved when using radial basis functions, which makes the localization of the defect correction even more appealing. A further remedy consists in using a fast multipole acceleration of the radial basis function reconstruction.

Several numerical examples will be given for illustration. All simulation results are based on FEn-iCS and implementation details for computational magnetics examples will be discussed in this context. Higher order Nédélec's elements are readily available, but more effort is needed to implement the nonlinearity, which is given by monotonic cubic splines in general. We also give details of the defect correction scheme and discuss computational costs. Current implementations can handle three-dimensional problems, but the defect correction scheme is limited to two-dimensions. Nevertheless, we sketch the extension to three-dimensional and time-dependent problems.

References

- [1] H.J. Stetter, The defect correction principle and discretization methods, *Numerische Mathematik*, 29 (1978), 425-443.
- [2] M. B. Giles, E. Süli, Adjoint methods for PDEs: a posteriori error analysis and postprocessing by duality, *Acta numerica*, 11 (2002), 145-236.
- [3] U. Römer, S. Schöps, H. De Gersem: A defect corrected finite element approach for the accurate evaluation of magnetic fields on unstructured grids, *Journal of Computational Physics*, 335 (2017), 688-699.
- [4] N.A. Pierce, M.B. Giles: Adjoint and defect error bounding and correction for functional estimates, *Journal of Computational Physics*, 200 (2004), 769-794.

Implementation of Diffuse Interface Models in FEniCS

Martin Řehoř, Charles University, rehor@karlin.mff.cuni.cz,
Jan Blechta, Charles University, blechta@karlin.mff.cuni.cz

Keywords: *diffuse interface models, Cahn-Hilliard equations, Navier-Stokes(-Fourier) equations*

Diffuse interface models have been successfully used as a tool for mathematical description of the motion of several immiscible fluids and their mutual interaction. Multicomponent flows of this type are of interest in many engineering applications.

The diffuse interface models usually couple the system of fourth-order partial differential equations of Cahn-Hilliard type together with incompressible Navier-Stokes equations. We have recently developed a variant of the model that is applicable also in a general non-isothermal setting. The coupling allows for consistent incorporation of surface tension effects into the model. The interface between the components is in this case treated as a thin layer across which the components are allowed to mix. The main challenge in numerical simulations is the very fine spatial resolution required for capturing the dynamics of the interface.

Several variants of diffuse interface models have been implemented and tested using the FEniCS project [1] and FENaPack [2], which is an open-source package implementing preconditioners for Navier-Stokes equations. A particular variant of the model has been applied in a practical problem of simulating the float glass process (Pilkington process), see [3].

References

- [1] M. ALNÆS, J. BLECHTA, J. HAKE, A. JOHANSSON, B. KEHLET, A. LOGG, C. RICHARDSON, J. RING, M. ROGNES, AND G. WELLS, *The FEniCS Project version 1.5*, Archive of Numerical Software, 3 (2015). URL: <http://journals.ub.uni-heidelberg.de/index.php/ans/article/view/20553>, doi:10.11588/ans.2015.100.20553.
- [2] J. BLECHTA AND M. ŘEHOŘ, *FENaPack - FEniCS Navier-Stokes preconditioning package*. URL: <https://github.com/blechta/fenapack>.
- [3] M. ŘEHOŘ, J. BLECHTA, AND O. SOUČEK, *On some practical issues concerning the implementation of Cahn-Hilliard-Navier-Stokes type models*, International Journal of Advances in Engineering Sciences and Applied Mathematics, 9 (2017), pp. 30-39. URL: <http://dx.doi.org/10.1007/s12572-016-0171-4>, doi:10.1007/s12572-016-0171-4.

Optimality Of Apriori Error Estimates Of Discontinuous Galerkin Method For Elliptic Problems With Nonlinear Newton Boundary Conditions

Miloslav Feistauer, Charles University, feist@karlin.mff.cuni.cz,

Filip Roskovec, Charles University, roskovec@gmail.com,

Anna-Margarete Sändig, IANS, Universität Stuttgart, saendig@ians.uni-stuttgart.de

Keywords: *a priori error estimates, discontinuous Galerkin method, nonlinear boundary conditions*

We are concerned with the discontinuous Galerkin method (dG) for the numerical solution of the Poisson problem with nonlinear Newton boundary condition in bounded two-dimensional polygonal domains. Such boundary condition models a polynomial behavior at the boundary which can be used in modeling of electrolysis of aluminium (see e.g. [1]). Similar nonlinear property of the solution on the boundary is observed also in radiation heat transfer problems or in nonlinear elasticity.

We consider the following *boundary value problem*: Find $u : \bar{\Omega} \rightarrow \mathbb{R}$ such that

$$\begin{aligned} -\Delta u &= f & \text{in } \Omega, \\ \frac{\partial u}{\partial n} + \kappa |u|^\alpha u &= \varphi & \text{on } \partial\Omega, \end{aligned} \tag{1}$$

where $f : \Omega \rightarrow \mathbb{R}$ and $\varphi : \partial\Omega \rightarrow \mathbb{R}$ are given functions and $\kappa > 0$, $\alpha \geq 0$ are given constants.

