

Dune-CurvedGrid – A Dune module for surface parametrization

Simon Praetorius¹ and Florian Stenger²

¹Technische Universität Dresden, Institut für Wissenschaftliches Rechnen, D-01062 Dresden, Germany (simon.praetorius@tu-dresden.de).

²Technische Universität Dresden, Institut für Wissenschaftliches Rechnen, D-01062 Dresden, Germany (florian.stenger@tu-dresden.de).

Received: Oct. 7th 2020; **final revision:** Oct. 20th 2021; **published:** Feb. 9th 2022.

Abstract: In this paper we introduce and describe an implementation of curved (surface) geometries within the DUNE framework for grid-based discretizations. Therefore, we employ the abstraction of geometries as element-functions bound to a grid element, and the abstraction of a grid as connectivity of elements together with a grid-function that can be localized to the elements to provide element local parametrizations of the curved surface.

1 Introduction

Numerical computations on curved surfaces are an important tool for studying physical phenomena in thin structures, on curved boundaries of domains, and on interfaces between two bulk regions. Those problems arise, e.g., in fluid dynamics when considering free-surface flows, in interfacial transport problems on biological membranes or fluid droplets, in mathematical geosciences, and in the physics of thin films with vanishing film thickness. All these applications require the discretization of partial differential equations (PDEs) on the embedded surface, involving quantities like curvature, surface measures, normal vectors, and covariant derivatives that need to be available by a numerical method. An overview about some of these applications can be found in [Bothe and Reusken \(2017\)](#); [Nestler et al. \(2017\)](#); [Grande et al. \(2018\)](#); [Jankuhn et al. \(2018\)](#); [Frieden and Schreiner \(2009\)](#).

In discretization methods based on finite elements or finite volumes one is often faced with the problem of representing the curved geometry by numerical grids and discretizing a (geometric) partial differential equation on this approximation of the actual smooth surface. One distinguishes between implicit surface representations using cut cells, level-sets or diffuse interfaces, see [Rätz and Voigt \(2006\)](#); [Lehrenfeld \(2016\)](#); [Olshanskii and Reusken \(2017\)](#); [Burman et al. \(2018\)](#), and explicit surface representations by triangulations and surface finite element or surface finite volume schemes. For the latter, the representation with piecewise flat elements is a widely used lowest order approximation that is easy to construct and to implement with standard software. Unfortunately, it has some drawbacks. The local flat approximation leads to vanishing curvature inside the elements and this can sometimes lead to non-converging numerical schemes, as shown

in [Heine \(2004\)](#) for the discrete mean-curvature vector and Weingarten map, in [Fritz \(2013\)](#) for a finite-element approximation of the Ricci curvature, and in [Hansbo et al. \(2019\)](#); [Hardering and Praetorius \(2021\)](#) for the discretization of a surface vector Laplacian¹. The numerical error involved in a discretization of PDEs on curved surfaces depends on two properties of the discretization, the representation of the objective function and the representation of the geometry. Thus, a higher-order scheme is only possible with also a higher-order description of the surface approximation.

Better approximations of the geometry than piecewise flat are necessary. Those can be found in piecewise polynomial approximations or even exact representations of the surface if an analytical description of the surface's geometry is available. Some numerical libraries for the implementation of partial-differential equation solvers, like, e.g., [Arndt et al. \(2021\)](#); [Anderson et al. \(2021\)](#); [Geuzaine and Remacle \(2009\)](#); [Johnen et al. \(2014\)](#); [Schöberl \(1997, 2014\)](#), allow for higher-order geometry parametrization. Others do not always provide a usable interface for such non-linear geometry transformation or representation of such curved grids and often do not provide the utilities necessary to implement such simulations efficiently.

The present work tries to fill this gap for the numerical library DUNE, [Bastian et al. \(2021\)](#), a framework for the discretization of grid-based numerical problems. This modular library is centered around the abstraction of a grid interface and the coupling to various different grid implementations for structured or unstructured grids, with or without adaptivity, sequential or parallel traversal, for volume and surface grids and additionally allowing to wrap any grid implementation to extend or modify its functionality. This concept of wrapped grids, called *meta-grids* in DUNE, is the basis of our implementation. We provide a wrapper around any grid implementation in DUNE, transforming a (piecewise) flat reference grid into a curved grid using a (non-linear) geometry transformation in the local elements while preserving the grid connectivity and element numbering.

All implementations discussed in this paper can be found in individual DUNE modules in publicly available source code repositories: [Praetorius and Stenger \(2020a\)](#); [Praetorius \(2020, 2019\)](#); [Praetorius and Stenger \(2020c\)](#). The code examples and results in this manuscript are summarized and made available in the repository [Praetorius and Stenger \(2020b\)](#).

1.1 Initial Example

For a motivation of the introduced functionality in DUNE-CURVEDGRID, we consider the polynomial approximation of a spherical surface by local polynomials of some given order.

C++ code

```

1 // 1. Construct a reference grid
2 std::unique_ptr refGrid =
3   Gmsh4Reader< FoamGrid<2,3> >::createGridFromFile("sphere.msh");
4
5 // 2. Define the geometry mapping
6 auto sphere = [](const auto& x) { return x / x.two_norm(); };
7 auto sphereGridFct = analyticDiscreteFunction(sphere, *refGrid, order);
8
9 // 3. Wrap the reference grid to build a curved grid
10 CurvedGrid grid{*refGrid, sphereGridFct};

```

The curved grid is build on top of a piecewise flat reference grid provided by a Gmsh² file and represented by a DUNE-FOAMGRID, [Sander et al. \(2017\)](#). Although this reference grid is created using the Gmsh4Reader which will be discussed in [subsection A.1](#), the grid could also be created by any other method³. The polynomial representation is given by interpolating an

¹In these three examples a stabilization technique or additional higher-order geometric knowledge allows to overcome these flat element limitations

²The Gmsh format is described in [Geuzaine and Remacle \(2009\)](#)

³For example DUNE-GRID provides it's own GmshReader which supports the older file format version 2.

analytic coordinate projection, `sphere`, i.e., a simple normalization of global coordinates, locally into a polynomial space of given polynomial order. This is achieved by the function wrapper `AnalyticDiscreteFunction`, that provides evaluation of values and derivatives of the given function in terms of its local discrete approximation, see [section 5](#).

The library provides grid-functions that can be used to represent geometry mappings. Both, analytical and discrete representations, are implemented and thus even evolving parametrizations as solutions of PDEs or other external descriptions are possible.

Finally, the reference grid together with the coordinate mapping build the curved grid. The `CurvedGrid` fulfills the `DUNE` grid requirements and can be used instead of any other regular `DUNE` grid. It provides the geometric mappings for grid elements and for element intersections, allowing for continuous and discontinuous discretization schemes build on top of this grid.

1.2 Structure of the Paper

In [section 2](#) the mathematical foundation is laid, describing smooth surfaces and their (polynomial) approximation, the grid and its geometry transforms as well as grid-functions. This section introduces the notation and defines the objects and mappings that are implemented in the library. The subsequent [section 3](#) introduces the interface for the geometry classes representing the local mapping of coordinates. This is followed by a section about the class interface for the `CurvedGrid` in [section 4](#) that implements the grid wrapper providing locally defined curved geometries. The description of the actual surface or its approximation requires grid-functions, introduced in [section 5](#), and projection mappings that are shown by examples in [section 6](#). This concludes the implementation aspects of the library. In [section 7](#) numerical validation of the implementation is given by analyzing known error bounds for geometric quantities like distance between surfaces, normal vectors, and mean curvature. This is followed by a selection of numerical examples of finite-element problems on curved surfaces.

2 Parametric Discrete Surfaces

Let $\Gamma \subset \mathbb{R}^{m+1}$ be an oriented, connected, and smooth m -dimensional manifold. Γ can be described in multiple ways, e.g., by a parametrization over a reference domain, by an implicit representation as level-set of a function, or by closest-point projection of coordinates on another manifold in a close neighborhood of Γ . All these descriptions have advantages and disadvantages and are summarized in [Dziuk and Elliott \(2013\)](#). While continuous descriptions allow to extract geometric measures and characteristics of the surface, like its metric or curvature, they are complicated to use in numerical computations. Hence, an approximation, or piecewise representation of the surface for local evaluation of quantities and data is desirable.

2.1 Reference Geometry

Such a representation might be given by a piecewise flat surface Γ_h , topologically equivalent to the smooth surface Γ . This reference surface is composed of finitely many regular and quasi-uniform (flat) m -dimensional elements of diameter h . The collection of these patches, typically simplices or hyper-cubes, is denoted by \mathcal{G}_h and is called the grid representation of Γ_h , with

$$\Gamma_h = \bigcup_{e \in \mathcal{G}_h} e \quad (1)$$

where e denotes an element of the grid. We assume that the patches do not overlap, i.e., for $e_1, e_2 \in \mathcal{G}_h$ we have that $\text{int}(e_1) \cap \text{int}(e_2) = \emptyset$ and if $e_1 \cap e_2 = I \neq \emptyset$ and $\dim(I) = m - 1$, it is called an intersection of e_1 and e_2 and is assumed to be a subset of an $(m - 1)$ -dimensional facet of e_1 and e_2 , respectively.

Each element of the grid $e \in \mathcal{G}_h$ is parametrized over a reference element $\hat{e} \subset \mathbb{R}^m$ by an invertible and differentiable mapping $\mu_e : \hat{e} \rightarrow e$, called the geometry mapping of e . Additionally, we assume that there exists a bijective mapping $\mathbf{X} : \Gamma_h \rightarrow \Gamma$, such that the smooth surface can be represented by the union of (non-overlapping) mapped patches, i.e.,

$$\Gamma = \bigcup_{e \in \mathcal{G}_h} \mathbf{X}(e) = \bigcup_{e \in \mathcal{G}_h} \mathbf{X}(\mu_e(\hat{e})) =: \bigcup_{e \in \mathcal{G}_h} \mathbf{X}_e(\hat{e}). \quad (2)$$

With this property, we call Γ_h the reference manifold or reference domain of Γ and the family $\{\mathbf{X}_e\}_{e \in \mathcal{G}_h}$ its reference parametrization.

2.2 (Higher-order) Approximations of the Manifold

The reference manifold Γ_h from the last section is not used directly for an approximate discretization of functions on Γ , since it does not necessarily approximate the smooth manifold well enough. It just provides a reference domain for the parametrization \mathbf{X} . For numerical computations and discretizations, we need another manifold in the proximate neighborhood of Γ . For a piecewise polynomial surface approximation, we follow the general notation of [Demlow \(2009\)](#).

