

## The DUNE framework: Basic concepts and recent developments



Peter Bastian<sup>a</sup>, Markus Blatt<sup>b</sup>, Andreas Dedner<sup>c</sup>, Nils-Arne Dreier<sup>d</sup>, Christian Engwer<sup>d</sup>, René Fritze<sup>d</sup>, Carsten Gräser<sup>e</sup>, Christoph Grüninger<sup>f</sup>, Dominic Kempf<sup>a</sup>, Robert Klöforn<sup>g,\*</sup>, Mario Ohlberger<sup>d</sup>, Oliver Sander<sup>h</sup>

<sup>a</sup> Heidelberg University, Germany

<sup>b</sup> Dr. Blatt HPC-Simulation-Software & Services, Germany

<sup>c</sup> University of Warwick, United Kingdom of Great Britain and Northern Ireland

<sup>d</sup> Applied Mathematics: Institute of Analysis and Numerics, University of Münster, Germany

<sup>e</sup> Freie Universität Berlin, Germany

<sup>f</sup> Universität Stuttgart, Germany

<sup>g</sup> NORCE Norwegian Research Centre AS, Norway

<sup>h</sup> Technische Universität Dresden, Germany

### ARTICLE INFO

#### Article history:

Available online 17 July 2020

### ABSTRACT

This paper presents the basic concepts and the module structure of the *Distributed and Unified Numerics Environment* and reflects on recent developments and general changes that happened since the release of the first DUNE version in 2007 and the main papers describing that state Bastian et al. (2008a, 2008b). This discussion is accompanied with a description of various advanced features, such as coupling of domains and cut cells, grid modifications such as adaptation and moving domains, high order discretizations and node level performance, non-smooth multigrid methods, and multiscale methods. A brief discussion on current and future development directions of the framework concludes the paper.

© 2020 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

The *Distributed and Unified Numerics Environment* DUNE<sup>1</sup> [1,2] is a free and open source software framework for the grid-based numerical solution of partial differential equations (PDEs) that has been developed for more than 15 years as a collaborative effort of several universities and research institutes. In its name, the term “distributed” refers to distributed development as well as distributed computing. The enormous importance of numerical methods for PDEs in applications has lead to the development of a large number of general (i.e. not restricted to a particular application) PDE software projects. Many of them will be presented in this special issue and an incomplete list includes AMDIS [3], deal.II [4], FEniCS [5], FreeFEM [6], HiFlow [7], Jaumin [8], MFEM [9], Netgen/NGSolve [10], PETSc [11], and UG4 [12].

The distinguishing feature of DUNE in this arena is its flexibility combined with efficiency. *The main goal of DUNE is to provide well-defined interfaces for the various components of a PDE solver for which then specialized implementations can be*

\* Corresponding author.

E-mail address: [robertk@posteo.org](mailto:robertk@posteo.org) (R. Klöforn).

<sup>1</sup> [www.dune-project.org](http://www.dune-project.org).

provided. DUNE is not build upon one single grid data structure or one sparse linear algebra implementation nor is the intention to focus on one specific discretization method only. All these components are meant to be exchangeable. This philosophy is based on a quote from *The Mythical Man-Month: Essays on Software Engineering* by Frederick Brooks [13, p. 102]:

Sometimes the strategic breakthrough will be a new algorithm, ...

Much more often, strategic breakthrough will come from redoing the representation of the data or tables. This is where the heart of a program lies.

This observation has lead to the design principles of DUNE stated in [2]:

- i. Keep a clear separation of data structures and algorithms by providing abstract interfaces algorithms can be built on. Provide different, special purpose implementations of these data structures.
- ii. Employ generic programming using templates in C++ to remove any overhead of these abstract interfaces at compile-time. This is very similar to the approach taken by the C++ standard template library (STL).
- iii. Do not reinvent the wheel. This approach allows us also to reuse existing legacy code from our own and other projects in one common platform.

Another aspect of flexibility in DUNE is to structure code as much as possible into separate modules with a clear dependence.

This paper is organized as follows. In Section 2 we describe the modular structure of DUNE and the ecosystem it provides. Section 3 describes the core modules and their concepts while Section 4 describes selected advanced features and illustrates the concepts with applications. The latter section is intended for selective reading depending on the interests of the reader. Section 5 concludes the paper with current development trends in DUNE.

## 2. The Dune ecosystem

The modular structure of DUNE is implemented by conceptually splitting the code into separate, interdependent libraries. These libraries are referred to as DUNE-modules (not to be confused with C++-modules and translation units). The DUNE project offers a common infrastructure for hosting, developing, building, and testing these modules. However, modules can also be maintained independently of the official DUNE infrastructure.

The so-called DUNE *core modules* are maintained by the DUNE developers and each stable release provides consistent releases of all core modules. These core modules are:

**DUNE-COMMON:** Basic infrastructure for all DUNE modules such as build scripts or dense linear algebra.

**DUNE-GEOMETRY:** Reference elements, element transformations, and quadrature rules.

**DUNE-GRID:** Abstract interface for general grids featuring any dimension, various element types, conforming and non-conforming, local hierarchical refinement, and parallel data decomposition. Two example implementations are provided.

**DUNE-ISTL:** Abstract interfaces and implementations for sparse linear algebra and iterative solvers on matrices with small dense blocks of sizes often known at compile time.

**DUNE-LOCALFUNCTIONS:** Finite elements on the reference element.

While the core modules build a common, agreed-upon foundation for the DUNE framework, higher level functionality based on the core modules is developed by independent groups, and concurrent implementations of some features focusing on different aspects exist. These additional modules can be grouped into the following categories:

**Grid modules** provide additional implementations of the grid interface including so-called meta grids, implementing additional functionality based on another grid implementation.

**Discretization modules** provide full-fledged implementations of finite element, finite volume, or other grid based discretization methods using the core modules. The most important ones are DUNE-FEM, DUNE-FUFEM, and DUNE-PDELAB.

**Extension modules** provide additional functionality interesting for all DUNE users which are not yet core modules. Examples currently are Python bindings, a generic implementation of grid functions, and support for system testing.

**Application modules** provide frameworks for applications. They make heavy use of the features provided by the DUNE ecosystem but do not intend to merge upstream as they provide application-specific physical laws, include third-party libraries, or implement methods outside of DUNE's scope. Examples are the porous media simulators OPM [14] and DuMu<sup>x</sup> [15], the FEM toolbox KASCADE7 [16], and the reduced basis method module DUNE-RB [17].

**User modules** are all other DUNE modules that usually provide applications implementations for specific research projects and also new development features that are not yet used by a large number of other DUNE users but over time may become extension modules.

Some of the modules that are not part of the DUNE core are designated as so-called *staging modules*. These are considered to be of wider interest and may be proposed to become part of the core in the future.

The development of the DUNE software takes place on the DUNE gitlab instance (<https://gitlab.dune-project.org>) where users can download and clone all openly available DUNE git repositories. They can create their own new projects, discuss issues, and open merge requests to contribute to the code base. The merge requests are reviewed by DUNE developers and others who want to contribute to the development. For each commit to the core and extension modules continuous integration tests are run to ensure code stability.

### 3. DUNE core modules and re-usable concepts

In this section we want to give an overview of the central components of DUNE, as they are offered through the core modules. These modules are described as they are in DUNE version 2.7, released in January 2020. Focus of the core modules is on those components modeling mathematical abstractions needed in a finite element method. We will discuss in detail the DUNE-GRID and DUNE-ISTL modules, explain the basic ideas of the DUNE-LOCALFUNCTIONS and DUNE-FUNCTIONS module, and discuss how the recently added Python support provided by the DUNE-PYTHON module works. While DUNE-COMMON offers central infrastructure and foundation classes, its main purpose is that of a generic C++ toolbox and we will only briefly introduce it, when discussing the build system and infrastructure in general.

#### 3.1. The DUNE grid interface – DUNE-GRID

The primary object of interest when solving partial differential equations (PDEs) are functions

$$f : \Omega \rightarrow R,$$

where the domain  $\Omega$  is a (piecewise) differentiable  $d$ -manifold embedded in  $\mathbb{R}^w$ ,  $w \geq d$ , and  $R = \mathbb{R}^m$  or  $R = \mathbb{C}^m$  is the range. In grid-based numerical methods for the solution of PDEs the domain  $\Omega$  is partitioned into a finite number of open, bounded, connected, and nonoverlapping subdomains  $\Omega_e$ ,  $e \in E$ ,  $E$  the set of elements, satisfying

$$\bigcup_{e \in E} \overline{\Omega_e} = \overline{\Omega} \quad \text{and} \quad \Omega_e \cap \Omega_{e'} = \emptyset \text{ for } e \neq e'.$$

This partitioning serves three separate but related tasks:

- i. *Description of the manifold.* Each element  $e$  provides a diffeomorphism (invertible and differentiable map)  $\mu_e : \hat{\Omega}_e \rightarrow \Omega_e$  from its reference domain  $\hat{\Omega}_e \subset \mathbb{R}^d$  to the subdomain  $\Omega_e \subset \mathbb{R}^w$ . It is assumed that the maps  $\mu_e$  are continuous and invertible up to the boundary  $\partial\hat{\Omega}_e$ . Together these maps give a piecewise description of the manifold.
- ii. *Computation of integrals.* Integrals can be computed by partitioning and transformation of integrals  $\int_{\Omega} f(x) dx = \sum_{e \in E} \int_{\hat{\Omega}_e} f(\mu_e(\hat{x})) d\mu_e(\hat{x})$ . Typically, the reference domains  $\hat{\Omega}_e$  have simple shape that is amenable to numerical quadrature.
- iii. *Approximation of functions.* Complicated functions can be approximated subdomain by subdomain for example by multivariate polynomials  $p_e(x)$  on each subdomain  $\Omega_e$ .

The goal of DUNE-GRID is to provide a C++ interface to describe such subdivisions, from now on called a “grid”, in a generic way. Additionally, approximation of functions (Task iii) requires further information to associate data with the constituents of a grid. The grid interface can handle arbitrary dimension  $d$  (although naive grid-based methods become inefficient for larger  $d$ ), arbitrary world dimension  $w$  as well as different types of elements, local grid refinement, and parallel processing.

##### 3.1.1. Grid entities and topological properties

Our aim is to separate the description of grids into a *geometrical part*, mainly the maps  $\mu_e$  introduced above, and a *topological part* describing how the elements of the grid are constructed hierarchically from lower-dimensional objects and how the grid elements are glued together to form the grid.

The topological description can be understood recursively over the dimension  $d$ . In a one-dimensional grid, the elements are edges connecting two vertices and two neighboring elements share a common vertex. In the combinatorial description of a grid the position of a vertex is not important but the fact that two edges share a vertex is. In a two-dimensional grid the elements might be triangles and quadrilaterals which are made up of three or four edges, respectively. Elements could also be polygons with any number of edges. If the grid is *conforming*, neighboring elements share a common edge with two vertices or at least one vertex if adjacent.

In a three-dimensional grid elements might be tetrahedra or hexahedra consisting of triangular or quadrilateral faces, or other types up to very general polyhedra.

In order to facilitate a dimension-independent description of a grid we call its constituents *entities*. An entity  $e$  has a dimension  $\dim(e)$ , where the dimension of a vertex is 0, the dimension of an edge is 1, and so on. In a  $d$ -dimensional grid the highest dimension of any entity is  $d$  and we define the *codimension* of an entity as

$$\text{codim}(e) = d - \dim(e).$$

We introduce the *subentity relation*  $\subseteq$  with  $e' \subseteq e$  if  $e' = e$  or  $e'$  is an entity contained in  $e$ , e.g. a face of a hexahedron. The set  $U(e) = \{e' : e' \subseteq e\}$  denotes all subentities of  $e$ . The *type* of an entity  $\text{type}(e)$  is characterized by the graph  $(U(e), \subseteq)$  being isomorphic to a specific reference entity  $\hat{e} \in \hat{E}$  (the set of all reference entities).

A  $d$ -dimensional grid is now given by all its entities  $E^c$  of codimension  $0 \leq c \leq d$ . Entities of each set  $E^c$  are represented by a different C++ type depending on the codimension  $c$  as a template parameter. In particular we call  $E^d$  the set of vertices,  $E^{d-1}$  the set of edges,  $E^1$  the set of facets, and  $E^0$  the set of elements. Grids of mixed dimension are not allowed, i.e. for every  $e' \in E^c$ ,  $c > 0$  there exists  $e \in E^0$  such that  $e' \subseteq e$ . We refer to [1,18] for more details on formal properties of a grid.

DUNE provides several implementations of grids all implementing the DUNE-GRID interface. Algorithms can be written generically to operate on different grid implementations. We now provide some code snippets to illustrate the DUNE-GRID interface. First we instantiate a grid:

```
const int dim = 4;
using Grid = Dune::YaspGrid<dim>;
Dune::FieldVector<double,dim> length; for (auto& l : length) l=1.0;
std::array<int,dim> nCells; for (auto& c : nCells) c=4;
Grid grid(length,nCells);
```

Here we selected the YaspGrid implementation providing a  $d$ -dimensional structured, parallel grid. The dimension is set to 4 and given as a template parameter to the YaspGrid class. Then arguments for the constructor are prepared, which are the length of the domain per coordinate direction and the number of elements per direction. Finally, a grid object is instantiated. Construction is implementation specific. Other grid implementations might read a coarse grid from a file.

Grids can be refined in a hierarchic manner, meaning that elements are subdivided into several smaller elements. The element to be refined is kept in the grid and remains accessible. More details on local grid refinement are provided in Section 4.1 below. The following code snippet refines all elements once and then provides access to the most refined elements in a so-called GridView:

```
grid.globalRefine(1);
auto gv = grid.leafGridView();
```

A GridView object provides read-only access to the entities of all codimensions in the view. Iterating over entities of a certain codimension is done by the following snippet using a range-based for loop:

```
const int codim = 2;
for (const auto& e : entities(gv,Dune::Codim<codim>{}))
    if (!e.type().isCube()) std::cout << "no cube " << std::endl;
```

In the loop body the type of the entity is accessed and tested for being a cube (here of dimension  $2=4-2$ ). Access via more traditional names is also possible:

```
for (const auto& e : elements(gv)) assert(e.codim()==0);
for (const auto& e : vertices(gv)) assert(e.codim()==dim);
for (const auto& e : edges(gv)) assert(e.codim()==dim-1);
for (const auto& e : facets(gv)) assert(e.codim()==1);
```

Range-based for loops for iterating over entities have been introduced with release 2.4 in 2015. Entities of codimension 0, also called elements, provide an extended range of methods. For example it is possible to access subentities of all codimensions that are contained in a given element:

```
for (const auto& e : elements(gv))
    for (unsigned int i=0; i<e.subEntities(codim); ++i)
        auto v = e.template subEntity<codim>(i);
```

This corresponds to iterating over  $U(e) \cap E^c$  for a given  $e \in E^0$ .

### 3.1.2. Geometric aspects

Geometric information is provided for  $e \in E^c$  by a map  $\mu_e : \hat{\Omega}_e \rightarrow \Omega_e$ , where  $\hat{\Omega}_e$  is the domain associated with the reference entity  $\hat{e}$  of  $e$  and  $\Omega_e$  is its image on the manifold  $\Omega$ . Usually  $\hat{\Omega}_e$  is one of the usual shapes (simplex, cube, prism, pyramid) where numerical quadrature formulas are available. However, the grid interface also supports arbitrary polygonal elements. In that case no maps  $\mu_e$  are provided and only the measure and the barycenter of each entity is available. Additionally, the geometry of intersections  $\Omega_e \cap \Omega_{e'}$  with  $d-1$ -dimensional measure for  $e, e' \in E^0$  is provided as well.

Working with geometric aspects of a grid requires working with positions, e.g.  $x \in \hat{\Omega}_e$ , functions, such as  $\mu_e$ , or matrices. In DUNE these follow the DenseVector and DenseMatrix interface and the most common implementations are the class templates FieldVector and FieldMatrix providing vectors and matrices with compile-time known size built on any data type having the operations of a field. Here are some examples (using dimension 3):

```
Dune::FieldVector<double,3> x({1.0,2.0,3.0}); // construct a vector
Dune::FieldVector<double,3> y(x);
y *= 1.0/3.0; // scale by scalar value
double s = x*y; // scalar product
double norm = x.two_norm(); // compute Euclidean norm
```




Fig. 1. Maps related to an interior intersection.

```
Dune::FieldMatrix<double,3,3> A({{1,0,0},{0,1,0},{0,0,1}});
A.mv(x,y); // y = Ax
A.usmv(0.5,x,y); // y += 0.5*Ax
```

An entity  $e$  (of any codimension) offers the method `geometry()` returning (a reference to) a geometry object which provides, among other things, the map  $\mu_e : \hat{\Omega}_e \rightarrow \Omega_e$  mapping a local coordinate in its reference domain  $\hat{\Omega}_e$  to a global coordinate in  $\Omega_e$ . Additional methods provide the barycenter of  $\Omega_e$  and the volume of  $\Omega_e$ , for example. They are used in the following code snippet to approximate the integral over a given function using the midpoint rule:

```
auto u = [](const auto& x){return std::exp(x.two_norm());};
double integral=0.0;
for (const auto& e : elements(gv))
    integral += u(e.geometry().center())*e.geometry().volume();
```

For more accurate integration DUNE provides a variety of quadrature rules which can be selected depending on the reference element and quadrature order. Each rule is a container of quadrature points having a position and a weight. The code snippet below computes the integral over a given function with fifth order quadrature rule on any grid in any dimension. It illustrates the use of the `global()` method on the geometry which evaluates the map  $\mu_e$  for a given (quadrature) point. The method `integrationElement()` on the geometry provides the measure arising in the transformation formula of the integral.

```
double integral = 0.0;
using QR = Dune::QuadratureRules<Grid::ctype,Grid::dimension>;
for (const auto& e : elements(gv)) {
    auto geo = e.geometry();
    const auto& quadrature = QR::rule(geo.type(),5);
    for (const auto& qp : quadrature)
        integral += u(geo.global(qp.position()))
            *geo.integrationElement(qp.position())*qp.weight();
}
```

An intersection  $I$  describes the intersection  $\Omega_I = \partial\Omega_{i(I)} \cap \partial\Omega_{o(I)}$  of two elements  $i(I)$  and  $o(I)$  in  $E^0$ . Intersections can be visited from each of the two elements involved. The element from which  $I$  is visited is called the *inside* element  $i(I)$  and the other one is called the *outside* element  $o(I)$ . Note that  $I$  is *not* necessarily an entity of codimension 1 in the grid. In this way DUNE-GRID allows for nonconforming grids. In a conforming grid, however, every intersection corresponds to a codimension 1 entity. For an intersection three maps are provided:

$$\mu_I : \hat{\Omega}_I \rightarrow \Omega_I, \quad \eta_{I,i(I)} = \hat{\Omega}_I \rightarrow \hat{\Omega}_{i(I)}, \quad \eta_{I,o(I)} = \hat{\Omega}_I \rightarrow \hat{\Omega}_{o(I)}.$$

The first map describes the domain  $\Omega_I$  by a map from a corresponding reference element. The second two maps describe the embedding of the intersection into the reference elements of the inside and outside element, see Fig. 1, such that

$$\mu_I(\hat{x}) = \mu_{i(I)}(\eta_{I,i(I)}(\hat{x})) = \mu_{o(I)}(\eta_{I,o(I)}(\hat{x})).$$

Intersections  $\Omega_I = \partial\Omega_{i(I)} \cap \partial\Omega$  with the domain boundary are treated in the same way except that the outside element is omitted.

As an example consider the approximative computation of the elementwise divergence of a vector field  $\text{div}_e = \int_{\Omega_e} \nabla \cdot f(x)dx = \int_{\partial\Omega_e} f \cdot nds$  for all elements  $e \in E^0$ . Using again the midpoint rule for simplicity this is achieved by the following snippet:

```
auto f = [](const auto& x){return x;};
for (const auto& e : elements(gv)) {
```

```

double divergence=0.0;
for (const auto& I : intersections(gv,e)) {
    auto geo = I.geometry();
    divergence += f(geo.center())*I.centerUnitOuterNormal()
                *geo.volume();
}
}

```

### 3.1.3. Attaching data to a grid

In grid-based methods data, such as degrees of freedom in the finite element method, is associated with geometric entities and stored in containers, such as vectors, external to the grid. To that end, the grid provides an index for each entity that can be used to access random-access containers. Often there is only data for entities of a certain codimension and geometrical type (identified by its reference entity). Therefore we consider subsets of entities having the same codimension and reference entity

$$E^{c,\hat{e}} = \{e \in E^c : e \text{ has reference entity } \hat{e}\}.$$

The grid provides bijective maps

$$\text{index}_{c,\hat{e}} : E^{c,\hat{e}} \rightarrow \{0, \dots, |E^{c,\hat{e}}| - 1\}$$

enumerating all the entities in  $E^{c,\hat{e}}$  consecutively and starting with zero. In simple cases where only one data item is to be stored for each entity of a given codimension and geometric type this can be used directly to store data in a vector as shown in the following example:

```

auto& indexset = gv.indexSet();
Dune::GeometryType gt(Dune::GeometryTypes::cube(dim)); // encodes (c,ê)
std::vector<double> volumes(indexset.size(gt)); // allocate container
for (const auto& e : elements(gv))
    volumes[indexset.index(e)] = e.geometry().volume();

```

Here, the volumes of the elements in a single element type grid are stored in a vector. Note that a `GeometryType` object encodes both, the dimension and the geometric type, e.g. simplex or cube. In more complicated situations an index map for entities of different codimensions and/or geometry types needs to be composed of several of the simple maps. This leaves the layout of degrees of freedom in a vector under user control and allows realization of different blocking strategies. DUNE-GRID offers several classes for this purpose, such as `MCMGMapper` which can map entities of multiple codimensions and multiple geometry types to a consecutive index.