In paper [2], the dG discretization of problem (1) is analyzed and theoretical a priori error estimates for the error are derived. The discretization error is measured by the so called dG norm, which consists of the $H^1(\Omega)$ -norm augmented by additional contributions over the edges of the mesh elements penalizing the discontinuity of the discrete dG solution.

We assume that the exact solution $u \in H^s(\Omega)$, $s > 2$, and the dG discretization with discontinuous piecewise polynomial functions of degree p is used. When the solution u is non-zero at least on a part of the boundary with positive $d - 1$ measure, it is proven in [3] that the dG method reaches the optimal rate of convergence measured in the dG norm $O(h^r)$, where $r = \min\{p, s - 1\}$. When $u \equiv 0$ on the whole boundary then the theoretical convergence rate is decreased to $O(h^{\frac{r}{\alpha+1}})$ due to the nonlinearity of the problem.

We present experiments, computed by the FEniCS software, performed in order to verify the optimality of these theoretical bounds and convergence rates. The decrease of the error both for problems with smooth solution and for irregular solutions in non-convex domains is examined. We measure the dependence of the experimental order of convergence on the regularity of the solution and also on the parameter α of the nonlinearity of the problem.

It seems that the convergence behaves similarly to the theoretical estimates in these experiments. Only, in the case when $u|_{\partial\Omega} = 0$, we observed that the $L^2(\Omega)$ -norm was the major part of the error decreasing as $O(h^{\frac{r+1}{\alpha+1}})$ while the $H^1(\Omega)$ -seminorm decreased with the optimal order r . Hence, it seems that the error estimate of order $O(h^{\frac{r}{\alpha+1}})$ proved in [2] is not optimal. An optimal error estimate based on the analysis of the $L^2(\Omega)$ norm is our goal for further work.

References

- [1] Feistauer M., Kalis H., Rokyta M.: Mathematical modelling of an electrolysis process. *Comment. Math. Univ. Carolinae* 30 1989, 465–477.
- [2] Feistauer, M., Roskovec, F., Sändig, A. M.: Discontinuous Galerkin Method for an Elliptic Problem with Nonlinear Newton Boundary Conditions in a Polygon, *IMA Journal of Numerical Analysis*, 2016 (submitted)
- [3] Bartoš, O.: Discontinuous Galerkin method for the solution of boundary-value problems in non-smooth domains, *Master thesis*, Charles University, Faculty of Mathematics and Physics, 2017 (in preparation)

Loop fusion for finite element assembly in PyOP2

Fabio Luporini, Imperial College London, f.luporini12@imperial.ac.uk,
Tianjiao Sun, Imperial College London, tianjiao.sun14@imperial.ac.uk

Keywords: *finite element assembly, performance optimization, loop fusion*

We discuss loop fusion techniques implemented in PyOP2 which speeds up the global assembly of linear and bilinear forms in Firedrake. The techniques involve direct fusion, applied to mixed finite elements, and indirect fusion, applied to Discontinuous Galerkin finite elements. We will also present preliminary results with several representative assembly kernels.

References

- [1] Florian Rathgeber, David A. Ham, Lawrence Mitchell, Michael Lange, Fabio Luporini, Andrew T. T. Mcrae, Gheorghe-Teodor Bercea, Graham R. Markall, and Paul H. J. Kelly: Firedrake: automating the finite element method by composing abstractions, *ACM Trans. Math. Softw.*, 2016.
- [2] Florian Rathgeber, Graham R. Markall, Lawrence Mitchell, Nicolas Lorient, David A. Ham, Carlo Bertolli, Paul H. J. Kelly: PyOP2: A High-Level Framework for Performance-Portable Simulations on Unstructured Meshes, *High Performance Computing, Networking Storage and Analysis, SC Companion*, 2012, 1116-1123.

UFL and the Automatic Derivation of Weak and Strong Shape Derivatives

Stephan Schmidt, University of Würzburg, stephan.schmidt@uni-wuerzburg.de

Keywords: *UFL, Tree Transformations, Shape optimization, Newton methods*

A fully automatic framework for conducting shape optimization using UFL and FEniCS is discussed. Such problems encompass a wide variety of problems such as finding drag optimal shapes in fluids or reconstructing inclusions by measuring reflections of waves within the area of non-destructive testing. Because the physics are usually modeled via partial differential equations (PDEs), shape optimization usually forms a special sub-class of problems within PDE constrained optimization, where the domain of the PDE is the design unknown.

Thus, the directional derivative of the objective J with respect to perturbations V of the input mesh Ω is necessary. Under sufficient regularity assumptions, tangential calculus can be used to find a surface representation of such derivatives. However, this usually requires applying the divergence theorem on surfaces.

The main goal of this talk is to present an automatic derivation of the shape derivative by first formulating the problem in the Uniform Form Language (UFL) and then processing the expression by a semantic analysis, which automatically applies the formal differentiation rules of shape calculus. In addition to applying the divergence theorem on surfaces, more complex transformations need to be applied when dependencies on the curvature or normal are present or when higher order derivatives are desired. An example of a UFL-tree representation of such a derivative with dependency on the normal n is shown in Figure 1.