Therefore, let $\mathbf{X}^k := \mathbb{I}_h^k \mathbf{X} \in \mathbb{P}_k(e)$ be a k th-order polynomial (Lagrange) interpolation of the mapping \mathbf{X} on the element e of the reference manifold, with Lagrange nodes sitting on the smooth surface Γ . $\mathbb{P}_k(e)$ denotes the space of polynomials on e of degree at most k and \mathbb{I} the (componentwise) Lagrange interpolation operator. The interpolation can be expressed in terms of local basis functions in the reference element \hat{e} , by

$$\mathbf{X}^k(\mu_e(\hat{x})) := \mathbf{X}_e^k(\hat{x}) = \sum_{j=1}^{n_k} \xi^j \cdot \phi_j(\hat{x}), \quad \text{for } \hat{x} \in \hat{e}, \quad (3)$$

with $\xi^j := \mathbf{X}_e(\hat{x}^j) \in \Gamma$ for $\{\hat{x}^j\}_{j=1 \dots n_k}$ local Lagrange nodes, the corresponding local Lagrange basis functions $\{\phi_j\}_{j=1 \dots n_k}$, and n_k the number of local basis functions of $\mathbb{P}_k(\hat{e})$. The nodes and basis functions fulfill the nodal interpolation property $\phi_i(\hat{x}^j) = \delta_{ij}$. Thus, the mapped Lagrange nodes ξ^j sit on the smooth surface Γ , see [Figure 1](#) for an illustration.

Then, the k th-order approximation Γ_h^k of the smooth surface Γ can be obtained by the union of (non-overlapping) mapped elements, mapped by \mathbf{X}^k :

$$\Gamma_h^k = \bigcup_{e \in \mathcal{G}_h} \mathbf{X}^k(e) = \bigcup_{e \in \mathcal{G}_h} \mathbf{X}^k(\mu_e(\hat{e})) = \bigcup_{e \in \mathcal{G}_h} \mathbf{X}_e^k(\hat{e}). \quad (4)$$

In case of $k = 1$ we speak of a piecewise flat or polyhedral surface grid Γ_h^1 . Since the mappings \mathbf{X}_e^k are locally smooth, we obtain a piecewise differentiable manifold.

2.3 The Grid, Entities, and Intersections

The reference manifold Γ_h is composed of elements e that build the grid \mathcal{G}_h , or vice versa, the grid defines the manifold. The (higher-order) mapped elements and the elements mapped to the smooth manifold also form grids, namely

$$\mathcal{G}_h^k := \{\mathbf{X}^k(e)\}_{e \in \mathcal{G}_h} \quad \text{and} \quad \mathcal{G} := \{\mathbf{X}(e)\}_{e \in \mathcal{G}_h}. \quad (5)$$

Since, we assume that all manifolds Γ , Γ_h , and Γ_h^k have the same dimension, also the grids are composed of elements e of that same dimension m . We speak of a conforming grid if the surface is continuous and the non-empty intersection of each two elements e_1 and e_2 is an l -dimensional facet of both elements, called sub-entity $s \leq e_1$ and $s \leq e_2$. We say s has co-dimension $c = m - l$. In conforming grids the sub-entities of co-dimension one are the intersections I of two elements.




Figure 1: Lagrange parametrization of order $k = 2$ with Lagrange nodes on vertices and edges.

Corresponding to the reference element \hat{e} there is a reference element $\hat{s} \subset \mathbb{R}^l$ of the sub-entity s of e . The relation between the geometries of \hat{s} and \hat{e} is given by the invertible and differentiable mapping $\eta_{s,e} : \hat{s} \rightarrow \hat{e}$, called the local-geometry mapping between sub-entity and element. With this, the parametrization of the real sub-entity s is given by the chain of η and μ , i.e., $s = \mu_e(\eta_{s,e}(\hat{s}))$, see Figure 2 for an illustration. If this parametrization is equivalent to the direct mapping from the reference element \hat{s} to s , i.e., if it holds $\mu_e \circ \eta_{s,e} = \mu_s$ for all sub-entities s of e , we call the grid *twist-free*, see Dedner and Nolte (2012). This property is assumed to hold at least for intersections, in the following. Examples of twist-free grids are Dune::OneDGrid, Dune::YaspGrid, and Dune::ALUGrid.




Figure 2: Coordinate mappings μ and η between reference element, flat element, and curved element. An additional mapping $X_{s,e}$ can be defined by chaining of η , μ and X , i.e. $X_{s,e} = X \circ \mu_e \circ \eta_{s,e}$.

2.4 Grid-functions and Element-functions

When discussing the mapping X or its polynomial variants X^k , we often define it by its local representation X_e and X_e^k , respectively, with e the element it is defined on. Those functions, defined via their local element variant, are called *grid-functions* in the following and are directly connected to their local variant, called *element-functions*. The evaluation of the grid-function in global coordinates $x \in \mathbb{R}^n$ might be an expensive operation, whereas the evaluation in the corresponding local coordinate \hat{x} can be easily defined. An example is the evaluation of a discrete function by linear combination of evaluations of local basis functions inside the grid elements.

In general, we denote by $f_e : \hat{e} \rightarrow \mathcal{R}$, $f_e = f \circ \mu_e$ the element-function bound to the element e . It is associated to the grid-function $f : e \rightarrow \mathcal{R}$ with range \mathcal{R} . If f is smooth or at least differentiable

inside the element e , we denote by $Df : e \rightarrow L(e, \mathcal{R})$ its derivative as linear mapping, often represented as a matrix $\mathbb{R}^{\dim(\mathcal{R}) \times n}$. The corresponding localized derivative $(Df)_e$ is then given by

$$(Df)_e : \hat{e} \rightarrow L(e, \mathcal{R}), \quad (Df)_e = Df \circ \mu_e, \quad (6)$$

with the same linear mapping in the range as the global Df . This notation follows Engwer et al. (2017). If instead just a local Jacobian $D(f_e) : \hat{e} \rightarrow L(\hat{e}, \mathcal{R})$ is available or requested, we have the relation $D(f_e) = (Df)_e \cdot D(\mu_e)$.

Note that the geometry mapping \mathbf{X} is a grid-function with element-function \mathbf{X}_e . Also, the parametrized geometry mappings \mathbf{X}^k and \mathbf{X}_e^k are grid-function and element-function, respectively. These mappings are differentiable, by differentiating their local basis functions ϕ_j , i.e.,

$$D(\mathbf{X}_e^k) = \sum_j^{n_k} \xi^j \otimes D(\phi_j). \quad (7)$$

In the next section we want to introduce the implementation of the element geometry mappings \mathbf{X}_e and \mathbf{X}_e^k and then in the subsequent section a wrapper to transform the reference grid \mathcal{G}_h into \mathcal{G} or \mathcal{G}_h^k using the geometry mappings and its global variants \mathbf{X} and \mathbf{X}^k from above.

3 CURVEDGEOMETRY and the DUNE geometry interface

A *Geometry* is a mapping from local coordinates in \mathbb{R}^l to global coordinates in \mathbb{R}^n , where the local coordinates are in the coordinate system of an entity e which this geometry belongs to. The entity can be an element of the grid, or a sub-entity of any co-dimension. This flexibility requires the geometry parametrization to be evaluable in different coordinate systems and also its derivatives to be available for the corresponding coordinate transformations.

Depending on how the parametrization of the geometry is given, different implementations are provided. The `ParametrizedGeometry` expects the mapping \mathbf{X}_e and a local finite-element and constructs the local interpolation \mathbf{X}_e^k internally, whereas the `LocalFunctionGeometry` wraps a given \mathbf{X}_e or \mathbf{X}_e^k directly.

3.1 Geometry with Local Interpolation

The first implementation, `ParametrizedGeometry`, expects only a callable function that maps entity local coordinates to global coordinates. This mapping is internally interpolated into a local finite-element space, e.g., local Lagrange functions, that allows to evaluate values and Jacobians of the parametrization from linear combinations of evaluated local basis functions and its derivatives.

C++ code

```
1 // <dune/curvedgeometry/parametrizedgeometry.hh>
2
3 template <class LocalFiniteElement, int coorddim, class Traits = (...)>
4 class ParametrizedGeometry;
```

The template parameters are defined by

LocalFiniteElement A class representing a local finite-element in the sense of `DUNE-LOCALFUNCTIONS`.

coorddim Dimension $n \geq l$ of the global coordinates this geometry maps into.

Traits (optional) Parameters for internal optimization of operations.

This geometry mapping directly corresponds to the description of the discrete geometry \mathbf{X}^k where the input function represents the mapping \mathbf{X} combined with the local-to-global mapping $\mu_e \circ \eta_{s,e}$. Thus, the actual input is $\mathbf{X}_{s,e}$. In the geometry, the local interpolation $\mathbf{X}_{s,e}^k$ is represented by the interpolation coefficients $\{\mathbf{X}_{s,e}(\hat{x}^j)\}_j$, i.e., Lagrange nodes on the surface Γ , and the set of local basis functions $\{\phi_j\}$ associated to these nodes, see (3). The Jacobian of the geometry mapping can thus be provided by evaluating the gradients of the local basis functions and its linear combination with the stored coefficients, see (7).

The `LocalFiniteElement` parameter here is the crucial input characterizing which type of local basis functions and local interpolation to use for calculating and representing the (Lagrange) nodes. `DUNE-LOCALFUNCTIONS` provides various implementations of local finite-elements, like Lagrange functions on all supported geometry types. The corresponding local finite-element can be obtained either by an explicit instantiation if the geometry type is known and identical for all elements, or by using a local finite-element cache. The latter provides the local finite-element of one kind for all geometry types by type-erasure or variadic visitors, see the example below.

3.2 Geometry with Differentiable Parametrization

The second implementation, `LocalFunctionGeometry`, expects a mapping for coordinates and additionally the Jacobian of that mapping so that the geometry Jacobian can be represented directly by the function. We expect this mapping function to be compatible with a `ElementFunction` interface, to be defined below. An element-function f_e typically can be evaluated only in element local coordinates of the element e , denoted by \mathbb{R}^m , but not in $\text{codim} > 0$ entity-local coordinates, denoted by \mathbb{R}^l . In order to allow the geometry to be defined also for these entities, or even element intersections, this geometry implementation is parametrized additionally with a `LocalGeometry` coordinate transform η . This coordinate transform maps the entity-local coordinates to the grid-element local coordinates where the element-function can be evaluated in. Thus, the geometry mapping is a chaining $\mathbb{R}^l \rightarrow \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $l \leq m \leq n$, that is, $\mathbf{X} = f_e \circ \eta = f \circ \mu_e \circ \eta$, where f is the global grid-function associated to the element-function f_e that is bound to an element e .