When a grid is modified through adaptive refinement, coarsening, or load balancing in the parallel case, the index maps may change as they are required to be consecutive and zero-starting. In order to store and access data reliably when the grid is modified each geometric entity is equipped with a global id:

$$\text{globalid} : \bigcup_{c=0}^d E^c \rightarrow \mathbb{I}$$

where  $\mathbb{I}$  is a set of unique identifiers. The map `globalid` is injective and persistent, i.e. `globalid(e)` does not change under grid modification when entity  $e$  is in the old and the new grid, and `globalid(e)` is not used when  $e$  was in the old grid and is not in the new grid (note that global ids may be used again after the next grid modification). There are very weak assumptions on the ids provided by `globalid`. They do not need to be consecutive, actually they do not even need to be numbers. Here is how element volumes would be stored in a map:

```

const auto& globalidset = gv.grid().globalIdSet();
using GlobalId = Grid::GlobalIdSet::IdType;
std::map<GlobalId, double> volumes2;
for (const auto& e : elements(gv))
    volumes2[globalidset.id(e)] = e.geometry().volume();

```

The type `GlobalId` represents  $\mathbb{I}$  and must be sortable and hashable. This requirement is necessary to be able to store data for example in an `std::map` or `std::unordered_map`. For example, `YaspGrid` uses the `bigunsignedint` class from DUNE-COMMON that implements arbitrarily large unsigned integers, while DUNE-UGGrid uses a `std::uint_least64_t` which is stored in each entity. In DUNE-ALUGRID the `GlobalId` of an element is computed from the macro element's unique vertex ids, codimension, and refinement information.

The typical use case would be to store data in vectors and use an `indexset` while the grid is in *read-only* state and to copy only the necessary data to a map using `globalidset` when the grid is being modified. Since using a `std::map` may not be the most efficient way to store data, a utility class `PersistentContainer<Grid, T>` exists, that implements the strategy outlined above for arbitrary types  $T$ . To allow for optimization, this class can be specialized by the grid implementation using structural information to optimize performance.






Fig. 2. Level and leaf grid views.

### 3.1.4. Grid refinement and coarsening

Adaptive mesh refinement using a posteriori error estimation is an established and powerful technique to reduce the computational effort in the numerical solution of PDEs, see e.g. [19]. DUNE-GRID supports the typical *estimate-mark-refine* paradigm as illustrated by the following code example:

```
const int dim = 2;
using Grid = Dune::UGGrid<dim>;
auto pgrid = std::shared_ptr<Grid>(
    Dune::GmshReader<Grid>::read("circle.msh"));
auto h = [](const auto& x)
    {auto d=x.two_norm(); return 1E-6*(1-d)+0.01*d;};
for (int i=0; i<15; ++i) {
    auto gv = pgrid->leafGridView();
    for (const auto& e : elements(gv)) {
        auto diameter=std::sqrt(e.geometry().volume()/M_PI);
        if (diameter>h(e.geometry().center())) pgrid->mark(1,e);
    }
    pgrid->preAdapt();
    pgrid->adapt();
    pgrid->postAdapt();
}
```

Here the UGGrid implementation is used in dimension 2. In each refinement iteration those elements with a diameter larger than a desired value given by the function `h` are marked for refinement. The method `adapt()` actually modifies the grid, while `preAdapt()` determines grid entities which might be deleted and `postAdapt()` clears the information about new grid entities. In between `preAdapt()` and `adapt()` data from the old grid needs to be stored using persistent global ids and in between `adapt()` and `postAdapt()` this data is transferred to the new grid. In order to identify elements that may be affected by grid coarsening and refinement the element offers two methods. The method `mightVanish()`, typically used between `preAdapt()` and `adapt()`, returns `true` if the entity might vanish during the grid modifications carried out in `adapt()`. Afterwards, the method `isNew()` returns `true` if an element was newly created during the previous `adapt()` call. How an element is refined when it is marked for refinement is specific to the implementation. Some implementations offer several ways to refine an element. Furthermore some grids may refine non-marked elements in an implementation specific way to ensure certain mesh properties like, e.g., conformity. For implementation of data restriction and prolongation a `geometryInFather()` method provides geometrical mapping between parent and children elements.

Grid refinement is hierarchic in all currently available DUNE-GRID implementations. Each entity is associated with a *grid level*. After construction of a grid object all its entities are on level 0. When an entity is refined the entities resulting from this refinement, also called its direct children, are added on the next higher level. Each level-0-element and all its descendants form a tree. All entities on level  $l$  are the entities of the level  $l$  grid view. All entities that are not refined are the entities of the leaf grid view. This is illustrated in Fig. 2. The following code snippet traverses all vertices on all levels of the grid using a `levelGridView`:

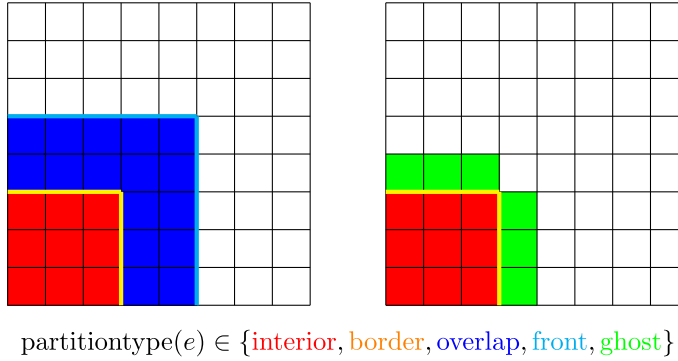
```
for (int l=0; l<=pgrid->maxLevel(); l++)
    for (const auto& v : vertices(pgrid->levelGridView(l)))
        assert( v.level() == l); // check level consistency
```

Each `GridView` provides its own `IndexSet` and so allows to associate data with entities of a single level or with all entities in the leaf view.

### 3.1.5. Parallelization

Parallelization in DUNE-GRID is based on three concepts: i. data decomposition, ii. message passing paradigm and iii. single-program-multiple-data (SPMD) style programming. As for the refinement rules in grid adaptation the data decomposition is implementation specific but must adhere to certain rules:

- i. The decomposition of codimension 0 entities  $E^0$  into sets  $E^{0,r}$  assigned to process rank  $r$  form a (possibly overlapping) partitioning  $\bigcup_{i=0}^{p-1} E^{c,r} = E^c$ .



**Fig. 3.** Two types of data decomposition implemented by YaspGrid (left) and UGGrid/ALUGrid (right). The colored entities show the entities of one rank. Sub-entities of elements assume the partition type of the element unless those sub-entities are located on a border between different partition types of interior, overlap or ghost.

ii. When process  $r$  has a codimension 0 entity  $e$  then it also stores all its subentities, i.e.  $e \in E^{0,r} \wedge E^c \ni f \subseteq e \Rightarrow f \in E^{c,r}$  for  $c > 0$ .

iii. Each entity is assigned a partition type attribute

$$\text{partitiontype}(e) \in \{\text{interior}, \text{border}, \text{overlap}, \text{front}, \text{ghost}\}$$

with the following semantics:

iii.a. Codimension 0 entities may only have the partition types interior, overlap, or ghost. The interior codimension 0 entities  $E^{0,r,\text{interior}} = \{e \in E^{0,r} : \text{partitiontype}(e) = \text{interior}\}$  form a nonoverlapping partitioning of  $E^0$ . Codimension 0 entities with partition type overlap can be used like regular entities whereas those with partition type ghost only provide a limited functionality (e.g. intersections may not be provided).

iii.b. The partition type of entities with codimension  $c > 0$  is derived from the codimension 0 entities they are contained in. For any entity  $f \in E^c$ ,  $c > 0$ , set  $\Sigma^0(f) = \{e \in E^0 : f \subseteq e\}$  and  $\Sigma^{0,r}(f) = \Sigma^0(f) \cap E^{0,r}$ . If  $\Sigma^{0,r}(f) \subseteq E^{0,r,\text{interior}}$  and  $\Sigma^{0,r}(f) = \Sigma^0(f)$  then  $f$  is interior, else if  $\Sigma^{0,r}(f) \cap E^{0,r,\text{interior}} \neq \emptyset$  then  $f$  is border, else if  $\Sigma^{0,r}(f)$  contains only overlap entities and  $\Sigma^{0,r}(f) = \Sigma^0(f)$  then  $f$  is overlap, else if  $\Sigma^{0,r}(f)$  contains overlap entities then  $f$  is front, else  $f$  is ghost.

Two examples of typical data decomposition models are shown in Fig. 3. Variant (a) on the left with interior/overlap codimension 0 entities is implemented by YaspGrid, variant (b) on the right with the interior/ghost model is implemented by UGGrid and ALUGrid.

To illustrate SPMD style programming we consider a simple example. Grid instantiation is done by all processes  $r$  with identical arguments and each stores its respective grid partition  $E^{c,r}$ .

```

const int dim = 2;
using Grid = Dune::YaspGrid<dim>;
Dune::FieldVector<double,dim> length; for (auto& l:length) l=1.0;
std::array<int,dim> nCells; for (auto& c : nCells) c=20;
Grid grid(length,nCells,std::bitset<dim>(OULL),1);
auto gv = grid.leafGridView();

```

Here, the third constructor argument of the grid controls periodic boundary conditions and the last argument sets the amount of overlap in codimension 0 entities.

Parallel computation of an integral over a function using the midpoint rule is illustrated by the following code snippet:

```

auto u = [](const auto& x){ return std::exp(x.two_norm()); };
double integral=0.0;
for (const auto& e : elements(gv,Dune::Partitions::interior))
    integral += u(e.geometry().center())*e.geometry().volume();
integral = gv.comm().sum(integral);

```

In the range-based for loop we specify in addition that iteration is restricted to interior elements. Thus, each element of the grid is visited exactly once. After each process has computed the integral on its elements a global sum (allreduce) produces the result which is now known by each process.

Data on overlapping entities  $E^{c,r} \cap E^{c,s}$  stored by two processes  $r \neq s$  can be communicated with the abstraction `CommDataHandleIF` describing which information is sent for each entity and how it is processed. Communication is



then initiated by the method `communicate()` on a `GridView`. Although all current parallel grid implementation use the message passing interface (MPI) in their implementation, nowhere the user has to make explicit MPI calls. Thus, an implementation could also use shared memory access to implement the DUNE-GRID parallel functionality. Alternatively, multithreading can be used within a single process by iterating over grid entities in parallel. This has been implemented in the EXA-DUNE project [20,21] or in DUNE-FEM [22] but a common interface concept is not yet part of DUNE core functionality.

### 3.1.6. List of grid implementations

The following list gives an overview of existing implementations of the DUNE-GRID interface and their properties and the DUNE module these are implemented in. Where noted, the implementation wraps access to an external library. A complete list can be found on the DUNE web page <https://dune-project.org/doc/grids/>.

**AlbertaGrid** (`dune-grid`) Provides simplicial grids in two and three dimensions with bisection refinement based on the ALBERTA software [23].

**ALUGrid** (`dune-alugrid`) Provides a parallel unstructured grid in two and three dimensions using either simplices or cubes. Refinement is nonconforming for simplices and cubes. Conforming refinement based on bisection is supported for simplices only [24].

**CurvilinearGrid** (`dune-curvilineargrid`) Provides a parallel simplicial grid [25] supporting curvilinear grids read from Gmsh [26] input.

**CpGrid** (`opm-grid`) Provides an implementation of a corner point grid, a nonconforming hexahedral grid, which is the standard in the oil industry <https://opm-project.org/>.

**FoamGrid** (`dune-foamgrid`) Provides one and two-dimensional grids embedded in three-dimensional space including non-manifold grids with branches [27].

**OneDGrid** (`dune-grid`) Provides an adaptive one-dimensional grid.

**PolygonGrid** (`dune-polygongrid`) A grid with polygonal cells (2d only).

**UGGrid** (`dune-grid`) Provides a parallel, unstructured grid with mixed element types (triangles and quadrilaterals in two, tetrahedra, pyramids, prisms, and hexahedra in three dimensions) and local refinement. Based on the UG library [28].

**YaspGrid** (`dune-grid`) A parallel, structured grid in arbitrary dimension using cubes. Supports non-equidistant mesh spacing and periodic boundaries.

Metagrids use one or more implementations of the DUNE-GRID interface to provide either a new implementation of the DUNE-GRID interface or new functionality all together. Here are examples:

**GeometryGrid** (`dune-grid`) Takes any grid and replaces the geometries of all entities  $e$  by the concatenation  $\mu_{geo} \circ \mu_e$  where  $\mu_{geo}$  is a user-defined mapping, see Section 4.1.2.

**PrismGrid** (`dune-metagrid`) Takes any grid of dimension  $d$  and extends it by a structured grid in direction  $d + 1$  [29].

**GridGlue** (`dune-grid-glue`) Takes two grids and provides a projection of one on the other as a set of intersections [30,31], see Section 4.2.1.

**MultiDomainGrid** (`dune-multidomaingrid`) Takes a grid and provides possibly overlapping sets of elements as individual grids [32], see Section 4.2.2.

**SubGrid** (`dune-subgrid`) Takes a grid and provides a subset of its entities as a new grid [33].

**IdentityGrid** (`dune-grid`) Wraps all classes of one grid implementation in new classes that delegate to the existing implementation. This can serve as an ideal base to write a new metagrid.

**CartesianGrid** (`dune-metagrid`) Takes a unstructured quadrilateral or hexahedral grid (e.g. ALUGrid or UGGrid) and replaces the geometry implementation with a strictly Cartesian geometry implementation for performance improvements.

**FilteredGrid** (`dune-metagrid`) Takes any grid and applies a binary filter to the entity sets for codimension 0 provided by the grid view of the given grid.

**SphereGrid** (`dune-metagrid`) A meta grid that provides the correct spherical mapping for geometries and normals for underlying spherical grids.

### 3.1.7. Major developments in the DUNE-GRID interface

Version 1.0 of the DUNE-GRID module was released on December 20, 2007. Since then a number of improvements were introduced, including the following:

- Methods `center()`, `volume()`, `centerUnitOuterNormal()` on `Geometry` and `Intersection` were introduced to support FV methods on polygonal and polyhedral grids.
- `GridFactory` provides an interface for portably creating initial meshes. `GmshReader` uses that to import grids generated with `gmsh`.
- `EntitySeed` replaced `EntityPointer`; this allows the grid to free the memory occupied by the entity and to recreate the entity from the seed.
- The DUNE-GEOMETRY module was introduced as a separate module to provide reference elements, geometry mappings, and quadrature formulas independent of DUNE-GRID.
- The automatic type deduction using `auto` makes using heavily template-based libraries such as DUNE more convenient to use.
- Initially, the DUNE-GRID interface tried to avoid copying objects for performance reasons. Many methods returned const references to internal data and disallowed copying. With copy elision becoming standard, copyable lightweight entities and intersections were introduced. Given an `Entity` with codimension  $c$  to obtain the geometry one would write:

```
using Geometry = typename Grid::template Codim< c >::Geometry;
const Geometry& geo = entity.geometry();
```

whereas in newer DUNE versions one can simply write:

```
const auto geo = entity.geometry();
```

using both, the automatic type deduction and the fact that objects are copyable. A performance comparison discussing the references vs. copyable grid objects can be found in [34,35]. In order to save on memory when storing entities the entity seed concept was introduced.

- Range-based for loops for entities and intersections made iteration over grid entities very convenient. With newer DUNE versions this is simply:

```
for( const auto& element : Dune::elements( gridView ) )
    const auto geo = element.geometry();
```

- `UGGrid` and `ALUGrid` became dune modules instead of being external libraries. This way they can be downloaded and installed like other dune modules.

As the DUNE grid interface has been adapted only slightly, it proved to work for a wide audience. Looking back, the separation of both topology and geometry, and mesh and data were good principles. Further, having entities as a view was a successful choice. Having DUNE split up in modules helped to keep the grid interface separated from other concerns. Some interface changes resulted from extensive advancements of C++11 and the subsequent standards; many are described in [34]. Other changes turned out to make the interface easier to use or to enable different methods on top of the grid interface.

## 3.2. The template library for iterative solvers – DUNE-ISTL

DUNE-ISTL is the linear algebra library of DUNE. It consists of two main components. First it offers a collection of different vector and matrix classes. Second it features different solvers and preconditioners. While the grid interface consists of fine grained interfaces and relies heavily on static polymorphism, the abstraction in DUNE-ISTL uses a combination of dynamic and static polymorphism.

### 3.2.1. Concepts behind the DUNE-ISTL interfaces

A major design decision in DUNE-ISTL was influenced by the observation, that linear solvers can significantly benefit from inherent structure of PDE discretizations. For example a discontinuous Galerkin (DG) discretization leads to a block structured matrix for certain orderings of the unknowns, the same holds for coupled diffusion reaction systems with many components. Making use of this structure often allows to improve convergence of the linear solver, reduce memory consumption and improve memory throughput.

DUNE-ISTL offers different vector and matrix implementations and many of these can be nested. The whole interface is fully templated w.r.t. the underlying scalar data type (double, float, etc.), also called field type. Examples of such nested types are:

`Dune::BlockVector<Dune::FieldVector<std::complex<double>,2>>` a dynamic block vector, which consists of  $N$  blocks of vectors with static size 2 over the field of the complex numbers, i.e. a vector in  $(\mathbb{C}^2)^N$ .

`Dune::BCRSMMatrix<Dune::FieldMatrix<float,27,27>>` a sparse block matrix with dense  $27 \times 27$  matrices as its entries. The dense matrices use a low precision `float` representation. The whole matrix represents a linear mapping  $(\mathbb{R}^{27 \times 27})^{N \times M}$ .

`Dune::BCRSMMatrix<Dune::BCRSMMatrix<double>>` a (sparse) matrix whose entries are sparse matrices with scalar entries. These might for example arise from a Taylor–Hood [36] discretization of the Navier–Stokes equations, where we obtain a  $4 \times 4$  block matrix of sparse matrices.

It is not necessary to use the same field type for matrices and vectors, as the library allows for mixed-precision setups with automatic conversions and determination of the correct return type of numeric operations. In order to allow for an efficient access to individual entries in matrices or vectors, these matrix/vector interfaces are static and make use of compile-time polymorphism, mainly by using duck typing<sup>2</sup> like in the STL.<sup>3</sup>

Solvers on the other hand can be formulated at a very high level of abstraction and are a perfect candidates for dynamic polymorphism. DUNE-ISTL defines abstract interfaces for operators, scalar products, solvers, and preconditioners. A solver, like `LoopSolver`, `CGSolver`, and similar is parameterized with the operator, possibly a preconditioner, and usually the standard euclidean scalar product. The particular implementations, as well as the interface, are strongly typed on the underlying vector types, but it is possible to mix and shuffle different solvers and preconditioners dynamically at runtime. While linear operators are most often stored as matrices, the interface only requires that an operator can be applied to a vector and thus also allows for implementing on-the-fly operators for implicit methods; this drastically reduces the memory consumption and allows for increased arithmetic intensity and thus overcomes performance limitations due to slow memory access. The benefit of on-the-fly operators is highlighted in Section 4.5. It should be noted that many strong preconditioners, like the `Dune::Amg::AMG` have stricter requirements on the operator and need access to the full matrix.

The interface design also offers a simple way to introduce parallel solvers. The parallelization of DUNE-ISTL's data structures and solvers differs significantly from that of other libraries like PETSc. While many libraries rely on a globally consecutive numbering of unknowns, we only require a locally consecutive numbering, which allows for fast access into local containers. Global consistency is then achieved by choosing appropriate parallel operators, preconditioners, and scalar products. Note that the linear solvers do not need any information about parallel data distribution, as they only rely on operator (and preconditioner) applications and scalar product computations, which are hidden under the afore introduced high level abstractions. This allows for a fully transparent switch between sequential and parallel computations.

### 3.2.2. A brief history

The development of DUNE-ISTL began nearly in parallel with DUNE-GRID around the year 2004 and it was included in the 1.0 release of DUNE in 2007. The serial implementation was presented in [37]. One goal was to make DUNE-ISTL usable standalone without DUNE-GRID. Hence a powerful abstraction of parallel iterative solvers based on the concept of parallel index sets was developed as described in [38]. As a first showcase of it an aggregation based parallel algebraic multigrid method for continuous and discontinuous Galerkin discretizations of heterogeneous elliptic problems was added to the library, see [39,40]. It was one of the first solvers scaling to nearly 295,000 processors on a problem with 150 billion unknowns, see [41].

### 3.2.3. Feature overview and recent developments

The afore outlined concepts are implemented using several templated C++ structures. Linear operators, that do not do their computations on the fly, will often use an underlying matrix representation. DUNE-ISTL offers dense matrices, both either of dynamic size or size known already at compile time, as well as several sparse (block) matrices. For a comprehensive list see Table 1. The corresponding vector classes can be found in Table 2.

The most important building blocks of the iterative solvers in DUNE-ISTL are the preconditioners. Together with the scalar product and linear operator they govern whether a solver will be serial/sequential only or capable of running in parallel. To mark sequential solvers the convention is that their name starts with `Seq`. Using the idea of inexact block Jacobi methods or Schwarz type methods, the `BlockPreconditioner` allows to turn any sequential into a parallel preconditioner, given information about the parallel data decomposition. Such so-called hybrid preconditioners are commonly used in parallel (algebraic) multigrid methods, see [42]. A list of preconditioners provided by DUNE-ISTL is in Table 3. The third column indicates whether a preconditioner is sequential (s), parallel (p), or both. For simple preconditioners, that do not need to store a decomposition, a recursion level can be given to the class. Those are marked with “yes” in the last column. The level given to the class indicates where the inversion on the matrix block happens. For a `BCRSMMatrix<FieldMatrix<double,n,m>>` a level of 0 will lead to the inversion of the scalar values inside of the small dense matrices whereas a level of 1 would invert the `FieldMatrix`. The latter variant, which leads to a block preconditioner, is the default. All of the listed preconditioners can be used in the iterative solvers provided by DUNE-ISTL. Table 4 contains a list of these together with the direct solvers. The latter are only wrappers to existing well established libraries.

<sup>2</sup> Used characteristics rather than actual type for algorithms: “if it walks like a duck and quacks like a duck, it is a duck”.

<sup>3</sup> Standard template library.

**Table 1**

Matrix types in DUNE-ISTL, the first three matrix types cannot be used as a block matrices.

