

Function space bases in the dune-functions module

Christian Engwer¹, Carsten Gräser², Steffen Müthing³, and
Oliver Sander⁴

¹Universität Münster, Institute for Computational und Applied Mathematics,
christian.engwer@uni-muenster.de

²Freie Universität Berlin, Institut für Mathematik, graeser@mi.fu-berlin.de

³Universität Heidelberg, Institut für Wissenschaftliches Rechnen,
steffen.muething@iwr.uni-heidelberg.de

⁴TU Dresden, Institute for Numerical Mathematics, oliver.sander@tu-dresden.de

June 26, 2018

This work is licensed under a
Creative Commons Attribution-NoDerivatives 4.0 International License.



The full license text is available here:
<https://creativecommons.org/licenses/by-nd/4.0/legalcode>

The `dune-functions` DUNE module provides interfaces for functions and function space bases. It forms one abstraction level above grids, shape functions, and linear algebra, and provides infrastructure for full discretization frameworks like `dune-pdelab` and `dune-fem`. This document describes the function space bases provided by `dune-functions`. These are based on an abstract description of bases for product spaces as trees of simpler bases. From this description, many different numberings of degrees of freedom by multi-indices can be derived in a natural way. We describe the abstract concepts, document the programmer interface, and give a complete example program that solves the stationary Stokes equation using Taylor–Hood elements.

Introduction

The core modules of the DUNE software system focus on low-level infrastructure for implementations of simulation algorithms for partial differential equations. Modules like `dune-grid` and `dune-istl` provide programmer interfaces (APIs) to finite element grids and sparse linear algebra, respectively, but little more. Actual finite element functions only appear in the `dune-localfunctions` module, which deals with discrete function spaces on single grid elements exclusively.

On top of these core modules, various other modules in the DUNE ecosystem implement finite element and finite volume assemblers and solvers, and the corresponding discrete function spaces. The most prominent ones are `dune-pdelab`¹ and `dune-fem`,² but smaller ones like `dune-fufem`³ exist as well. The functionality of these modules overlaps to a considerable extent, even though each such module has a different focus.

The `dune-functions` module was written to partially overcome this fragmentation, and to unify parts of the competing implementations. It picks a well-defined aspect of finite element assembly—finite element spaces and functions—and, in the DUNE spirit, provides abstract interfaces that try to be both extremely flexibly and efficient. The hope is that other implementations of the same functionality eventually replace their implementations by a dependence on `dune-functions`. Indeed, at the time of writing at least `dune-pdelab` and `dune-fufem` are in the process of migrating, and have stated their clear intention to complete this migration eventually.

Of the two parts of `dune-functions` functionality, the APIs for discrete and closed-form functions have already been described in a separate paper [6]. The present document focuses on spaces of discrete functions. However, the central concept is not the function space itself, but rather the *basis* of the function space. This is because even though finite element spaces play a central role in theoretical considerations of the finite element method, actual computations use coefficient vectors, which are defined with respect to particular bases. Also, for various finite element spaces, more than one basis is used in practice. For example, the space of second-order Lagrangian finite elements is used both with the nodal (Lagrange) basis [3], and with the hierarchical basis [1]. Discontinuous Galerkin spaces can be described in terms of Lagrange bases, monomial bases, Legendre bases, and more [7]. It is therefore important to be able to distinguish these different representations of the same space in the application code. For these reasons, the main `dune-functions` interface represents a basis of a discrete function space, and not the space itself.

Finite element function space bases frequently exhibit a fair amount of structure. In particular, vector-valued and mixed finite element spaces can be written as products of simpler spaces. Even more, such spaces have a natural structure as a tree, with scalar-valued or otherwise irreducible spaces forming the leaves, and products forming the inner nodes. The `dune-functions` module allows to systematically construct new bases by multiplication of existing bases. The resulting tree structure is reproduced as type information in the code. This tree construction of finite element spaces has first

¹<https://dune-project.org/modules/dune-pdelab>

²<https://dune-project.org/modules/dune-fem>

³<https://dune-project.org/modules/dune-fufem>

been systematically worked out in [11].

For the basis functions in such a non-trivial tree structure, there is no single canonical way to index them. Keeping all degrees of freedom in a single standard array would require indexing by a contiguous, zero-starting set of natural numbers. On the other hand, from the tree structure of the basis follows a natural indexing by multi-indices, which can be used to address nested vector and matrix data types, like the ones provided by `dune-istl`. Closer inspection reveals that these two possibilities are just two extreme cases of a wider scale of indexing rules. The `dune-functions` module therefore provides a systematic way to construct such rules. While some of them are somewhat contrived, many others really are useful in applications.

This document describes Version 2.6 of the `dune-functions` module. The module is hosted on the DUNE project homepage www.dune-project.org. Installation instructions and an up-to-date class documentation can be found there.