C++ code

```

1 // <dune/curvedgeometry/localfunctiongeometry.hh>
2
3 template <class ElementFunction, class LocalGeometry, class Traits = (...)>
4 class LocalFunctionGeometry;
5
6 template <class ElementFunction, class ctype, int dim>
7 using ElementLocalFunctionGeometry = LocalFunctionGeometry<ElementFunction,
8   DefaultLocalGeometry<ctype, dim, dim>, LocalFunctionGeometryTraits<ctype> >;

```

This geometry is parametrized with the types `ElementFunction`, `LocalGeometry`, and `Traits` that fulfill the following requirements:

LocalGeometry represents a geometric mapping from an entity of $\text{codim } c$ to the element with $\text{codim } 0$. The geometry is bound to the domain element whereas the `ElementFunction` can be bound to the range element of this geometric mapping. Thus, it is a differentiable function $\eta : \mathbb{R}^l \rightarrow \mathbb{R}^m$ with $l = m - c$ that fulfills a reduced `DUNE::GEOMETRY` concept, i.e., there is a local-to-global mapping from the coordinate system of the $\text{codim}-c$ entity to the element geometry. Let η be of type `LocalGeometry` and \hat{x} of type `LocalCoordinate`, then the expression $\eta(\hat{x}) \in \mathbb{R}^m$ results in a type that is the domain type of the `LocalFunction`.

The derivative of this parametrization must be accessible by evaluating the geometry method `\eta.jacobianTransposed(\hat{x})` that returns the transposed of the Jacobian of η , that is, $D^T \eta$ with $D\eta : \mathbb{R}^l \rightarrow L(\mathbb{R}^n, \mathbb{R}^l) \cong \mathbb{R}^{n \times l}$.

ElementFunction represents a differentiable mapping $f_e : \mathbb{R}^m \rightarrow \mathbb{R}^n$ with given derivative $(Df)_e : \mathbb{R}^m \rightarrow L(\mathbb{R}^n, \mathbb{R}^n) \cong \mathbb{R}^{n \times n}$. It is required that f_e is a model of the concept *Callable*, i.e., let x be of type `LocalGeometry::GlobalCoordinate`, then the expression $f_e(x)$ must result in a valid type denoted by the `GlobalCoordinate` of the `LocalFunctionGeometry`. f_e must be *differentiable*, i.e., there must exist a function `derivative(f)` whose return type is another model of the *ElementFunction* concept. It returns the global derivative $D(f)_e$ of the grid-function f associated to f_e .

Traits (optional) is a class holding parameters for the implementation of the geometry, like the tolerance and iteration limit for a Newton solver implementing the global-to-local function. Additionally it allows to specify element properties that cannot be deduced from the `ElementFunction` or `LocalGeometry` directly, like the `GeometryType` of the element, if there is only one.

The definition of the `ElementFunction` follows the definition of localized functions in Engwer et al. (2017). Especially, the definition of the range of the derivative of the element-function as derivative w.r.t. global coordinates is taken from there. The final `jacobianTransposed` of the geometry mapping \mathbf{X} is thus given by the chaining $D^\top \mathbf{X} = D^\top \eta \cdot (D^\top \mu_e \circ \eta) \cdot (Df \circ \mu_e \circ \eta)^\top$. Note that the mappings \mathbf{X} , η , and μ_e are geometry mappings and thus provide only the transposed of their Jacobians, whereas the localized function f_e provides the non-transposed Jacobian derivative. In order to provide the transposed of the final geometry Jacobian $D\mathbf{X}$, it needs to be a linear map that is *transposable*, i.e. it can be applied in a transposed fashion to a vector, by implementing the method `mtv()`, a method representing the transposed matrix times vector multiplication operation, or it must be representable as a matrix directly.

3.3 Examples for the Usage of Local Geometries

The two geometry implementations can now be used directly to parametrize a surface while traversing a flat reference grid \mathcal{G}_h . The following examples show the wrapping of a flat geometry into a `LocalFunctionGeometry` and `ParametrizedGeometry`, respectively.

At first, we introduce a differentiable mapping `torus : $\mathbb{R}^2 \rightarrow \mathbb{R}^3$` representing a torus parametrization.

C++ code

```

1 struct Torus {
2     double const r1 = 2.0, r2 = 1.0;
3
4     auto operator() (FieldVector<double,2> const& u) const {
5         return FieldVector<double,3>{
6             (r1 + r2*std::cos(u[0])) * std::cos(u[1]),
7             (r1 + r2*std::cos(u[0])) * std::sin(u[1]),
8             r2*std::sin(u[0])
9         };
10    }
11
12    friend auto derivative (Torus t) {
13        return [r1=t.r1,r2=t.r2](FieldVector<double,2> const& u) {
14            return FieldMatrix<double,3,2>{
15                {-r2*std::sin(u[0])*std::cos(u[1]), -(r1 + r2*std::cos(u[0]))*std::sin(u[1])},
16                {-r2*std::sin(u[0])*std::sin(u[1]), (r1 + r2*std::cos(u[0]))*std::cos(u[1])},
17                { r2*std::cos(u[0]),                0.0}
18            };
19        };
20    }
21 };

```

Second, we have to define a reference grid that provides the actual elements e , its topological connectivity, and the mapping \mathbf{X} for the curved geometry. The reference grid is a simple structured grid `YaspGrid`.

C++ code

```

1 // Construct a reference grid as a (periodic) 2d rectangular domain
2 auto refGrid = YaspGrid<2>{{2*M_PI,2*M_PI}, {8,16}, std::bitset<2>("11")};
3
4 // Define the geometry mapping
5 auto torus = Torus{};
6 auto torusGridFct = analyticGridFunction<YaspGrid<2>>(torus);

```

The mapping torus is transformed into a differentiable grid-function `torusGridFct` using the wrapper `AnalyticGridFunction`, see [section 5](#), that is provided by the library.

For the construction of a `LocalFunctionGeometry`, we have to provide an element-function of that grid-function and a `LocalGeometry`. In case of wrapping the grid element geometry, i.e., `codim` is zero, this local-geometry mapping is the identity. An efficient implementation is given by the class `DefaultLocalGeometry`.

C++ code

```

1 // Define an element-function from the grid-function
2 auto torusElemFct = localFunction(torusGridFct);
3 auto localGeometry = DefaultLocalGeometry<double,2,2>{};
4
5 // traverse the reference grid
6 for (const auto& e : elements(refGrid.leafGridView()))
7 {
8     // bind the element-function to the grid element
9     torusElemFct.bind(e);
10
11     // construct the LocalFunctionGeometry from ElementFunction and LocalGeometry
12     auto localFctGeometry
13         = LocalFunctionGeometry{e.type(), torusElemFct, localGeometry};
14
15     // (optionally) unbind from the element, i.e., free memory and unset variables
16     torusElemFct.unbind();
17 }

```

An element-function must be bound to an element (and optionally unbound at the end of usage). The type supports class-template argument deduction and if the geometry is constructed on the grid element, the `LocalGeometry` argument can even be omitted, defaulting to `DefaultLocalGeometry` in this case:

C++ code

```

1 auto localFctGeometry = LocalFunctionGeometry{referenceElement(e), torusElemFct};

```

Note, it is necessary to pass the element type as `Dune::ReferenceElement`, in order to allow the deduction of the element dimension.

Similarly, we can construct a `ParametrizedGeometry` (parametrized by Lagrange local finite-elements) by using the torus function from above. Therefore, we have to either use the element-function wrapper or have to construct the local-to-global mapping from reference element coordinates to element coordinates directly.

A local finite-element can be provided by using a local finite-element cache, or by explicit instantiation. Both variants are shown in the example below.

C++ code

```

1 // <dune/localfunctions/lagrange/lagrangelfecache.hh>
2 // <dune/localfunctions/lagrange/lagrangecube.hh>
3 ...
4 LagrangeLocalFiniteElementCache<double, double, 2, order> lfeCache;
5
6 // traverse the reference grid
7 for (const auto& e : elements(refGrid.leafGridView()))

```

```

8 {
9 // projection from local coordinates
10 auto X_e = [&torus,geo=e.geometry()](const auto& local) {
11     return torus(geo.global(local));
12 };
13
14 // construct the ParametrizedGeometry from lfe cache
15 auto curvedGeometry = ParametrizedGeometry{e.type(), lfeCache.get(e.type()), X_e};
16
17 // construct the ParametrizedGeometry from local finite-element
18 auto lfe = LagrangeCubeLocalFiniteElement<double, double, 2, order>{};
19 auto curvedGeometry2 = ParametrizedGeometry{e.type(), lfe, X_e};
20 }

```

4 CURVEDGRID and the DUNE grid interface

Instead of wrapping the geometries manually while traversing the flat reference grid, the whole grid can be wrapped. This allows to return the wrapped geometry directly in a call to `e.geometry()` instead of the flat element geometry while preserving the grid topology and element connectivity given by the wrapped reference grid. The library provides such a grid wrapper with the class `CurvedGrid` which is an implementation of \mathcal{G}_h^k or \mathcal{G} , depending on the element parametrization provided.

The class signature is given by

C++ code

```

1 // <dune/curvedgrid/grid.hh>
2
3 template <class RefGrid, class GridFunction, bool useInterpolation = false>
4 class CurvedGrid;
5
6 // constructor using Lagrange geometry interpolation
7 template <class RefGrid, class GridFunction>
8 CurvedGrid<RefGrid, GridFunction, true>
9     ::CurvedGrid(const RefGrid&, const GridFunction&, int order);
10
11 // constructor using LocalFunctionGeometry
12 template <class RefGrid, class GridFunction>
13 CurvedGrid<RefGrid, GridFunction, false>
14     ::CurvedGrid(const RefGrid&, const GridFunction&);

```

with template parameters

RefGrid The reference grid the curved grid is based on

GridFunction the type of a grid-function associated with a reference grid

useInterpolation (optional) if true, uses Lagrange ParametrizedGeometry, otherwise construct a geometry of type LocalFunctionGeometry

This class allows to locally construct both the `LocalFunctionGeometry` and the `ParametrizedGeometry`, depending on the properties of the grid-function passed to the grid wrapper and the `useInterpolation` parameter given to the grid. If the latter is true, it is assumed that a local interpolation should be constructed of the passed grid-function and thus a `ParametrizedGeometry` with Lagrange local finite-element is used as local geometry parametrization. Otherwise, if `useInterpolation` is false and the grid-function is locally differentiable, a `LocalFunctionGeometry` is used.

