# Sixth deal.II Users and Developers Workshop

Monday, July 23, 2018 - Friday, July 27, 2018

SISSA Main Building

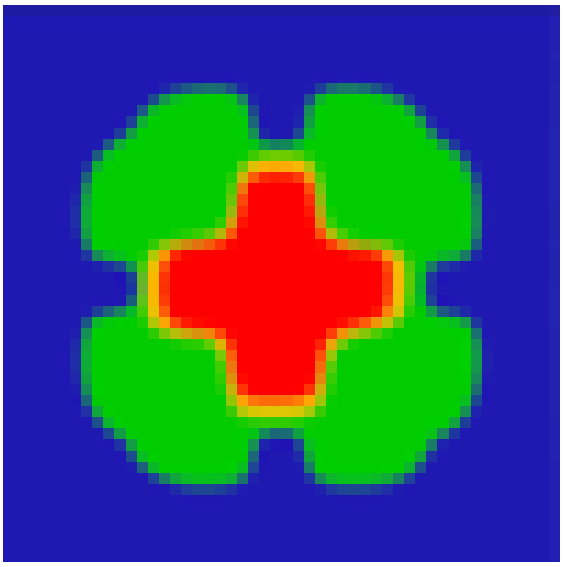

# Book of Abstracts

# Contents

**Keynotes** / **0**

## Welcome remarks. Past and future of deal.II

**Author:** Wolfgang Bangerth[1]

[1] *Colorado State University*

**Corresponding Author:** bangerth@colostate.edu

deal.II is now 20 years old. I'll give a brief overview of how it started, how it has developed, and where I see it going in the future. This includes both the technical side and, in particular, the social side: deal.II is not just a software project, but also a community of developers and users.

**Keynotes** / **8**

## Present and Future of adaptive, parallel, geometric multigrid in deal.II

**Author:** Timo Heister[1]

**Co-authors:** Guido Kanschat [2]; Martin Kronbichler [3]; Thomas Clevenger [1]

[1] *Clemson University*

[2] *Heidelberg University*

[3] *TU Munich*

**Corresponding Author:** heister@clemson.edu

Multigrid methods are the only known approach to efficiently solve large linear systems for solving PDEs by having runtime cost proportional to the number of unknowns. Algebraic multigrid methods construct a hierarchy of problems from the system matrix, while geometric multigrid methods use the mesh hierarchy and system matrices from each level of this hierarchy.

While algebraic multigrid methods have been successful in the past and are still the method of choice for large problems in deal.II, they have several disadvantages over geometric multigrid methods when implemented in a matrix-free version: First, scaling to more than a few thousand cores is limited due to the large setup cost. Second, the low arithmetic intensity does not lend itself to modern architectures. Third, implementation of multithreading is difficult. Fourth, the matrices consume large amounts of system memory.

Here we present the current status of adaptive, parallel, geometric multigrid how it is implemented in deal.II. Several numerical tests demonstrate the approach. Finally, we will close with future plans.

**Contributed talks** / **13**

## Accelerating augmented and deflated Krylov space methods for convection-diffusion problems

**Author:** Giuseppe Pitton[1]

**Co-author:** Luca Heltai [1]

[1] *SISSA*

**Corresponding Author:** gpitton@sissa.it

In this talk I will recall some basic notions of augmented and deflated Krylov space methods for the iterative solution of linear systems.
Then I will discuss a few strategies to apply these techniques to the solution of linear systems coming from nonlinear scalar convection-diffusion equations.
A simple back-of-the-envelope analysis allows to argue that in some cases it may be convenient to exploit alternative recycling strategies based on the SVD selection of previous solutions, and this conclusion is corroborated by numerical tests based on Finite Element and Spectral Element discretisations.