Contents

1	Function space bases	4
1.1	Trees of function spaces	5
1.2	Trees of function space bases	6
1.3	Indexing basis functions by multi-indices	7
1.4	Strategy-based construction of multi-indices	11
1.5	Localization to single grid elements	16
2	Programmer interface for function space bases	17
2.1	The interface for a global function space basis	18
2.2	The user interface for a localized basis	20
2.3	The user interface of the tree of local bases	22
2.4	Multi-indices	24
3	Constructing trees of function space bases	27
3.1	Basis implementations provided by <code>dune-functions</code>	28
3.2	Combining bases into trees	29
4	Treating subtrees as separate bases	33
5	Combining global bases and coefficient vectors	34
5.1	Vector backends	35
5.2	Interpreting coefficient vectors as finite element functions	36
5.3	Interpolation	37
6	Example: Solving the Stokes equation with <code>dune-functions</code>	40
6.1	The Stokes equation	40
6.2	The driven-cavity benchmark	41
6.3	Implementation	42
	References	53

1 Function space bases

Before we can explain the programmer interface for bases of discrete function spaces in Chapter 2, we need to say a few words about how these bases can be endowed with an abstract tree structure. Readers who are only interested in finite element spaces of scalar-valued functions may try to proceed directly to Chapter 2. They should only know that whenever a local finite element tree is mentioned there, this tree consists of a single node only, which is the local finite element basis. Similarly, for a scalar finite element space the tree of multi-indices used to index the basis functions simply represents a contiguous, zero-starting set of natural numbers.

1.1 Trees of function spaces

Throughout this paper we assume that we have a single fixed domain Ω , and all function spaces that we consider are defined on this domain. The focus is on spaces of functions that are piecewise polynomial with respect to a grid, but that is not actually required yet.

For a set R we denote by $R^\Omega := \{f : \Omega \rightarrow R\}$ the set of all functions mapping from Ω to R . For domains $\Omega \subset \mathbb{R}^d$ we write $P_k(\Omega) \subset \mathbb{R}^\Omega$ for the space of all scalar-valued continuous piecewise polynomials of degree at most k on Ω with respect to some given triangulation. We will omit the domain if it can be inferred from the context.

Considering the different finite element spaces that appear in the literature, there are some that we will call *irreducible*. By this term we mean all bases of scalar-valued functions, but also others like the Raviart–Thomas basis that cannot easily be written as a combination of simpler bases. Many other finite element spaces arise naturally as a combination of simpler ones. There are primarily two ways how two vector spaces V and W can be combined to form a new one: sums and products.⁴

For sums, both spaces need to have the same range space R , and thus both be subspaces of R^Ω . Then the vector space sum

$$V + W := \{v + w : v \in V, w \in W\}$$

in R^Ω will have that same range space. For example, a P_2 -space can be viewed as a P_1 -space plus a hierarchical extension spanned by bubble functions [1]. XFEM spaces [10] are constructed by adding particular weighted Heaviside functions to a basic space to capture free discontinuities. The `dune-functions` module does not currently support constructing sums of finite element bases, but this may be added in later versions.

The second way to construct finite element spaces from simpler ones uses Cartesian products. Let $V \subset (\mathbb{R}^{r_1})^\Omega$ and $W \subset (\mathbb{R}^{r_2})^\Omega$ be two function spaces. Then we define the product of V and W as

$$V \times W := \{(v, w) : v \in V, w \in W\}.$$

Functions from this space take values in $\mathbb{R}^{r_1} \times \mathbb{R}^{r_2} = \mathbb{R}^{r_1+r_2}$. It should be noted that the Cartesian product of vector spaces must not be confused with the tensor product of these spaces. Rather, the k -th power of a single space can be viewed as the tensor product of that space with \mathbb{R}^k , i.e.,

$$(V)^k = \underbrace{V \times \cdots \times V}_{k\text{-times}} = \mathbb{R}^k \otimes V.$$

The product operation allows to build vector-valued and mixed finite element spaces of arbitrary complexity. For example, the space of first-order Lagrangian finite elements with values in \mathbb{R}^3 can be seen as the product $P_1 \times P_1 \times P_1$. The lowest-order Taylor–Hood element is the product $P_2 \times P_2 \times P_2 \times P_1$ of $P_2 \times P_2 \times P_2$ for the velocities

⁴While these are also called internal and external sums, respectively, we stick to the terminology *sum* and *product* in the following.

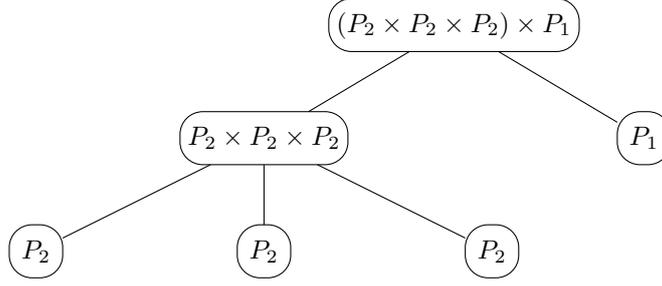


Figure 1: Function space tree of the Taylor–Hood space $(P_2 \times P_2 \times P_2) \times P_1$

with P_1 for the pressure. More factor bases can be included easily, if necessary. We call such products of spaces *composite spaces*.

In the Taylor–Hood space, the triple $P_2 \times P_2 \times P_2$ forms a semantic unit—it contains the components of a velocity field. The associativity of the product allows to write the Taylor–Hood space as $(P_2 \times P_2 \times P_2) \times P_1$, which makes the semantic relationship clearer. Grouped expressions of this type are conveniently visualized as tree structures. This suggests to interpret composite finite element spaces as tree structures. In these structures, leaf nodes represent scalar or otherwise irreducible spaces, and inner nodes represent products of their children. Subtrees then represent composite finite element spaces. Figure 1 shows the Taylor–Hood finite element space in such a tree representation. Note that in this document all trees are *rooted* and *ordered*, i.e., they have a dedicated root note, and the children of each node have a fixed given ordering. Based on this child ordering we associate to each child the corresponding zero-based index.