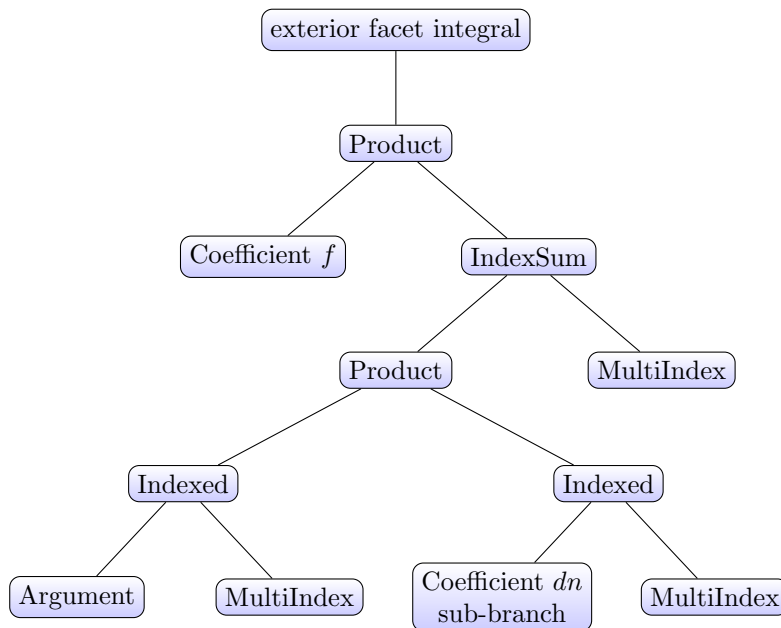


Figure 1: UFL tree representation of the surface form of a shape derivative. The shape derivative dn of the normal n can then be combined via independently processed subtrees.

Having a UFL-tree as an output means the result of the automatic differentiation procedure can immediately and automatically be discretized into a full numerical solution framework via FEniCS. The talk features examples of automatically created Newton-like solvers for problems in drag minimization in fluid flow and wave reconstructions. Because there are surface representations of such derivatives, this differentiation procedure also nicely interfaces with FEniCS and its capability to solve PDEs on manifolds, which is augmented with different approaches to compute curvature and other surface intrinsic quantities.

From medical images to uncertainty quantification with FEniCS and 3D Slicer

Vahid Shahrabi, University of Luxembourg, vahid.shahrabi.001@student.uni.lu,
Davide Baroli, University of Luxembourg, davide.baroli@uni.lu,
Jack S. Hale, University of Luxembourg, jack.hale@uni.lu,
Stéphane P.A. Bordas, University of Luxembourg, stephane.bordas@uni.lu

Keywords: computational biomechanics pipeline, segmentation, mesh smoothing, bone augmentation operation, FEniCS Project, uncertainty quantification, model-order reduction.

It is increasingly recognised that simulation tools can be useful in aiding physicians making critical medical decisions.


In this work we propose an image to finite element analysis pipeline for patient-specific simulation and demonstrate its effectiveness by simulating a Kyphoplasty operation on a fractured vertebra. Computing the von-Mises stress allows us to identify the region of the risk of fracture after the bone augmentation.

The pipeline consists a pre-processing of medical images segmentation using the open-source software 3DSlicer [Fedorov et al. 2012]. Then, the tetrahedral mesh quality is enhanced by modern smoothing technique within PyMesh [Zhou et al. 2016].

Finally, we use the FEniCS to compactly express and efficiently solve the elasticity model using the finite element method.

In addition, we enhance the classical pipeline with two cutting-edge numerical approaches: model reduction and uncertainty quantification. We feel it is increasingly important to understand the sensitivity of biomechanical models to underlying uncertainties in parameters, e.g. boundary conditions, material properties. However, the cost of such an analysis is computationally expensive, because it demands that the same model be run with many different realisations of the parameters.

To overcome this issue and obtain a near-real-time simulation, we introduce the model reduction technique.



References

- [Fedorov et al. 2012] Fedorov, Andriy and Beichel, Reinhard and Kalpathy-Cramer, Jayashree and Finet, Julien and Fillion-Robin, Jean-Christophe and Pujol, Sonia and Bauer, Christian and Jennings, Dominique and Fennessy, Fiona and Sonka, Milan and Buatti, John and Aylward, Stephen and Miller, James V. and Pieper, Steve and Kikinis, Ron. (2012) 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*, 30(9):1323–1341.
- [Zhou et al. 2016] Qingnan Zhou. PyMesh Geometry Processing Library for Python PyMesh 0.1 documentation manuscript.

Functional-type two-sided a posteriori error estimates for high-order discretizations

Satyendra Tomar, University of Luxembourg, satyendra.tomar@uni.lu,
Jack S. Hale, University of Luxembourg, jack.hale@uni.lu,
Stéphane P.A. Bordas, University of Luxembourg,
University of Western Australia, Cardiff University, stephane.bordas@uni.lu.

Keywords: *a posteriori error estimates, Poisson, higher-order, sharp, reliability.*

We present reliable and sharp two-sided a posteriori estimates for high-order discretisations of the Poisson problem using functional-type arguments. The resulting scheme is implemented in FEniCS.