4.1 Examples for the Usage of the Grid Wrapper

In the following examples we construct both a wrapper using the `ParametrizedGeometry` and the `LocalFunctionGeometry`.

At first, we construct the grid by wrapping a callable representing the geometry mapping X .

C++ code

```

1 // Construct a reference grid
2 auto refGrid = Gmsh4Reader< AlbertaGrid<2,3> >::createGridFromFile("sphere.msh");
3
4 // Define the geometry mapping
5 auto sphere = [](const auto& x) { return x / x.two_norm(); };
6
7 // Wrap the reference grid to build a curved grid
8 CurvedGrid grid{*refGrid, sphere, order};

```

In this example, a grid-function is automatically constructed from the callable `sphere` as an instance of `AnalyticGridFunction`. If class template-argument deduction from C++17 cannot be used, a generator function `curvedSurfaceGrid(*refGrid, sphere)` is provided.

In the second example we construct a grid-function first which either uses a local interpolation internally, or is given as a differentiable function as above.

C++ code

```

1 // Define a discrete grid-function on the reference grid
2 auto gridFct = discreteGridViewFunction<3>(refGrid->leafGridView(), order);
3
4 // Interpolate the parametrization into the grid-function
5 Functions::interpolate(gridFct.basis(), gridFct.coefficients(), sphere);
6
7 // Wrap the reference grid to build a curved grid
8 CurvedGrid grid{*refGrid, gridFct};

```

Here, in the example we use a generic discrete function with range type `FieldVector<double, 3>`, that is represented by a `DUNE-FUNCTIONS` global basis.

C++ code

```

1 power<3>(lagrange(order), blockedInterleaved())

```

This represents a product basis (factory) composed of three times the same Lagrange basis with an ordering of the global indices in groups of the components. The basis is stored together with a coefficient vector inside the `DiscreteGridViewFunction`, see [section 5](#). Note that this grid-function is associated to a `GridView` and not the whole grid, since the global basis is bound to a `GridView`. Each time the (reference) grid changes, e.g., by local refinement or parallel load balancing, the grid-function must be updated as well, using

C++ code

```

1 gridFct.update(refGrid->leafGridView());

```

5 Grid-functions and Parametrizations

In order to construct the `ParametrizedGeometry` or `LocalFunctionGeometry` and thus the curved grid, parametrizations in form of mappings X_e or element-functions must be provided. Various grid-functions are implemented in `DUNE-CURVEDGRID` to simplify the construction and to act as reference implementations:

AnalyticGridFunction Implementation of a grid-function that can be bound to any entity in the grid given by a *Callable*, mapping global coordinates of Γ_h to a range type. This range type defines the global coordinates in the curved geometry. If the callable is *differentiable* so is the grid-function. It can thus be used in the *LocalFunctionGeometry*.

The *AnalyticGridFunction* can be constructed by

C++ code

```

1 // <dune/curvedgrid/gridfunctions/analyticgridfunction.hh>
2
3 template <class Grid, class Functor>
4 class AnalyticGridFunction;
5
6 template <class Grid, class Functor>
7 auto analyticGridFunction (Functor&& functor)
8   -> AnalyticGridFunction<Grid, std::decay_t<Functor>>;

```

AnalyticDiscreteFunction Similarly to *AnalyticGridFunction* this grid-function is constructed from a *Callable*, mapping global coordinates to global coordinates, but the mapping is locally interpolated by means of a Lagrange basis. Thus, this grid-function does not represent an exact geometry but an approximation. Moreover, it provides derivatives by differentiating the local basis functions, see (7). It can be used to parametrize *LocalFunctionGeometry*.

The *AnalyticDiscreteFunction* can be constructed by

C++ code

```

1 // <dune/curvedgrid/gridfunctions/analyticdiscretetfunction.hh>
2
3 template <class Grid, class Functor, int order = -1>
4 class AnalyticDiscreteFunction;
5
6 template <class Functor, class Grid>
7 auto analyticDiscreteFunction (Functor&& functor, const Grid&, int order)
8   -> AnalyticDiscreteFunction<Grid, std::decay_t<Functor>>;
9
10 template <int order, class Functor, class Grid>
11 auto analyticDiscreteFunction (Functor&& functor, const Grid&)
12   -> AnalyticDiscreteFunction<Grid, std::decay_t<Functor>, order>;

```

Note that this grid-function requires `DUNE-FUNCTIONS` as module dependency.

DiscreteGridViewFunction This grid-function is restricted to a specific *GridView* of the grid and is build by a set of global basis functions and a coefficient vector, both stored inside this grid-function. It can be used as *LocalFunctionGeometry*, since the basis functions provide a derivative and thus the grid-function is differentiable. Additionally, the coefficient vector, i.e., the vector of Lagrange nodes on the smooth surface Γ , can be modified and thus evolving grids can be parametrized easily.

The *DiscreteGridViewFunction* can be constructed by

C++ code

```

1 // <dune/curvedgrid/gridfunctions/discretetgridviewfunction.hh>
2
3 template <class GridView,
4         int components = GridView::dimensionworld,
5         int ORDER = -1,
6         class T = double>
7 class DiscreteGridViewFunction;
8
9 template <int components, int ORDER = -1, class T = double, class GridView>
10 auto discreteGridViewFunction (const GridView& gridView, int order = ORDER)
11   -> DiscreteGridViewFunction<GridView, components, ORDER, T>;

```

Note, the template parameter `ORDER == -1` means, the polynomial order must be given as constructor argument or as final function argument in `discreteGridViewFunction`. Otherwise, if `ORDER >= 0`, a static polynomial order is considered. In case the template parameter and the function argument are both negative, an error is thrown.

This grid-function is not as general as the `DiscreteGlobalBasisFunction` of `DUNE-FUNCTIONS`, i.e., the range type is fixed to `FieldVector<T, components>` and the global basis is explicitly defined as `power<components>(lagrange(order), blockedInterleaved())`, see [Engwer et al. \(2018\)](#). But it defines the necessary derivative and is implemented as a grid-function that includes the coefficient vector. Note that this grid-function requires `DUNE-FUNCTIONS` as module dependency.

6 Geometries

In order to test PDE discretizations or in benchmarks, geometry parametrizations for simple shapes must be available. A common example is the sphere parametrization used in all the examples above. But, additionally, shapes with less symmetry might be of interest for benchmarks and numerical validation. We have implemented the geometries of the sphere, ellipsoid, and torus as simple parametrizable shapes. Those three geometries can be obtained by

C++ code

```

1 // <dune/curvedgrid/geometries/sphere.hh>
2
3 template <int dim, class ctype = double>
4 class SphereProjection;
5
6 // sphere radius r
7 template <class Grid, class T>
8 auto sphereGridFunction (T r)
9 {
10     auto sphere = SphereProjection<Grid::dimensionworld, T>{r};
11     return analyticGridFunction<Grid>(sphere);
12 }

```

for the sphere parametrization.

C++ code

```

1 // <dune/curvedgrid/geometries/ellipsoid.hh>
2
3 template <class ctype = double>
4 class EllipsoidProjection;
5
6 // major axis a, b, and c
7 template <class Grid, class T>
8 auto ellipsoidGridFunction (T a, T b, T c)
9 {
10     auto ellipsoid = EllipsoidProjection<T>{a,b,c};
11     return analyticGridFunction<Grid>(ellipsoid);
12 }

```

for the ellipsoid parametrization.

C++ code

```

1 // <dune/curvedgrid/geometries/torus.hh>
2
3 template <class ctype = double>
4 class TorusProjection;
5
6 // Outer radius R and inner radius r
7 template <class Grid, class T>
8 auto torusGridFunction (T R, T r)

```

```

9 {
10     auto torus = TorusProjection<T>{R,r};
11     return analyticGridFunction<Grid>(torus);
12 }

```

for the torus parametrization.

For all three a projection and a corresponding grid-function that just wraps the callable into an `AnalyticGridFunction` is provided.

Additionally, two geometry parametrizations are implemented for an explicit or implicit surface representation. That is, a representation as high-resolution (piecewise) flat grid, or as zero-level set of an implicit function. For both representations the corresponding coordinate projection is implemented which is required to construct the grid-function for the curved geometries.

6.1 Projection to High-Resolution Surface Grid

The explicit surface representation is based on a flat grid that approximates the smooth surface with higher-resolution than the target grid we want to construct from a reference grid. This would allow to run simulations on a coarse grid, while the surface is only given by a very fine grid. Additionally, it allows to construct higher-order geometries for a surface that is given only with piecewise flat geometries.

C++ code

```

1 // <dune/curvedgrid/geometries/explicitsurface.hh>
2
3 template <class ctype = double>
4 class ExplicitSurfaceProjection;
5
6 // Constructor with grid and an option to activate caching
7 template <class ctype>
8     template <class Grid>
9     ExplicitSurfaceProjection<ctype>
10    ::ExplicitSurfaceProjection (const Grid& grid, bool cached = true);

```

The input to construct the `ExplicitSurfaceProjection` is a (surface) grid representing the high-resolution (piecewise) flat surface. Internally, the vertices of the grid are stored in a fast search tree, a KDTree implementation based on [Blanco and Rai \(2014\)](#) which supports nearest neighbor search. Each time the projection is evaluated for a global coordinate x the closest vertex in the high-resolution surface grid is searched. Afterwards the adjacent elements are considered. For each of them the closest point to x is determined by orthogonal projection. The closest found point is then returned.

This approach only works well if the high-resolution surface grid is of sufficient quality, i.e., no overly acute-angled elements occur. Otherwise the adjacent elements of the closest vertex don't need to contain the actual closest point to x . When choosing a method for creating the grid a Delaunay-triangulation, for instance, is a good candidate. Grids not fulfilling the element-quality condition can typically be adapted without losing their important features by using the meshing tool `meshconv`, [Stenger \(2020\)](#).

The decision to only consider the elements adjacent to the unique closest vertex is a compromise made for performance-reasons. By extending the nearest neighbor search in the KDTree to several closest vertices and considering adjacent elements to all of them one could allow for lower quality surface grids at the cost of sacrificing performance.