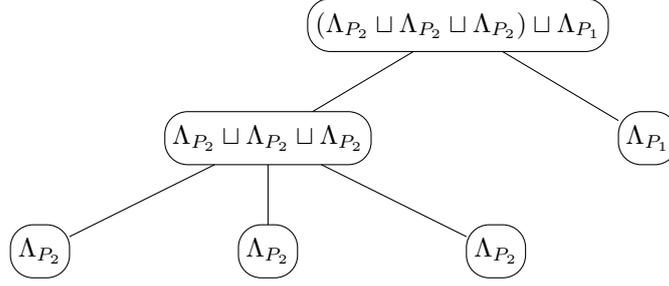
While the inner tree nodes may initially appear like useless artifacts of the tree representation, they are often extremely useful because we can treat the subtrees rooted in those nodes as individual trees in their own right. This often allows to reuse existing algorithms that expect to operate on those subtrees in more complex settings.

1.2 Trees of function space bases

The multiplication of finite-dimensional spaces naturally induces a corresponding operation on bases of such spaces. We introduce a generalized tensor product notation: Consider linear ranges R_0, \dots, R_{m-1} of function spaces $R_0^\Omega, \dots, R_{m-1}^\Omega$, and the i -th canonical basis vector \mathbf{e}_i in \mathbb{R}^m . Then

$$\mathbf{e}_i \otimes f := (0, \dots, 0, \underbrace{f}_{i\text{-th entry}}, 0, \dots, 0) \in \prod_{j=0}^{m-1} (R_j^\Omega) = \left(\prod_{j=0}^{m-1} R_j \right)^\Omega,$$

where 0 in the j -th position denotes the zero-function in R_j^Ω . Let Λ_i be a function space basis of the space $V_i = \text{span } \Lambda_i$ for $i = 0, \dots, m-1$. Then a natural basis Λ of

Figure 2: Function space basis tree of the Taylor-Hood space $(P_2 \times P_2 \times P_2) \times P_1$

the product space

$$V_0 \times \cdots \times V_{m-1} = \prod_{i=0}^{m-1} V_i = \prod_{i=0}^{m-1} \text{span } \Lambda_i$$

is given by

$$\Lambda = \Lambda_0 \sqcup \cdots \sqcup \Lambda_{m-1} = \bigsqcup_{i=0}^{m-1} \Lambda_i := \bigcup_{i=0}^{m-1} \mathbf{e}_i \otimes \Lambda_i. \quad (1)$$

The product $\mathbf{e}_i \otimes \Lambda_i$ is to be understood element-wise, and the “disjoint union” symbol \sqcup is used here as a simple short-hand notation for (1) and not to be understood as an associative binary operation. Using this new notation we have

$$\text{span } \Lambda = \text{span}(\Lambda_0 \sqcup \cdots \sqcup \Lambda_{m-1}) = (\text{span } \Lambda_0) \times \cdots \times (\text{span } \Lambda_{m-1}).$$

Similarly to the case of function spaces, bases can be interpreted as trees. If we associate a basis Λ_V to each space V in the function space tree, then the induced natural function space basis tree is obtained by simply replacing V by Λ_V in each node. For the Taylor-Hood basis this leads to the tree depicted in Figure 2.

1.3 Indexing basis functions by multi-indices

To work with the basis of a finite element space, the basis functions need to be indexed. Indexing the basis functions is what allows to address the corresponding vector and matrix coefficients in suitable vector and matrix data structures. In simple cases, indexing means simply enumerating the basis functions with natural numbers, but for many applications hierarchically structured matrix and vector data structures are more natural or efficient. This leads to the idea of hierarchically structured multi-indices.

Definition 1 (Multi-indices). *A tuple $I \in \mathbb{N}_0^k$ for some $k \in \mathbb{N}_0$ is called a multi-index of length k , and we write $|I| := k$. The set of all multi-indices is denoted by $\mathcal{N} = \bigcup_{k \in \mathbb{N}_0} \mathbb{N}_0^k$.*

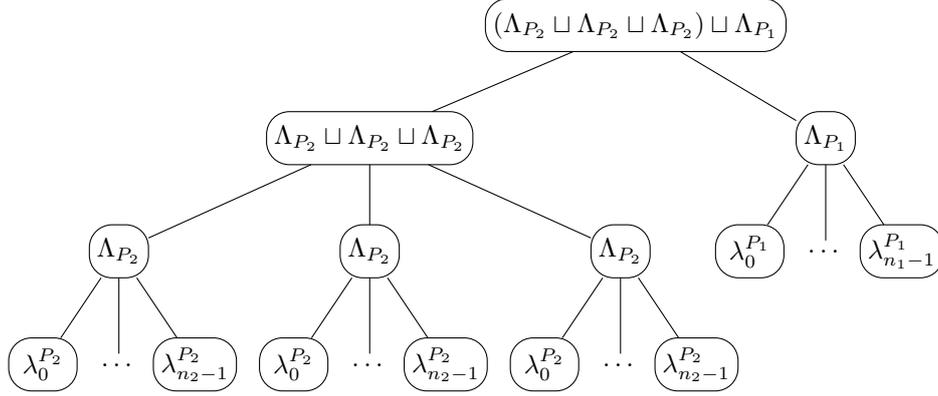


Figure 3: Tree of basis vectors for the Taylor–Hood basis

To establish some structure in a set of multi-indices it is convenient to consider prefixes.