Multi-Level Monte Carlo for Large-scale Vibrations Problems Using FEniCS and podS

H. Juliette T. Unwin, University of Cambridge, hjtu2@cam.ac.uk,
Garth N. Wells, University of Cambridge, gnw20@cam.ac.uk,
Nathan Sime, University of Cambridge, njcs4@cam.ac.uk,
Hugo Hadfield, University of Cambridge, hh409@cam.ac.uk

Keywords: *uncertainty quantification, multi-level Monte Carlo, eigenvalue problem, random vibrations*

We consider multi-level Monte Carlo methods (MLMC) [1] for accelerating the computation of stochastic eigenvalue problems arising from PDEs in the presence of uncertainty. Eigenvalue problems are important in a number of applications where randomness in the underlying system has a significant impact on the response, such as structural vibration. In these systems, designing to coincide with or avoid natural frequencies is important.

Monte Carlo (MC) methods can quantify the impact of uncertainties on eigenvalue problems in a simple way, but are considered to be too expensive in many practical cases. We propose employing MLMC methods to accelerate the computation of the statistics of eigenvalue problems to make MC techniques viable. MLMC methods compute the expectation and variance of the quantity of interest, Q , by solving the governing equation on grids of varying fidelity. The expectation of Q at the finest level, L , is calculated using the following telescopic sum:

$$\mathbb{E}[Q_L] = \mathbb{E}[Q_0] + \sum_{l=1}^L \mathbb{E}[Q_l - Q_{l-1}] \quad (1)$$

where increasing l represents levels of increasing mesh refinement. More expensive ‘finer grids’ are sampled less often than coarser grids because the variance of the ‘finer’ levels is smaller than the variance between the ‘coarser’ levels. This achieves equivalent accuracy as MC methods but at a much lower cost. The total variance of the estimator can also be calculated by summing the variance of the estimator at each level.

We consider uncertainty in our systems from random coefficients in the governing equations. We solve the eigenvalue problem using a discrete realisation of the random field and add point masses to the weak form of the equation. This is similar to the approach used in other random vibration research, e.g. [2]. We discuss the development of the FEniCS *PointSource* code, which now enables correct addition of the point sources to mass matrices.

We use the SLEPc eigenvalue solver combined with the FEniCS libraries to compute the natural frequencies for each random realisation and schedule and run the MLMC algorithm using the podS library in parallel on HPC and Microsoft Cloud facilities.

We argue that the demonstrated efficiency gains of the multi-level method make the approach competitive with alternatives to MC for eigenvalue problems. This enables us to compute quantities of interest when the underlying system has many degrees of freedom without having to make assumptions that reduce the accuracy of the result.

References

- [1] M. Giles: Multilevel Monte Carlo methods, *Acta Numerica*, 24 (2015), 259-328.
- [2] R. S. Langley and A. W. M. Brown: The ensemble statistics of the energy of a random system subjected to harmonic excitation, *Journal of Sound and Vibration*, 275 (2004), 823-846.

Multiscale model reduction using GMsFEM for applied problems in heterogeneous domains

Maria Vasilyeva,
North-Eastern Federal University, Russia and Texas A & M University, USA,
vasilyevadotmdotv@gmail.com

Keywords: *multiscale model reduction, generalized multiscale finite element method, heterogeneous domain, fractured media, perforated domain, approximation, FEniCS*

Mathematical modeling and numerical solution of the applied problems in highly heterogeneous media is a actual, necessary and interesting task. A classical numerical method for solving such problems should use a computational grid that resolve all small heterogeneities. For such problems, a numerical homogenization or multiscale methods are used. For the processes in a periodic heterogeneous media, where the length of the period is small in comparison with other parameters of the problem, the asymptotic homogenization methods is used to construct coarse-scale approximation and calculate effective properties that takes into account small scale heterogeneities. Nowadays, a multiscale methods are becoming popular, for example, heterogeneous multiscale method (HMM), multiscale finite element method (MsFEM), variational multiscale method (VMS) and others. Multiscale methods should combine the simplicity and efficiency of a coarse-scale models, and the accuracy of microscale approximations.