Figure 3: Coarse grid (left) used as reference domain Γ_h for parametrization with closest-point projection to fine grid Γ_h^{fine} (right) of Stanford-bunny geometry. The Coarse grid is obtained by feature-preserving coarsening of the fine-grid, see [Stenger \(2020\)](#); [Valette and Chassery \(2004\)](#).

6.2 Projection to Zero-Level set

If the surface is given implicitly as the zero-level set of a higher-order function, the closest-point projection must be calculated iteratively using either a Newton-method or a fixed-point iteration. The essential property that is used in these algorithms, is that the normal vector of the surface in a proximate neighborhood is given by the normalized gradient of the implicit function.

Thus, the input to the `ImplicitSurfaceProjection` is a differentiable function ψ providing the surface as $\Gamma = \{x \in \mathbb{R}^n : \psi(x) = 0\}$ with $n_\Gamma(x) = \nabla\psi(x)/\|\nabla\psi(x)\|$ at $x \in \Gamma$.

Given an initial guess for the projected point x_0 , the authors of [Persson \(2004\)](#); [Nitschke \(2014\)](#) describe a scheme to iteratively compute better guesses of the projection of x_0 to Γ by approximating the closest-point property $\mathbf{X}(x) = x - d(x)n_\Gamma(\mathbf{X}(x))$ with a representation of the approximate distance $d(x) \approx \psi(x)/\|\nabla\psi(x)\|$ and the normal vector representation from above:

$$x_{i+1} = x_i - \nabla\psi(x_i) \frac{\psi(x_i)}{\|\nabla\psi(x_i)\|^2}, \quad \widehat{\text{err}}_i = \frac{|\psi(x_i)|}{\|\nabla\psi(x_i)\|}. \quad (8)$$

This scheme applies this relation iteratively, eventually converging to a point on Γ near the closest-point $\mathbf{X}(x)$. It is implemented in the class

C++ code

```

1 // <dune/curvedgrid/geometries/implicitsurface.hh>
2
3 template <class Functor>
4 class SimpleImplicitSurfaceProjection;
5
6 template <class Functor>
7 SimpleImplicitSurfaceProjection<Functor>
8   ::SimpleImplicitSurfaceProjection (const Functor& psi, int maxIter = 10);
```

where the maximal number of iterations in the iterative scheme is given by `maxIter`.

An improved version of that scheme that takes the iterate from the simple scheme as initial approximation of the closest-point projection and uses this point to get a better estimate for the actual distance of x to Γ is proposed in [Demlow and Dziuk \(2007\)](#):

$$\begin{aligned} \tilde{x}_{i+1} &= x_i - \nabla\psi(x_i) \frac{\psi(x_i)}{\|\nabla\psi(x_i)\|^2}, \quad \text{dist} = \text{sign}(\psi(x_0))\|\tilde{x}_{i+1} - x_0\| \\ x_{i+1} &= x_0 - \nabla\psi(\tilde{x}_{i+1}) \frac{\text{dist}}{\|\nabla\psi(\tilde{x}_{i+1})\|}, \quad \text{err}_i = \widehat{\text{err}}_i + \left\| \frac{\nabla\psi(x_i)}{\|\nabla\psi(x_i)\|} \pm \frac{(x_i - x_0)}{\|x_i - x_0\|} \right\| \end{aligned} \quad (9)$$

The computational demand is higher than in the simple scheme, but it converges to the actual closest point on Γ . This scheme is implemented in the class

C++ code

```

1 // <dune/curvedgrid/geometries/implicitsurface.hh>
2
3 template <class Functor>
4 class ImplicitSurfaceProjection;
5
6 template <class Functor>
7 ImplicitSurfaceProjection<Functor>
8 ::ImplicitSurfaceProjection (const Functor& psi, int maxIter = 10);

```

6.2.1 Example of an application of the iterative scheme We consider a surface with genus two, given by the function

$$\psi(x, y, z) = 2y(y^2 - 3x^2)(1 - z^2) + (x^2 + y^2)^2 - (9z^2 - 1)(1 - z^2).$$

A reference surface can be obtained by extracting the zero-level set contour, e.g., by using the tool ParaView, see [Ahrens et al. \(2005\)](#), by a surface Delaunay triangulation combined with a surface projection, see [Persson and Strang \(2004\)](#); [Persson \(2004\)](#), or by reconstructing the implicitly defined surface using some fast marching algorithm, see [Engwer and Nüßing \(2017\)](#). We followed the first approach, combined with a coarsening of the obtained surface grid using [Stenger \(2020\)](#), see also [Valette and Chassery \(2004\)](#) for a similar approach.

Applying the implicit projection methods from above to such a coarse reference grid results in very fast convergence of both schemes to a machine epsilon tolerance. If no closest-point property is explicitly required, the simple iterative scheme (8) performs faster than the full closest-point iterative projection scheme (9).




Figure 4: Surface extracted from the implicit description as zero-level set of ψ , using an implicit projection method for the higher-order surface approximation. The colored plane illustrates a cut through the function ψ . The table in (a) shows the maximal number of iterations i necessary to reach a pointwise tolerance $e_i < \sqrt{\epsilon} \approx 1.5 \cdot 10^{-8}$ with ϵ the machine epsilon of double floating-point numbers of either the error $e_i = \widehat{\text{err}}_i$ that approximates the distance of x_i to the surface for all projected points x , or the error $e_i = \text{err}_i$ in the closest-point property of the iterates.

7 Numerical examples

In order to verify the implementation and to test different geometric representations, we have first analyzed the difference between the discrete surface and the continuous surface. This numerical verification considers the difference between the smooth surface quantities, the closest-point projection \mathbf{X} , the surface normal n , and the mean curvature H , in the $L^\infty(\Gamma_h)$ norm. In [Demlow \(2009\)](#); [Hansbo et al. \(2019\)](#), upon many others, the following estimates are shown:

Proposition 1 For h small enough, we have the estimates

$$\|\mathbf{X} - \mathbf{X}^k\|_{L^\infty(\Gamma_h)} \leq Ch^{k+1}, \quad \|n \circ \mathbf{X} - n_h^k\|_{L^\infty(\Gamma_h^k)} \leq Ch^k, \quad \|H \circ \mathbf{X} - H_h^k\|_{L^\infty(e)} \leq Ch^{k-1} \quad (10)$$

for $e \in \mathcal{G}_h^k$, with C a generic constant independent of the mesh parameter h .

We show for three smooth geometries, the unit sphere, an ellipsoid with major axis (1, 1.25, 0.75), and torus with the two radii (2, 1), the convergence in the L^2 -norm on the discrete surface that follows from the assertion in the L^∞ -norm. Therefore, we have first created a reference grid, then interpolated the surface parametrization into the element geometries with order k and finally, evaluated the three quantities by iterating over the reference surface. The error norms are shown in Figure 5. Note, the L^∞ -norm is only approximated by computing the maximum in all quadrature points of the elements that are also used for computing the L^2 -norms.

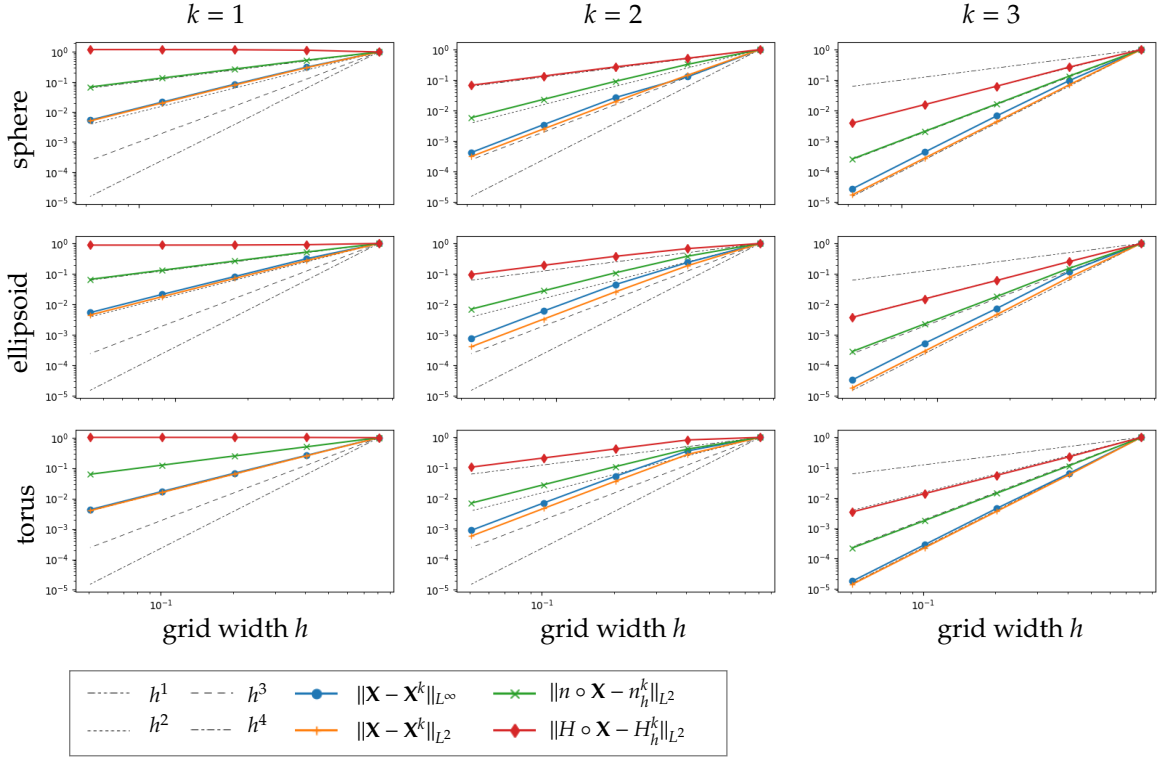


Figure 5: Geometric error norms, normalized by error of largest grid-size, evaluated for three different geometries, the unit sphere, an ellipsoid, and a torus, for $k = 1, 2$, and 3 . In various dashed lines, the ideal convergence lines h^p are shown.

7.1 Surface Vector Helmholtz equation

We consider the vector Helmholtz equation for a test of the surface parametrization. The corresponding intrinsic surface formulation reads: Find the tangential-vector field $u \in H_{\tan}^1(\Gamma, TT)$ such that

$$\left(\nabla_\Gamma u, \nabla_\Gamma v\right)_\Gamma + (u, v)_\Gamma = (f, v)_\Gamma \quad \forall v \in H_{\tan}^1(\Gamma, TT), \quad (11)$$

with ∇_Γ the covariant derivative of the vector fields, $(\cdot, \cdot)_\Gamma$ the generic L^2 -inner product on Γ , and f a given tangential vector field.