Definition 2 (Multi-index prefixes).

1. If $I \in \mathcal{N}$ takes the form $I = (I^0, I^1)$ for $I^0, I^1 \in \mathcal{N}$, then we call I^0 a prefix of I . If additionally $|I^1| > 0$, then we call I^0 a strict prefix of I .
2. For $I, I^0 \in \mathcal{N}$ and a set $\mathcal{M} \subset \mathcal{N}$:
 - a) We write $I = (I^0, \dots)$, if I^0 is a prefix of I ,
 - b) we write $I = (I^0, \bullet, \dots)$, if I^0 is a strict prefix of I ,
 - c) we write $(I^0, \dots) \in \mathcal{M}$, if I^0 is a prefix of some $I \in \mathcal{M}$,
 - d) we write $(I^0, \bullet, \dots) \in \mathcal{M}$, if I^0 is a strict prefix of some $I \in \mathcal{M}$.

It is important to note that the multi-indices from a given set do not necessarily all have the same length. For an example, Figure 3 illustrates the set of all basis functions by extending the basis tree of Figure 2 by leaf nodes for individual basis functions. A possible indexing of the basis functions of the Taylor–Hood basis Λ_{TH} then uses multi-indices of the form $(0, i, j)$ for velocity components, and $(1, k)$ for pressure components. For the velocity multi-indices $(0, i, j)$, the $i = 0, \dots, 2$ determines the component of the velocity vector field, and the $j = 0, \dots, n_2 - 1 := |\Lambda_{P_2}| - 1$ determines the number of the scalar P_2 basis function that determines this component. For the pressure multi-indices $(0, k)$ the $k = 0, \dots, n_1 - 1 := |\Lambda_{P_1}| - 1$ determines the number of the P_1 basis function for the scalar P_1 function that determines the pressure.

It is evident that the complete set of these multi-indices can again be associated to a rooted tree. In this tree, the multi-indices correspond to the leaf nodes, their strict prefixes correspond to interior nodes, and the multi-index digits labeling the edges are

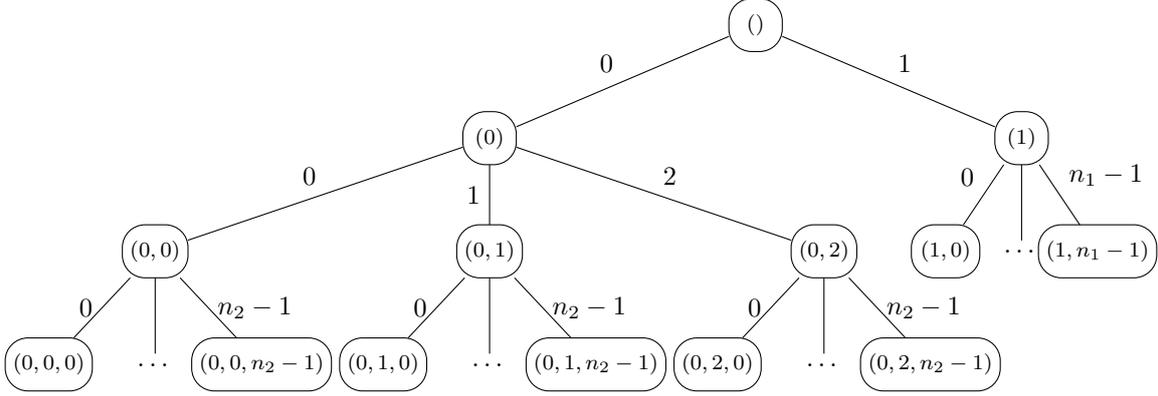


Figure 4: Index tree for the Taylor–Hood basis inherited from the basis tree

the indices of the children within the ordered tree. Prefixes can be interpreted as paths from the root to a given node.

This latter fact can be seen as the defining property of index trees. Indeed, a set of multi-indices (together with all its strict prefixes) forms a tree as long as it is consistent in the sense that the multi-indices can be viewed as the paths to the leaves in an ordered tree. That is, the children of each node are enumerated using consecutive zero-based indices and paths to the leaves (i.e., the multi-indices) are built by concatenating those indices starting from the root and ending in a leaf. Since the full structure of this tree is encoded in the multi-indices associated to the leaves we will—by a slight abuse of notation—call the set of multi-indices itself a tree from now on.

Definition 3. A set $\mathcal{I} \subset \mathcal{N}$ is called an index tree if for any $(I, i, \dots) \in \mathcal{I}$ there are also $(I, 0, \dots), (I, 1, \dots), \dots, (I, i - 1, \dots) \in \mathcal{I}$, but $I \notin \mathcal{I}$.

The index tree for the example indexing of the Taylor–Hood basis given above is shown in Figure 4.

Definition 4. Let $(I, \dots) \in \mathcal{I}$, i.e., I is a prefix of multi-indices in \mathcal{I} . Then the size of \mathcal{I} relative to I is given by

$$\deg_{\mathcal{I}}^+[I] := \max\{k : \exists(I, k, \dots) \in \mathcal{I}\} + 1. \quad (2)$$

In terms of the ordered tree associated with \mathcal{I} this corresponds to the out-degree of I , i.e., the number of direct children of the node indexed by I .