**Contributed talks** / **11**

## Coupling a temperature field to level set advection in a Stefan phase change problem

**Author:** Sean McGovern[1]

[1] *Forschungszentrum Juelich*

**Corresponding Author:** s.mcgovern@fz-juelich.de

Given the context of distributed parallelism, the purpose of this contribution is to consider a solution procedure to the classical Stefan phase change problem, an example of a free boundary problem. In the Stefan problem, the interface between water and ice moves through a domain as the ice melts due to a heat source. As an example of an interface problem, we choose to apply the level set method to implicitly track the evolution of the ice/water interface. The diffusive temperature equation must be solved over the whole domain while an evolving interface separates regions endowed with different material properties, i.e. specific heat of water and ice. The Gibbs-Thomson relation gives the discontinuity in the temperature field at the interface and the interface velocity is given by the jump in the heat flux.

To solve this kind of coupled nonlinear problem, we can 1) evolve the interface and 2) solve the temperature equation iteratively. The first is a hyperbolic advection equation and the second is parabolic. As is well-known, these require different solution tactics. We consider a setup for which an analytical solution exists.
The approach remains pedagogical as I would like to package this as a possible tutorial step.

**Contributed talks** / **18**

## Deal.II in cosmic ray propagation

**Authors:** Alberto Sartori[1]; Jiaxin Wang[1]; Luca Heltai[1]; Piero Ullio[2]

[1] *SISSA*

[2] *Astroparticle Physics Group, SISSA*

**Corresponding Author:** jiwang@sissa.it

Cosmic ray propagation in fluid approach can be described as a PDE system,

which lives in 6D spatial-spectral space.

In our project, we attempt to build up a solver for such a system,

by combining independent triangulation in two domains (spatial and spectral).

In this talk, I will present why and how we build this framework with Kronecker product.

**Keynotes** / **1**

## Integration and application of symbolic and automatic differentiation frameworks in deal.II

**Author:** Jean-Paul Pelteret[1]

**Co-authors:** Andrew McBride [2]; Isuru Fernando [3]; Paul Steinmann [4]

[1] *Chair of Applied Mechanics, Friedrich-Alexander University Erlangen-Nuremberg*

[2] *University of Glasgow*

[3] *University of Illinois*

[4] *Friedrich-Alexander University Erlangen-Nuremberg*

**Corresponding Author:** jean-paul.pelteret@fau.de

In recent years there has been a notable growth in research among fields involving coupled media, and in multi-physics problems in general. As the complexity of the physics describing the material loading and response increases, so does the time cost and sophistication of the computational tools used to analyse these problems both in the academic and industrial settings. Associated with this are further challenges, one example being that potential errors in implementation and subsequent numerical instabilities could be misinterpreted as physically plausible material instabilities.

Tools that automate differentiation are one mechanism to tackle the challenges of constitutive law and finite element formulation prototyping, development and validation. In this work we discuss the implementation of frameworks to perform assisted differentiation within the deal.II finite element library. We outline and discuss some of the challenges surrounding the integration of two algorithmic (automatic) and symbolic differentiation libraries within the pre-established tensor and finite element frameworks. Lastly, we discuss applications in the simulation of magnetorheological elastomers; these are smart, field-responsive composite materials that are of increasing interest in numerous industries. We demonstrate and benchmark the current implementation in the context of two problems, namely (i) a rate-dependent magnetic field-response constitutive law and (ii) homogenisation of representative volume element for a rate-independent magneto-rheological elastomer.

**Keynotes** / **10**

## CFL: A domain-specific language for simplifying integration kernels

**Author:** Daniel Arndt[1]

**Co-author:** Guido Kanschat [2]

[1] *University Heidelberg*

[2] *Heidelberg University*

**Corresponding Author:** daniel.arndt@iwr.uni-heidelberg.de

There are several different ways of describing bilinear forms in deal.II. All of these require to describe the assembly loop inside the program. This talk discusses an approach for providing the assembly

operator from the outside using a domain-specific language.
In particular, this easily allows changing the backend used in the end. CFL also provides the possibility to generate the Newton linearization from a residual operator described in this language.