For the discretization of this equation we follow the ideas of [Nestler et al. \(2017, 2019\)](#); [Hansbo et al. \(2019\)](#); [Gross et al. \(2018\)](#); [Jankuhn and Reusken \(2020\)](#) and represent the vector field in an embedding space – in this case the Euclidean space $T\Gamma \cong \mathbb{R}^3$ – and transform the corresponding covariant derivatives into the embedding space. By allowing the vector field to also have non-tangential components, the equation can be written as a coupled system of scalar-valued equations with a penalization term to enforce tangentiality: Find the vector field $u \in [H^1(\Gamma, \mathbb{R})]^3$ such that

$$\left(\nabla_{\Gamma} P u, \nabla_{\Gamma} P v\right)_{\Gamma} + (P u, P v)_{\Gamma} + \omega(n \cdot u, n \cdot v)_{\Gamma} = (f, P v)_{\Gamma} \quad \forall v \in [H^1_{\text{tan}}(\Gamma, \mathbb{R})]^3, \quad (12)$$

with $\omega \gg 0$ a penalization parameter, $P = \text{Id} - n \otimes n$ the tangential projection operator w.r.t. the surface normal vector n . For extended vector fields u , the surface covariant derivative can be expressed in terms of the Euclidean derivative ∇ in the ambient space, by $\nabla_{\Gamma} P u = P \nabla(P u) P = P(u \otimes \nabla_{\Gamma}) - (n \otimes \nabla_{\Gamma})(n \cdot u)$. The expression $u \otimes \nabla_{\Gamma}$ means the componentwise surface gradient.

In order to discretize this equation, we approximate Γ by Γ_h^k and $H^1(\Gamma, \mathbb{R})$ by $V_{h,k}^r$ the Lagrange finite-element space of order r , given by

$$V_{h,k}^r := \left\{v \in C^0(\Gamma_h^k) : v \circ \mathbf{X}_e^k \in \mathbb{P}_r \quad \forall e \in \mathcal{G}_h\right\}$$

Following the analysis of [Hansbo et al. \(2019\)](#) and [Jankuhn and Reusken \(2020\)](#); [Gross et al. \(2018\)](#) the normal vector involved in the covariant derivative and mass-matrix term can have the same approximation order as the geometry, but the normals involved in the penalty term should be at least one order better. We denote this “better” normal by \tilde{n} . Additionally, the scaling for the penalization factor should be of order h^{-2} , thus we take $\omega = \beta h^{-2}$ with $\beta = 10$ in the numerical experiments below. Note, in [Hardering and Praetorius \(2021\)](#) it is shown that when interested in the tangential part of the solution only, a better choice for the approximation is $\tilde{n} \equiv n_h$ the normal vector of the discrete surface with $\omega \simeq h^{-1}$.

The resulting discrete variational formulation reads: Find the vector field $u_h \in [V_{h,k}^r]^3$ such that

$$\left(\nabla_{\Gamma_h^k} P_h u_h, \nabla_{\Gamma_h^k} P_h v_h\right)_{\Gamma_h^k} + (P_h u_h, P_h v_h)_{\Gamma_h^k} + \omega(\tilde{n}_h \cdot u_h, \tilde{n}_h \cdot v_h)_{\Gamma_h^k} = (f \circ \mathbf{X}^k, P_h v_h)_{\Gamma_h^k} \quad \forall v \in [V_{h,k}^r]^3. \quad (13)$$

where the inner product and derivatives have to be understood elementwise and locally, i.e., $(A, B)_{\Gamma_h^k} = \sum_{e \in \mathcal{G}_h^k} \int_e A : B \, d\Gamma$. The challenge hereby is to integrate over the parametrized geometries and to provide normal vectors of differing approximation order. Since we have (10), the higher-order normals can be obtained by constructing a local geometry of parametrization order $k + 1$.

C++ code

```

1  template <int order>
2  using LFE_t = LagrangeSimplexLocalFiniteElement<double, double, 2, order>;
3
4  // traverse the reference grid
5  for (const auto& e : elements(refGrid->leafGridView()))
6  {
7    // projection from local coordinates
8    auto X_e = [&sphere, geo=e.geometry()](const auto& local) {
9      return sphere(geo.global(local));
10   };
11
12   // construct the CurvedGeometries from local parametrization
13   ParametrizedGeometry geometry(e.type(), LFE_t<k>{}, X_e);
14   ParametrizedGeometry higherOrderGeometry(e.type(), LFE_t<k+1>{}, X_e);
15
16   const auto& quadRule = QuadratureRules<double, 2>::rule(e.type(), quad_order);
17   for (const auto& qp : quadRule) {
18     // integration element
19     double dS = geometry.integrationElement(qp.position()) * qp.weight();
20     // surface normal n_h
21     auto nh = geometry.normal(qp.position());

```

```

22 // higher-order surface normal n~_h
23 auto nh2 = higherOrderGeometry.normal(qp.position());
24
25 // ...
26 }
27 }

```

Thus, instead of traversing the CurvedGrid, one could iterate over the reference grid \mathcal{G}_h instead and locally construct the curved geometries of order k and $k + 1$ to obtain the surface element and the surface normal vectors.

7.1.1 Vector Fields on Spherical Geometry Following the example in Nestler et al. (2019) we define an exact solution $u^* := \text{rot}_n(xyz)$ of (11) and construct⁴ the corresponding load-vector function $f := -\text{div}_\Gamma \nabla_\Gamma u^* + u^*$. A coarse grid of the sphere is explicitly provided using a Gmsh mesh with nearly equal sized elements.

In the following numerical test the geometry parametrization and function parametrization take the same polynomial order, i.e., $r = k$.

level	grid width h	error ($k = 1$)	eoc	error ($k = 2$)	eoc	error ($k = 3$)	eoc
0	$7.8462 \cdot 10^{-1}$	$3.0270 \cdot 10^{-1}$	—	$2.3687 \cdot 10^{-2}$	—	$1.7959 \cdot 10^{-2}$	—
1	$3.9937 \cdot 10^{-1}$	$9.2247 \cdot 10^{-2}$	1.809	$3.7851 \cdot 10^{-3}$	2.713	$1.1312 \cdot 10^{-3}$	4.094
2	$2.0496 \cdot 10^{-1}$	$2.4538 \cdot 10^{-2}$	2.000	$4.7189 \cdot 10^{-4}$	3.121	$7.3444 \cdot 10^{-5}$	4.099
3	$1.0322 \cdot 10^{-1}$	$6.2489 \cdot 10^{-3}$	1.998	$5.6328 \cdot 10^{-5}$	3.099	$4.6436 \cdot 10^{-6}$	4.025
4	$5.1721 \cdot 10^{-2}$	$1.5703 \cdot 10^{-3}$	2.000	$6.7794 \cdot 10^{-6}$	3.064	$2.9136 \cdot 10^{-7}$	4.007
5	$2.5878 \cdot 10^{-2}$	$3.9311 \cdot 10^{-4}$	2.000	$8.2829 \cdot 10^{-7}$	3.036	$1.8234 \cdot 10^{-8}$	4.002

Table 1: L^2 -error of linear ($k = 1$), quadratic ($k = 2$), and cubic ($k = 3$) iso-parametric finite elements for the vector Helmholtz equation with experimental order of convergence (eoc) in the grid width h that tends towards 2, 3, and 4, respectively.




Figure 6: Solution of the vector Helmholtz equation on the sphere and on the ellipsoid.

7.2 Moving Grids

Let $\Gamma_h \subset \mathbb{R}^3$ be a smooth closed and stationary reference surface. A map $\mathbf{X} : \Gamma_h \times [0, T] \rightarrow \mathbb{R}^3$ then defines a parametrization of a family of surfaces $\Gamma(t) \subset \mathbb{R}^3$ over this reference manifold:

$$\Gamma(t) = \{\mathbf{X}(x, t) : x \in \Gamma_h\}. \quad (14)$$

⁴The corresponding symbolic computations are done using sympy.

The evolution of this family of surfaces is characterized by its velocity $v(X, t) \in \mathbb{R}^3$ at each point $X = X(x, t) \in \Gamma(t)$,

$$\partial_t X(x, t) = v(X(x, t), t). \quad (15)$$

We consider the surface evolution driven by its mean curvature, the geometric *mean curvature flow* flow, see [Deckelnick et al. \(2005\)](#); [Kovács et al. \(2019\)](#); [Dziuk \(1990\)](#). Therefore, we introduce $H := \text{tr}(\kappa)$ the mean curvature of the surface with extended Weingarten map $\kappa = -n \otimes \nabla_\Gamma$, and the surface evolution $v = -Hn$.

Utilizing the geometric identity $\Delta_\Gamma X = -Hn$, see, e.g., [Dziuk \(1990\)](#); [Dziuk and Elliott \(2013\)](#), a weak formulation of the evolution law can be written: For all $t \in [0, T]$, find $\mathbf{X}(\cdot, t) \in [H^1(\Gamma_h)]^3$ such that

$$\int_{\Gamma(t)} \partial_t x(t) \cdot y \, d\Gamma(t) = - \int_{\Gamma(t)} \sum_i \nabla_{\Gamma(t)} x^i(t) \cdot \nabla_{\Gamma(t)} y^i \, d\Gamma(t) \quad \forall y \in [H^1(\Gamma(t))]^3, \quad (16)$$

with the surface identity $x(t) = X(\cdot, t)$. Note that on the right-hand side of that equation, we find the componentwise surface gradient of the parametrization.

In order to discretize this equation, we introduce a splitting of the time-interval $[0, T]$ into discrete time steps $0 < t_0 < t_1 < \dots < t_N = T$ with generic time step width $\tau = t_s - t_{s-1}$ and denote by $\mathbf{X}_s \cong \mathbf{X}(\cdot, t_s)$ the parametrization at time step t_s . Correspondingly, we denote by $\Gamma_s = \mathbf{X}_s(\Gamma_h)$ the surface at that time step. Since the grid-function \mathbf{X}_s is parametrized over the reference surface Γ_h we replace the integration over Γ_s by an integration over the reference surface using a transformation of the surface elements $d\Gamma_s \rightarrow d\Gamma_h$.