Using the idea of multi-index trees, an indexing of a function space basis is an injective map from the leaf nodes of a tree of basis functions to the leaves of an index tree.

Definition 5. Let M be a finite set and $\iota : M \rightarrow \mathcal{N}$ an injective map whose range $\iota(M)$ forms an index tree. Then ι is called an index map for M . The index map is called uniform if additionally $\iota(M) \subset \mathbb{N}_0^k$ for some $k \in \mathbb{N}$, and flat if $\iota(M) \subset \mathbb{N}_0$.

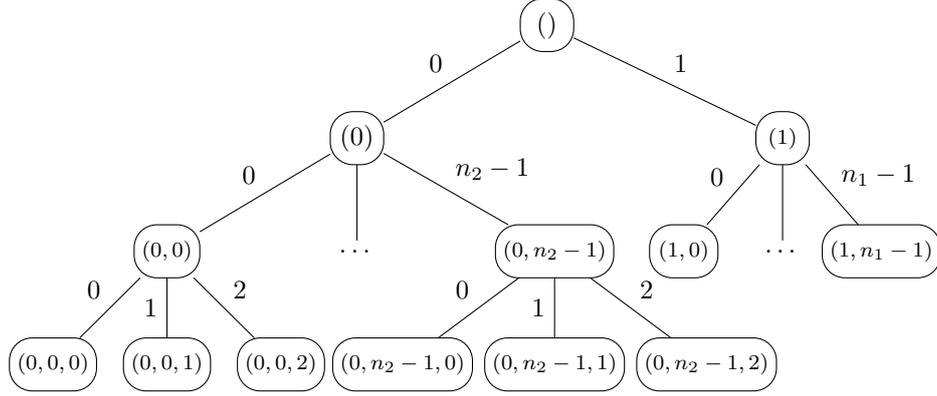


Figure 5: Index tree for Taylor–Hood with blocking of velocity components

Continuing the Taylor–Hood example, if all basis functions $\Lambda_{\text{TH}} = \{\lambda_I\}$ of the whole finite element tree are indexed by multi-indices of the above given form, and if X is a coefficient vector that has a compatible hierarchical structure, then a finite element function (v_h, p_h) with velocity v_h and pressure p_h defined by the coefficient vector X is given by

$$(v_h, p_h) = \sum_{i=0}^2 \sum_{j=0}^{n_2-1} X_{(0,i,j)} \lambda_{(0,i,j)} + \sum_{k=0}^{n_1-1} X_{(1,k)} \lambda_{(1,k)}, \quad (3)$$

with basis functions

$$\lambda_{(0,i,j)} = \mathbf{e}_0 \otimes (\mathbf{e}_i \otimes \lambda_j^{P_2}), \quad i = 0, 1, 2, \quad \text{and} \quad \lambda_{(1,k)} = \mathbf{e}_1 \otimes \lambda_k^{P_1}.$$

Introducing the corresponding index map $\iota : \Lambda_{\text{TH}} \rightarrow \mathcal{N}$ with $\iota(\lambda_I) = I$ on the set Λ_{TH} of all basis functions we can write this in compact form as

$$(v_h, p_h) = \sum_{\lambda \in \Lambda_{\text{TH}}} X_{\iota(\lambda)} \lambda = \sum_{I \in \iota(\Lambda_{\text{TH}})} X_I \lambda_I.$$

Alternatively the individual velocity and pressure fields v_h and p_h are given by

$$v_h = \sum_{i=0}^2 \sum_{j=0}^{n_2-1} X_{(0,i,j)} (\mathbf{e}_i \otimes \lambda_j^{P_2}), \quad p_h = \sum_{k=0}^{n_1-1} X_{(1,k)} \lambda_k^{P_1}.$$

In the previous example, the index tree was isomorphic to the basis function tree depicted in Figure 3. However, one may also be interested in constructing multi-indices that do not mimic the structure of the basis function tree: For example, to increase data locality in assembled matrices for the Taylor–Hood basis it may be preferable to group

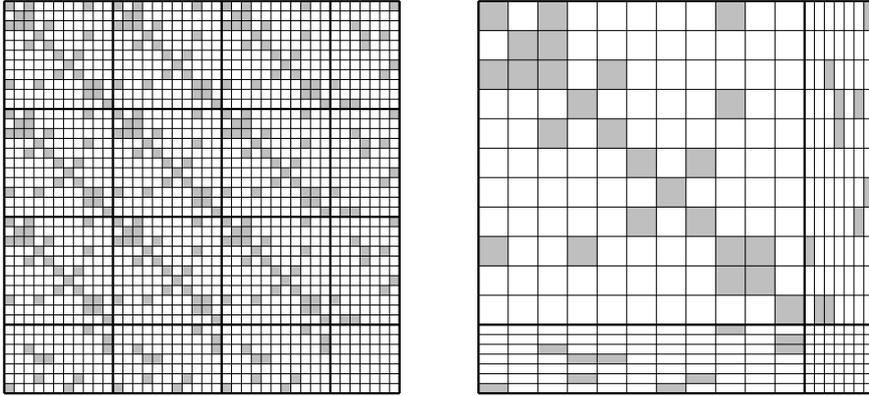


Figure 6: Two matrix occupation patterns for different indexings of the Taylor–Hood bases. Left: Corresponding to the index tree of Figure 4. Right: Corresponding to the index tree of Figure 5.

all velocity degrees of freedom corresponding to a single P_2 basis function together, i.e., to use the index $(0, j, i)$ for the j -th P_2 basis function for the i -th component. The corresponding alternative index tree is shown in Figure 5. Figure 6 shows the corresponding layouts of a hierarchical stiffness matrix.