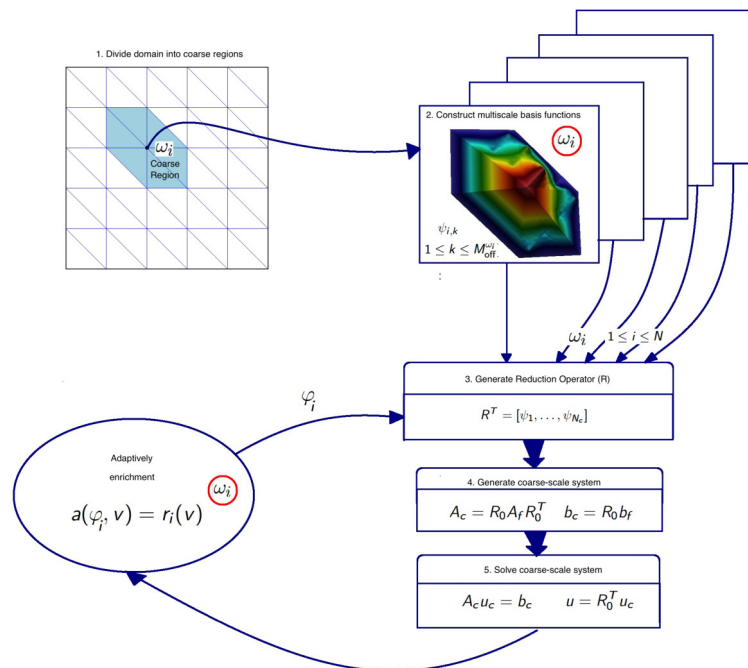


Figure 1: Multiscale model reduction using GMsFEM

In this work, we construct a reduced order model using the generalized multiscale finite element method (GMsFEM). This method involves two basic steps: (1) the construction of multiscale basic functions that take into account small scale heterogeneities in the local domains and (2) the construction of the coarse scale approximation [1]. For the construction of the multiscale basic functions we solve spectral problems in local domains. Spectral problems help to identify the most important characteristics of the solution. In contrast to the available techniques, this method allows to avoid the limitations associated with idealization and limitations on the applicability of the method [2]. This method also more general technology that takes into account the different scale processes. The construction of basic functions occurs

independently for each local domain, doesn't require the exchange of information between processors, and has a high parallelization efficiency. Using constructed multiscale basic functions, we construct a mathematical model on a coarse grid that allows to significantly reduce the solution time, the amount of used memory, and can be used to perform calculations for a given configuration of heterogeneous properties.

We consider several applied problems with different types of heterogeneities. As first problem, we consider gas filtration in the fractured poroelastic medium and present multiscale method for solution of the problem on a coarse grid [3]. Next, we consider wave propagation in the elastic fractured media [4], and finally present multiscale method for the solution of the pore-scale electrochemical processes in the lithium-ion batteries.

References

- [1] Y. Efendiev, J. Galvis, and T. Hou, Generalized multiscale finite element methods, *Journal of Computational Physics*, 251, pp. 116-135, 2013.
- [2] E. T. Chung, Y. Efendiev, G. Li, and M. Vasilyeva, Generalized multiscale finite element method for problems in perforated heterogeneous domains, *Applicable Analysis*, 255 (2015), pp. 1-15.
- [3] I Y Akkutlu, Y. Efendiev, and M. Vasilyeva. Multiscale model reduction for shale gas transport in fractured media. *Computational Geosciences* , 20(5), pp. 953973, 2016.
- [4] E. T. Chung, Y. Efendiev, R. L. Gibson, and M. Vasilyeva. A generalized multiscale finite element method for elastic wave propagation in fractured media. *GEM-International Journal on Geomathematics*, 7(2), pp. 163-182, 2016.

Shape Optimization with Geometric Constraints

Florian Wechsung, University of Oxford, wechsung@maths.ox.ac.uk,
 Patrick Farrell, University of Oxford, patrick.farrell@maths.ox.ac.uk

Keywords: *shape optimization, optimal control, Moreau-Yosida regularization, optimization in function spaces*

Shape optimization has received significant interest from both a theoretical and an applied point of view over the last decades. The approaches used can roughly be categorized into those based on using a parametrization for the shape or its deformation and then applying an optimization algorithm to the discretized problem (discretize-then-optimize) and those that formulate the problem as optimization over an infinite-dimensional space of shapes or deformations and then discretize afterwards (optimize-then-discretize). We follow the latter approach and search for diffeomorphisms $T \in [W^{1,\infty}(\Omega)]^d$ that deform an initial shape Ω . The optimization problem then reads as follows

$$\underset{T \in \mathcal{X} \cap K}{\text{minimize}} \quad J(T(\Omega)). \quad (1)$$

We choose $\mathcal{X} = \{T \in [W^{1,\infty}(\Omega)]^d : T \text{ is a } W^{1,\infty}\text{-diffeomorphism}\}$ as the space of admissible deformations as Lipschitz regularity of the domain is needed by many problems in which a PDE constraint is included in the optimization problem.

Furthermore we want to include certain geometric constraints; this is represented by the requirement $T \in K$. A classical example for a geometric constraint that is often considered is volume/mass conservation, i.e. $K = \{T \in \mathcal{X} : \text{vol}(T(\Omega)) = \text{vol}(\Omega)\}$. In our work we investigate constraints of the form

$$K = \{T \in \mathcal{X} : T(\partial\Omega) \subset C\} \quad (2)$$

for some convex set $C \in \mathbb{R}^d$. A classical application where this is relevant is wing design in Formula 1, where the teams are given bounding boxes in which the wing needs to be contained.




Figure 1: Shape at different stages of an optimization subject to box constraints.