For the discretization in space, we introduce the finite-element space V_h^r of Lagrange finite-elements on Γ_h ,

$$V_h^r := \left\{ v \in C^0(\Gamma_h) : v \circ \mathbf{X}_e \in \mathbb{P}_r \, \forall e \in \mathcal{G}_h \right\}.$$

Then we get the discrete variational formulation by simple Euler discretization in time: Let \mathbf{X}_0 be a given initial parametrization. For all $s = 1, \dots, N$, find $\mathbf{X}_s \in [V_h^r]^3$ such that

$$\int_{\Gamma_h} (\mathbf{X}_s - \mathbf{X}_{s-1}) \cdot \mathbf{Y} \, d\Gamma_{s-1} = - \int_{\Gamma_h} \tau \sum_i \nabla_{\Gamma_{s-1}} X_s^i \cdot \nabla_{\Gamma_{s-1}} Y^i \, d\Gamma_{s-1} \quad \forall \mathbf{Y} \in [V_h^r]^3. \quad (17)$$

So, while traversing the reference grid, we need the geometry of the curved grid from the last time step $\Gamma_{s-1} = \mathbf{X}_{s-1}(\Gamma_h)$. This is given by the grid-function \mathbf{X}_{s-1} :

Initially we construct a `DiscreteGridViewFunction`:

C++ code

```
1 auto X = discreteGridViewFunction(refGrid->leafGridView(), order);
2 auto X_e = localFunction(X);
3
4 // interpolate the initial surface parametrization
5 auto perturbedSphere = [](auto const& x) { return ...; };
6 Functions::interpolate(X.basis(), X.coefficients(), perturbedSphere);
```

This grid-function additionally provides a global basis that can be localized to an element:

C++ code

```
1 auto localView = X.basis().localView();
2
3 // traverse the reference grid
4 for (const auto& e : elements(X.basis().gridView()))
5 {
6 // bind the local function to the element
7 X_e.bind(e);
8 }
```

```

9 // bind the localized basis to the element
10 localView.bind(e);
11
12 // the localized basis provides a local finite-element
13 auto const& localFE = localView.tree().child(0).finiteElement();
14 auto const& localBasis = localFE.localBasis();
15
16 // ...
17 }

```

On each element of the reference grid, we can construct a curved geometry. This can be used to obtain the integration element and the transform of the local gradients of the local basis-functions to the actual domain of the curved element.

C++ code

```

1 LocalFunctionGeometry geometry(referenceElement(e), X_e);
2
3 const auto& quadRule = QuadratureRules<double,2>::rule(e.type(), quad_order);
4 for (const auto& qp : quadRule) {
5 // integration element dG_{s-1}
6 double dS = geometry.integrationElement(qp.position()) * qp.weight();
7
8 // the inverse of the transposed geometry Jacobian
9 auto Jtinv = geometry.jacobianInverseTransposed(qp.position());
10
11 // evaluate the local basis Jacobians in the quadrature point
12 std::vector<FieldMatrix<double,1,2>> shapeGradients;
13 localBasis.evaluateJacobian(qp.position(), shapeGradients);
14
15 // transform the local basis Jacobians to the real element
16 std::vector<FieldVector<double,3>> gradients(shapeGradients.size());
17 for (std::size_t i = 0; i < gradients.size(); ++i)
18 Jtinv.mv(shapeGradients[i][0], gradients[i]);
19 }

```

The evolution of a perturbed initial sphere can be found in [Figure 7](#). It starts the evolution by smoothing high curvature regions while continuously shrinking the surface. Eventually the surface gets sphere-like with a radius that tends to zero.




Figure 7: Mean-curvature flow of a perturbed spherical surface with parametrization of polynomial order 2 at four different time-steps in the evolution.

8 Conclusion and Outlook

We have implemented parametrized and curved geometries for the DUNE framework by wrapping grid-functions or differentiable functions into the Geometry interface defined by the DUNE-GRID module. Additionally, we have implemented a wrapper for flat grids providing a curved geometry on traversal, while preserving the grid topology and element connectivity.

It is shown in several examples how these wrappers provide high flexibility while preserving simple usage patterns. In a numerical study we have verified the implementation by showing classical geometry error bounds to be achieved.

The DUNE modules `DUNE-CURVEDGEOMETRY` and `DUNE-CURVEDGRID` provide not only the geometry and grid wrappers but also utility functions to simplify the work with curved geometries. These utilities include some reference geometries and surface projections, as well as grid-functions for various requirements. A long-term goal is to integrate these DUNE modules into the core functionality of DUNE.

The implementation of the DUNE modules is neither restricted to only surface parametrizations nor to grids without boundaries. The transformation of 1d, 2d, or 3d geometries is implemented and handling of boundary parametrized intersections is included. A future work is the development and application of more grid-functions for the geometry parametrization, e.g., based on a b-spline basis or p-adaptive Lobatto functions. The latter may allow for boundary adapted high-order parametrizations with inner elements described as affine mappings. Another topic for further studies is the differentiability of the element-functions used in the grid parametrization. Automatic or numeric differentiation of the projections \mathbf{X} or \mathbf{X}_e , as well as implicit differentiation of the levelset function ψ , could be a way to allow for exact geometry parametrizations of more surfaces.

Acknowledgments

This work was supported by German Research Foundation (DFG), Research Unit *Vector- and Tensor-Valued Surface PDEs* (FOR 3013)

A Appendix

A.1 Input and output of curved geometries

In all the examples above, a reference grid \mathcal{G}_h is provided by reading a surface mesh from file. DUNE provides a multitude of grid file readers, but lacks support for a reader that can read curved geometries directly. This would allow to not only start from a reference grid with an analytical projection, but to provide a discrete representation of the curved surface from the beginning. Many meshing tools allow to directly construct such curved meshes and provide a file format that is able to represent the additional nodes for a parametrization. We have implemented two readers, the `VtkReader` for the VTK file format and a `Gmsh4Reader` for the GMSH file format. Both associated meshing and visualization tools, ParaView, [Ahrens et al. \(2005\)](#), and GMSH, [Geuzaine and Remacle \(2009\)](#), allow to design curved geometries and to export these in the mentioned file formats.

In addition to file readers, the result of a numerical computation must be exported to allow visualization and postprocessing. Our tool of choice is ParaView, supporting the VTK file format also for curved geometries. We have implemented a grid and data writer for this file format.

The module `DUNE-VTK` provides grid readers and writers with flexible input and output policies in the VTK file format, while the module `DUNE-GMSH4` provides grid readers for the Gmsh4 file format. In the following code snippets we show the VTK reader and writer, a corresponding Gmsh reader works analogously.

A.1.1 File Readers for Curved Grids When reading a higher-order grid from file, we need to construct both the reference grid and the parametrization. The reference grid could be obtained by evaluating the higher-order grid representation in the element's corner vertices, whereas the parametrization must be extracted from the additional Lagrange nodes stored in the file.

In order to read these nodes and to construct an element connectivity, an input policy called *grid-creator* must be provided. It creates local vertex coordinates and element indices from the fields read from file and passes those to a `GridFactory` to create the actual grid. The grid-creator for VTK and Gmsh4 files with parametrized grid elements is called `LagrangeGridCreator`.

The grid itself does not contain the additional Lagrange nodes, but a parametrization or coordinate mapping describing the higher-order geometries. Thus, one needs to associate these nodes to a local Lagrange basis. We provide a grid-function representation of the higher-order geometries parametrized over the extracted reference grid. This grid-function is represented by the *grid-creator* itself.

C++ code

```

1 // <dune/vtk/vtkreader.hh>
2 // <dune/vtk/gridcreators/lagrangegridcreator.hh>
3
4 using Grid = FoamGrid<2,3>;
5 using Creator = Vtk::LagrangeGridCreator<Grid>;
6 auto grid = VtkReader<Grid, Creator>::createGridFromFile("filename.vtu");

```

To extract the parametrization in addition to the reference grid, we need to obtain the grid-creator directly, that acts as a grid-function after reading from the file:

C++ code

```

1 using Grid = FoamGrid<2,3>;
2 GridFactory<Grid> factory;
3 Vtk::LagrangeGridCreator creator{factory};
4 VtkReader reader{creator};
5 reader.read("filename.vtu");
6
7 // construct the reference grid
8 auto grid = factory.createGrid();

```

Thus, the grid-function can be used to fill any other storage to represent the geometry, e.g., by interpolating into a `DiscreteGridViewFunction`, or can be used directly for the parametrization of the `CurvedGrid`:

C++ code

```

1 CurvedGrid curvedGrid{*grid, creator, creator.order()};

```

A.1.2 VTK Writer VTK supports higher order cell types including Lagrange parametrizations of cells since version 9. This allows to directly write the curved geometries to files. A corresponding output policy, called *data-collector*, is added to support these cell types. This data-collector is responsible for transforming a grid-view into a list of point coordinates and a connectivity table. Additionally, it collects values associated to the point coordinates, if data should be written to the VTK file.

For writing higher-order Lagrange parametrized grids, in addition to the corner vertices of grid elements internal Lagrange nodes are written to the file. The connectivity table collects these nodes in a specific order so that they can be associated to Lagrange basis functions. The corresponding data-collector is called `LagrangeDataCollector`:

C++ code

```

1 // <dune/vtk/datacollectors/lagrangedatacollector.hh>
2
3 namespace Vtk {
4     template <class GridView, int ORDER = -1>
5     class LagrangeDataCollector;
6 }

```

The template parameter `GridView` represents the type of the grid-view that shall be written and the second (optional) parameter `ORDER` represents the Lagrange polynomial order of the cell parametrization. This second parameter is optional since also runtime polynomial order is supported, by passing the polynomial order parameter to the constructor instead.

C++ code

```

1 LagrangeDataCollector (const GridView& gridView, int order = ORDER)

```

In case no constructor parameter for the order is given, the template `ORDER` parameter is used as default value. Note that either in the template parameter or in the constructor argument a positive value for order must be given.

The corresponding writer object can be instantiated by either passing a data-collector object or by letting the writer construct it with the given grid-view object.

C++ code

```

1 using DataCollector = Vtk::LagrangeDataCollector<GridView,4>;
2 using Writer = VtkUnstructuredGridWriter<GridView, DataCollector>;
3
4 // a) default construct the data-collector with the passed gridView
5 Writer vtkWriter1{gridView};
6
7 // b) construct the data-collector before and pass it to the writer
8 DataCollector dataCollector{gridView};
9 Writer vtkWriter2(dataCollector);

```

Note, we are using an unstructured-grid writer to generate a `.vtu` file that represents the grid.