**Contributed talks / 14**

## Fluid-Structure Interaction: simulating the cavitation phenomenon

**Author:** Dario Abbondanza[1]

**Co-authors:** Carlo Massimo Casciola [1]; Francesco Magaletti [1]; Mirko Gallo [1]

[1] *Sapienza University of Rome*

**Corresponding Author:** dario.abbondanza@uniroma1.it

It has been recently shown that a Diffuse Interface model is able to fully describe the dynamics of multiphase flows. In particular, the possibility of deriving a set of equations addressing the dynamic of the entire domain of simulation, makes it ideal to describe phenomena such as vapor bubble dynamics in a surrounding liquid. The model is able to capture surface tension effects, phase changes (condensation and evaporation) and even transitions to supercritical conditions, bubble collapse stages and the related shock-wave emissions. The diffuse interface model is intrinsically multiscale, enabling the study of all the characteristic size of the bubbles (from nm up to mm of radius) and the typical frequencies of evolution (of order 1-100 MHz). Moreover, the numerical simulations gives access to all the fields, anywhere in the domain, in particular predicting the pressure and thermal loads on the nearby solid walls.
A main application we are interested in is the cavitation phenomenon (implosion of vapor bubbles near solid surfaces), in order to take advantage of the high mechanical stresses (for biomedical applications, e.g.) or prevent them (to reduce or prevent structural damages). A full understanding of this complex phenomena requires, on one hand, a full description of the fluid part, and on the other hand, a suitable modeling of the solid dynamics with particular attention to plasticity and damage analysis. In this context, the use of the deal.II library could help thanks to its ability in automatically refining the mesh (i.e. in solving the smallest scales involved in the dynamics), and thanks to the large implementation possibility, already offered from the structural point of view, in terms of plasticity models.

**Contributed talks / 16**

## Model Order Reduction and Computational Homogenisation of Magnetorheological Elastomers

**Author:** Benjamin Brands[1]

**Co-authors:** Denis Davydov [2]; Julia Mergheim [1]; Paul Steinmann [1]

[1] *Friedrich-Alexander-Universität Erlangen-Nürnberg*

[2] *Friedrich-Alexander University Erlangen-Nuremberg*

**Corresponding Author:** benjamin.brands@fau.de

The mechanical response of magnetorheological composites is highly affected by an applied magnetic field. Since a generally valid constitutive law does not exist for such heterogeneous materials, multiscale techniques like computational homogenisation are commonly used to approximate effective macroscopic properties. In our approach the macroscopic quantities at a material point of a magnetorheological elastomer are derived from the response of the underlying micro-structure, where the constitutive law is known, using first-order homogenisation.

The computational cost of this nested solution scheme known as the $FE^2$ method prohibits the simulation of complex macroscopic problems. To mitigate the computational bottleneck the FE models on the microscale are replaced by reduced-order models (ROMs). In projection-based ROM the governing equations are projected onto the reduced basis, which is an approximation of the solution manifold of the parametrised partial differential equations (pPDE). The reduced basis is commonly constructed using previously computed solutions of the pPDE, e.g. by applying proper orthogonal decomposition or the reduced basis method.

We will present our approach for the construction and computation of the reduced-order models on the microscale. Through various numerical examples the accuracy and time savings of the reduced models will be discussed.

**Contributed talks / 4**

## Numerical (micro)aquarium using the deal.II library

**Author:** Nicola Giuliani[1]

[1] *SISSA*

**Corresponding Author:** ngiuliani@sissa.it

During the last 8 years the mathLab laboratory at SISSA has analysed the swimming mechanisms of micro-organisms (unicellular beings as protists, algae or bacteria) both from a theoretical and an experimental point view. In this work we present a way to derive an accurate and reliable numerical setting for the resolution of general swimming mechanism. Such a framework can be seen as a numerical microscope.

A general swimming motion (both for micro than macro organism) can be seen as a Fluid Structure Interaction problem.
Given the characteristic length-scale of the problem (100 microns at most) the fluid part of the problem is well described by the Stokes equations (in particular we use the Boundary Integral Equations for the Stokes system). For what concerns the structural part of the system the typical swimming mechanisms (flagellar movements, body deformations) involve large domain deformations which pose severe mesh quality problems to finite element or finite volume solvers, and make Boundary Element Method an ideal discretisation strategy for these FSI problems, requiring only the discretisation of the domain boundaries.