In order to include this constraint, we begin by considering the indicator function $\phi : L^2(\partial\Omega) \rightarrow (-\infty, +\infty]$, defined via

$$\phi(T) = \begin{cases} 0, & \text{if } T(\partial\Omega) \subset C \\ +\infty, & \text{otherwise,} \end{cases} \quad (3)$$

and observe that (1) is equivalent to

$$\underset{T \in \mathcal{X}}{\text{minimize}} \quad J(T(\Omega)) + \phi(T). \quad (4)$$

In [Ito] it was shown that the Moreau-Yosida regularization of ϕ given by ϕ_c for $c > 0$ has the following useful properties:

1. $T \mapsto \phi_c(T, \lambda)$ is convex and Lipschitz-continuously differentiable;
2. for any $\lambda \in L^2(\partial\Omega)$, $\phi_c(T, \lambda) \uparrow \phi(T)$ as $c \rightarrow \infty$;

3. under some additional assumptions: T^* solves (4) if and only if (T^*, λ^*) are a solution of

$$\lambda^* = \phi'_c(T^*, \lambda^*) \quad (5)$$

and T^* minimizes

$$\underset{T \in \mathcal{X}}{\text{minimize}} \quad J(T(\Omega)) + \phi_c(T, \lambda^*). \quad (6)$$

This gives rise to three optimization approaches:

1. Find minimizers of $J(T(\Omega)) + \phi_c(T, 0)$ for increasingly large $c > 0$ (similar to the approach used in [Keu]).
2. Alternate between minimizing (6) and updating λ according to (5).
3. Apply a semi-smooth Newton method to (5) and (6) to solve for T^* and λ^* simultaneously.

We implement our methods in **Firedrake** and present synthetic examples without PDE constraint as seen in Figure 1 as well as the minimization of the Drag/Lift ratio of an airfoil, as seen in Figure 2, as an example for a PDE constrained problem.




Figure 2: Left: initial airfoil; middle: optimized airfoil without constraints; right: optimized airfoil with constraints.

References

- [1] M. Keuthen and M. Ulbrich: Moreau-Yosida regularization in shape optimization with geometric constraints, *Computational Optimization and Applications*, 62, 2014.
- [2] K. Ito and K. Kunisch: Augmented Lagrangian methods for nonsmooth, convex optimization in Hilbert spaces, *Nonlinear Analysis: Theory, Methods & Applications*, 41, 2000.

Grid Independence and Manufacturability in Microchannel Heat Sink Topology Optimization

Yang Zhou, KU Leuven, yang.zhou@kuleuven.be,
Paul Lacko, KU Leuven, paul.lacko@kuleuven.be,
Martine Baelmans, KU Leuven, martine.baelmans@kuleuven.be

Keywords: topology optimization, conjugate heat transfer, microchannel heat sinks, manufacturing constraints

The ongoing pursuit for higher power densities in electronic components such as CPUs, graphics cards, and other devices poses severe demands on cooling solutions. Air-cooled heat sinks are becoming increasingly inadequate and being replaced by liquid cooling counterparts. To explore the currently untapped potential of liquid-cooled microchannel heat sinks, topology optimization is used in this work. Using this methodology, novel tree-like microchannel layouts with superior performance over parallel channel layouts can be achieved automatically.

This work starts from earlier topology optimization studies for optimal cooling topologies for pressure-driven liquid-cooled microchannel heat sinks [1]. From a physical point of view, optimized designs should incorporate a large amount of very fine channels, since this is beneficial towards the heat transfer to the fluid. Additionally, for a fixed pressure drop over the heat sink, these channels should be fairly short to limit friction losses. Consequently, optimized designs include a combination of small channels with excellent heat transfer characteristics and larger channels acting as fluid distributors and collectors to limit the pressure drop, as seen in figure 1. In this figure, black zones correspond to solid walls and white to liquid channels. The left and right side are the fluid inlet and outlet respectively, and the top and bottom are fixed solid walls.




Figure 1: Example optimized heat sink design.

The optimization problem is solved on a discretized finite volume mesh based on the conjugate heat transfer equations. By introducing a material with imaginary properties into the simulation, material properties in every cell are updated during the optimization procedure. Within this procedure, the fictitious material is pushed towards either solid or fluid, in the final optimized design. This numerical solution procedure brings with it a number of issues, most notably, the appearance of channel dimensions that equal the coarseness or fineness of the design grid used in the optimization. Consequently, optimized designs exhibit a high degree of grid dependence. Moreover, on the one hand an accurate solution of state variables asks for a fine resolution, while on the other hand designs optimized on those fine grids are not manufacturable as a result.

In order to avoid such grid dependent designs and fabrication limits violation, regularization methods are employed. These methods originate from structural topology optimization research, where similar problems have been encountered [2]. Among multiple existing approaches, parameter space reduction [3] and mesh-independent filtering methods [4] are tested in this work.

We assume here that the microchannel heat sinks are fabricated using silicon etching techniques such as deep reactive ion etching. In order to ensure manufacturability, a lower threshold is imposed on the width of the smallest channels, equal to etching dimensions that lie comfortably within current technological limits. This threshold is subsequently translated into specific parameters used in each of the regularization methods. Optimized designs are deemed manufacturable if all feature sizes are above this selected lower threshold. Grid independence is verified by simulating the optimization problem with varying grid coarseness and comparing the resulting topologies.

Testing of the two regularization methods indicates that using parameter space reduction, grid independent optimized designs can be realized, and manufacturability can be guaranteed. However, the design space is limited by the coarseness of the design grid, resulting in features such as staircase designs. Mesh-independent filtering methods allow for smoother solid-liquid interfaces using finer design grids, but simulations show that although there is a moderate level grid independency in the optimized designs, the filtering methods fail to ensure manufacturability for any given design. Reasons for this undesirable behavior include the non-convexity of the optimization problem, which has a tendency to get stuck in different local minima, and the fact that the filter methods appear not to completely remove the high frequencies in the design.