Class	Implements
FieldMatrix	(small) matrix with size known at compile time
DiagonalMatrix	Storage optimal representation of a diagonal matrix with size known at compile time
ScaledIdentityMatrix	Storage optimal representation of a scaled identity matrix with size known at compile time
BCRSMatrix	(block) compressed row storage matrix
BDMatrix	(block) diagonal matrix
BTDMatrix	(block) tri-diagonal matrix
Matrix	Generic dynamic dense (block) matrix
MultiTypeBlockMatrix	Dense block matrix with differing block type known at compile time

**Table 2**

Vector types in DUNE-ISTL, the first vector type cannot be used as a block vector.

Class	Implements
FieldVector	(small) vector with size known at compile time
BVector	(block) vector, blocks have same size
VariableBlockVector	Block vector where each block may vary in size
MultiTypeBlockVector	Block vector with differing block type known at compile time

**Table 3**

Preconditioners in DUNE-ISTL.

Class	Implements	s/p	Recursive
Richardson	Richardson (multiply with a scalar)	s	No
SeqJac	Jacobi method	s	Yes
SeqSOR	Successive overrelaxation (SOR)	s	Yes
SeqSSOR	Symmetric SOR	s	Yes
SeqOverlappingSchwarz	Overlapping Schwarz for arbitrary subdomains	s	No
SeqILU	Incomplete LU decomposition	s	No
SeqILDL	Incomplete LDL decomposition	s	No
Pamg::AMG	Algebraic multigrid solver based on aggregation	s/p	No
BlockPreconditioner	Wraps sequential preconditioner to parallel hybrid one	p	No

**Table 4**

Iterative and direct solvers in DUNE-ISTL. Some of these solvers can handle non-static preconditioner, i.e. the preconditioner might change from iteration to iteration.

Class	Implements	Direct
LoopSolver	Simply applies preconditioner in each step	No
GradientSolver	Simple gradient solver	No
CGSolver	Conjugate gradient method	No
BiCGSTABSolver	Biconjugate gradient stabilized method	No
MINRESSolver	Minimal Residual method	No
RestartedGMResSolver	Restarted GMRes solver	No
RestartedFlexibleGMResSolver	Flexible restarted GMRes solver (for non-static preconditioners)	No
GeneralizedPCGSolver	Flexible conjugate gradient solver (for non-static preconditioners)	No
RestartedFCGSolver	Flexible conjugate gradient solver proposed by Notay (for non-static preconditioners)	No
CompleteFCGSolver	Flexible conjugate gradient method reusing old orthogonalizations when restarting	No
SuperLU	Wrapper for SuperLU library	Yes
UMFPack	Wrapper for UMFPack direct solver library	Yes

In recent time DUNE-ISTL has seen a lot of new development. Both, regarding usability, as well as feature-wise. We now briefly discuss some noteworthy improvements.

- i. From the start, DUNE-ISTL was designed to support nested vector and matrix structures. However, the nesting recursion always had to end in `FieldVector` and `FieldMatrix`, respectively. Scalar entries had to be written as vectors of length 1 or matrices of size  $1 \times 1$ . Exploiting modern C++ idioms now allows to support scalar values directly to end the recursion. In other words, it is now possible to write

```
Dune::BCRSMatrix<double>
```

instead of the lengthy

```
Dune::BCRSMatrix<FieldMatrix<double,1,1>>
```

Internally, this is implemented using the `Dune::IsNumber` traits class to recognize scalar types. Note that the indirections needed internally to implement the transparent use of scalar and blocked entries is completely optimized away by the compiler.

- ii. As discussed in the concepts section, operators, solvers, preconditioners, and scalar products offer only coarse grained interfaces. This allows to use dynamic polymorphism. To enable full exchangeability of these classes at runtime we introduced abstract base classes and now store shared pointers to these base classes. With this change it is now possible to configure the solvers at runtime. Additionally, most solvers can now be configured using a `Dune::ParameterTree` object, which holds configuration parameters for the whole program. A convenient solver factory is currently under development, which will complete these latest changes. For example the restarted GMRES solver was constructed as

```
Dune::RestartedGMRESolver<V> solver(op, preconditioner, reduction, restart, maxit, verbose);
```

where `reduction`, `restart`, `maxit`, and `verbose` are just scalar parameters, which the user usually wants to change often to tweak the solvers. Now these parameters can be specified in a section of an INI-style file like:

```
[GMRES]
reduction = 1e-8
maxit = 500
restart = 10
verbose = 0
```

This configuration is parsed into a `ParameterTree` object, which is passed to the constructor:

```
Dune::RestartedGMRESolver<V> solver(op, preconditioner, parametertree);
```

- iii. From a conceptual point of view DUNE-ISTL was designed to support vectors and matrices with varying block structure since the very first release. In practice, it took a very long time to actually fully support such constructs. Only since the new language features of C++11 are available it was possible to implement the classes `MultiTypeBlockVector` and `MultiTypeBlockMatrix` in a fully featured way. These classes implement dense block matrices and vectors with different block types in different entries. The user now can easily define matrix structures like

```
using namespace Dune;
using Row0 = MultiTypeBlockVector<
    Matrix<FieldMatrix<double, 3, 3>>,
    Matrix<FieldMatrix<double, 3, 1>>> >;
using Row1 = MultiTypeBlockVector<
    Matrix<FieldMatrix<double, 1, 3>>,
    Matrix<double>> >;

MultiTypeBlockMatrix<Row0, Row1> A;
```

Such a matrix type would be natural, e.g., for a Taylor–Hood discretization of a three-dimensional Stokes or Navier–Stokes problem, combining a velocity vector field with a scalar pressure.

- iv. With the DUNE 2.6 release an abstraction layer for SIMD-vectorized data types was introduced. This abstraction layer provides functions for transparently handling SIMD data types, as provided by libraries, e.g. `Vc`<sup>4</sup> [43,44] or `VectorClass` [45], and scalar data types, like `double` or `std::complex`. The layer consists of free-standing functions, for example `Simd::lane(int l, VT& v)`, where `v` is of vector-type `VT` and `Simd::lane` gives access to the `l`-th entry of the vector. Operators like `+` or `*` are overloaded and applied component-wise. The result of boolean expressions are also vectorized and return data types with `bool` as scalar type. To handle these values DUNE offers functions like `Simd::cond`, `Simd::allTrue`, or `Simd::anyTrue` for testing them. The `Simd::cond` function has the semantics of the ternary operator, which cannot be overloaded. This operator is necessary, as `if-else` expressions might lead to different branches in the different lanes, which contradicts the principle of vectorization.

Using this abstraction layer it is possible to construct a solver in DUNE-ISTL supporting multiple right hand sides. This is achieved by using the vectorized type as `field_type` in the data structures. For example using the type `Vec4d`, provided by `VectorClass`, the `Vector` type is constructed as:

```
Dune::BlockVector<Dune::FieldVector<Vec4d, 1>>
```

<sup>4</sup> Note that an interface similar to `Vc` is part of the C++ *Parallelism TS 2* standard.

It can be interpreted as a tall-skinny matrix in  $\mathbb{R}^{N \times 4}$ . Using these data types has multiple advantages:

- *Explicit use of vectorization instructions* – Modern CPUs provide SIMD-vectorization instructions, that can perform the same instruction on multiple data simultaneously. It is difficult for the compiler to make use of these instructions automatically. With the above approach we can make explicit use of the vectorization instructions.
- *Better utilization of memory bandwidth* – The application of the operator or the preconditioner is in most cases limited by the available memory bandwidth. This means the runtime of these operations depends on the amount of data that must be transferred from or to the memory. With our vectorization approach the matrix has to be loaded from memory only once for calculating  $k$  matrix–vector products.
- *Reduction of communication overhead* – On distributed systems the cost for sending a message is calculated as  $\alpha D + \beta$ , where  $D$  is the amount of data,  $\alpha$  is the bandwidth, and  $\beta$  is the latency. When using vectorized solvers,  $k$  messages are fused to a single message. Therefore the costs are reduced from  $k(\alpha D + \beta)$  to  $k\alpha D + \beta$ .
- *Block Krylov methods* – Block Krylov methods are Krylov methods for systems with multiple right hand sides. In every iteration the energy error is minimized in all search directions of all lanes. This improves the number of iterations, that are needed to achieve a certain residual reduction.

### 3.3. Finite element spaces on discretization grids

While DUNE focuses on grid-based discretization methods for PDEs, its modular design explicitly avoids any reference to ansatz functions for discretizations in the interfaces of grids and linear algebra of the modules discussed so far. Instead of this the corresponding interfaces and implementations are contained in separate modules. However, the DUNE infrastructure is not limited to finite element discretizations and, for example, a number of applications based on the finite volume method exist, for example DuMux [15] and the Open Porous Media Initiative [14], or higher order finite volume schemes on polyhedral grids [46] as well as a tensor product multigrid approach for grids with high aspect ratios in atmospheric flows [47].

#### 3.3.1. Local functions spaces

The DUNE-LOCALFUNCTIONS core module contains interfaces and implementations for ansatz functions on local reference domains. In terms of finite element discretizations, this corresponds to the finite elements defined on reference geometries. Again following the modular paradigm, this is done independently of any global structures like grids or linear algebra, such that the DUNE-LOCALFUNCTIONS module does not depend on the DUNE-GRID and DUNE-ISTL module. The central concept of the DUNE-LOCALFUNCTIONS module is the `LocalFiniteElement` which is defined along the lines of a *finite element* in terms of [48]. There, a finite element is a triple  $(\mathcal{D}, \Pi, \Sigma)$  of the local domain  $\mathcal{D}$ , a local function space  $\Pi$ , and a finite set of functionals  $\Sigma = \{\sigma_1, \dots, \sigma_n\}$  which induces a basis  $\lambda_1, \dots, \lambda_n$  on the local ansatz space by  $\sigma_i(\lambda_j) = \delta_{ij}$ .

Each `LocalFiniteElement` provides access to its polyhedral domain  $\mathcal{D}$  by exporting a `GeometryType`. The exact geometry of the type is defined in the DUNE-GEOMETRY module. The local basis functions  $\lambda_i$  and the functionals  $\sigma_i$  are provided by each `LocalFiniteElement` by exporting a `LocalBasis` and `LocalInterpolation` object, respectively. Finally, a `LocalFiniteElement` provides a `LocalCoefficients` object. The latter maps each basis function/functional to a triple  $(c, i, j)$  which identifies the basis function as the  $j$ th one tied to the  $i$ th codimension- $c$  face of  $\mathcal{D}$ .

#### 3.3.2. Global functions spaces

In contrast to local shape functions provided by DUNE-LOCALFUNCTIONS, a related infrastructure for global function spaces on a grid view – denoted *global finite element spaces* in the following – is not contained in the DUNE core so far. Instead several concurrent/complementary discretization modules, like DUNE-FEM, DUNE-PDELAB, DUNE-FUFEM, are used to provide their own implementations. To improve interoperability, an interface for global function spaces was developed as a joint effort in the staging module DUNE-FUNCTIONS. This intended to be a common foundation for higher level discretization modules.

It can often be useful to use different bases of the same global finite element space for different applications. For example, a discretization in the space  $P_2$  will often use a classical Lagrange basis of the second order polynomials on the elements, whereas hierarchical error estimators make use of the so called hierarchical  $P_2$  basis. As a consequence the DUNE-FUNCTIONS module does not use the global finite element space itself but its basis as the central abstraction. This is represented by the concept of a `GlobalBasis` in the interface.

Inspired by the DUNE-PDELAB module, DUNE-FUNCTIONS provides a flexible framework for global bases of hierarchically structured finite element product spaces. Such spaces arise, e.g. in non-isothermal phase field models, primal plasticity, mixed formulations of higher order PDEs, multi-physics problems, and many more applications. The central feature is a generic mechanism for the construction of bases for arbitrarily structured product spaces from existing elementary spaces. Within this construction, the global DOF indices can easily be customized according to the matrix/vector data structures suitable for the problem at hand, which may be flat, nested, or even multi-type containers.

In the following we will illustrate this using the  $k$ th order Taylor–Hood space  $P_{k+1}^3 \times P_k$  in  $\mathbb{R}^3$  for  $k \geq 1$  as example. Here  $P_j$  denotes the space of  $j$ th order continuous finite elements. Notice that the Taylor–Hood space provides a natural hierarchical structure: It is the product of the space  $P_{k+1}^3$  for the velocity with the space  $P_k$  for the pressure, where the



former is again the 3-fold product of the space  $P_{k+1}$  for the individual velocity components. Any such product space can be viewed as a tree of spaces, where the root denotes the full space (e.g.  $P_{k+1}^3 \times P_k$ ) inner nodes denote intermediate product spaces (e.g.  $P_{k+1}^3$ ), and the leaf nodes represent *elementary* spaces that are not considered as products themselves (e.g.  $P_{k+1}$  and  $P_k$ ).

The DUNE-FUNCTIONS module on the one hand defines an interface for such nested spaces and on the other hand provides implementations for a set of elementary spaces together with a mechanism for the convenient construction of product spaces. For example, the first order Taylor–Hood space on a given grid view can be constructed using

```
using namespace Dune::Functions;
using namespace Dune::Functions::BasisFactory;
auto basis = makeBasis(gridView,
    composite( power<3>( lagrange<2>()), lagrange<1>()));
```

Here, `lagrange<k>()` creates a descriptor of the Lagrange basis of  $P_k$  which is one of the pre-implemented elementary bases, `power<m>(...)` creates a descriptor of an  $m$ -fold product of a given basis, and `composite(...)` creates a descriptor of the product of an arbitrary family of (possibly different) bases. Finally, `makeBasis(...)` creates the global basis of the desired space on the given grid view. For other applications, `composite` and `power` can be nested in an arbitrary way, and the mechanism can be extended by implementing further elementary spaces providing a certain implementers interface.

The interface of a `GlobalBasis` is split into to several parts. All functionality that is related to the basis as a whole is directly provided by the basis, whereas all functionality that can be localized to grid elements is accessible by a so called `LocalView` obtained using `basis.localView()`. Binding a local view to a grid element using `localView.bind(element)` will then initialize (and possibly pre-compute and cache) the local properties. To avoid costly reallocation of internal data, one can rebind an existing `LocalView` to another element.

Once a `LocalView` is bound, it gives access to all non-vanishing basis functions on the bound-to element. Similar to the global basis, the localized basis forms a local tree which is accessible using `localView.tree()`. Its children can either be directly obtained using the `child(...)` method or traversed in a generic loop. Finally, shape functions can be accessed on each local leaf node in terms of a `LocalFiniteElement` (cf. Section 3.3.1).

The mapping of the shape functions to global indices is done in several stages. First, the shape functions of each leaf node have unique indices within their `LocalFiniteElement`. Next, the per-`LocalFiniteElement` indices of each leaf node can be mapped to per-`LocalView` indices using `leafNode.localIndex(i)`. The resulting local indices enumerate all basis functions on the bound-to element uniquely. Finally, the local per-`LocalView` indices can be mapped to globally unique indices using `localView.index(j)`. To give a full example, the global index of the  $i$ th shape function for the  $d$ th velocity component of the Taylor–Hood basis on a given element can be obtained using

```
using namespace Dune::Indices; // Use compile-time indices
auto localView = basis.localView(); // Create a LocalView
localView.bind(element); // Bind to a grid element
auto&& velocityNode = localView.child(_0, d); // Obtain leaf node of the
// d-th velocity component
auto localIndex = velocityNode.localIndex(i); // Obtain local index of
// i-th shape function
auto globalIndex = localView.index(localIndex); // Obtain global index
```

Here, we made use of the compile time index `Dune::Indices::_0` because direct children in a composite construction may have different types.

While all local indices are flat and zero-based, global indices can in general be multi-indices which allows to efficiently access hierarchically structured containers. The global multi-indices do in general form an index tree. The latter can be explored using `basis.size(prefix)` with a given prefix multi-index. This provides the size of the next digit following the prefix, or, equivalently, the number of direct children of the (possibly interior) node denoted by the prefix. Consistently, `basis.size()` provides the number of entries appearing in the first digit. In case of flat indices, this corresponds to the total number of basis functions.

The way shape functions are associated to indices can be influenced according to the needs of the used discretization, algebraic data structures, and algebraic solvers. In principle an elementary basis provides a pre-defined set of global indices. When defining more complex product space bases using `composite` and `power`, the indices provided by the respective direct children are combined in a customizable way. Possible strategies are, for example, to prepend or append the number of the child to the global index within the child, or to increment the global indices to get consecutive flat indices.

Additionally to the interfaces and implementations of global finite element function space bases, the DUNE-FUNCTIONS module provides utility functions for working with global bases. The most basic utilities are `subspaceBasis(basis, childIndices)`, which constructs a view of only those basis functions corresponding to a certain child in the ansatz tree, `makeDiscreteGlobalBasisFunction<Range>(basis, vector)`, which allows to construct the finite element function (with given range type) obtained by weighting the basis functions with the coefficients stored in a suitable

vector, and `interpolate(basis, vector, function)`, which computes the interpolation of a given function storing the result as coefficient vector with respect to a basis.

The following example interpolates a function into the pressure degrees of freedom only and later constructs the velocity vector field as a function. The latter can e.g. be used to write a subsampled representation in the VTK format.

```
// Interpolate f into vector x
auto f = [] (auto x) { return sin(x[0]) * sin(x[1]); };
interpolate(subspaceBasis(basis, _1), x, f);
// [ Do something to compute x here ]
using Range = FieldVector<double,dim>;
auto velocityFunction =
    makeDiscreteGlobalBasisFunction<Range>(subspaceBasis(basis, _0), x);
```

A detailed description of the `GlobalBasis` interface, the available elementary basis implementations, the mechanism to construct product spaces, the rule-based combination of global indices, and the basis-related utilities can be found in [49]. The type-erasure based polymorphic interface of global functions and localizable grid functions as e.g. implemented by `makeDiscreteGlobalBasisFunction` is described in [50].

### 3.4. Python interfaces for DUNE

Combining easy to use scripting languages with state-of-the-art numerical software has been a continuous effort in scientific computing for a long time. For solution of PDEs the pioneering work of the FEniCS team [5] inspired many others, e.g. [51,52] to also provide Python scripting for high performance PDE solvers usually coded in C++.

Starting with the 2.6 release in 2018, DUNE can also be used from within the Python scripting environment. The DUNE-PYTHON staging module provides i. a general concept for exporting realizations of polymorphic interfaces as used in many DUNE modules and ii. Python bindings for the central interfaces of the DUNE core modules described in this section. These bindings make rapid prototyping of new numerical algorithms easy since they can be implemented and tested within a scripting environment. Our aim was to keep the Python interfaces as close as possible to their C++ counterparts so that translating the resulting Python algorithms to C++ to maximize efficiency of production code is as painless as possible. Bindings are provided using [53].

We start with an example demonstrating these concepts. We revisit the examples given in Section 3.1.1 starting with the construction of a simple Cartesian grid in four space dimensions and the approximation of the integral of a function over that grid. The corresponding Python code is

```
from dune.grid import yaspGrid, cartesianDomain
from math import exp
dim = 4
lower = 4*[0.]
upper = 4*[1.]
nCells = 4*[4]
gv = yaspGrid(constructor=cartesianDomain(lower, upper, nCells))
u = lambda x: exp(x.two_norm)
integral = 0.0
for e in gv.elements:
    integral += u(e.geometry.center)*e.geometry.volume
```

A few changes were made to make the resulting code more Pythonic, i.e., the use of class attributes instead of class methods, but the only major change is that the function returning the grid object in fact returns the leaf grid view and not the hierarchic grid structure. Notice that the life time of this underlying grid is managed automatically by Python's internal reference counting mechanism. It can be obtained using a class attribute, i.e., to refine the grid globally

```
gv.hierarchicalGrid.globalRefine(1);
```

Other interface classes and their realizations have also been exported so that for example the more advanced quadrature rules used in the previous sections can also be used in Python:

```
from dune.geometry import quadratureRule
integral = 0.0
for e in gv.elements:
    geo = e.geometry
    quadrature = quadratureRule(e.type, 5)
    for qp in quadrature:
        integral += u(geo.toGlobal(qp.position)) \
            *geo.integrationElement(qp.position)*qp.weight
```

Again, the changes to the C++ code are mostly cosmetics or due to the restrictions imposed by the Python language.

While staying close to the original C++ interface facilitates rapid prototyping, it also can lead to a significant loss of efficiency. A very high level of efficiency was never a central target during the design of the Python bindings to DUNE – to achieve this, a straightforward mechanism is provided to call DUNE algorithms written in C++. Nevertheless, we made some changes to the interface and added a few extra features to improve the overall efficiency of the code. The two main strategies are to reduce the number of calls from Python to C++ by, for example, not returning single objects for a given index but iterable structures instead. The second strategy is to introduce an extended interface taking a vector of its arguments to allow for vectorization.