Alternatively, the case of indexing all basis functions from the Taylor–Hood basis with a single natural number can be represented by an index tree with $3n_2 + n_1$ leaf nodes all directly attached to a single root. Different variations of such a tree differ by how the degrees of freedom are ordered.

1.4 Strategy-based construction of multi-indices

Let Λ be the set of basis functions of a finite element basis tree. In principle, **dune-functions** allows any indexing scheme that is given by an index map, i.e., any map $\iota : \Lambda \rightarrow \mathcal{N}$ that is injective and whose range $\iota(\Lambda)$ is an index tree. In practice, out of this large set of maps, **dune-functions** allows to construct the most important ones systematically using certain transformation rules.

Consider a tree of function space bases in the sense of Section 1.2. We want to construct an indexing for this tree, that is an index tree \mathcal{I} and a bijection ι from the set of all basis functions Λ to the multi-indices in \mathcal{I} . The construction proceeds recursively. To describe it, we assume in the following that Λ is a node in the function space basis tree, i.e., it is the set of all basis functions corresponding to a node $V := \text{span } \Lambda$ in the function space tree.

To end the recursion, we assume that an index map $\iota : \Lambda \rightarrow \mathcal{N}$ is given if $V = \text{span } \Lambda$ is a leaf node of the function space tree. The most obvious choice would be a flat zero-based index of the basis functions of Λ . However, other choices are possible. For example, in case of a discontinuous finite element space, each basis function $\lambda \in \Lambda$

could also be associated to a two-digit multi-index $\iota(\lambda) = (i, k)$, where i is the index of the grid element that forms the support of λ , and k is the index of λ within this element.

For the actual recursion, if Λ is any non-leaf node in the function space basis tree, then it takes the form

$$\Lambda = \Lambda_0 \sqcup \cdots \sqcup \Lambda_{m-1} = \bigcup_{i=0}^{m-1} \mathbf{e}_i \otimes \Lambda_i,$$

where $\Lambda_0, \dots, \Lambda_{m-1}$ are the direct children of Λ , i.e., the sets of basis functions of the child spaces $\{\text{span } \Lambda_i\}_{i=0, \dots, m-1}$ of the product space

$$\text{span } \Lambda = \text{span}(\Lambda_0 \sqcup \cdots \sqcup \Lambda_{m-1}) = (\text{span } \Lambda_0) \times \cdots \times (\text{span } \Lambda_{m-1}).$$

For the recursive construction we assume that an index map $\iota_i : \Lambda_i \rightarrow \mathcal{N}$ on Λ_i is given for any $i = 0, \dots, m-1$. The task is to construct an index map $\iota : \Lambda \rightarrow \mathcal{N}$ from the maps ι_i . In the following we describe four strategies to achieve this; all have been implemented in **dune-functions**. When reading about these strategies, remember that any $\lambda \in \Lambda$ has a unique representation $\lambda = \mathbf{e}_i \otimes \hat{\lambda}$ for $i \in \{0, \dots, m-1\}$ and some $\hat{\lambda} \in \Lambda_i$. It will be necessary to distinguish the special case that all children Λ_i are identical.

Definition 6. *An inner node Λ will be called power node if all of its children Λ_i are identical and equipped with identical index maps ι_i . An inner node that is not a power node is called composite node.*

This definition is needed because some of the following strategies can only be applied to power nodes.

- **BlockedLexicographic:** This strategy prepends the child index to the multi-index within the child basis. That is, the index map $\iota : \Lambda \rightarrow \mathcal{N}$ is given by

$$\iota(\mathbf{e}_i \otimes \hat{\lambda}) = (i, \iota_i(\hat{\lambda})).$$

It is straightforward to show that ι is always an index map for Λ . To demonstrate the strategy the following table shows the multi-indices at inner nodes, when the basis functions of the subtrees $\Lambda_0, \Lambda_1, \dots$ are labeled by multi-indices $(I^0), (I^1), \dots$ for Λ_0 , $(K^0), (K^1), \dots$ for Λ_1 , and so on.

indices for Λ_0	indices for Λ_1	...	indices for Λ
$\iota_0(\hat{\lambda}_{0,0}) = (I^0)$			$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,0}) = (0, I^0)$
$\iota_0(\hat{\lambda}_{0,1}) = (I^1)$			$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,1}) = (0, I^1)$
	$\iota_1(\hat{\lambda}_{1,0}) = (K^0)$		$\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,0}) = (1, K^0)$
	$\iota_1(\hat{\lambda}_{1,1}) = (K^1)$		$\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,1}) = (1, K^1)$
	$\iota_1(\hat{\lambda}_{1,2}) = (K^2)$		$\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,2}) = (1, K^2)$
	