In conclusion, it appears that parameter space reduction performs favorably compared to mesh-independent filtering methods with respect to manufacturability of optimized designs and both methods display a similar level of grid independence. However, both methods have their own specific advantages and disadvantages.

References

- [1] Van Oevelen T., Optimal Heat Sink Design for Liquid Cooling of Electronics, *PhD thesis*, KU Leuven, 2014.
- [2] Sigmund O. and Maute K., Topology optimization approaches: A comparative review, *Structural and Multidisciplinary Optimization*, Vol. 48, No. 6. pp. 1031-1055, Dec. 2013.
- [3] Wang F., Lazarov B. S., and Sigmund O., On projection methods, convergence and robust formulations in topology optimization, *Structural and Multidisciplinary Optimization*, Vol. 43, No. 6, pp. 767-784, Jun. 2011.
- [4] Sigmund O., Morphology-based black and white filters for topology optimization, *Structural and Multidisciplinary Optimization*, Vol. 33, No. 4-5, pp. 401-424, Feb. 2007.

Simulating the Response of Bioprosthetic Heart Valve Tissues to Cyclic Loading in FEniCS

Will Zhang, University of Texas at Austin, willwiz@utexas.edu,
Michael S. Sacks, University of Texas at Austin, msacks@ices.utexas.edu

Keywords: *permanent set, time dependent simulation, soft tissue*

There are over 100,000 heart valve surgeries done each year in the U.S. costing over \$20 billion [1]. Bioprosthetic heart valves(BHV) are the most popular surgical replacements, but the life expectancy remains stuck at 10-15 years. One of the most crucial effects involved in the mechanical failure of BHVs is permanent set(PS); especially in the early stage (2-3 years). PS cause the geometry of BHVs to change permanently over time, resulting in extraneous stress on the leaflets. While this process does not damage the tissue, it accelerates the rate of failure. PS is a result of the scission-healing of exogenously crosslinked matrix. This allows the unloaded configuration of the matrix to evolve while convecting the underlying collagen fiber architecture(CFA). Thus, we developed a structural constitutive model for the exogenously crosslinked BHVs biomaterials based on the PS mechanism [2]. The results show that PS alone can capture and more importantly predict how the shape and mechanical response leaflet biomaterial change during the early stage.

In this work, we developed a numerical implementation of our PS model to simulate BHV leaflet tissues under cyclic loading (Fig. 1). The simulations were performed quasi-statically. At each time step, we simulate the loaded state using the elastic part of the PS model implemented in FEniCS. The time dependent updates to the material model is done using a finite difference approach and a full implementation of the PS constitutive model in Python and C++. To facilitate the Finite Element simulation, we developed a simplified material model based on the non-covariant Hencky strains presented by Criscione *et al.* [3] to match the mechanical response of the elastic part of the full model,

$$\Psi(\mathbf{C}) = \sum_{i,j,k} c_{ijk} \gamma_i \gamma_j \gamma_k \quad (1)$$

$$\gamma_i = \ln(\mathbf{m}^T \cdot \mathbf{F} \cdot \mathbf{m}), \gamma_j = \ln(\mathbf{s}^T \cdot \mathbf{F} \cdot \mathbf{s}), \gamma_k = (\mathbf{m}^T \cdot \mathbf{F} \cdot \mathbf{m})^{-1} \mathbf{m}^T \cdot \mathbf{F} \cdot \mathbf{s},$$

where \mathbf{F} is the applied deformation, \mathbf{m} is the preferred direction of the collagen fibers and \mathbf{s} is the cross-preferred direction, and c_{ijk} are the material constants. This is a necessary step as structural models have substantial computational cost. Finite Element simulation is ran at each time step to determine the loaded state, which is then added to the strain history. This is then used to determine the evolution of the reference configuration and material constant using the discretized PS model (Fig. 2A),

$$\mathbf{S} = \phi_m \text{Exp} \left[-k \sum_{i=1}^n a(i) \right] \bar{\mathbf{S}}(\mathbf{F}_{\text{PS}}, \mathbf{I}, \mathbf{C}) + \phi_m \sum_{j=1}^n k a(j) \text{Exp} \left[-k \left(\sum_{i=1}^n a(i) - \sum_{i=1}^j a(i) \right) \right] \bar{\mathbf{S}}(\mathbf{F}_{\text{PS}}, \mathbf{A}(i\Delta s), \mathbf{C}), \quad (2)$$

where ϕ_m is the total mass fraction, $\bar{\mathbf{S}}$ is the stress of the material relative to the strain history $\mathbf{A}(t)$, k is the rate constant of PS, \mathbf{F}_{PS} is the deformation due to PS, $a(i)$ is the step size in time as a function of the iteration i , and the exponential functions are the relative mass fraction of each part of the constrained mixture. The step size in time, $a(i)$, is accelerated based on the expected deformation due to PS (Fig. 2B), where less PS results in larger time step. The geometry of the Finite Element mesh is then updated, and a new set of material parameters for the simplified model is generated at each integration point. This process is then repeated for the next iteration until the end of the simulation. The model geometry is change depending on when simulation experimental studies on BHV leaflet tissue patches, or BHV leaflets or intact valves.