Consider, for example the C++ interface methods on the Geometry class `geometry.corners()` and `geometry.corner(i)` which return the number of corners of the elements and their coordinates in physical space, respectively. Using these methods, loops would read as follows:

```
auto center = geometry.corner(0);
for (std::size_t i=1; i<geometry.corners(); ++i)
    center += geometry.corner(i);
center /= geometry.corners();
```

To reduce the number of calls into C++, we decided to slightly change the semantics of method pairs of this type: the plural version now returns an iterable object, while the singular version still exists in its original form. So in Python the above code snippet can be written as follows:

```
corners = geometry.corners
center = corners[0]
for c in corners[1:]:
    center += c
center /= len(corners)
```

As discussed above, quadrature loops are an important ingredient of most grid based numerical schemes. As the code snippet given at the beginning of this section shows, this requires calling methods on the geometry for each point of the quadrature rule which again can lead to a significant performance penalty. To overcome this issue we provide vectorized versions of the methods on the geometry class so that the above example can be more efficiently implemented

```
import numpy
from dune.geometry import quadratureRule
u = lambda x: numpy.exp( numpy.sqrt( sum(x*x) ) )
integral = 0.0
for e in gv.elements:
    hatxs, hatws = quadratureRule(e.type, 5).get()
    weights = hatws * e.geometry.integrationElement(hatxs)
    integral += numpy.sum(u(hatxs) * weights, axis=-1)
```

The following list gives a short overview of changes and extensions we made to the DUNE interface while exporting it to Python:

- Since `global` is a keyword in Python we cannot export the `global` method on the Geometry directly. So we have exported `global` as `toGlobal` and for symmetry reasons `local` as `toLocal`.
- Some methods take compile-time static arguments, e.g., the `codimension` argument for `entity.subEntity<c>(i)`. These had to be turned into dynamic arguments, so in Python the `subEntity` is obtained via `entity.subEntity(i,c)`.
- In many places we replaced methods with properties, i.e., `entity.geometry` instead of `entity.geometry()`.
- Methods returning a `bool` specifying that other interface methods will return valid results are not exported (e.g. `neighbor` on the intersection class). Instead `None` is returned to specify a non valid call (e.g. `to outside`).
- Some of the C++ interfaces contain pairs of methods where the method with the *plural name* returns an integer (the *number of*) and the singular version takes an integer and returns the *ith* element. The plural version was turned to a range-returning method in Python as discussed above.
- In C++, free-standing functions can be found via argument-dependent lookup. As Python does not have such a concept, we converted those free-standing functions to methods or properties. Examples are `elements`, `entities`, `intersections`, or `localFunction`.
- A *grid* in DUNE-PYTHON is always the `LeafGridView` of the hierarchical grid. To work with the actual hierarchy, i.e., to refine the grid, use the `hierarchicalGrid` property. Level grid views can also be obtained from that hierarchical grid.
- In contrast to C++, partitions are exported as objects of their own. The interior partition, for example, can be accessed by

```
partition = grid.interiorPartition
```

The partition, in turn, also exports the method entities and the properties elements, facets, edges, and vertices.

- An MCMGMapper can be constructed using the mapper method on the GridView class passing in the Layout as argument. The mapper class has an additional call method taking an entity, which returns an array with the indices of all DoFs (degrees of freedom) attached to that entity. A list of DoF vectors based on the same mapper can be communicated using methods defined on the mapper itself and without having to define a DataHandle.

A big advantage of using the Python bindings for DUNE is that non performance critical tasks, e.g., pre- and postprocessing can be carried out within Python while the time critical code parts can be easily carried out in C++. To make it easy to call functions written in C++ from within a Python script, DUNE-PYTHON provides a simple mechanism. Let us assume for example that the above quadrature for  $e^{|\mathbf{x}|}$  was implemented in a C++ function `integral` contained in the header file `integral.hh` using the DUNE interface as described in Section 3.1.1:

```
template <class GridView>
double integral(const GridView &gv) {
    auto u = [] (const auto& x){return std::exp(x.two_norm());};
    double integral=0.0;
    for (const auto& e : elements(gv))
        integral += u(e.geometry().center())*e.geometry().volume();
    return integral;
}
```

We can now call this function from within Python using

```
integral = algorithm.run('integral', 'integral.hh', gv)
```

Note that the correct version of the template function `integral` will be exported using the C++ type of the `gv` argument, i.e., `YaspGrid<4>`.

With the mechanism provided in the DUNE-PYTHON module, numerical schemes can first be implemented and tested within Python and can then be translated to C++ to achieve a high level of efficiency. The resulting C++ functions can be easily called from within Python making it straightforward to simply replace parts of the Python code with their C++ counterparts.

In addition to the features described so far, the DUNE-PYTHON module provides general infrastructure for adding bindings to other DUNE modules. Details are given in [54]. We will demonstrate the power of this feature in Section 4.1 where we also use the domain specific language UFL [55] to describe PDE models.

### 3.5. Build system and testing

Starting with the 2.4 release, DUNE has transitioned from its Autotools build system to a new, CMake-based build system. This follows the general trend in the open source software community to use CMake. The framework is split into separate modules; each module is treated as a CMake project in itself, with the build system managing inter-module dependencies and propagation of configuration results. In order to simplify the inter-module management, there is a shell script called `dunecontrol` (part of DUNE-COMMON) that resolves dependencies and controls the build order.

In the CMake implementation of the DUNE build system, special emphasis has been put on testing. Testing has become increasingly important with the development model of the DUNE core modules being heavily based on Continuous Integration. In order to lower the entrance barrier for adding tests to a minimum, a one-line solution in the form of a CMake convenience function `dune_add_test` has been implemented. Further testing infrastructure has been provided in the module DUNE-TESTTOOLS [56], which allows the definition of system tests. These system tests describe samples of framework variability covering both compile-time and run-time variations.

More information on the DUNE CMake build system can be found in the Sphinx-generated documentation, which is available on the DUNE website.<sup>5</sup>

## 4. Selected advanced features with applications

After having discussed the central components of DUNE and their recent changes, we now want to highlight some advanced features. The following examples all showcase features of DUNE extension modules in conjunctions with the core modules.

### 4.1. Grid modification

In this section we discuss two mechanisms of modifying grids within the DUNE framework: *dynamic local grid adaptation* and *moving domains*. In particular, dynamic local grid adaptation is of interest for many scientific and

<sup>5</sup> <https://dune-project.org/buildsystem/>.

engineering applications due to the potential high computational cost savings. However, especially parallel dynamic local grid adaptation is technically challenging and not many PDE frameworks offer a seamless approach. We will demonstrate in this section how the general concepts described in Section 3.1 for grid views and adaptivity provided by the core modules are used to solve PDE problems on grids with dynamic refinement and coarsening. Especially important for these concepts is the separation of topology, geometry, and user data provided by the grid interface.

To support grid modification the DUNE-FEM module provides two specialized GridViews: AdaptiveGridView and GeometryGridView. Both are based on a given grid view, i.e., the LeafGridView, and replace certain aspects of the implementation. In the first case, the index set is replaced by an implementation that provides additional information that can be used to simplify data transfer during grid refinement and coarsening. In the second case the geometry of each element is replaced by a given grid function, e.g., by an analytic function or by some discrete function over the underlying grid view. The advantage of this *meta grid view* approach is that any algorithm based on a DUNE grid view can be used without change while for example the data transfer during grid modification can be transparently handled by specialized algorithms using features of the *meta* grid view.

#### 4.1.1. Dynamic local grid adaptation

A vast number of structured or Cartesian grid managers are available which support adaptive refinement.<sup>6</sup> There exist far fewer open source unstructured grid managers, supporting adaptivity, for example, deal.II [4] which is build on top of p4est [57] for parallel computations, or another recent development the FEMPAR [58] package. Both provide hexahedral grids with non-conforming refinement. Other very capable unstructured grid managers providing tetrahedral elements are, for example, AMDIS [3], FEniCS [5], HiFlow [7], or the “Flexible Distributed Mesh Database (FMDB)” [59], libMesh [60], and others.

As previously described in Section 3.1.4 the DUNE grid interface offers the possibility to dynamically refine and coarsen grids if the underlying grid implementation offers these capabilities. Currently, there are two implementations that support parallel dynamic grid adaptation including load balancing, UGGrid and ALUGrid. AlbertaGrid supports grid adaptation but cannot be used for parallel computations.

A variety of applications make use of the DUNE grid interface for adaptive computations. For example, adaptive discontinuous Galerkin computations of compressible flow, e.g. Euler equations [61] or atmospheric flow [62]. A number of applications focus on hp-adaptive schemes, e.g. for continuous Galerkin approximations of Poisson type problems [63], or discontinuous Galerkin approximations of two-phase flow in porous media [64–67] or conservation laws [68]. Other works consider, for example, the adaptive solution of the Cahn–Larché system using finite elements [69].

In this section we demonstrate the capabilities of the DUNE grid interface and its realizations making use of the Python bindings for the Dune module DUNE-FEM. We show only small parts of the Python code here, the full scripts are part of the tutorial [70].

To this end we solve the Barkley model, which is a system of reaction–diffusion equations modeling excitable media and oscillatory media. The model is often used as a qualitative model in pattern forming systems like the Belousov–Zhabotinsky reaction and other systems that are well described by the interaction of an activator and an inhibitor component [71].

In its simplest form the Barkley model is given by

$$\frac{\partial u}{\partial t} = \frac{1}{\varepsilon} f(u, v) + D \Delta u, \quad \frac{\partial v}{\partial t} = h(u, v),$$

with  $f(u, v) = u(1 - u)\left(u - \frac{v+b}{a}\right)$  and  $h(u, v) = u - v$ . Finally,  $\varepsilon = 0.02$ ,  $a = 0.75$ ,  $b = 0.02$ , and  $D = 0.01$  are chosen according to the web page [http://www.scholarpedia.org/article/Barkley\\_model](http://www.scholarpedia.org/article/Barkley_model) and [71]. To evolve the equations in time, we employ the carefully constructed linear time stepping scheme for this model described in the literature: let  $u^n, v^n$  be approximations of the solution at a time  $t^n$ . To compute approximations  $u^{n+1}, v^{n+1}$  at a later time  $t^{n+1} = t^n + \Delta t$  we replace the nonlinear function  $f(u, v)$  by  $-m(u^n, v^n)u^{n+1} + f_\varepsilon(u^n, v^n)$  where using  $U^*(V) := \frac{V+b}{a}$

$$m(U, V) := \begin{cases} (U - 1)(U - U^*(V)) & U < U^*(V) \\ U(U - U^*(V)) & U \geq U^*(V), \end{cases}$$

$$f_\varepsilon(U, V) := \begin{cases} 0 & U < U^*(V) \\ U(U - U^*(V)) & U \geq U^*(V). \end{cases}$$

Note that  $u, v$  are assumed to be between zero and one so  $m(u^n, v^n) > 0$ . We end up with a linear, positive definite elliptic operator defining the solution  $u^{n+1}$  given  $u^n, v^n$ . In the following we will use a conforming Lagrange approximation with quadratic basis functions. To handle possible nonconforming grids we add interior penalty DG terms as discussed in [63]. The slow reaction  $h(u, v)$  can be solved explicitly leading to a purely algebraic equation for  $v^{n+1}$ . The initial data is piecewise constant chosen in such a way that a spiral wave develops.

<sup>6</sup> See [http://math.boisestate.edu/~calhoun/www\\_personal/research/amr\\_software/](http://math.boisestate.edu/~calhoun/www_personal/research/amr_software/).

The model and initial conditions are easily provided using the Unified Form Language (UFL) [55]. First the problem data needs to be provided

```
dt,t = 0.1,0
spiral_a,spiral_b,spiral_eps,spiral_D = 0.75, 0.02, 0.02, 0.01
def spiral_h(u,v): return u - v
```

and the discrete space and functions constructed

```
space = lagrange( gridView, order=2 )
x = ufl.SpatialCoordinate(space)
iu = lambda s: ufl.conditional(s > 1.25, 1, 0 )
top = ufl.conditional( x[2] > 1.25,1,0)
initial_u = iu(x[1])*top + iu(2.5-x[1])*(1.0 - top)
initial_v = ufl.conditional(x[0]<1.25,0.5,0)
uh = space.interpolate( initial_u, name="u" )
vh = space.interpolate( initial_v, name="v" )
uh_n, vh_n = uh.copy(), vh.copy()
```

Now we use UFL to describe the PDE for the function  $u$  adding DG skeleton terms to take care of possible conforming intersections caused by local grid modification [63]:

```
u, phi = ufl.TrialFunction(space), ufl.TestFunction(space)
hT = ufl.MaxCellEdgeLength(space.cell())
hS = ufl.avg( ufl.MaxFacetEdgeLength(space.cell()) )
hs = ufl.MaxFacetEdgeLength(space.cell())('+')
n = ufl.FacetNormal(space.cell())
penalty = 5 * (order * (order+1)) * spiral_D
ustar = lambda v: (v+spiral_b)/spiral_a
source = lambda u1,u2,u3,v: -1/spiral_eps * u1*(1-u2)*(u3-ustar(v))
# main terms
xForm = inner(D_spiral*grad(u), grad(phi)) * dx
xForm += ufl.conditional(uh_n<ustar(vh_n),
    source(u,uh_n,uh_n,vh_n), source(uh_n,u,uh_n,vh_n)) * phi * dx
# dg terms
xForm -= ( inner( outer(jump(u), n('+')), avg(spiral_D*grad(phi))) \
    inner( avg(spiral_D*grad(u)), outer(jump(phi), n('+')) ) ) * dS
xForm += penalty/hS * inner(jump(u), jump(phi)) * dS
# adding time discretization
form = ( inner(u,phi) - inner(uh_n, phi) ) * dx + dt*xForm
```

For adaptation we use a residual based error estimator derived in [63] for a Discontinuous Galerkin (DG) approximation for the Poisson problem. The error estimator for an element  $E$  at a given time step is given by  $\int_E \eta_E(u^{n+1}) + \frac{1}{2} \int_{\partial E} \eta_{\partial E}(u^{n+1})$  with

$$\eta_E(u_h) = h_E^2 \left( \frac{u_h - u^n}{\Delta t} - f(u_h, v^{n+1}) + \nabla \cdot D \nabla u_h \right)^2,$$

$$\eta_{\partial E}(u_h) = h_E \left[ D \nabla u_h \cdot \nu \right]^2 + \llbracket u_h \rrbracket^2$$

where  $\llbracket \cdot \rrbracket$  is the jump of the given quantity over the boundary of  $E$  and  $\nu$  denotes the outward unit normal. To describe the estimator using UFL we rewrite it in the form

$$R(u_h, \varphi_h) := \int_{\Omega} \eta_h(u_h) \varphi_h + \int_{\partial \Omega} \eta_{\partial \Omega}(u_h) \{ \varphi_h \}$$

such that computing  $R(u^{n+1}, \chi_E) = \int_E \eta_E(u^{n+1}) + \frac{1}{2} \int_{\partial E} \eta_{\partial E}(u^{n+1})$  provides the estimator on each element where  $\chi_E$  is the characteristic function on  $E$ . The characteristic functions are the basis of the finite-volume space provided by DUNE-FEM so that  $R(\cdot, \cdot)$  can be defined using UFL as bilinear form over the solution space of  $u^{n+1}$  and the scalar finite volume space:

```
fvspace = dune.fem.space.finiteVolume(uh.space.grid)
estimate = fvspace.interpolate([0], name="estimate")
chi = ufl.TestFunction(fvspace)
residual = (u-uh_n)/dt - div(spiral_D*grad(u)) + source(u,u,u,vh)
estimator_ufl = hT**2 * residual**2 * chi * dx + \
    hS * inner( jump(spiral_D*grad(u)), n('+'))**2 * avg(chi) * dS + \
    1/hS * jump(u)**2 * avg(chi) * dS
estimator = dune.fem.operator.galerkin(estimator_ufl)
```

Now the grid can be modified according to the estimator within the time loop by i. applying the operator constructed above to the discrete solution  $u^{n+1}$  ii. marking all elements where the error indicator exceeds a given tolerance for






Fig. 4. The evolution of  $u$  for different times using non-conforming grid adaptation in 2d with quadrilaterals.




Fig. 5. The evolution of  $u$  for different times using conforming bisection grid adaptation in 2d.

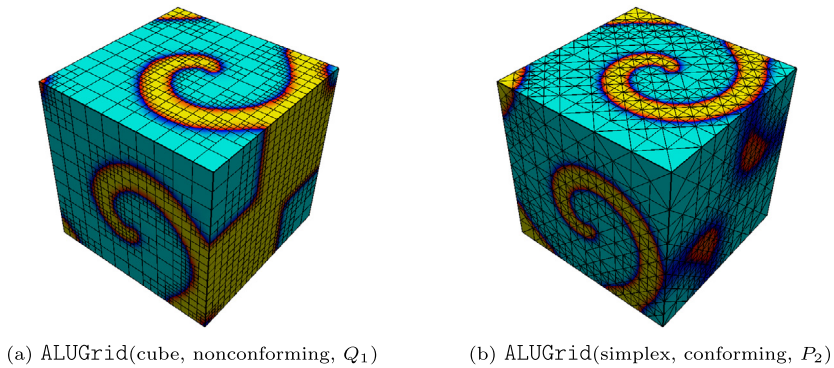


Fig. 6. The solution  $u$  at  $t = 11$  for (a) a non-conforming cube grid in 3d and (b) conforming bisection grid adaptation in 3d. The visual differences between the two solutions especially on the right face are caused by the spiral having rotated a fraction more on the cube grid compared to the spiral on the simplex grid. This is also noticeable on the top left corner. Looking at the front left face it is clear that the difference in angle of the spiral is quite small.

refinement and marking elements for coarsening with an indicator below a given threshold and finally iii. modifying the grid prolonging/restricting the data for  $u_h, v_h$  to the new element:

```
estimator(uh, estimate)
dune.fem.mark(estimate, maxTol, 0.1*maxTol, 0, maxLevel)
dune.fem.adapt([uh, vh])
```

where `maxTol` is some prescribed tolerance (in the following set to  $10^{-4}$ ). Note that only the data for  $u_h, v_h$  is retained during the adaptation process, the underlying storage for other discrete functions used in the simulation is resized as required but the values of the functions are not maintained.

Fig. 4 shows results for different times using a 2d quadrilateral grid with conforming, quadratic Lagrange basis functions. When using a conforming discrete space, the additional terms in the DG formulation vanish whenever basis functions are smooth across element intersection while these terms lead to a stabilization for nonconforming refinement/coarsening with hanging nodes. On a conforming mesh without hanging nodes the residual error estimator coincides with standard results known from the literature [23, and references therein].


Fig. 5 shows results using the quadratic Lagrange basis and a conforming simplicial grid with bisection refinement.

Fig. 6 shows the same example for 3d grids, using a bi-linear Lagrange basis for a non-conforming hexahedral grid in Fig. 6a and using a quadratic Lagrange basis on a conforming simplicial grid with bisection refinement in Fig. 6b.

Details on the available load balancing algorithms and parallel performance studies for the DUNE-ALUGRID package can be found in [24,72], and [73], and for UGGrid in [28].

#### 4.1.2. Moving grids

In this section we touch on another important topic for modern scientific computing: moving domains. Typically this is supported by moving nodes in the computational grid. In DUNE this can be done in a very elegant way. The



**Fig. 7.** Surface evolution towards a sphere using the `ALUGrid<2,3,conforming>` grid implementation. The color coding reflects the magnitude of the surface velocity  $U$ .

presence of an abstract grid interface allows the construction of meta grids where only parts of the grid implementation are re-implemented and, in addition, the original grid implementation stays untouched. Thus meta grids provide a very sophisticated way of adding features to the complete feature stack and keeping the code base modular. In DUNE-GRID one can use the meta grid `GeometryGrid` (see also Section 3.1.6) which allows to move nodes of the grid by providing an overloaded `Geometry` implementation. Another, slightly easier way, is to only overload geometries of grid views which is, for example, done in DUNE-FEM.

Both approaches re-implement the reference geometry mapping. In `GeometryGrid` an external vector of nodes providing the positions of the computational grid is used while for `GeometryGridView` a grid function, i.e., a function which is evaluated on each entity given reference coordinates, is used to provide a mapping for the coordinates. The advantage of both approaches is, that the implementation of the numerical algorithm does not need to change at all. The new grid or grid view follows the same interface as the original implementation. A moving grid can now be realized by modifying this grid function.

To demonstrate this feature of DUNE we solve a mean curvature flow problem which is a specific example of a geometric evolution equation where the evolution is governed by the mean curvature  $H$ . One real-life example of this is in how soap films change over time, although it can also be applied to other problems such as image processing. Assume we are given a reference surface  $\bar{\Gamma}$  such that we can write the evolving surface in the form  $\Gamma_t = X(t, \bar{\Gamma})$ . It is now possible to show that the vector valued function  $X = X(t, \bar{x})$  with  $\bar{x} \in \bar{\Gamma}$  satisfies

$$\frac{\partial}{\partial t} X = -H(X)v(X),$$

where  $H$  is the mean curvature of  $\Gamma_t$  and  $v$  is its outward pointing normal.

We use the following time discrete approximation as suggest in [74]

$$\int_{\Gamma^n} (U^{n+1} - x) \cdot \varphi \, d\sigma + \Delta t \int_{\Gamma^n} \nabla_{\Gamma^n} U^{n+1} : \nabla_{\Gamma^n} \varphi \, d\sigma = 0.$$

Here  $U^n$  parametrizes  $\Gamma^{n+1} \approx \Gamma_{n+1}$  over  $\Gamma^n := \Gamma_t^n$  and  $\Delta t$  is the time step.

In the example used here, the work flow can be set up as follows. First one creates a reference grid and a corresponding quadratic Lagrange finite element space to represent the geometry of the mapped grid.

```
refView = leafGridView("sphere.dgf", dimgrid=2, dimworld=3)
refView.hierarchicalGrid.globalRefine( 2 )
space = solutionSpace(refView, dimRange=refView.dimWorld, order=2)
```

Then, a deformation function is projected onto this Lagrange space

```
x = ufl.SpatialCoordinate(space)
positions = space.interpolate(
    x*(1 + 0.5*sin(2*pi*x[0])*x[1])*cos(pi*x[2])), name="position")
```

Using this grid function, a `GeometryGridView` can be created that uses these new coordinates to represent the grid geometries. This grid view is then used to create the solution space.

```
gridView = geometryGridView(positions)
space = lagrange(gridView, dimRange=gridView.dimWorld, order=2)
```

In each step of the time loop the coordinate positions can be updated, for example, by assigning the values from the computed solution of the mean curvature flow.

```
positions.dofVector.assign(uh.dofVector)
```

In Fig. 7 the evolution of the surface is presented.

Other successful applications of this meta grid concept for moving domains can be found, for example, in [75] where the compressible Navier–Stokes equations are solved in a moving domain and in [29] where free surface shallow water flow is considered.

#### 4.2. Grid coupling and complex domains

In recent years DUNE has gained support for different strategies to handle couplings of PDEs on different subdomains. One can distinguish three different approaches to describe and handle such different domains involved in a multi-physics simulation. As an important side effect, the last approach also provides support for domains with complex shapes.