We present a series of tools that allows the resolution of swimming problem starting from the analysis of an experimental video. In particular we isolate the shape changes of the swimmer and we use this informations as input data to recover the rigid movements of the overall organism. We use an OpenSource 3D modelling tool (Blender) to generate the computational mesh and we use an efficient MPI parallel solver (based on deal.II) for Stokes problems exploiting a collocation Boundary Element Method. The BEMStokes library, released under LGPL license, solve both 2D and 3D problem and it is rigorously tested using continuous integration over more than 120 different test cases.

The resulting softwares constitutes what we call a “numerical aquarium” since it allows for an efficient and reliable study of complex swimming organisms complementing the analysis that can be obtained through a real microscope.

**Keynotes / 2**

## Supports for GPU in deal.II

**Author:** Bruno Turcksin[1]

[1] *ORNL*

**Corresponding Author:** bruno.turcksin@gmail.com

GPUs are extensively used in supercomputers both in Europe with Piz Daint and in the United States with Titan, Sierra, and Summit. The main reason is that they offer great floating points performance for a given power envelope. GPUs are also the de facto standard architectue for neural networks whose usage has been increasing quite rapidly in the last few years. It is therefore very advantageous if other applications can use the same infrastructure as the one used for AI. However, progamming on GPU is not the same as programming on CPU. GPUs offer a lot more parallelism but GPU cores are a lot less powerful than CPU cores.

In this talk, we will discuss how to use deal.II on GPUs for matrix-based and matrix-free applications. For matrix-based applications, we rely on cuSPARSE and cuSOLVER for operations on sparse matrices and for some solvers. Matrix-free methods rely on the fact that many modern codes are using Krylov methods as their solver of choice. These solvers do not require the matrix of the system but only the result of a sparse matrix-vector multiplication. This operation is memory bound both on CPU and GPU. Matrix-free methods reduce drastically the memory requirement and thus, allow for a better utilization of the hardware. Moreover, the extensive use of tensor products make these methods extremely interesting for high-order finite elements. High-order finite elements are attractive not only because of the convergence properties but also because they introduce structure even on unstructured meshes. Finally, the lesser memory requirement of matrix-free algorithms is very advantageous for GPUs since they do not have a lot of memory per core. We will show some results and discuss our path forward.

**Keynotes** / **5**

## Creative freedom for computational mesh generation in deal.II

**Author:** Nicola Giuliani[1]

**Co-author:** Luca Heltai [1]

[1] *SISSA*

**Corresponding Author:** ngiuliani@sissa.it

Many modern 3D modelling tools offer the possibility of creating complex geometries in a very straightforward way (what is usually called creative freedom). Such softwares are extensively used in computer graphics (video-gaming, animation studios) and they are able to produce smooth surfaces exploiting the subdivision surface modelling which was derived by Edwin Catmull and Jim Clark in 1978 as a generalisation of bi-cubic uniform B-spline surfaces to arbitrary topology. We believe that a bridge between such computer graphics techniques and scientific Finite Element computations can open new perspectives for the deal.II users community.
We present the possibility to use such tools to generate a complex geometry to be integrated as a computational mesh inside the deal.II library.

To demonstrate the possibilities of our approach we use Blender, an Open Source python based software which has emerged as one of the most successful 3D modelling tool. We believe that its success is due both to its outstanding possibilities and spreading, in particular we stress that Blender

- is Open Source software,
- allows for straightforward geometry creation using Graphic User Interface,
- is highly scriptable using a Python 3 based Text User Interface,
- has a large user community (tutorials, manuals, youtube channel, wikipedia).

We use the Open Asset Import Library (ASSIMP) to convert the standard Blender format into a deal.II geometry through the interface that is already inside the library. These passages allow deal.II users to generate complex three dimensional surfaces in a very controllable way.

We prove the capabilities of our approach in different framework varying from Boundary Element Method for Stokes flows, to the discretisation of surface operators on non trivial geometries.