- **BlockedInterleaved:** This strategy is only well-defined for power nodes. It appends the child index to the multi-index within the child basis. That is, the index map $\iota : \Lambda \rightarrow \mathcal{N}$ is given by

$$\iota(\mathbf{e}_i \otimes \hat{\lambda}) = (\iota_i(\hat{\lambda}), i).$$

An example is given in the following table:

indices for Λ_0	indices for Λ_1	...	indices for Λ
$\iota_0(\hat{\lambda}_{0,0}) = (I^0)$	$\iota_1(\hat{\lambda}_{1,0}) = (I^0)$		$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,0}) = (I^0, 0)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,0}) = (I^0, 1)$
$\iota_0(\hat{\lambda}_{0,1}) = (I^1)$	$\iota_1(\hat{\lambda}_{1,1}) = (I^1)$ $\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,1}) = (I^1, 0)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,1}) = (I^1, 1)$
$\iota_0(\hat{\lambda}_{0,2}) = (I^2)$	$\iota_1(\hat{\lambda}_{1,2}) = (I^2)$ $\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,2}) = (I^2, 0)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,2}) = (I^2, 1)$
	

To see that this strategy does not work for general composite nodes, consider $\iota_0(\Lambda_0) = \{0\}$ and $\iota_1(\Lambda_1) = \{(0, 0)\}$. Then $\iota(\Lambda) = \{(0, 0), (0, 0, 1)\}$ which is not an index tree.

Unlike the previous two strategies, the following two do not introduce new multi-index digits. Such strategies are called *flat*.

- **FlatLexicographic:** This strategy merges the roots of all index tree $\iota_i(\Lambda_i)$ into a single new one. Assume that we split the multi-index $\iota_i(\hat{\lambda})$ according to

$$\iota_i(\hat{\lambda}) = (i_0, I), \quad (4)$$

where $i_0 \in \mathbb{N}_0$ is the first digit. The index map $\iota : \Lambda \rightarrow \mathcal{N}$ is then given by

$$\iota(\mathbf{e}_i \otimes \hat{\lambda}) = (L_i + i_0, I),$$

where the offset L_i for the first digit is computed by

$$L_i = \sum_{j=0}^{i-1} \deg_{\iota_j(\Lambda_j)}^+[\()].$$

This construction offsets the first digits of the multi-indices of all basis functions from Λ_j with $j > 0$ such that they form a consecutive sequence. This guarantees that ι is always an index map for Λ . An example is given in the following table:

indices for Λ_0	indices for Λ_1	...	indices for Λ
$\iota_0(\hat{\lambda}_{0,0}) = (0, I^0)$ $\iota_0(\hat{\lambda}_{0,1}) = (1, I^1)$	$\iota_1(\hat{\lambda}_{1,0}) = (0, K^0)$ $\iota_1(\hat{\lambda}_{1,1}) = (0, K^1)$ $\iota_1(\hat{\lambda}_{1,2}) = (1, K^2)$...	$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,0}) = (0, I^0)$ $\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,1}) = (1, I^1)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,0}) = (2, K^0)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,1}) = (2, K^1)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,2}) = (3, K^2)$
	

The digit zero deliberately appears twice in the column for Λ_1 , to demonstrate that a consecutive first digit is not required.

- **FlatInterleaved:** This strategy again only works for power nodes. It also merges the roots of all child index trees $\iota_i(\Lambda_i)$ into a single one, but it interleaves the children. Again using the splitting $\iota_i(\hat{\lambda}) = (i_0, I)$ introduced in (4), the index map $\iota : \Lambda \rightarrow \mathcal{N}$ is given by

$$\iota(\mathbf{e}_i \otimes \hat{\lambda}) = (i_0 m + i, I),$$

where the fixed stride m is given by the number of children of Λ . The following table shows an example:

indices for Λ_0	indices for Λ_1	...	indices for Λ
$\iota_0(\hat{\lambda}_{0,0}) = (0, I^0)$	$\iota_1(\hat{\lambda}_{1,0}) = (0, I^0)$		$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,0}) = (0, I^0)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,0}) = (1, I^0)$
	
$\iota_0(\hat{\lambda}_{0,1}) = (1, I^1)$	$\iota_1(\hat{\lambda}_{1,1}) = (1, I^1)$		$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,1}) = (m + 0, I^1)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,1}) = (m + 1, I^1)$
	
$\iota_0(\hat{\lambda}_{0,2}) = (2, I^2)$	$\iota_1(\hat{\lambda}_{1,2}) = (2, I^2)$		$\iota(\mathbf{e}_0 \otimes \hat{\lambda}_{0,2}) = (2m + 0, I^2)$ $\iota(\mathbf{e}_1 \otimes \hat{\lambda}_{1,2}) = (2m + 1, I^2)$
	

Again, for this interleaved strategy, ι may not be an index map for general composite nodes.

These four strategies are offered by `dune-functions`, but there are others that are sometimes useful. Experimentally, `dune-functions` therefore also provides a way to use self-implemented custom rules.