- i. **Coupling of individual grids:** In the first approach, each subdomain is treated as a separate grid, and meshed individually (see Fig. 8). The challenge is then the construction of the coupling interfaces, i.e., the geometrical relationships at common subdomain boundaries. As it is natural to construct nonconforming interfaces in this way, coupling between the subdomains will in general require some type of weak coupling, like the mortar method [76], penalty methods [77,78], or flux-based coupling [79].
- ii. **Partition of a single grid:** In contrast, one may construct a single host grid that resolves the boundaries of all subdomains. The subdomain meshes are then defined as subsets of elements of the host grid (see Fig. 9). While the construction of the coupling interface is straightforward, generating the initial mesh is an involved process, if the subdomains have complicated shapes. As the coupling interfaces are conforming (as long as the host grid is), it is possible to enforce coupling conditions in strong form.
- iii. **Cut-cell grids:** The third approach is similar to the second one, and again involves a host grid. However, it is more flexible because this time the host grid can be arbitrary, and does not have to resolve the boundaries of the subdomains (see Fig. 10). Instead, subdomain grids are constructed by intersecting the elements with the independent subdomain geometry, typically described as the 0-level set of a given function. This results in so-called cut-cells, which are fragments of host grid elements. Correspondingly, the coupling interfaces are constructed by intersecting host grid elements with the subdomain boundary. It is important to note that the shapes of the cut-cells can be arbitrary and the resulting cut-cell grids are not necessarily shape-regular. As a consequence, it is not possible to employ standard discretization techniques. Instead, a range of different methods like the unfitted discontinuous Galerkin (UDG) method [80,81] and the CutFEM method [82] have been developed for cut-cell grids.

All three concepts for handling complex domains are available as special DUNE modules.

##### 4.2.1. DUNE-GRID-GLUE – Coupling of individual grids

When coupling simulations on separate grids, the main challenge is the construction of coupling operators, as these require detailed neighborhood information between cells in different meshes. The DUNE-GRID-GLUE module [30,31], available from <https://dune-project.org/modules/dune-grid-glue>, provides infrastructure to construct such relations efficiently. Neighborhood relationships are described by the concept of `RemoteIntersections`, which are closely related to the `Intersections` known from the DUNE-GRID module (Section 3.1.2): Both keep references to the two elements that make up the intersection, they store the shape of the intersection in world space, and the local shapes of the intersection when embedded into one or the other of the two elements. However, a `RemoteIntersection` is more general than its DUNE-GRID cousin: For example, the two elements do not have to be part of the same grid object, or even the same grid implementation. Also, there is no requirement for the two elements to have the same dimension. This allows mixed-dimensional couplings like the one in [83].

Constructing the set of remote intersections for a pair of grids first requires the selection of two coupling patches. These are two sets of entities that are known to be involved in the coupling, like a contact boundary, or the overlap between two overlapping grids. Coupling entities can have any codimension. In principle all entities of a given codimension could always be selected as coupling entities, but it is usually easy and more efficient to preselect the relevant ones.


There are several algorithms for constructing the set of remote intersections for a given pair of coupling patches. Assuming that both patches consist of roughly  $N$  coupling entities, the naive algorithm will require  $O(N^2)$  operations. This is too expensive for many situations. Gander and Japhet [84] proposed an advancing front algorithm with expected linear complexity, which, however, slows down considerably when the coupling patches consist of many connected components, or contain too many entities not actually involved in the coupling. Both algorithms are available in DUNE-GRID-GLUE. A third approach using a spatial data structure and a run-time of  $O(N \log N)$  still awaits implementation.

A particular challenge arises in the case of parallel grids, as the partitioning of both grids is also unrelated. DUNE-GRID-GLUE can also compute the set of `RemoteIntersection` in parallel codes, using additional communication. For details on the algorithm and how to handle couplings in the parallel case we refer to [31].


As an example we implement the assembly of mortar mass matrices using DUNE-GRID-GLUE. Let  $\Omega$  be a domain in  $\mathbb{R}^d$ , split into two parts  $\Omega_1, \Omega_2$  by a hypersurface  $\Gamma$ , as in Fig. 8. On  $\Omega$  we consider an elliptic PDE for a scalar function  $u$ , subject to the continuity conditions

$$u_1 = u_2, \quad \langle \nabla u_1, \mathbf{n} \rangle = \langle \nabla u_2, \mathbf{n} \rangle \quad \text{on } \Gamma,$$


where  $u_1$  and  $u_2$  are the restrictions of  $u$  to the subdomains  $\Omega_1$  and  $\Omega_2$ , respectively, and  $\mathbf{n}$  is a unit normal of  $\Gamma$ .



**Fig. 8.** Coupling of two unrelated meshes via a merged grid: Intersecting the coupling patches yields a set of remote intersections, which can be used to evaluate the coupling conditions.



**Fig. 9.** Partition of a given host mesh into subdomains.



**Fig. 10.** Construction of two cut-cell subdomain grids from a Cartesian background grid and a level-set geometry representation: cut-cells are constructed by intersecting a background cell with the zero-iso-surface of the level-set.

For a variationally consistent discretization, the mortar methods discretize the weak form of the continuity condition

$$\int_{\Omega \cap \Gamma} (u_1|_T - u_2|_T) w \, ds = 0, \tag{1}$$

which has to hold for a space of test functions  $w$  defined on the common subdomain boundary. Let  $\Omega_1$  and  $\Omega_2$  be discretized by two independent grids, and let  $\{\theta_i^1\}_{i=1}^{n_1}$  and  $\{\theta_i^2\}_{i=1}^{n_2}$  be nodal basis functions for these grids, respectively. We use the nonzero restrictions of the  $\{\theta_i^1\}$  on  $\Gamma$  to discretize the test function space. Then (1) has the algebraic form

$$M_1 \bar{u}_1 - M_2 \bar{u}_2 = 0,$$

with mortar mass matrices

$$M_1 \in \mathbb{R}^{n_{\Gamma,1} \times n_{\Gamma,1}}, \quad (M_1)_{ij} = \int_{\Omega \cap \Gamma} \theta_i^1 \theta_j^1 \, ds$$

$$M_2 \in \mathbb{R}^{n_{\Gamma,1} \times n_{\Gamma,2}}, \quad (M_2)_{ij} = \int_{\Omega \cap \Gamma} \theta_i^1 \theta_j^2 \, ds.$$

The numbers  $n_{\Gamma,1}$  and  $n_{\Gamma,2}$  denote the numbers of degrees of freedom on the interface  $\Omega \cap \Gamma$ . Assembling these matrices is not easy, because  $M_2$  involves shape functions from two different grids.

For the implementation, assume that the grids on  $\Omega_1$  and  $\Omega_2$  are available as two DUNE grid view objects `gridView1` and `gridView2`, of types `GridView1` and `GridView2`, respectively. The code first constructs the coupling patches, i.e., those parts of the boundaries of  $\Omega_1$ ,  $\Omega_2$  that are on the interface  $\Gamma$ . These are represented in DUNE-GRID-GLUE by objects called `Extractors`. Since we are coupling on the grid boundaries – which have codimension 1 – we need two `Codim1Extractors`:

```

using Extractor1 = GridGlue::Codim1Extractor<GridView1>;
using Extractor2 = GridGlue::Codim1Extractor<GridView2>;

VerticalFacetPredicate<GridView1> facetPredicate1;
VerticalFacetPredicate<GridView2> facetPredicate2;

auto domEx = std::make_shared<Extractor1>(gridView1, facetPredicate1);
auto tarEx = std::make_shared<Extractor2>(gridView2, facetPredicate2);

```

The extractors receive the information on what part of the boundary to use by two predicate objects `facetPredicate1` and `facetPredicate2`. Both implement a method

```

bool contains(const typename GridView::Traits::template Codim<0>::Entity& element,
             unsigned int facet) const

```

that returns true if the `facet`-th face of the element given in `element` is part of the coupling boundary  $\Gamma$ . For the example we use the hyperplane  $\Gamma \subset \mathbb{R}^d$  of all points with first coordinate equal to zero. Then the complete predicate class is

```

template <class GridView>
struct VerticalFacetPredicate
{
    bool operator()(const typename GridView::template Codim<0>::Entity& element,
                  unsigned int facet) const
    {
        const int dim = GridView::dimension;
        const auto& refElement = Dune::ReferenceElements<double, dim>::general(element.type());

        // Return true if all vertices of the facet
        // have the coordinate (numerically) equal to zero
        for (const auto& c : refElement.subEntities(facet, 1, dim))
            if ( std::abs(element.geometry().corner(c)[0]) > 1e-6 )
                return false;

        return true;
    }
};

```

Next, we need to compute the set of remote intersections from the two coupling patches. The different algorithms for this mentioned above are implemented in objects called “mergers”. The most appropriate one for the mortar example is called `ContactMerge`, and it implements the advancing front algorithm of Gander and Japhet.<sup>7</sup> The entire code to construct the remote intersections for the two trace grids at the interface  $\Gamma$  is

```

using GlueType = GridGlue::GridGlue<Extractor1, Extractor2>;

// Backend for the computation of the remote intersections
auto merger = std::make_shared<GridGlue::ContactMerge<dim, double>>();
GlueType gridGlue(domEx, tarEx, merger);

gridGlue.build();

```

The `gridGlue` object is a container for the remote intersections. These can now be used to compute the two mass matrices  $M_1$  and  $M_2$ . Let `mortarMatrix1` and `mortarMatrix2` be two objects of some (deliberately) unspecified matrix type. We assume that both are initialized and all entries are set to zero. The nodal bases  $\{\theta_i^1\}_{i=1}^{n_1}$  and  $\{\theta_i^2\}_{i=1}^{n_2}$  are represented by two DUNE-FUNCTIONS bases. The mortar assembly loop is much like the loop for a regular mass matrix

```

auto nonmortarView = nonmortarBasis.localView();
auto mortarView = mortarBasis.localView();

for (const auto& intersection : intersections(gridGlue))
{
    nonmortarView.bind(intersection.inside());
    mortarView.bind(intersection.outside());

    const auto& nonmortarFiniteElement = nonmortarView.tree().finiteElement();

```

<sup>7</sup> It is called `ContactMerge` because it can also handle the case where the two subdomains are separated by a physical gap, which is common in contact problems.

```

const auto& mortarFiniteElement = mortarView.tree().finiteElement();
const auto& testFiniteElement  = nonmortarView.tree().finiteElement();

// Select a quadrature rule: Use order = 2 just for simplicity
int quadOrder = 2;
const auto& quad = QuadratureRules<double, dim-1>::rule(intersection.type(), quadOrder);

// Loop over all quadrature points
for (const auto& quadPoint : quad)
{
    // compute integration element of overlap
    double integrationElement =
        intersection.geometry().integrationElement(quadPoint.position());

    // quadrature point positions on the reference element
    FieldVector<double, dim> nonmortarQuadPos =
        intersection.geometryInInside().global(quadPoint.position());
    FieldVector<double, dim> mortarQuadPos =
        intersection.geometryInOutside().global(quadPoint.position());

    // evaluate all shapefunctions at the quadrature point
    std::vector<FieldVector<double, 1> >
        nonmortarValues, mortarValues, testValues;

    nonmortarFiniteElement.localBasis()
        .evaluateFunction(nonmortarQuadPos, nonmortarValues);
    mortarFiniteElement.localBasis()
        .evaluateFunction(mortarQuadPos, mortarValues);
    testFiniteElement.localBasis()
        .evaluateFunction(nonmortarQuadPos, testValues);

    // Loop over all shape functions of the test space
    for (size_t i=0; i<testFiniteElement.size(); i++)
    {
        auto testIdx = nonmortarView.index(i);

        // loop over all shape functions on the nonmortar side
        for (size_t j=0; j<nonmortarFiniteElement.localBasis().size(); j++)
        {
            auto nonmortarIdx = nonmortarView.index(j);

            mortarMatrix1[testIdx][nonmortarIdx] += integrationElement *
                quadPoint.weight() * testValues[i] * nonmortarValues[j];
        }

        // loop over all shape functions on the mortar side
        for (size_t j=0; j<mortarFiniteElement.size(); j++)
        {
            auto mortarIdx = mortarView.index(j);

            mortarMatrix2[testIdx][mortarIdx] += integrationElement *
                quadPoint.weight() * testValues[i] * mortarValues[j];
        }
    }
}
}
}
}

```

After these loops, the objects `mortarMatrix1` and `mortarMatrix2` contain the matrices  $M_1$  and  $M_2$ , respectively.

The problem gets more complicated when  $\Gamma$  is not a hyperplane. The approximation of a non-planar  $\Gamma$  by unrelated grids will lead to “holes” at the interface, and the jump  $u_1|_{\Gamma} - u_2|_{\Gamma}$  is not well-defined anymore. This situation is usually dealt with by identifying  $\Gamma_1^h$  and  $\Gamma_2^h$  with a homeomorphism  $\Phi$ , and replacing the second mass matrix by

$$M_2 \in \mathbb{R}^{n_{\Gamma,1} \times n_{\Gamma,2}}, \quad (M_2)_{ij} = \int_{\Gamma_1^h} \theta_i^2(\theta_j^1 \circ \Phi) ds.$$

Only few changes have to be done to the code to implement this. First of all, the vertical predicate class has to be exchanged for something that correctly finds the curved coupling boundaries. Then, setting up extractor and `GridGlue` objects remains unchanged. The extra magic needed to handle the mapping  $\Phi$  is completely concealed in the `ContactMerge` implementation, which does not rely on  $\Gamma_1^h$  and  $\Gamma_2^h$  being identical. Instead, if there is a gap between them, a projection in normal direction is computed automatically and used for  $\Phi$ .



#### 4.2.2. DUNE-MULTIDOMAINGRID — Using element subsets as subdomains

The second approach to the coupling of subdomains is implemented in the DUNE-MULTIDOMAINGRID module, available at <https://dune-project.org/modules/dune-multidomaingrid>. This module allows to structure a given host grid into different subdomains. It is implemented in terms of two cooperating grid implementations MultiDomainGrid and SubDomainGrid: MultiDomainGrid is a meta grid that wraps a given host grid and extends it with an interface for setting up and accessing subdomains. It also stores all data required to manage the subdomains. The individual subdomains are exposed as SubDomainGrid instances, which are lightweight objects that combine the information from the host grid and the associated MultiDomainGrid. SubDomainGrid objects present a subdomain as a regular DUNE grid. A MultiDomainGrid inherits all capabilities of the underlying grid, including features like  $h$ -adaptivity and MPI parallelism. Extensions of the official grid interface allow to obtain the associate entities in the fundamental mesh and the corresponding indices in both grids.

A fair share of the ideas from DUNE-MULTIDOMAINGRID were incorporated in the coupling capabilities of DuMu<sup>x</sup> 3 [15].

#### 4.2.3. DUNE-TPMC — Assembly of cut-cell discretizations

The main challenge for cut-cell approaches is the construction of appropriate quadrature rules to evaluate integrals over the cut-cell and its boundary. We assume that the domain is given implicitly as a discrete level set function  $\Phi_h$ , s.t.  $\Phi(x) < 0$  if  $x \in \Omega^{(i)}$ . The goal is now to compute a polygonal representation of the cut-cell and a decomposition into sub-elements, such that standard quadrature can be applied on each sub-element. This allows to evaluate weak forms on the actual domain, its boundary, and the internal skeleton (when employing DG methods).

The DUNE-TPMC library implements a *topology preserving marching cubes* (TPMC) algorithm [85], assuming that  $\Phi_h$  is given as a piecewise multilinear scalar function (i.e. a  $P^1$  or  $Q^1$  function). The fundamental idea in this case is the same as that of the classical marching cubes algorithm, known from computer graphics. Given the sign of the vertex values the library identifies the topology of the cut-cell. In certain ambiguous cases additional points in the interior of the cell need to be evaluated. From the topological case the actual decomposition is retrieved from a lookup table and mapped according to the real function values.

*Evaluating integrals over a cut-cell domain using DUNE-TPMC.* We look at a simple example to learn how to work with cut-cell domains. As stated, the technical challenge regarding cut-cell methods is the construction of quadrature rules. We consider a circular domain of radius 1 in 2d and compute the area using numerical quadrature. The scalar function  $\Phi : x \in \mathbb{R}^d \rightarrow |x|_2 - 1$  describes the domain boundary as the isosurface  $\Phi = 0$  and the interior as  $\Phi < 0$ .

```
using namespace Dune;

double R = 1.0;
auto phi = [R](FieldVector<ctype,dim> x){ return x.two_norm()-R; };
```

After having set up a grid, we iterate over a given gridview gv, compute the  $Q_1$  representation of  $\Phi$  (or better to say the vertex values in an element e)

```
double volume = 0.0;
std::vector<ctype> values;
for (const auto & e : elements(gv))
{
    const auto & g = e.geometry();
    // fill vertex values
    values.resize(g.corners());
    for (std::size_t i = 0; i < g.corners(); i++)
        values[i] = phi(g.corner(i));
    volume += localVolume(values,g);
}
```

We now compute the local volume by quadrature over the cut-cell  $e|_{\Phi < 0}$ . In order to evaluate the integral we use the TpmcRefinement and construct snippets, for which we can use standard quadrature rules:

```
template<typename Geometry>
double localVolume(std::vector<double> values, const Geometry & g)
{
    double volume = 0.0;
    // calculate tpmc refinement
    TpmcRefinement<ctype,dim> refinement(values);
    // sum over inside domain
    for (const auto & snippet : refinement.volume(tpmc::InteriorDomain))
    {
        // get zero-order quadrature rule
        const QuadratureRule<double,dim>& quad =
            QuadratureRules<double,dim>::rule(snippet.type(),0);
```

```

// sum over snippets
for (size_t i=0; i<quad.size(); i++)
  volume += quad[i].weight()
           * snippet.integrationElement(quad[i].position())
           * g.integrationElement(snippet.global(quad[i].position()));
}
}

```

This gives us a convergent integral, approximating  $\pi$ . Unsurprisingly we obtain an  $O(h^2)$  convergence of the quadrature error, as the geometry is approximated as a polygonal domain.

### 4.3. Non-smooth multigrid

Various interesting PDEs from application fields such as computational mechanics or phase-field modeling can be written as nonsmooth convex minimization problems with certain separability properties. For such problems, the module DUNE-TNNMG offers an implementation of the Truncated Nonsmooth Newton Multigrid (TNNMG) algorithm [86,87].

#### 4.3.1. The truncated nonsmooth Newton multigrid algorithm

TNNMG operates at the algebraic level of PDE problems. Let  $\mathbb{R}^N$  be endowed with a block structure

$$\mathbb{R}^N = \prod_{i=1}^m \mathbb{R}^{N_i},$$

and call  $R_i: \mathbb{R}^N \rightarrow \mathbb{R}^{N_i}$  the canonical restriction to the  $i$ th block. Typically, the factor spaces  $\mathbb{R}^{N_i}$  will have small dimension, but the number of factors  $m$  is expected to be large. A strictly convex and coercive objective functional  $J: \mathbb{R}^N \rightarrow \mathbb{R} \cup \{\infty\}$  is called block-separable if it has the form

$$J(v) = J_0(v) + \sum_{i=1}^m \varphi_i(R_i v), \quad (2)$$

with a convex  $C^2$  functional  $J_0: \mathbb{R}^N \rightarrow \mathbb{R}$ , and convex, proper, lower semi-continuous functionals  $\varphi_i: \mathbb{R}^{N_i} \rightarrow \mathbb{R} \cup \{\infty\}$ .

Given such a functional  $J$ , the TNNMG method alternates between a nonlinear smoothing step and a damped inexact Newton correction. The smoother solves local minimization problems

$$\tilde{v}^k = \underset{\tilde{v} \in \tilde{v}^{k-1} + V_k}{\operatorname{arg\,min}} J(\tilde{v}) \quad \text{for all } k = 1, \dots, m, \quad (3)$$

in the subspaces  $V_k \subset \mathbb{R}^N$  of all vectors that have zero entries everywhere outside of the  $k$ th block. The inexact Newton step typically consists of a single multigrid iteration for the linearized problem, but other choices are possible as well.

For this method global convergence has been shown even when using only inexact local smoothers [87]. In practice it is observed that the method degenerates to a multigrid method after a finite number of steps, and hence multigrid convergence rates are achieved asymptotically [86].

The DUNE-TNNMG module, available from <https://git.imp.fu-berlin.de/agnumpde/dune-tnnmg>, offers an implementation of the TNNMG algorithm in the context of DUNE. The coupling to DUNE is very loose – as TNNMG operates on functionals in  $\mathbb{R}^N$  only, there is no need for it to know about grids, finite element spaces, etc.<sup>8</sup> The DUNE-TNNMG module therefore only depends on DUNE-ISTL and DUNE-SOLVERS.

#### 4.3.2. Numerical example: small-strain primal elastoplasticity

The theory of elastoplasticity describes the behavior of solid objects that can undergo both temporary (elastic) and permanent (plastic) deformation. In its simplest (primal) form, its variables are a vector field  $\mathbf{u}: \Omega \rightarrow \mathbb{R}^d$  of displacements, and a matrix field  $\mathbf{p}: \Omega \rightarrow \operatorname{Sym}_0^{d \times d}$  of plastic strains. These strains are assumed to be symmetric and trace-free [88]. Displacements  $\mathbf{u}$  live in the Sobolev space  $H^1(\Omega, \mathbb{R}^d)$ , and (in theories without strain gradients) plastic strains live in the larger space  $L^2(\Omega, \operatorname{Sym}_0^{d \times d})$ . Therefore, the easiest space discretization employs continuous piecewise linear finite elements for the displacement  $\mathbf{u}$ , and piecewise constant plastic strains  $\mathbf{p}$ .

Implicit time discretization of the quasistatic model leads to a sequence of spatial problems [88,89]. These can be written as minimization problems

$$J(u, p) := \frac{1}{2} (u^T \ p^T) A \begin{pmatrix} u \\ p \end{pmatrix} - b^T \begin{pmatrix} u \\ p \end{pmatrix} + \sigma_c \sum_{i=1}^{n_2} \int_{\Omega} \theta_i(x) dx \cdot \|p_i\|_F, \quad (4)$$

which do not involve a time step size because the model is rate-independent. Here,  $u$  and  $p$  are the finite element coefficients of the displacement and plastic strains, respectively, and  $A$  is a symmetric positive definite matrix. The number

<sup>8</sup> The only exception to this are the multigrid transfer operators. These require access to finite element bases, but are not central to TNNMG.