**Contributed talks** / **17**

## Non matching grids in a distributed setting

**Author:** Giovanni Alzetta[1]

**Co-author:** Luca Heltai [1]

[1] *SISSA*

**Corresponding Author:** galzetta@sissa.it

I am beginning this talk with a quick recall of the general difficulties faced when using multiple grids on a distributed memory setting and some possible solutions.
As a model problem I am presenting briefly the Lagrange-multiplier method of step-60, where a harmonic equation with values prescribed at the boundary and on a sub-domain is solved, with a particular focus on the coupling matrix used. After discussing how a partial solution making use of bounding boxes was implemented in deal.II 9.0, I am discussing its current limitations and intended developements.

**Contributed talks** / **19**

## Multilevel solver for discontinuous viscosity Stokes problem

**Author:** Michal Wichrowski[1]

[1] *Instytut Podstawowych Problemów Techniki*

The goal of my work is to develop robust scalable solver for linear system arising from finite element discretization of Stokes problem with strongly variable viscosity. Such linear systems appear in fluid-structure problems.
Symmetric positive define problems with discontinuous coefficient can be solved effectively by multigrid preconditioned Krylov subspace methods. Here we extend this idea on saddle point problems by using multilevel preconditioned GMRes. The multilevel method developed by Breass and Sarazin and further refined by Zulehner is applied. A multigrid algorithm works on the saddle point problem by applying constrained smoother. The relaxation allows the iterative procedure to remain in divergence free subspace while smoothing out the error in both pressure and velocity. Similarly to other multilevel methods, the algorithm presented here does not require explicitly storing the matrix and thus, deal.II matrix-free framework can be used.

**Contributed talks** / **9**

## Fast Tensor Product Schwarz Smoothers

**Author:** Julius Witte[1]

**Co-authors:** Daniel Arndt [1]; Guido Kanschat [1]

[1] *Heidelberg University*

**Corresponding Author:** julius.witte@iwr.uni-heidelberg.de

The goal of the ExaDG project is to provide efficient finite element solvers relying on deal.II's MatrixFree framework based on higher order DG methods. However, the matrix-free operators restrict the choice of preconditioners. A powerful tool is offered by a geometric multigrid algorithm having a favorable linear complexity. The quality of multigrid preconditioners is crucially dependent on the choice of the smoother. Overlapping Schwarz smoothers using local spaces supported on vertex patches were introduced by Arnold, Falk, and Winther and provide robust convergence in particular for divergence-constrained problems and almost incompressible elasticity. In classic finite element codes, the local stiffness matrices are extracted from the global system matrix. Therefore, such smoothers are expensive to compute and store, especially for higher order DG discretizations. Exploiting again the tensor product finite elements on the subproblem level offers a remedy to both and fits well into a matrix-free framework. Thus, we approximate the global bilinear form by separable problems in the smoother.

First, we show how to reduce the complexity of inverting and storing subspace problems by means of the fast diagonalization method. Second, a general tensor product setting is presented to illustrate the sum-factorization approaches used in finite element related matrix-vector products resulting from the global forward as well as local inverse problems. Finally, the efficiency of our proposed fast diagonalized Schwarz smoothers is underlined by numerical experiments and we close with a brief overview of future plans.

**Keynotes** / **7**

## Algorithms and data structures for matrix-free finite element operators with MPI-parallel sparse multi-vectors

**Author:** Denis Davydov[1]

**Co-authors:** Martin Kronbichler [2]; Paul Steinmann [1]

[1] *Friedrich-Alexander University Erlangen-Nuremberg*

[2] *Technical University of Munich*

**Corresponding Author:** davydden@gmail.com

Traditional solution approaches for problems in quantum mechanics scale as $\mathcal{O}(N^3)$, where $N$ is the number of electrons. Various methods have been proposed to address this issue and obtain linear scaling $\mathcal{O}(N)$. One promising formulation is the direct minimization of energy. Such methods take advantage of physical localization of the solution, namely that the solution can be sought in terms of non-orthogonal orbitals with local support. This is often called the near-sightedness principle of matter.

In this talk we present numerically efficient implementation of sparse parallel vectors within the deal.II open-source finite element library suitable for matrix-free operator evaluation. Based on the a-priori chosen support for each vector, we develop algorithms and data structures to perform (i) matrix-free sparse matrix multi-vector products (SpMM) (ii) projection of an operator onto a sparse sub-space (inner products) (iii) post-multiplication with a matrix. Strong and weak scaling results are reported for a typical benchmark problem using quadratic and quartic finite element bases.