To validate the numerical implementation, we simulated the experimental studies we used to establish the PS constitutive model form, where exogenously crosslinked bovine pericardium patches are cycled under strain control [4] or stress control [5], and with the preferred collagen fiber direction or orthogonal to the preferred collagen fiber direction. Next we performed parametric simulations under different loading conditions such as uniaxial extension, biaxial extension, simple shear, pure shear, and combinations of each, stress vs strain boundary conditions, heterogeneous vs homogeneous collagen fiber architecture, aligned vs spread fiber splays, collagen fiber crimp, and direction of collagen fiber alignment.

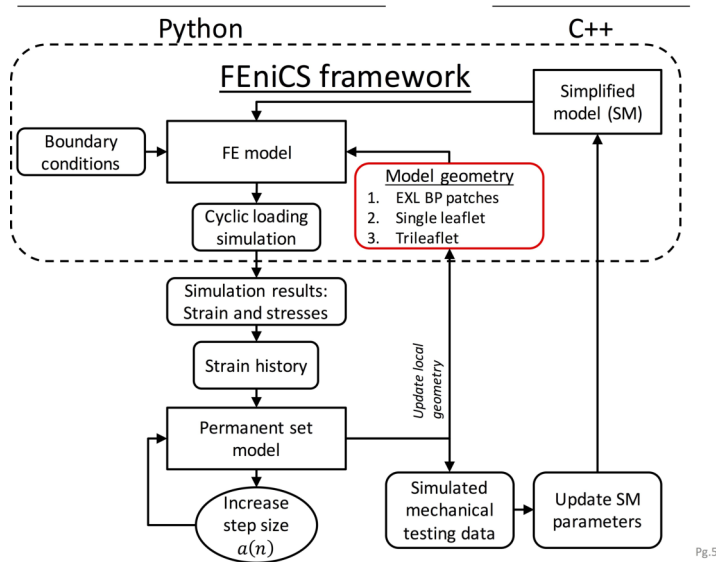


Figure 1: Numerical implementation using FEniCS, python and C++

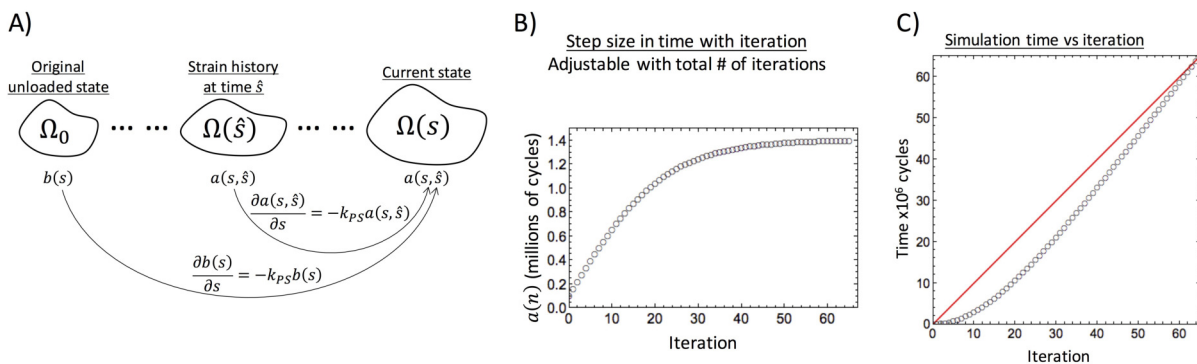


Figure 2: A) Updating the strain energy function using the constrained mixture approach and the transfer of mass fraction between reference states. B) Accelerating the time step size based on the expected size of the PS deformation. C) An example of the iteration vs simulated time.

We have developed a time dependent numerical implementation of our PS model using the FEniCS framework, and performed parametric studies to examine different loading conditions and collagen fiber architecture on the effects of PS. The results show the potential impact of PS on BHVs and we moved one step closer to simulating intact BHVs under cyclic loading using the FEniCS framework.

References

- [1] Mozaffarian, D.; Benjamin, E. J.; Go, A. S., *et al.* American Heart Association Statistics, C. & Stroke Statistics, S. Heart Disease and Stroke Statistics-2016 Update: A Report From the American Heart Association *Circulation*, 2016.
- [2] Zhang, W. and Sacks, M.S. Modeling the Response of Exogenously Crosslinked Tissue to Cyclic Loading: The Effects of Permanent Set, *Journal of the Mechanical Behavior of Biomedical Materials*, submitted.
- [3] Criscione, J. C.; Sacks, M. S. and Hunter, W. C. Experimentally tractable, pseudo-elastic constitutive law for biomembranes: I. Theory *J Biomech Eng*, 2003, 125.
- [4] Sun, W.; Sacks, M.; Fulchiero, G.; Lovekamp, J.; Vyavahare, N. and Scott, M. Response of heterograft heart valve biomaterials to moderate cyclic loading. *J Biomed Mater Res*, 2004.
- [5] Sellaro, T.L.; Hildebrand, D.; Lu, Q.; Vyavahare, N.; Scott, M. and Sacks, M. S. Effects of collagen fiber orientation on the response of biologically derived soft tissue biomaterials to cyclic loading. *J Biomed Mater Res A*, 2007.