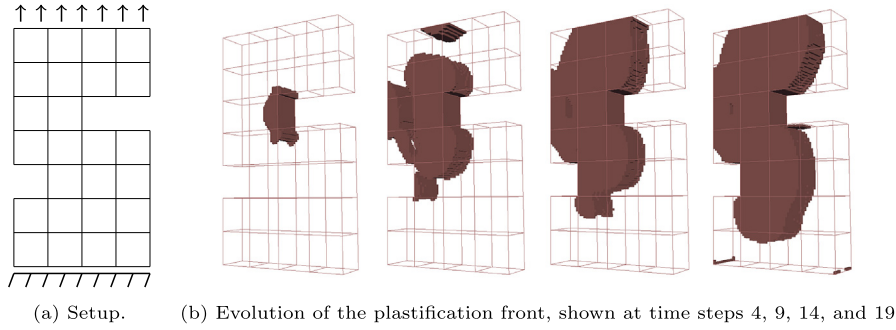


Fig. 11. Setup and numerical results for the three-dimensional benchmark from [91].

$\sigma_c$  is the yield stress, and  $b$  is the load vector. The functions  $\theta_1, \dots, \theta_{n_2}$  are the canonical basis functions of the space of piecewise constant functions, and  $\|\cdot\|_F : \text{Sym}_0^{d \times d} \rightarrow \mathbb{R}$  is the Frobenius norm. In the implementation, trace free symmetric matrices  $p_i \in \text{Sym}_0^{d \times d}$  are represented by vectors of length  $\frac{1}{2}(d+1)d - 1$ .

By comparing (4) to (2), one can see that the increment functional (4) has the required form [89]. By a result of [90], the local nonsmooth minimization problems (3) can be solved exactly.

The implementation used in [89] employs several of the recent hybrid features of DUNE-FUNCTIONS and DUNE-ISTL. The pair of finite element spaces for displacements and plastic strains for a three-dimensional problem forms the tree  $(P_1)^3 \times (P_0)^5$  (where we have identified  $\text{Sym}_0^{3 \times 3}$  with  $\mathbb{R}^5$ ). This tree can be constructed by

```
constexpr size_t nPlasticStrainComponents = dim*(dim+1)/2-1;
auto plasticityBasis = makeBasis(gridView, composite(
    power<dim>(lagrange<1>()), // Deformation basis
    power<nPlasticStrainComponents>(lagrange<0>()) // Plastic strain basis
));
```

The corresponding linear algebra data structures must combine block vectors with block size 3 and block vectors with block size 5. Hence the vector data type definition is

```
using DisplacementVector = BlockVector<FieldVector<double, dim> >;
using PlasticStrainVector = BlockVector<
    FieldVector<double, nPlasticStrainComponents> >;
using Vector = MultiTypeBlockVector<DisplacementVector,
    PlasticStrainVector>;
```

and this is the type used by DUNE-TNNMG for the iterates. The corresponding matrix type combines four BCRSMATRIX objects of different block sizes in a single MultiTypeBlockMatrix, and the multigrid solver operates directly on this type of matrix.

We show a numerical example of the TNNMG solver for a three-dimensional test problem. Note that in this case, the space  $\text{Sym}_0^{d \times d}$  is 5-dimensional, and therefore isomorphic to  $\mathbb{R}^5$ . Let  $\Omega$  be the domain depicted in Fig. 11a, with bounding box  $(0, 4) \times (0, 1) \times (0, 7)$ . We clamp the object at  $\Gamma_D = (0, 4) \times (0, 1) \times \{0\}$ , and apply a time-dependent normal load

$$\langle l(t), \mathbf{u} \rangle = 20t \int_{\Gamma_N} \mathbf{u} \cdot \mathbf{e}_3 \, ds$$

on  $\Gamma_N = (0, 4) \times (0, 1) \times \{7\}$ . The material parameters are taken from [91]. Fig. 11b shows the evolution of the plastification front at several points in time. See [89] for more detailed numerical results and performance measurements.

#### 4.4. Multiscale methods

There has been a tremendous development of numerical multiscale methods in the last two decades including the multiscale finite element method (MsFEM) [92–94], the heterogeneous multiscale method (HMM) [95–97], the variational multiscale method (VMM) [98–100] or the local orthogonal decomposition (LOD) [101–103]. More recently, extensions to parameterized multiscale problems have been presented, such as the localized multiscale reduced basis method (LRBMS) [104–106] or the generalized multiscale finite element method (GMsFEM) [107–109]. In [110,111] we have demonstrated that most of these methods can be cast into a general abstract framework that may then be used for the design of a common implementation framework for multiscale methods, which has been realized in the DUNE-MULTISCALE module [112]. In the following, we concentrate on an efficient parallelized implementation of MsFEM within the DUNE software framework.

4.4.1. Multiscale model problem

As a model problem we consider heterogeneous diffusion. Given a domain  $\Omega \subset \mathbb{R}^n$ ,  $n \in \mathbb{N}_{>0}$  with a polygonal boundary, an elliptic diffusion tensor  $A^\epsilon \in (L^\infty(\Omega))^{n \times n}$  with microscopic features, and an  $f \in L^2(\Omega)$  we define our model problem as

$$\text{find } u^\epsilon \in \dot{H}^1(\Omega) : \int_{\Omega} A^\epsilon \nabla u^\epsilon \cdot \nabla \Phi = \int_{\Omega} f \Phi \quad \forall \Phi \in \dot{H}^1(\Omega) \tag{5}$$

with  $\dot{H}^1(\Omega) := \overline{C^\infty(\Omega)}^{\|\cdot\|_{H^1(\Omega)}}$ .

For the discretization of Eq. (5) we require a regular partition  $\mathcal{T}_H$  of  $\Omega$  with elements  $T$  and a nested refinement  $\mathcal{T}_h$  of  $\mathcal{T}_H$  with elements  $t$  and choose associated piece-wise linear finite element spaces  $U_H := S_0^1(\mathcal{T}_H) \subset U_h := S_0^1(\mathcal{T}_h) \subset \dot{H}^1(\Omega)$ .

We assume that  $U_h$  is sufficiently accurate. By  $A_h^\epsilon$  we denote a suitable piecewise polynomial approximation of  $A^\epsilon$ , and for  $T \in \mathcal{T}_H$ , we call  $U(T)$  an *admissible environment* of  $T$ , if it is connected, if  $T \subset U(T) \subset \Omega$  and if it is the union of elements of  $\mathcal{T}_h$ . Admissible environments will be used for oversampling. In particular  $T$  is an admissible environment of itself.

The MsFEM in Petrov–Galerkin formulation with oversampling is defined in the following. The typical construction of an explicit multiscale finite element basis is already indirectly incorporated in the method. Also note that for  $U(T) = T$  we obtain the MsFEM without oversampling.

Let now  $U_H = \{U(T) \mid T \in \mathcal{T}_H\}$  denote a set of admissible environments of elements of  $\mathcal{T}_H$ . We call  $\mathcal{R}_h^\epsilon(u_H) \in U_h \subset \dot{H}^1(\Omega)$  the MsFEM-approximation of  $u^\epsilon$ , if  $u_H \in U_H$  solves:

$$\sum_{T \in \mathcal{T}_H} \int_T A_h^\epsilon \nabla \mathcal{R}_h^\epsilon(u_H) \cdot \nabla \Phi_H = \int_{\Omega} f \Phi_H \quad \forall \Phi_H \in U_H.$$

For  $\Phi_H \in U_H$ , the reconstruction  $\mathcal{R}_h^\epsilon(\Phi_H)$  is defined by  $\mathcal{R}_h^\epsilon(\Phi_H)|_T := \tilde{Q}_h^\epsilon(\Phi_H) + \Phi_H$ , where  $\tilde{Q}_h^\epsilon(\Phi_H)$  is obtained in the following way: First we solve for  $Q_{h,T}^\epsilon(\Phi_H) \in \dot{U}_h(U(T))$  with

$$\int_{U(T)} A_h^\epsilon (\nabla \Phi_H + \nabla Q_{h,T}^\epsilon(\Phi_H)) \cdot \nabla \phi_h = 0 \quad \forall \phi_h \in \dot{U}_h(U(T)) \tag{6}$$

for all  $T \in \mathcal{T}_H$ , where  $\dot{U}_h(U(T))$  is the underlying fine scale finite element space on  $U(T)$  with zero boundary values on  $\partial U(T)$ . Since we are interested in a globally continuous approximation, i.e.  $\mathcal{R}_h^\epsilon(u_H) \in U_h \subset \dot{H}^1(\Omega)$ , we still need a conforming projection  $P_{H,h}$  which maps the discontinuous parts  $Q_{h,T}^\epsilon(\Phi_H)|_T$  to an element of  $U_h$ . Therefore, if

$$P_{H,h} : \{\phi_h \in L^2(\Omega) \mid \phi_h \in U_h(T) \forall T \in \mathcal{T}_H\} \longrightarrow U_h$$

denotes such a projection, we define

$$\tilde{Q}_h^\epsilon(\Phi_H) := P_{H,h} \left( \sum_{T \in \mathcal{T}_H} \chi_T Q_{h,T}^\epsilon(\Phi_H) \right)$$

with indicator function  $\chi_T$ .

For a more detailed discussion and analysis of this method we refer to [94].

4.4.2. Implementation and parallelization

Our implementation of the general framework for multiscale methods (DUNE-MULTISCALE, [112]) is an effort birthed from the EXA-DUNE project [20,113,114] and is built using the DUNE Generic Discretization Toolbox (DUNE-GDT, [115]) and DUNE-XT [116] as well as the DUNE core modules described in Section 3.

To maximize CPU utilization we employ multi-threading to dynamically load balance work items inside one CPU without expensive memory transfer or cross-node communication. This effectively reduces the communication/overlap region of the coarse grid in a scenario with a fixed number of available cores. Within DUNE we decided to use Intel's Thread Building Blocks (TBB) library as our multithreading abstraction.

Let us now consider an abstract compute cluster that is comprised of a set of processors  $\mathcal{P}$ , where a set of cores  $C_{P_i} = \{C_{P_i}^j\}$  is associated with each  $P_i \in \mathcal{P}$  and a collection of threads  $t_{C_j} = \{t_{C_j}^k\}$ . For simplicity, we assume here that  $j = k$  across  $\mathcal{P}$ .

Since we are interested in globally continuous solutions in  $U_H$ , we require an overlapping distribution  $\mathcal{T}_{H,P_i} \subset \mathcal{T}_H$  where cells can be present on multiple  $P_i$ . Furthermore, we denote by  $\mathcal{I}_i \subset \mathcal{T}_{H,P_i}$  the set of inner elements, if for all  $T_H \in \mathcal{I}_i \Rightarrow T_H \notin \mathcal{I}_j$  for all  $i, j$  with  $i \neq j$ . The first important step in the multiscale algorithm is to solve the cell corrector problems (6) for all  $U(T_H)$ ,  $T_H \in \mathcal{I}_i$ . These are truly locally solvable in the sense of data independence with respect to neighboring coarse cells. We build upon extensions to the DUNE-GRID module made within EXA-DUNE, presented in [117], that allow us to partition a given GridView into connected ranges of cells. The assembler was modified to use TBB such that different threads iterate over different of these partitions in parallel (Fig. 12).

For each  $T_H$  we create a new structured `Dune::YaspGrid` to cover  $U(T_H)$ . Next we need to obtain  $Q_{h,T}^\epsilon(\Phi_H)$  for all  $J$  coarse scale basis functions. After discretization this actually means assembling only one linear system




Fig. 12. Non-overlapping hybrid macro grid distribution of  $\mathcal{T}_{\mathcal{H}}$  for  $\mathcal{P} = P_0, \dots, P_3$  with the hatched area symbolizing sub-distribution over  $t_{\mathcal{G}}$  and zoomed fine scale sub-structure of  $U_{h,T}$  for  $U(T) = T$  (left). Overlapping macro grid distribution of  $\mathcal{T}_{\mathcal{H}}$  for  $\mathcal{P} = P_0, \dots, P_3$  (right).




Fig. 13. Strong scaling absolute runtimes (left) and speedup (right) for the MsFEM from 1792 MPI ranks with roughly 500 coarse cells per rank, up to 14336 ranks with around 60 cells per rank. Performed on a full 512-node island of the Phase 2 partition of the SuperMUC Petascale System in Garching with 28 ranks per node.

matrix and  $J$  different right hand sides. The assembly handled by `Dune::GDT::SystemAssembler`, which is parameterized by an elliptic operator `GDT::Operators::EllipticCG` and corresponding right hand side functionals `GDT::LocalFunctional::Codim0Integral`. The `SystemAssembler` performs loop-fusion by merging cell-local operations of any number of input functors. This allows to perform the complete assembly in one single sweep over the grid, using a configurable amount of thread-parallelism.

Since the cell problems usually only contain up to about 100,000 elements it is especially efficient to factorize the assembled system matrix once and then backsolve for all right hand sides. For this we employ the `UMFPACK`[118] direct solver from the `SuiteSparse` library<sup>9</sup> and its abstraction through `DUNE-ISTL`. Another layer of abstraction on top of that in `DUNE-XT` allows us to switch to an iterative solver at run-time, should we exceed the suitability constraints of the direct solver.

After applying the projections  $P_{H,h}$  to get  $\tilde{Q}_h^\epsilon(\Phi_H)$ , we discretize Eq. (6) which yields a linear system in the standard way. Since this is a system with degrees of freedom (DoF) distributed across all  $P_i$  we need to select an appropriate iterative solver. Here we use the implementation of the bi-conjugate gradient stabilized method (`BICGSTAB`) in `DUNE-ISTL`, preconditioned by an Algebraic Multigrid (`AMG`) solver, see Section 3.2. We note that the application of the linear solver for the coarse system is the only step in our algorithm that requires global communication. This allows the overall algorithm to scale with high parallel efficiency in setups with few coarse grid cells per rank, where a distributed iterative solver cannot be expected to scale with its runtime dominated by communication overhead. We demonstrate this case in Fig. 13.

While the `DUNE-MULTISCALE` module can function as a standalone application to apply the method to a given problem, it is also possible to use it as a library in other modules as well (see for example `DUNE-MLMC`[119]). Run-time parameters like the problem to solve, oversampling size, micro and macro grid resolution, number of threads per process, etc. are read from a structured INI-style file or passed as a `XT::Common::Configuration` object. New problems with associated data functions and computational domains can easily be added by defining them in a new header file. The central library routine to apply the method to a given problem, with nulled solution and prepared grid setup is very concise as it follows the mathematical abstractions discussed above.

```
void Elliptic_MsFEM_Solver::apply(
    DMP::ProblemContainer& problem,
    const CommonTraits::SpaceType& coarse_space,
    std::unique_ptr<LocalSolutionProxy>& solution,
```

<sup>9</sup> <http://faculty.cse.tamu.edu/davis/suitesparse.html>.

```

LocalGridList& localgrid_list) const
{
  CommonTraits::DiscreteFunctionType coarse_msfem_solution(coarse_space, "Coarse Part MsFEM
  Solution");
  LocalProblemSolver(problem, coarse_space, localgrid_list).solve_for_all_cells();
  CoarseScaleOperator elliptic_msfem_op(problem, coarse_space, localgrid_list);
  elliptic_msfem_op.apply_inverse(coarse_msfem_solution);
  //! projection and summation
  identify_fine_scale_part(problem, localgrid_list, coarse_msfem_solution, coarse_space,
    solution);
  solution->add(coarse_msfem_solution);
}

```

#### 4.5. Sum-factorization for high order discretizations to improve node level performance

In this last example we showcase how DUNE is used to develop HPC simulation code for modern hardware architectures. We discuss some prevalent trends in hardware development and how they affect finite element software. Then a matrix-free solution technique for high order discretizations is presented and its node level performance on recent architectures is shown. This work was implemented in DUNE-PDELAB and was originally developed within the Exa-DUNE project. The complexity of the performance engineering efforts have led to a reimplementaion and continued development in DUNE-CODEGEN.

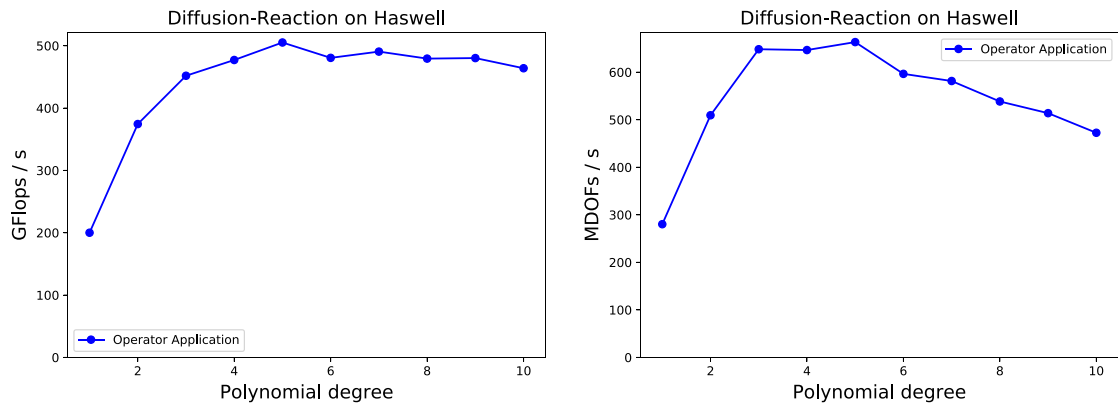
With the end of frequency scaling, performance increases on current hardware rely on an ever-growing amount of parallelism in modern architectures. This includes a drastic increase of CPU floating point performance through instruction level parallelism (SIMD vectorization, superscalar execution, fused multiplication and addition). However, memory bandwidth has not kept up with these gains, severely restricting the performance of established numerical codes and leaving them unable to saturate the floating point hardware. Developers need to both reconsider their choice of algorithms, as well as adapt their implementations in order to overcome this barrier. E.g. in traditional FEM implementations, the system matrix is assembled in memory and the sparse linear system is solved with efficient solvers based on sparse matrices. Optimal complexity solvers scale linearly in the number of unknowns. Despite their optimal complexity, these schemes cannot leverage the capabilities of modern HPC systems as they rely on sparse matrix vector products of the assembled system matrix, which have very low arithmetic intensity and are therefore inherently memory-bound.

One possible approach to leverage the capabilities of current hardware is to directly implement the application of the sparse matrix on a vector. This direct implementation shifts the arithmetic intensity into the compute-bound regime of modern CPUs. Other software projects are pursuing similar ideas for highly performant simulation codes on modern architectures, e.g. libceed [120] and deal.ii [121]. Given an optimal complexity algorithm on suitable discretizations, it is possible to compute the matrix–vector product faster than the entries of an assembled system matrix that can be loaded from main memory. Such optimal complexity algorithms make use of a technique called sum factorization [122] which exploits the tensor product structure of finite element basis functions and quadrature formulas. Given polynomial degree  $k$  and minimal quadrature order, it allows to reduce the computational complexity of one operator application from  $\mathcal{O}(k^{2d})$  to  $\mathcal{O}(k^{d+1})$  by rewriting the evaluation of finite element functions as a  $d$  sequence of tensor contractions. To compute local contributions of the operator it is necessary to have access to the 1d shape functions and quadrature rule that was used in the tensor-product construction of the 2d or 3d variants. Although this optimal complexity algorithm cannot use 3d shape functions, the implementation is still hard-coded, but uses 1d shape functions from DUNE-LOCALFUNCTIONS. By this the implementation can still be fairly generic and easily switch between different polynomial degrees and polynomial representation (e.g. Lagrange- or Legendre-Polynomials).

In order to fully exploit the CPU's floating point capabilities, an implementation needs to maximize its use of SIMD instructions. In our experience, general purpose compilers are not capable to sufficiently autovectorize this type of code, especially as the loop bounds of tensor contractions depend on the polynomial degree  $k$  and are thus not necessarily an integer multiple of the SIMD width. Explicit SIMD vectorization is a challenging task that requires both an algorithmic idea of how to group instructions and possibly rearrange loops as well as a technical realization. In the following we apply a vectorization strategy developed in [123]: Batches of several sum factorization kernels arising from the evaluation of finite element functions and their gradients are parallelized using SIMD instructions. In order to achieve portability between different instruction sets, code is written using a SIMD abstraction layer [45]. This however requires the innermost loops of finite element assembly to be rewritten using SIMD types. With DUNE-PDELAB's abstraction of a local operator, these loops are typically located in user code. This led to the development of DUNE-CODEGEN, which will be further described in Section 5.

Fig. 14 shows node level performance numbers for a Discontinuous Galerkin finite element operator for the diffusion reaction equation on an Intel Haswell node. The measurements use MPI to saturate the node and make extensive use of SIMD instructions which lead to a performance of roughly 40% of the theoretical peak performance of 1.17 TFlops/s ( $10^{12}$  floating point operations per second) on this 32 core node. Discontinuous Galerkin discretizations benefit best from





**Fig. 14.** Performance measurements on an Intel Haswell node for a matrix-free application of a convection–diffusion DG operator on an axis-parallel, structured grid: On the left side, the machine utilization in GFlops/s ( $10^9$  floating point operations per second) is shown. The theoretical peak performance of this Haswell node is 1.17 TFlops/s. On the right hand side, the degree of freedom throughput is measured in degrees of freedom per second.

this compute-bound algorithm, as they allow to minimize memory transfers by omitting the costly setup of element-local data structured, operating directly on suitably blocked global data structures instead. A dedicated assembler for DG operators, `Dune::PDELab::FastDGAssembler`, is now available in DUNE-PDELAB. It does not gather/scatter data from global memory into element-local data structures, but just uses views onto the global data. By this it avoids unnecessary copy operations and index calculations. This assembler is essential to achieve the presented node level performance, but can also be beneficial for traditional DG implementations.

It is worth noting that iterative solvers based on this kind of matrix-free operator evaluation require the design of preconditioners that preserve the low memory bandwidth requirements while ensuring good convergence behavior, as the per-iteration speedup would otherwise be lost to a much higher number of solver iterations. We studied matrix-free preconditioning techniques for Discontinuous Galerkin problems in [124]. This matrix-free solution technology has been used for an advanced application with the Navier–Stokes equations in [125].