Figure 8: Three different approximations of the sphere visualized using the VTK writer with ParaView. Shown are the element edges and the Lagrange nodes. Left: reference grid, Center: Lagrange parametrization with polynomial order $k = 4$, Right: Lagrange parametrization with polynomial order $k = 1$ and two grid refinements.

References

J. Ahrens, B. Geveci, and C. Law. *ParaView: An End-User Tool for Large Data Visualization*. Visualization Handbook. Elsevier, 2005. ISBN 978-0123875822.

- R. Anderson, J. Andrej, A. Barker, J. Bramwell, J.-S. Camier, J. Cervený, V. Dobrev, Y. Dudouit, A. Fisher, T. Kolev, W. Pazner, M. Stowell, V. Tomov, I. Akkerman, J. Dahm, D. Medina, and S. Zampini. MFEM: A modular finite element methods library. *Computers & Mathematics with Applications*, 81:42–74, jan 2021. doi: 10.1016/j.camwa.2020.06.009.
- D. Arndt, W. Bangerth, D. Davydov, T. Heister, L. Heltai, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, and D. Wells. The deal.II finite element library: Design, features, and insights. *Journal of Numerical Mathematics*, 81:407–422, jan 2021. doi: 10.1016/j.camwa.2020.02.022. URL <https://dealii.org>.
- P. Bastian, M. Blatt, A. Dedner, N.-A. Dreier, C. Engwer, R. Fritze, C. Gräser, D. Kempf, R. Klöforn, M. Oehlberger, and O. Sander. The dune framework: Basic concepts and recent developments. *Computers & Mathematics with Applications*, 81:75–112, 2021. ISSN 0898-1221. doi: 10.1016/j.camwa.2020.06.007.
- J. L. Blanco and P. K. Rai. nanoflann: a C++ header-only fork of FLANN, a library for nearest neighbor (NN) with kd-trees, 2014. URL <https://github.com/jlblancoc/nanoflann>.
- D. Bothe and A. Reusken, editors. *Transport Processes at Fluidic Interfaces*. Springer International Publishing, 2017. doi: 10.1007/978-3-319-56602-3.
- E. Burman, P. Hansbo, M. G. Larson, and A. Massing. Cut finite element methods for partial differential equations on embedded manifolds of arbitrary codimensions. *ESAIM: Mathematical Modelling and Numerical Analysis*, 52(6):2247–2282, 2018. doi: 10.1051/m2an/2018038.
- K. Deckelnick, G. Dziuk, and C. M. Elliott. Computation of geometric partial differential equations and mean curvature flow. *Acta Numerica*, 14:139–232, 2005. doi: 10.1017/S0962492904000224.
- A. Dedner and M. Nolte. Construction of local finite element spaces using the generic reference elements. In A. Dedner, B. Flemisch, and R. Klöforn, editors, *Advances in DUNE*, pages 3–16, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-28589-9. doi: 10.1007/978-3-642-28589-9_1.
- A. Demlow. Higher-order finite element methods and pointwise error estimates for elliptic problems on surfaces. *SIAM Journal on Numerical Analysis*, 47(2):805–827, 2009. doi: 10.1137/070708135.
- A. Demlow and G. Dziuk. An adaptive finite element method for the Laplace–Beltrami operator on implicitly defined surfaces. *SIAM Journal on Numerical Analysis*, 45(1):421–442, 2007. doi: 10.1137/050642873.
- G. Dziuk. An algorithm for evolutionary surfaces. *Numerische Mathematik*, 58:603–611, 1990. doi: 10.1007/BF01385643.
- G. Dziuk and C. M. Elliott. Finite element methods for surface pdes. *Acta Numerica*, 22:289–396, 2013. doi: 10.1017/S0962492913000056.
- C. Engwer and A. Nüßing. Geometric reconstruction of implicitly defined surfaces and domains with topological guarantees. *ACM Transactions on Mathematical Software*, 44(2), 2017. doi: 10.1145/3104989.
- C. Engwer, C. Gräser, S. Müthing, and O. Sander. The interface for functions in the dune-functions module. *Archive of Numerical Software*, 5(1):95–110, 2017. doi: 10.11588/ans.2017.1.27683.
- C. Engwer, C. Gräser, S. Müthing, and O. Sander. Function space bases in the dune-functions module. arXiv, 2018, 1806.09545.
- W. Freeden and M. Schreiner. *Spherical Functions of Mathematical Geosciences*. Springer Berlin Heidelberg, 2009. doi: 10.1007/978-3-540-85112-7.

- H. Fritz. Isoparametric finite element approximation of Ricci curvature. *IMA Journal of Numerical Analysis*, 33(4):1265–1290, 2013. doi: 10.1093/imanum/drs037.
- C. Geuzaine and J.-F. Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11): 1309–1331, 2009. doi: 10.1002/nme.2579. URL <https://gmsh.info>.
- J. Grande, C. Lehrenfeld, and A. Reusken. Analysis of a high-order trace finite element method for PDEs on level set surfaces. *SIAM Journal on Numerical Analysis*, 56(1):228–255, 2018. doi: 10.1137/16M1102203.
- S. Gross, T. Jankuhn, M. A. Olshanskii, and A. Reusken. A trace finite element method for vector-Laplacians on surfaces. *SIAM Journal on Numerical Analysis*, 56(4):2406–2429, 2018. doi: 10.1137/17M1146038.
- P. Hansbo, M. G. Larson, and K. Larsson. Analysis of finite element methods for vector Laplacians on surfaces. *IMA Journal of Numerical Analysis*, 04 2019. doi: 10.1093/imanum/drz018.
- H. Hardering and S. Praetorius. Tangential errors of tensor surface finite elements. arXiv, Jun 2021, 2106.01000. arxiv:2106.01000.
- C. J. Heine. Isoparametric finite element approximation of curvature on hypersurfaces. Technical report, Abteilung für Angewandte Mathematik, Universität Freiburg, Hermann-Herder-Straße 10, 79104 Freiburg i.Br., Germany, 2004.
- T. Jankuhn and A. Reusken. Trace finite element methods for surface vector-Laplace equations. 41(1):48–83, may 2020. doi: 10.1093/imanum/drz062.
- T. Jankuhn, M. Olshanskii, and A. Reusken. Incompressible fluid problems on embedded surfaces: modeling and variational formulations. *Interfaces and Free Boundaries*, 20(3):353–377, 2018. doi: 10.4171/IFB/405.
- A. Johnen, J.-F. Remacle, and C. Geuzaine. Geometrical validity of high-order triangular finite elements. *Engineering with Computers*, 30(3):375–382, 2014. doi: 10.1007/s00366-012-0305-7.
- B. Kovács, B. Li, and C. Lubich. A convergent evolving finite element algorithm for mean curvature flow of closed surfaces. *Numerische Mathematik*, 143:797–853, 2019. doi: 10.1007/s00211-019-01074-2.
- C. Lehrenfeld. High order unfitted finite element methods on level set domains using isoparametric mappings. *Computer Methods in Applied Mechanics and Engineering*, 300:716–733, 2016. doi: 10.1016/j.cma.2015.12.005.
- M. Nestler, I. Nitschke, S. Praetorius, and A. Voigt. Orientational order on surfaces: The coupling of topology, geometry, and dynamics. *Journal of Nonlinear Science*, 28(1):147–191, 2017. doi: 10.1007/s00332-017-9405-2.
- M. Nestler, I. Nitschke, and A. Voigt. A finite element approach for vector- and tensor-valued surface pdes. *Journal of Computational Physics*, 389:48–61, 2019. doi: 10.1016/j.jcp.2019.03.006.
- I. Nitschke. Diskretes Äußeres Kalkül (DEC) auf Oberflächen ohne Rand. Master’s thesis, Technische Universität Dresden, Dresden, 2014.
- M. A. Olshanskii and A. Reusken. Trace finite element methods for PDEs on surfaces. In *Lecture Notes in Computational Science and Engineering*, pages 211–258. Springer International Publishing, 2017. doi: 10.1007/978-3-319-71431-8_7.
- P.-O. Persson. *Mesh Generation for Implicit Geometries*. PhD thesis, Department of Mathematics, MIT, Dec 2004.

- P.-O. Persson and G. Strang. A simple mesh generator in MATLAB. *SIAM Review*, 46(2):329–345, 2004. doi: 10.1137/S0036144503429121.
- S. Praetorius. Dune-Vtk – grid reader and writer for the vtk file format, 2019. URL <https://gitlab.dune-project.org/extensions/dune-vtk>.
- S. Praetorius. Dune-CurvedGeometry – parametrizations of curved geometries, 2020. URL <https://gitlab.mn.tu-dresden.de/iwr/dune-curvedgeometry>.
- S. Praetorius and F. Stenger. Dune-CurvedGrid – meta grid for wrapping element geometries into a curved geometries, 2020a. URL <https://gitlab.mn.tu-dresden.de/iwr/dune-curvedgrid>.
- S. Praetorius and F. Stenger. Example codes of this manuscript collected in a dune module, 2020b. URL <https://gitlab.mn.tu-dresden.de/iwr/dune-curvedgrid-examples>.
- S. Praetorius and F. Stenger. Dune-Gmsh4 – grid reader and writer for the gmsh-4 file format, 2020c. URL <https://gitlab.mn.tu-dresden.de/iwr/dune-gmsh4>.
- A. Rätz and A. Voigt. PDE’s on surfaces — a diffuse interface approach. *Communications in Mathematical Sciences*, 4(3):575–590, 09 2006. URL <https://projecteuclid.org:443/euclid.cms/1175797557>.
- O. Sander, T. Koch, N. Schröder, and B. Flemisch. The DUNE-FOAMGRID implementation for surface and network grids. *Archive of Numerical Software*, 5(1):217–244, 2017. doi: 10.11588/ans.2017.1.28490.
- J. Schöberl. NETGEN an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and Visualization in Science*, 1(1):41–52, Jul 1997. doi: 10.1007/s007910050004.
- J. Schöberl. C++11 implementation of finite elements in NGSolve. Technical report, Institute for analysis and scientific computing, Vienna University of Technology, 2014. URL <https://ngsolve.org>.
- F. Stenger. meshconv: a tool for various mesh-conversions and mesh-transformations., 2020. URL <https://gitlab.mn.tu-dresden.de/iwr/meshconv>. v3.20.
- S. Valette and J.-M. Chassery. Approximated centroidal Voronoi diagrams for uniform polygonal mesh coarsening. *Computer Graphics Forum*, 23(3):381–389, 2004. doi: 10.1111/j.1467-8659.2004.00769.x.