**Keynotes** / **6**

## Numerical Approximation of the Integral Fractional Laplacian

**Author:** Wenyu Lei[1]

[1] *Texas A&M University*

**Corresponding Author:** wenyu@math.tamu.edu

In this talk, we propose a new nonconforming finite element algorithm to approximate a linear problem involving the integral fractional Laplacian. The method is based on the Dunford-Taylor integral representation of the bilinear form corresponding to the variational formulation. We obtain the approximated solution using the deal.II finite element library with a non-uniform mesh setting. Finally, as an application, we consider an obstacle problem involving the integral fractional Laplacian and several numerical simulations are also provided.

**Contributed talks** / **21**

## deal.II interface with CAD date structures: description, results and open issues

**Author:** Andrea Mola[1]

[1] *SISSA*

To exploit the automated mesh generation and adaptive refinement capabilities offered by deal.II, the user must provide its program with a suitable description of the domain geometry. In particular, it is fundamental that the geometries of boundary and internal surfaces are available to the triangulation so that the nodes generated across each mesh refinement are placed onto such surfaces, in a way that preserves mesh quality. This task is currently carried out by the manifold classes. A wide variety of manifold descriptors for the most common analytical geometries is available in the library. Along with such useful instruments, a set of manifold descriptors which import and interrogate CAD files has been implemented to treat also arbitrary geometries. All such manifolds are based on wrappers of functions contained in the OpenCASCADE library, that are included in the corresponding deal.II namespace. This contribution will discuss such manifold descriptors and their working principles, mentioning relevant details of their implementation and providing examples and results of their applications. The discussion will then move to a series of open issues which currently limit the OpenCASCADE manifolds computational efficiency.

**Keynotes** / **12**

## An overview of some new features of deal.II 9.0

**Author:** Luca Heltai[1]

[1] *SISSA*

**Corresponding Author:** luca.heltai@sissa.it

Version 9.0.0 of the deal.II library offers many new features and interfaces with several new external libraries.

In this talk I will provide an overview of some of the new features that I have contributed in adding to the library, including (in no particular order)

- Improved support for curved geometries
- Non-standard quadrature rules
- Improved support for user-defined run-time parameters
- Support for GMSH

- Support for Assimp
- Support for NanoFlann
- Support for SUNDIALS

**3**

## Closing remarks

**Authors:** Luca Heltai[1]; Timo Heister[2]; Wolfgang Bangerth[3]

[1] *SISSA*

[2] *Clemson University*

[3] *Colorado State University*

**Corresponding Authors:** bangerth@colostate.edu, heister@clemson.edu, luca.heltai@sissa.it

Some conclusions and some plans for future developments and opportunities. In particular:

- How to attract funds to deal.II
- How to attract new users
- Open discussions on the best platforms for Forums/Discussions/Mailinglists
- Anything else

**20**

## deal.II interface with CAD date structures: description, results and open issues

**Author:** Andrea Mola[1]

[1] *SISSA*

**Corresponding Author:** amola@sissa.it

To exploit the automated mesh generation and adaptive refinement capabilities offered by deal.II, the user must provide its program with a suitable description of the domain geometry. In particular, it is fundamental that the geometries of boundary and internal surfaces are available to the triangulation so that the nodes generated across each mesh refinement are placed onto such surfaces, in a way that preserves mesh quality. This task is currently carried out by the manifold classes. A wide variety of manifold descriptors for the most common analytical geometries is available in the library. Along with such useful instruments, a set of manifold descriptors which import and interrogate CAD files has been implemented to treat also arbitrary geometries. All such manifolds are based on wrappers of functions constrained in the OpenCASCADE library, that are included in the corresponding deal.II namespace. This contribution will discuss such manifold descriptors and their working principles, mentioning relevant details of their implementation and providing examples and results of their applications. The discussion will then move to a series of open issues which currently limit the OpenCASCADE manifolds computational efficiency.