## 5. Development trends in DUNE

DUNE, and especially its grid interface, have proven themselves. Use cases range from personal laptops to TOP500 super computers, from mathematical and computer science methodologies to engineering application, and from bachelor thesis to research of Fortune 500 corporations.

As we laid out, DUNE's structure and its interfaces remained stable over the time. The modular structure of DUNE is sometimes criticized, as it might lead to different implementations to solve the same problem. We still believe that the decision was right, as it allows to experiment with new features and make them publicly available to the community without compromising the stability of the core. Other projects like FEniCS have taken similar steps. In our experience a high granularity of the software is useful.

DUNE remains under constant development and new features are added regularly. We briefly want to highlight four topics that are subject of current research and development.

### 5.1. Asynchronous communication

The communication overhead is expected to be an even greater problem, in future HPC systems, as the numbers of processes will increase. Therefore, it is necessary to use asynchronous communication. A first attempt to establish asynchronous communication in DUNE was demonstrated in [22,24].

With the MPI 3.0 standard, an official interface for asynchronous communication was established. Based on this standard, as part of the EXA-DUNE project, we are currently developing high-level abstractions for DUNE for such asynchronous communication, following the *future-promise* concept which is also used in the STL library. An `MPIFuture` object encapsulates the `MPI_Request` as well as the corresponding memory buffer. Furthermore, it provides methods to check for the state of the communication and access the result.

Thereby the fault-tolerance with respect to soft- and hard-faults that occur on remote processes is improved as well. We are following the recommended way of handling failures by throwing exceptions. Unfortunately, this concept integrates poorly with MPI. An approach how to propagate exceptions through the entire system and handle them properly, using the ULFM functionality proposed in [126,127], can be found in [128].

Methods like pipelined CG [129] overlap global communication and operator application to hide communication costs. Such asynchronous solvers will be incorporated in DUNE-ISTL, along with the described software infrastructure.

## 5.2. Thread parallelism

Modern HPC systems exhibit different levels of concurrency. Many numerical codes are now adopting the MPI+X paradigm, meaning that they use internode parallelism via MPI and intranode parallelism, i.e. threads, via some other interface. While early works were based on OpenMP and pthreads, for example in [22], the upcoming interface changes in DUNE will be based on the Intel Thread Building Blocks (TBB) to handle threads. Up to now the core modules do not use multi-threading directly, but the consensus on a single library ensures interoperability among different DUNE extension modules.

In the EXA-DUNE project several numerical components like assemblers or specialized linear solvers have been developed using TBB. As many developments of EXA-DUNE are proof of concepts, these cannot be merged into the core modules immediately, but we plan to port the most promising approaches to mainline DUNE. Noteworthy features include mesh partitioning into entity ranges per thread, as it is used in the MS-FEM code in Section 4.4, the block-SELL-C- $\sigma$  matrix format [130] (an extension of the work of [131]) and a task-based DG-assembler for DUNE-PDELAB.

## 5.3. C++ and Python

Combining easy to use scripting languages with state-of-the-art numerical software has been a continuous effort in scientific computing for a long time. While much of the development of mathematical algorithms still happens in Matlab, there is increasing use of Python for such efforts, also in other scientific disciplines. For solution of PDEs the pioneering work of the FEniCS team [5] inspired many others, e.g. [51,52] to also provide Python scripting for high performance PDE solvers usually coded in C++. As discussed in Section 3.4, DUNE provides Python-bindings for central components like meshes, shape functions, and linear algebra. DUNE-PYTHON also provides infrastructure for exporting static polymorphic interfaces to Python using just in time compilation and without introducing a virtual layer and thus not leading to any performance losses when objects are passed between different C++ components through the Python layer. Bindings are now being added to a number of modules like the DUNE-GRID-GLUE module discussed in Section 4.2.1 and further modules will follow soon.

## 5.4. DSLs and code-generation

Code-generation techniques allow to use scripting languages, while maintaining high efficiency. Using a domain-specific language (DSL), the FEniCS project first introduced a code generator to automatically generate efficient discretization code in Python. The *Unified Form Language* UFL [5,55] is an embedded Python DSL for describing a PDE problem in weak form. UFL is now used by several projects, in particular Firedrake [52]. We also started adopting this input in several places in DUNE. For example, UFL can now be used for the generating model descriptions for DUNE-FEM [70] as demonstrated in Section 4.1. Another effort is the currently developed DUNE-CODEGEN module, which tries to make performance optimization developed in the EXA-DUNE project accessible to the DUNE community.

In Section 4.5 we highlighted how highly tuned matrix-free higher-order kernels can achieve 40% peak performance on modern architectures. While DUNE offers the necessary flexibility, this kind of optimizations is hard to implement for average users. To overcome this issue and improve sustainability, we introduced a code generation toolchain in [132], using UFL as our input language. From this DSL, a header file containing a performance-optimized `LocalOperator` class is generated. The `LocalOperator` interface is DUNE-PDELAB's abstraction for local integration kernels. The design decisions for this code generation toolchain are discussed in detail in [133]. This toolchain achieves near-optimal performance by applying structural transformations to an intermediate representation based on [134]. A search space of SIMD vectorization strategies is explored from within the code generator through an autotuning procedure. This work now leads to the development of DUNE-CODEGEN, which also offers other optimizations, like block-structured meshes, similar to the concepts described in [135], or extruded meshes, like in [136,137]. This is an ongoing effort and is still in early development.

## CRedit authorship contribution statement

**Peter Bastian:** Conceptualization, Writing - original draft, Supervision, Project administration, Funding acquisition. **Markus Blatt:** Conceptualization, Software, Writing - original draft. **Andreas Dedner:** Conceptualization, Methodology, Software, Validation, Investigation, Writing - original draft, Supervision, Visualization. **Nils-Arne Dreier:** Methodology, Software, Investigation, Writing - original draft. **Christian Engwer:** Conceptualization, Software, Investigation, Writing - original draft, Supervision, Project administration, Funding acquisition. **René Fritze:** Investigation, Software, Writing - original draft, Visualization. **Carsten Gräser:** Conceptualization, Software, Investigation, Writing - original draft, Supervision. **Christoph Grüninger:** Software, Writing - review & editing. **Dominic Kempf:** Investigation, Software, Writing - original draft. **Robert Klöforn:** Conceptualization, Methodology, Software, Validation, Investigation, Writing - original draft, Writing - review & editing, Supervision, Project administration. **Mario Ohlberger:** Supervision, Project administration, Funding acquisition, Writing - original draft, Conceptualization. **Oliver Sander:** Conceptualization, Software, Investigation, Writing - original draft, Supervision, Visualization.

## Acknowledgments

We thank the DUNE users and contributors for their continuous support, as only a vivid community unfolds the power of open source. We would like to point out the essential contributions of those DUNE core developers that are not authors of this paper: Ansgar Burchardt, Jorrit Fahlke, Christoph Gersbacher, Steffen Müthing, and Martin Nolte. The file LICENSE.md within every DUNE module attributes the work of numerous more contributors.

Robert Klöforn acknowledges the support of the Research Council of Norway through the INTPART project INSPIRE (274883). Peter Bastian, Nils-Arne Dreier, Christian Engwer, René Fritze, and Mario Ohlberger acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under SPP 1648: Software for Exascale Computing through the project EXA-DUNE – Flexible PDE Solvers, Numerical Methods, and Applications under contract numbers Ba1498/10-2, EN 1042/2-2, and OH98/5-2. Christian Engwer, Dominic Kempf, and Peter Bastian also acknowledge funding through the BMBF project HPC<sup>2</sup>SE under reference number 01H16003A. Nils-Arne Dreier, Christian Engwer, René Fritze, and Mario Ohlberger acknowledge funding by the Deutsche Forschungsgemeinschaft under Germany's Excellence Strategy EXC 2044-390685587, Mathematics Münster: Dynamics – Geometry – Structure.

## References

- [1] P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. Klöforn, M. Ohlberger, O. Sander, A generic grid interface for parallel and adaptive scientific computing. part I: Abstract framework, *Computing* 82 (2–3) (2008) 103–119, <http://dx.doi.org/10.1007/s00607-008-0003-x>.
- [2] P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. Klöforn, R. Kornhuber, M. Ohlberger, O. Sander, A generic grid interface for parallel and adaptive scientific computing. part II: Implementation and tests in DUNE, *Computing* 82 (2–3) (2008) 121–138, <http://dx.doi.org/10.1007/s00607-008-0004-9>.
- [3] S. Vey, A. Voigt, AMDiS: adaptive multidimensional simulations, *Comput. Vis. Sci.* 10 (1) (2007) 57–67, <http://dx.doi.org/10.1007/s00791-006-0048-3>.
- [4] D. Arndt, W. Bangerth, T.C. Clevenger, D. Davydov, M. Fehling, D. Garcia-Sanchez, G. Harper, T. Heister, L. Heltai, M. Kronbichler, R.M. Kynch, M. Maier, J.-P. Pelletier, B. Turcksin, D. Wells, The deal.ii library, version 9.1, *J. Numer. Math.* (2019) <http://dx.doi.org/10.1515/jnma-2019-0064>, online–first.
- [5] A. Logg, K.-A. Mardal, G. Wells, *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*, Springer Publishing Company, Incorporated, 2012.
- [6] F. Hecht, New development in FreeFem++, *J. Numer. Math.* 20 (3–4) (2012) 251–265, URL <https://freefem.org/>.
- [7] S. Gawlok, P. Gerstner, S. Haupt, V. Heuveline, J. Kratzke, P. Lösel, K. Mang, M. Schmidtbreick, N. Schoch, N. Schween, J. Schwegler, C. Song, M. Wlotzka, HiFlow3 – Technical Report on Release 2.0, Preprint Series of the Engineering Mathematics and Computing Lab (EMCL) 0 (06), 2017, <http://dx.doi.org/10.11588/emclpp.2017.06.42879>.
- [8] Q. Liu, Z. Mo, A. Zhang, Z. Yang, JAUMIN: a programming framework for large-scale numerical simulation on unstructured meshes, *CCF Trans. High Perform. Comput.* 1 (1) (2019) 35–48, <http://dx.doi.org/10.1007/s42514-019-00001-z>.
- [9] T. Kolev, V. Dobrev, MFEM: Modular finite element methods library, 2010, <http://dx.doi.org/10.11578/dc.2017.1025.1248>.
- [10] Netgen/NGSolve: high performance multiphysics finite element software, <https://ngsolve.org/>.
- [11] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, D. Karpeyev, D. Kaushik, M.G. Knepley, D.A. May, L.C. McInnes, R.T. Mills, T. Munson, K. Rupp, P. Sanan, B.F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc Web page, 2019, URL <https://www.mcs.anl.gov/petsc>.
- [12] A. Vogel, S. Reiter, M. Rupp, A. Nägel, G. Wittum, UG 4: A novel flexible software system for simulating PDE based models on high performance computers, *Comput. Vis. Sci.* 16 (4) (2013) 165–179, <http://dx.doi.org/10.1007/s00791-014-0232-9>.
- [13] F. Brooks, *The Mythical Man-Month: Essays on Software Engineering*, Addison-Wesley, 1975, URL <https://archive.org/details/mythicalmanmonth00fred>.
- [14] A.F. Rasmussen, T.H. Sandve, K. Bao, A. Lauser, J. Hove, B. Skaflestad, R. Klöforn, M. Blatt, A.B. Rustad, O. Sævareid, et al., The open porous media flow reservoir simulator, *Comput. Math. Appl.* (2020) <http://dx.doi.org/10.1016/j.camwa.2020.05.014>.
- [15] T. Koch, D. Gläser, K. Weishaupt, S. Ackermann, M. Beck, B. Becker, S. Burbulla, H. Class, E. Coltman, S. Emmert, T. Fetzer, C. Grüniger, K. Heck, J. Hommel, T. Kurz, M. Lipp, F. Mohammadi, S. Scherrer, M. Schneider, G. Seitz, L. Stadler, M. Utz, F. Weinhardt, B. Flemisch, *Dumux 3 – an open-source simulator for solving flow and transport problems in porous media with a focus on model coupling*, *Comput. Math. Appl.* (2020) <http://dx.doi.org/10.1016/j.camwa.2020.02.012>.
- [16] S. Götschel, M. Weiser, A. Schiela, Solving optimal control problems with the Kaskade 7 finite element toolbox, in: *Advances in DUNE*, Springer, 2012, pp. 101–112.
- [17] M. Drohmann, B. Haasdonk, S. Kaulmann, M. Ohlberger, A software framework for reduced basis methods using dune-RB and rbmatlab, in: A. Dedner, B. Flemisch, R. Klöforn (Eds.), *Advances in DUNE*, Springer, 2012, pp. 77–88, [http://dx.doi.org/10.1007/978-3-642-28589-9\\_6](http://dx.doi.org/10.1007/978-3-642-28589-9_6).
- [18] C.T. Lee, J.B. Moody, R.E. Amaro, J.A. McCammon, M.J. Holst, The implementation of the colored abstract simplicial complex and its application to mesh generation, *ACM Trans. Math. Software* 45 (3) (2019) <http://dx.doi.org/10.1145/3321515>.
- [19] M. Ainsworth, J. Oden, A posteriori error estimation in finite element analysis, *Comput. Methods Appl. Mech. Engrg.* 142 (1) (1997) 1–88, [http://dx.doi.org/10.1016/S0045-7825\(96\)01107-3](http://dx.doi.org/10.1016/S0045-7825(96)01107-3).
- [20] P. Bastian, C. Engwer, D. Göddeke, O. Iliiev, O. Ippisch, M. Ohlberger, S. Turek, J. Fahlke, S. Kaulmann, S. Müthing, D. Ribbrock, EXA-DUNE: Flexible PDE solvers, numerical methods and applications, in: Lopes, et al. (Eds.), *Euro-Par 2014: Parallel Processing Workshops. Euro-Par 2014 International Workshops, Porto, Portugal, August 25–26, 2014, Revised Selected Papers, Part II*, in: *Lecture Notes in Computer Science*, vol. 8806, Springer, 2014, pp. 530–541, [http://dx.doi.org/10.1007/978-3-319-14313-2\\_45](http://dx.doi.org/10.1007/978-3-319-14313-2_45).
- [21] P. Bastian, C. Engwer, J. Fahlke, M. Geveler, D. Göddeke, O. Iliiev, O. Ippisch, R. Milk, J. Mohring, S. Müthing, M. Ohlberger, D. Ribbrock, S. Turek, *Hardware-based efficiency advances in the EXA-DUNE project*, in: *Software for Exascale Computing - SPPEXA 2013-2015*, in: *Lecture Notes in Computational Science and Engineering*, Springer Verlag, 2016, pp. 3–23.
- [22] R. Klöforn, Efficient matrix-free implementation of discontinuous Galerkin methods for compressible flow problems, in: A. Handlovicova, et al. (Eds.), *Proceedings of the ALGORITHM, 2012*, pp. 11–21, URL <http://www.iam.fmph.uniba.sk/algorithm2012/zbornik/2Kloefkorn.pdf>.
- [23] A. Schmidt, K. Siebert, *Design of Adaptive Finite Element Software – The Finite Element Toolbox ALBERTA*, Springer, 2005, URL <http://www.alberta-fem.de/>.
- [24] M. Alkämper, A. Dedner, R. Klöforn, M. Nolte, The DUNE-ALUGrid module, *Arch. Numer. Softw.* 4 (1) (2016) 1–28, <http://dx.doi.org/10.11588/ans.2016.1.23252>.

- [25] A. Fomins, B. Oswald, Dune-CurvilinearGrid: Parallel dune grid manager for unstructured Tetrahedral curvilinear meshes, 2016, arXiv e-prints [arXiv:1612.02967](https://arxiv.org/abs/1612.02967).
- [26] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities, *Internat. J. Numer. Methods Engrg.* 79 (11) (2009) 1309–1331, <http://dx.doi.org/10.1002/nme.2579>.
- [27] O. Sander, T. Koch, N. Schröder, B. Flemisch, The Dune FoamGrid implementation for surface and network grids, *Arch. Numer. Softw.* 5 (1) (2017) 217–244.
- [28] P. Bastian, K. Birken, K. Johannsen, S. Lang, N. Neuß, H. Rentz-Reichert, C. Wieners, UG – A flexible software toolbox for solving partial differential equations, *Comput. Vis. Sci.* 1 (1) (1997) 27–40, <http://dx.doi.org/10.1007/s007910050003>.
- [29] C. Gersbacher, The Dune-PrismGrid module, in: A. Dedner, B. Flemisch, R. Klöforn (Eds.), *Advances in DUNE*, Berlin, Heidelberg, 2012, pp. 33–44, [http://dx.doi.org/10.1007/978-3-642-28589-9\\_3](http://dx.doi.org/10.1007/978-3-642-28589-9_3).
- [30] P. Bastian, G. Buse, O. Sander, Infrastructure for the coupling of dune grids, in: G. Kreiss, P. Lötstedt, A. Målqvist, M. Neytcheva (Eds.), *Numerical Mathematics and Advanced Applications 2009*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 107–114.
- [31] C. Engwer, S. Müthing, Concepts for flexible parallel multi-domain simulations, in: *Domain Decomposition Methods in Science and Engineering XXII*, Springer, 2016, pp. 187–195.
- [32] S. Müthing, A flexible framework for multi physics and multi domain PDE simulations, (Ph.D. thesis), Universität Stuttgart, 2015, <http://dx.doi.org/10.18419/opus-3620>.
- [33] C. Gräser, O. Sander, The dune-subgrid module and some applications, *Computing* 86 (4) (2009) 269, <http://dx.doi.org/10.1007/s00607-009-0067-2>.
- [34] M. Blatt, A. Burchardt, A. Dedner, C. Engwer, J. Fahlke, B. Flemisch, C. Gersbacher, C. Gräser, F. Gruber, C. Grüniger, D. Kempf, R. Klöforn, T. Malkmus, S. Müthing, M. Nolte, M. Piatkowski, O. Sander, The distributed and unified numerics environment, Version 2.4, *Arch. Numer. Softw.* 4 (100) (2016) 13–29, <http://dx.doi.org/10.11588/ans.2016.100.26526>.
- [35] R. Klöforn, M. Nolte, Performance pitfalls in the dune grid interface, in: A. Dedner, B. Flemisch, R. Klöforn (Eds.), *Advances in DUNE*, Springer Berlin Heidelberg, 2012, pp. 45–58, [http://dx.doi.org/10.1007/978-3-642-28589-9\\_4](http://dx.doi.org/10.1007/978-3-642-28589-9_4).
- [36] H. Elman, D. Silvester, A. Wathen, *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*, second ed., Oxford University Press, 2014.
- [37] M. Blatt, P. Bastian, The iterative solver template library, in: B. Kågström, E. Elmroth, J. Dongarra, J. Waśniewski (Eds.), *Applied Parallel Computing. State of the Art in Scientific Computing*, in: *Lecture Notes in Computer Science*, vol. 4699, Springer, 2007, pp. 666–675.
- [38] M. Blatt, P. Bastian, On the generic parallelisation of iterative solvers for the finite element method, *Int. J. Comput. Sci. Engrg.* 4 (1) (2008) 56–69, <http://dx.doi.org/10.1504/IJCSE.2008.021112>.
- [39] M. Blatt, A parallel algebraic multigrid method for elliptic problems with highly discontinuous coefficients, (Ph.D. thesis), Universität Heidelberg, 2010.
- [40] P. Bastian, M. Blatt, R. Scheichl, Algebraic multigrid for discontinuous Galerkin discretizations of heterogeneous elliptic problems, *Numer. Linear Algebra Appl.* 2 (19) (2012) 367–388.
- [41] O. Ippisch, M. Blatt, Scalability test of  $\mu\phi$  and the parallel algebraic multigrid solver of dune-istl, in: *JÜLich Blue Gene/P Extreme Scaling Workshop*, No. FZJ-JSC-IB-2011-02, JÜLich Supercomputing Centre, 2011, pp. 21–26, doi: 2128/7309.
- [42] U.M. Yang, On the use of relaxation parameters in hybrid smoothers, *Numer. Linear Algebra Appl.* 11 (2–3) (2004) 155–172.
- [43] M. Kretz, Extending C++ for explicit data-parallel programming via SIMD vector types, (Ph.D. thesis), Goethe University Frankfurt am Main, 2015.
- [44] M. Kretz, V. Lindenstruth, Vc: A c++ library for explicit vectorization, *Softw. - Pract. Exp.* 42 (11) (2012) 1409–1430, <http://dx.doi.org/10.1002/spe.1149>.
- [45] A. Fog, C++ vector class library, 2013, URL <http://www.agner.org/optimize/vectorclass.pdf>.
- [46] R. Klöforn, A. Kvaschuk, M. Nolte, Comparison of linear reconstructions for second-order finite volume schemes on polyhedral grids, *Comput. Geosci.* 21 (5) (2017) 909–919, <http://dx.doi.org/10.1007/s10596-017-9658-8>.
- [47] A. Dedner, E. Müller, R. Scheichl, Efficient multigrid preconditioners for atmospheric flow simulations at high aspect ratio, *Internat. J. Numer. Methods Fluids* 80 (1) (2016) 76–102, <http://dx.doi.org/10.1002/flid.4072>.
- [48] P.G. Ciarlet, *The finite element method for elliptic problems*, Vol. 40, SIAM, 2002.
- [49] C. Engwer, C. Gräser, S. Müthing, O. Sander, Function space bases in the dune-functions module, 2018, ArXiv e-prints [arXiv:1806.09545](https://arxiv.org/abs/1806.09545).
- [50] C. Engwer, C. Gräser, S. Müthing, O. Sander, The interface for functions in the dune-functions module, *Arch. Numer. Softw.* 5 (1) (2017) 95–109, <http://dx.doi.org/10.11588/ans.2017.1.27683>, [arXiv:1512.06136](https://arxiv.org/abs/1512.06136).
- [51] L.D. Dalcin, R.R. Paz, P.A. Kler, A. Cosimo, Parallel distributed computing using python, *Adv. Water Resour.* 34 (9) (2011) 1124–1139, <http://dx.doi.org/10.1016/j.advwatres.2011.04.013>.
- [52] F. Rathgeber, D.A. Ham, L. Mitchell, M. Lange, F. Luporini, A.T.T. McRae, G.-T. Bercea, G.R. Markall, P.H.J. Kelly, Firedrake: Automating the finite element method by composing abstractions, *ACM Trans. Math. Software* 43 (3) (2016) 24:1–24:27, <http://dx.doi.org/10.1145/2998441>.
- [53] W. Jakob, J. Rhineland, D. Moldovan, Pybind11 – Seamless operability between C++11 and Python, 2017, URL <https://github.com/pybind/pybind11>.
- [54] A. Dedner, M. Nolte, The Dune Python Module, 2018, ArXiv e-prints [arXiv:1807.05252](https://arxiv.org/abs/1807.05252).
- [55] M.S. Alnæs, A. Logg, K.B. Ølgaard, M.E. Rognes, G.N. Wells, Unified form language: A domain-specific language for weak formulations of partial differential equations, *ACM Trans. Math. Software* 40 (2) (2014) 9:1–9:37, <http://dx.doi.org/10.1145/2566630>.
- [56] D. Kempf, T. Koch, System testing in scientific numerical software frameworks using the example of DUNE, *Arch. Numer. Softw.* 5 (1) (2017) 151–168.
- [57] C. Burstedde, L. Wilcox, O. Ghattas, P4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.* 33 (3) (2011) 1103–1133, <http://dx.doi.org/10.1137/100791634>.
- [58] S. Badia, A.F. Martín, J. Principe, FEMPAR: An object-oriented parallel finite element framework, *Arch. Comput. Methods Eng.* 25 (2) (2018) 195–271, <http://dx.doi.org/10.1007/s11831-017-9244-1>.
- [59] T. Xie, S. Seol, M. Shephard, Generic components for petascale adaptive unstructured mesh-based simulations, *Eng. Comput.* 30 (1) (2014) 79–95, <http://dx.doi.org/10.1007/s00366-012-0288-4>.
- [60] B. Kirk, J. Peterson, R. Stogne, G. Carey, Libmesh: A c++ library for parallel adaptive mesh refinement/coarsening simulations, *Eng. Comput.* 22 (3–4) (2006) 237–254, <http://dx.doi.org/10.1007/s00366-006-0049-3>.
- [61] A. Dedner, R. Klöforn, A generic stabilization approach for higher order discontinuous Galerkin methods for convection dominated problems, *J. Sci. Comput.* 47 (3) (2011) 365–388, <http://dx.doi.org/10.1007/s10915-010-9448-0>.
- [62] D. Schuster, S. Brdar, M. Baldauf, A. Dedner, R. Klöforn, D. Kröner, On discontinuous Galerkin approach for atmospheric flow in the mesoscale with and without moisture, *Meteorol. Z.* 23 (4) (2014) 449–464, <http://dx.doi.org/10.1127/0941-2948/2014/0565>.
- [63] A. Dedner, R. Klöforn, M. Kränkel, Continuous finite-elements on non-conforming grids using discontinuous Galerkin stabilization, in: J. Fuhrmann, et al. (Eds.), *Finite Volumes for Complex Applications VII*, in: *Springer Proceedings in Mathematics & Statistics*, vol. 77, Springer, 2014, pp. 207–215, [http://dx.doi.org/10.1007/978-3-319-05684-5\\_19](http://dx.doi.org/10.1007/978-3-319-05684-5_19).



- [64] A. Dedner, B. Kane, R. Klöforn, M. Nolte, Python framework for hp-adaptive discontinuous Galerkin methods for two-phase flow in porous media, *Appl. Math. Model.* 67 (2019) 179–200, <http://dx.doi.org/10.1016/j.apm.2018.10.013>.
- [65] B. Kane, R. Klöforn, C. Gersbacher, Hp-adaptive discontinuous Galerkin methods for porous media flow, in: C. Cancès, P. Omnes (Eds.), *Finite Volumes for Complex Applications VIII - Hyperbolic, Elliptic and Parabolic Problems: FVCA 8*, Lille, France, June 2017, Springer International Publishing, Cham, 2017, pp. 447–456, [http://dx.doi.org/10.1007/978-3-319-57394-6\\_47](http://dx.doi.org/10.1007/978-3-319-57394-6_47).
- [66] B. Kane, Adaptive higher order discontinuous Galerkin methods for porous-media multi-phase flow with strong heterogeneities, (Dissertation), Universität Stuttgart, 2018, <http://dx.doi.org/10.18419/opus-9863>.
- [67] R. Klöforn, D. Kröner, M. Ohlberger, Parallel adaptive simulation of PEM fuel cells, in: H.-J. Krebs, W. Jäger (Eds.), *Mathematics – Key Technology for the Future*, Springer, 2008, pp. 235–249, [http://dx.doi.org/10.1007/978-3-540-77203-3\\_16](http://dx.doi.org/10.1007/978-3-540-77203-3_16).
- [68] C. Gersbacher, Higher-order discontinuous finite element methods and dynamic model adaptation for hyperbolic systems of conservation laws, (Dissertation), Albert-Ludwigs Universität Freiburg, 2017, <http://dx.doi.org/10.6094/unifr/12838>.
- [69] C. Gräser, R. Kornhuber, U. Sack, Numerical simulation of coarsening in binary solder alloys, *Comput. Mater. Sci.* 93 (2014) 221–233, <http://dx.doi.org/10.1016/j.commatsci.2014.06.010>.
- [70] A. Dedner, R. Klöforn, M. Nolte, Python Bindings for the DUNE-FEM module, Zenodo, 2020, <http://dx.doi.org/10.5281/zenodo.3706994>.
- [71] D. Barkley, A model for fast computer simulation of waves in excitable media, *Physica* 49 (1991) 61–70.
- [72] M. Alkämper, F. Gaspoz, R. Klöforn, A weak compatibility condition for newest vertex bisection in any dimension, *SIAM J. Sci. Comput.* 40 (6) (2018) A3853–A3872, <http://dx.doi.org/10.1137/17M1156137>.
- [73] M. Alkämper, R. Klöforn, Distributed newest vertex bisection, *J. Parallel Distrib. Comput.* 104 (2017) 1–11, <http://dx.doi.org/10.1016/j.jpdc.2016.12.003>.
- [74] K. Deckelnick, G. Dziuk, C.M. Elliott, Computation of geometric partial differential equations and mean curvature flow, *Acta Numer.* 14 (2005) 139–232.
- [75] R. Klöforn, M. Nolte, Solving the reactive compressible Navier-Stokes equations in a moving domain, in: K. Binder, G. Münster, M. Kremer (Eds.), *NIC Symposium 2014 - Proceedings*, Vol. 47, John von Neumann Institute for Computing Jülich, 2014, pp. 353–362, doi: 2128/5919.
- [76] C. Bernardi, Y. Maday, A. Patera, Domain decomposition by the mortar element method, in: H. Kaper, M. Garbey, G. Pieper (Eds.), *Asymptotic and Numerical Methods for Partial Differential Equations with Critical Parameters*, in: NATO ASI Series (Series C: Mathematical and Physical Sciences), vol. 384, Springer, 1993, pp. 269–286, [http://dx.doi.org/10.1007/978-94-011-1810-1\\_17](http://dx.doi.org/10.1007/978-94-011-1810-1_17).
- [77] R. Becker, P. Hansbo, R. Stenberg, A finite element method for domain decomposition with non-matching grids, *ESAIM Math. Model. Numer. Anal.* 37 (2) (2003) 209–225.
- [78] R.D. Lazarov, J.E. Pasciak, J. Schöberl, P.S. Vassilevski, Almost optimal interior penalty discontinuous approximations of symmetric elliptic problems on non-matching grids, *Numer. Math.* 96 (2) (2003) 295–315.
- [79] M.J. Gander, C. Japhet, Y. Maday, F. Nataf, A new cement to glue nonconforming grids with robin interface conditions: the finite element case, in: *Domain Decomposition Methods in Science and Engineering*, Springer, 2005, pp. 259–266.
- [80] P. Bastian, C. Engwer, An unfitted finite element method using discontinuous Galerkin, *Internat. J. Numer. Methods Engrg.* 79 (2009) 1557–1576.
- [81] C. Engwer, F. Heimann, Dune-UDG: A cut-cell framework for unfitted discontinuous Galerkin methods, in: *Advances in DUNE*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 89–100.
- [82] E. Burman, S. Claus, P. Hansbo, M.G. Larson, A. Massing, CutFEM: Discretizing geometry and partial differential equations, *Intern. J. Numer. Methods Engrg.* 104 (2015) 472–501.
- [83] T. Koch, K. Heck, N. Schröder, H. Class, R. Helmig, A new simulation framework for soil-root interaction, evaporation, root growth, and solute transport, *Vadose Zone J.* (2018) <http://dx.doi.org/10.2136/vzj2017.12.0210>.
- [84] M.J. Gander, C. Japhet, Algorithm 932: PANG: software for nonmatching grid projections in 2d and 3d with linear complexity, *ACM Trans. Math. Softw.* 40 (1) (2013) 6.
- [85] C. Engwer, A. Nüßing, Geometric reconstruction of implicitly defined surfaces and domains with topological guarantees, *ACM Trans. Math. Softw.* 44 (2) (2017) 14.
- [86] C. Gräser, R. Kornhuber, Multigrid methods for obstacle problems, *J. Comput. Math.* 27 (1) (2009) 1–44.
- [87] C. Gräser, O. Sander, Truncated nonsmooth Newton multigrid methods for block-separable minimization problems, *IMA J. Numer. Anal.* 39 (1) (2019) 454–481, <http://dx.doi.org/10.1093/imanum/dry073>.
- [88] W. Han, B.D. Reddy, *Plasticity*, second ed., Springer, 2013.
- [89] O. Sander, Solving primal plasticity increment problems in the time of a single predictor–corrector iteration, 2017, ArXiv e-prints [arXiv:1707.03733](https://arxiv.org/abs/1707.03733).
- [90] J. Albery, C. Carstensen, D. Zarrabi, Adaptive numerical analysis in primal elastoplasticity with hardening, *Comput. Methods Appl. Mech. Engrg.* 171 (1999) 175–204.
- [91] P. Neff, A. Sydow, C. Wieners, Numerical approximation of incremental infinitesimal gradient plasticity, *Internat. J. Numer. Methods Engrg.* 77 (2009) 414–436.
- [92] T.Y. Hou, X. Wu, A multiscale finite element method for elliptic problems in composite materials and porous media, *J. Comput. Phys.* 134 (1) (1997) 169–189, <http://dx.doi.org/10.1006/jcph.1997.5682>.
- [93] Y. Efendiev, T.Y. Hou, *Multiscale finite element methods*, in: *Surveys and Tutorials in the Applied Mathematical Sciences*, vol. 4, Springer, New York, 2009, p. xii+234, Theory and applications.
- [94] P. Henning, M. Ohlberger, B. Schweizer, An adaptive multiscale finite element method, *Multiscale Model. Simul.* 12 (3) (2014) 1078–1107, <http://dx.doi.org/10.1137/120886856>.
- [95] W. E, B. Engquist, The heterogeneous multiscale methods, *Commun. Math. Sci.* 1 (1) (2003) 87–132.
- [96] M. Ohlberger, A posteriori error estimates for the heterogeneous multiscale finite element method for elliptic homogenization problems, *Multiscale Model. Simul.* 4 (1) (2005) 88–114, <http://dx.doi.org/10.1137/040605229>.
- [97] A. Abdulle, On a priori error analysis of fully discrete heterogeneous multiscale FEM, *Multiscale Model. Simul.* 4 (2) (2005) 447–459, <http://dx.doi.org/10.1137/040607137>.
- [98] T.J.R. Hughes, Multiscale phenomena: Green's functions, the Dirichlet-to-Neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods, *Comput. Methods Appl. Mech. Engrg.* 127 (1–4) (1995) 387–401.
- [99] T.J.R. Hughes, G.R. Feijóo, L. Mazzei, J.-B. Quinicy, The variational multiscale method - a paradigm for computational mechanics, *Comput. Methods Appl. Mech. Engrg.* 166 (1–2) (1998) 3–24.
- [100] M.G. Larson, A. Malqvist, Adaptive variational multiscale methods based on a posteriori error estimation: duality techniques for elliptic problems, in: *Multiscale Methods in Science and Engineering*, in: *Lect. Notes Comput. Sci. Eng.*, vol. 44, Springer, Berlin, 2005, pp. 181–193.
- [101] A. Malqvist, D. Peterseim, Localization of elliptic multiscale problems, *Math. Comp.* 83 (290) (2014) 2583–2603, <http://dx.doi.org/10.1090/S0025-5718-2014-02868-8>.
- [102] P. Henning, A. Malqvist, D. Peterseim, A localized orthogonal decomposition method for semi-linear elliptic problems, *ESAIM Math. Model. Numer. Anal.* 48 (5) (2014) 1331–1349, <http://dx.doi.org/10.1051/m2an/2013141>.

- [103] C. Engwer, P. Henning, A. Målqvist, D. Peterseim, Efficient implementation of the localized orthogonal decomposition method, *Comput. Methods Appl. Mech. Engrg.* 350 (2019) 123–153.
- [104] F. Albrecht, B. Haasdonk, S. Kaulmann, M. Ohlberger, The localized reduced basis multiscale method, *Proc. ALGORITHMY* (2012) 393–403.
- [105] M. Ohlberger, F. Schindler, Error control for the localized reduced basis multiscale method with adaptive on-line enrichment, *SIAM J. Sci. Comput.* 37 (6) (2015) A2865–A2895, <http://dx.doi.org/10.1137/151003660>.
- [106] M. Ohlberger, S. Rave, F. Schindler, True error control for the localized reduced basis method for parabolic problems, in: *Model Reduction of Parametrized Systems*, in: MS&A. Model. Simul. Appl., vol. 17, Springer, Cham, 2017, pp. 169–182.
- [107] Y. Efendiev, J. Galvis, T.Y. Hou, Generalized multiscale finite element methods (gmsfem), *J. Comput. Phys.* 251 (2013) 116–135, <http://dx.doi.org/10.1016/j.jcp.2013.04.045>.
- [108] E.T. Chung, Y. Efendiev, G. Li, An adaptive GMSFEM for high-contrast flow problems, *J. Comput. Phys.* 273 (2014) 54–76, <http://dx.doi.org/10.1016/j.jcp.2014.05.007>.
- [109] E.T. Chung, Y. Efendiev, W.T. Leung, An adaptive generalized multiscale discontinuous Galerkin method for high-contrast flow problems, *Multiscale Model. Simul.* 16 (3) (2018) 1227–1257, <http://dx.doi.org/10.1137/140986189>.
- [110] M. Ohlberger, Error control based model reduction for multiscale problems, in: *Proceedings of the Conference ALGORITHMY*, 2015, pp. 1–10, URL <http://www.iam.fmph.uniba.sk/amuc/ojs/index.php/algorithmy/article/view/310>.
- [111] P. Henning, M. Ohlberger, On the implementation of a heterogeneous multiscale finite element method for nonlinear elliptic problems, in: *Advances in DUNE*, Springer, Berlin, 2012, pp. 143–155.
- [112] R. Milk, S. Kaulmann, *DUNE multiscale*, 2015, <http://dx.doi.org/10.5281/zenodo.16560>.
- [113] P. Bastian, C. Engwer, J. Fahlke, M. Geveler, D. Göddeke, O. Iliev, O. Ippisch, R. Milk, J. Mohring, S. Müthing, M. Ohlberger, D. Ribbrock, S. Turek, Advances concerning multiscale methods and uncertainty quantification in EXA-DUNE, in: *Software for Exascale Computing - SPPEXA 2013-2015*, in: *Lecture Notes in Computational Science and Engineering*, Springer Verlag, 2016, pp. 25–43.
- [114] P. Bastian, M. Altenbernd, N. Dreier, C. Engwer, J. Fahlke, R. Fritze, M. Geveler, D. Göddeke, O. Iliev, O. Ippisch, J. Mohring, J. Müthing, M. Ohlberger, D. Ribbrock, N. Shegunov, S. Turek, EXA-DUNE – Flexible PDE Solvers, Numerical Methods and Applications, in: Bungartz, H.-J. and Nagel, W.E., *Software for Exascale Computing - SPPEXA 2016-2018*, Springer Lecture Notes in Computational Science and Engineering.
- [115] F. Schindler, R. Milk, *DUNE generic discretization toolbox*, 2015, <http://dx.doi.org/10.5281/zenodo.16563>.
- [116] R. Milk, F. Schindler, T. Leibner, Extending DUNE: The dune-xt modules, *Arch. Numer. Softw.* 5 (1) (2017) 193–216, <http://dx.doi.org/10.11588/ans.2017.1.27720>.
- [117] C. Engwer, J. Fahlke, Scalable hybrid parallelization strategies for the DUNE grid interface, in: *Numerical Mathematics and Advanced Applications: Proceedings of ENUMATH 2013*, in: *Lecture Notes in Computational Science and Engineering*, vol. 103, 2014, pp. 583–590.
- [118] T.A. Davis, Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method, *ACM Trans. Math. Software* 30 (2) (2004) 196–199, <http://dx.doi.org/10.1145/992200.992206>.
- [119] R. Milk, J. Mohring, *DUNE-mimc (SPPEXA AnPleMeet '16)*, 2015, <http://dx.doi.org/10.5281/zenodo.34412>.
- [120] P. Fischer, M. Min, T. Rathnayake, S. Dutta, T. Kolev, V. Dobrev, J.-S. Camier, M. Kronbichler, T. Warburton, K. Swirydowicz, J. Brown, Scalability of high-performance PDE solvers, 2020, [arXiv:2004.06722](https://arxiv.org/abs/2004.06722).
- [121] M. Kronbichler, K. Kormann, Fast matrix-free evaluation of discontinuous Galerkin finite element operators, *ACM Trans. Math. Software* 45 (3) (2019) <http://dx.doi.org/10.1145/3325864>.
- [122] S.A. Orszag, Spectral methods for problems in complex geometries, *J. Comput. Phys.* 37 (1) (1980) 70–92, [http://dx.doi.org/10.1016/0021-9991\(80\)90005-4](http://dx.doi.org/10.1016/0021-9991(80)90005-4).
- [123] S. Müthing, M. Piatkowski, P. Bastian, High-performance implementation of matrix-free high-order discontinuous Galerkin methods, *Int. J. High Perform. Comput. Appl.* (2018) Accepted to [arXiv:1711.10885](https://arxiv.org/abs/1711.10885).
- [124] P. Bastian, E.H. Müller, S. Müthing, M. Piatkowski, Matrix-free multigrid block-preconditioners for higher order discontinuous Galerkin discretisations, *J. Comput. Phys.* (2019).
- [125] M. Piatkowski, S. Müthing, P. Bastian, A stable and high-order accurate discontinuous Galerkin based splitting method for the incompressible Navier–Stokes equations, *J. Comput. Phys.* 356 (2018) 220–239.
- [126] G. Bosilca, A. Bouteiller, A. Guermouche, T. Herault, Y. Robert, P. Sens, J. Dongarra, Failure detection and propagation in HPC systems, in: *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 2016, pp. 312–322.
- [127] W. Bland, A. Bouteiller, T. Herault, J. Hursey, G. Bosilca, J.J. Dongarra, An evaluation of user-level failure mitigation support in MPI, in: *European MPI Users' Group Meeting*, Springer, 2012, pp. 193–203.
- [128] C. Engwer, M. Altenbernd, N.-A. Dreier, D. Göddeke, A high-level c++ approach to manage local errors, asynchrony and faults in an MPI application, in: *2018 26th EuroMicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, IEEE, 2018, pp. 714–721, [arXiv:1804.04481](https://arxiv.org/abs/1804.04481).
- [129] P. Ghysels, W. Vanroose, Hiding global synchronization latency in the preconditioned conjugate gradient algorithm, *Parallel Comput.* 40 (7) (2014) 224–238.
- [130] S. Müthing, D. Ribbrock, D. Göddeke, Integrating multi-threading and accelerators into DUNE-istl, in: *Proceedings of ENUMATH 2013*, Vol. 103, Springer, 2014, pp. 601–609, <http://dx.doi.org/10.1007/978-3-319-10705-959>.
- [131] M. Kreutzer, G. Hager, G. Wellein, H. Fehske, A.R. Bishop, A unified sparse matrix data format for modern processors with wide SIMD units, *SIAM J. Sci. Comput.* 36 (5) (2014) C401–C423, <http://dx.doi.org/10.1137/130930352>.
- [132] D. Kempf, R. Heß, S. Müthing, P. Bastian, Automatic code generation for high-performance discontinuous Galerkin methods on modern architectures, 2018, [ArXiv e-prints arXiv:1812.08075](https://arxiv.org/abs/1812.08075).
- [133] D. Kempf, P. Bastian, An HPC perspective on generative programming, in: *Proceedings of the 14th International Workshop on Software Engineering for Science*, IEEE Press, 2019, pp. 9–16.
- [134] A. Klöckner, Loo.py: Transformation-based code generation for GPUs and CPUs, in: *Proceedings of ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming*, in: *ARRAY'14*, ACM, New York, NY, USA, 2014, pp. 82:82–82:87, <http://dx.doi.org/10.1145/2627373.2627387>.
- [135] B. Bergen, T. Gradl, F. Hulsemann, U. Rude, A massively parallel multigrid method for finite elements, *Comput. Sci. Eng.* 8 (6) (2006) 56–62.
- [136] A.E. MacDonald, J. Middlecoff, T. Henderson, J.-L. Lee, A general method for modeling on irregular grids, *Int. J. High Perform. Comput. Appl.* 25 (4) (2011) 392–403.
- [137] G. Bercea, A.T.T. McRae, D.A. Ham, L. Mitchell, F. Rathgeber, L. Nardi, F. Luperini, P.H.J. Kelly, A structure-exploiting numbering algorithm for finite elements on extruded meshes, and its performance evaluation in firedrake, *Geosci. Model Dev.* 9 (10) (2016) 3803–3815, <http://dx.doi.org/10.5194/gmd-9-3803-2016>.