To further illustrate the four index transformation strategies, we return to the Taylor–Hood example. While the indexing schemes proposed for this example so far were introduced in an ad-hoc way, we will now systematically apply the above given strategies. Recall that the Taylor–Hood basis is denoted by

$$\Lambda_{\text{TH}} = (\Lambda_{P_2} \sqcup \Lambda_{P_2} \sqcup \Lambda_{P_2}) \sqcup \Lambda_{P_1}.$$

For the bases $\Lambda_{P_1}, \Lambda_{P_2}$ of the elementary spaces P_1, P_2 we consider fixed given flat index maps

$$\iota_{P_1}(\Lambda_{P_1}) \rightarrow \mathbb{N}_0, \quad \iota_{P_2}(\Lambda_{P_2}) \rightarrow \mathbb{N}_0.$$

These are typically constructed by enumerating the grid entities the basis functions are associated to. Then the interior product space basis

$$\Lambda_V = \Lambda_{P_2} \sqcup \Lambda_{P_2} \sqcup \Lambda_{P_2}$$

	BL(BL)	BL(BI)	BL(FL)	BL(FI)	FL(BL)	FL(BI)	FL(FL)	FL(FI)
$v_{x_0,0}$	(0, 0, 0)	(0, 0, 0)	(0, 0)	(0, 0 + 0)	(0, 0)	(0, 0)	(0)	(0 + 0)
$v_{x_0,1}$	(0, 0, 1)	(0, 1, 0)	(0, 1)	(0, 3 + 0)	(0, 1)	(1, 0)	(1)	(3 + 0)
$v_{x_0,2}$	(0, 0, 2)	(0, 2, 0)	(0, 2)	(0, 6 + 0)	(0, 2)	(2, 0)	(2)	(6 + 0)
$v_{x_0,3}$	(0, 0, 3)	(0, 3, 0)	(0, 3)	(0, 9 + 0)	(0, 3)	(3, 0)	(3)	(9 + 0)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$v_{x_1,0}$	(0, 1, 0)	(0, 0, 1)	(0, $n_2 + 0$)	(0, 0 + 1)	(1, 0)	(0, 1)	($n_2 + 0$)	(0 + 1)
$v_{x_1,1}$	(0, 1, 1)	(0, 1, 1)	(0, $n_2 + 1$)	(0, 3 + 1)	(1, 1)	(1, 1)	($n_2 + 1$)	(3 + 1)
$v_{x_1,2}$	(0, 1, 2)	(0, 2, 1)	(0, $n_2 + 2$)	(0, 6 + 1)	(1, 2)	(2, 1)	($n_2 + 2$)	(6 + 1)
$v_{x_1,3}$	(0, 1, 3)	(0, 3, 1)	(0, $n_2 + 3$)	(0, 9 + 1)	(1, 3)	(3, 1)	($n_2 + 3$)	(9 + 1)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$v_{x_2,0}$	(0, 2, 0)	(0, 0, 2)	(0, $2n_2 + 0$)	(0, 0 + 2)	(2, 0)	(0, 2)	($2n_2 + 0$)	(0 + 2)
$v_{x_2,1}$	(0, 2, 1)	(0, 1, 2)	(0, $2n_2 + 1$)	(0, 3 + 2)	(2, 1)	(1, 2)	($2n_2 + 1$)	(3 + 2)
$v_{x_2,2}$	(0, 2, 2)	(0, 2, 2)	(0, $2n_2 + 2$)	(0, 6 + 2)	(2, 2)	(2, 2)	($2n_2 + 2$)	(6 + 2)
$v_{x_2,3}$	(0, 2, 3)	(0, 3, 2)	(0, $2n_2 + 3$)	(0, 9 + 2)	(2, 3)	(3, 2)	($2n_2 + 3$)	(9 + 2)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
p_0	(1, 0)	(1, 0)	(1, 0)	(1, 0)	(3 + 0)	($n_2 + 0$)	($3n_2 + 0$)	($3n_2 + 0$)
p_1	(1, 1)	(1, 1)	(1, 1)	(1, 1)	(3 + 1)	($n_2 + 1$)	($3n_2 + 1$)	($3n_2 + 1$)
p_2	(1, 2)	(1, 2)	(1, 2)	(1, 2)	(3 + 2)	($n_2 + 2$)	($3n_2 + 2$)	($3n_2 + 2$)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

Table 1: Different indexing strategies for the Taylor–Hood basis functions

together with the index map ι_{P_2} is a power node in the sense of Definition 6, while the tree root

$$\Lambda_{\text{TH}} = \Lambda_V \sqcup \Lambda_{P_1}$$

is a composite node. The basis functions for the k -th component of the velocity are denoted by

$$v_{x_k,i} = \mathbf{e}_0 \otimes (\mathbf{e}_k \otimes \lambda_i^{P_2})$$

where $i = 0, \dots, n_2 - 1$ for $n_2 = |\Lambda_{P_2}| = \dim P_2$ whereas the basis functions for the pressure are denoted by

$$p_j = \mathbf{e}_1 \otimes \lambda_j^{P_1}$$

where $j = 0, \dots, n_1 - 1$ for $n_1 = |\Lambda_{P_1}| = \dim P_1$.

As two of the above given strategies can be used for composite nodes, while all four can be applied to power nodes we obtain eight different index maps for the Taylor–Hood basis Λ_{TH} . They are listed in Table 1, where the label $X(Y)$ means that strategy X is used for the outer product and strategy Y for the inner product. For X and Y we use the abbreviations BL (BlockedLexicographic), BI (BlockedInterleaved), FL (FlatLexicographic), and FI (FlatInterleaved). Notice that the index maps depicted in Figure 4 and Figure 5 are reproduced for the strategies BL(BL) and BL(BI), respectively.

1.5 Localization to single grid elements

For the most part, access to finite element bases happens element by element. It is therefore important to consider the restrictions of bases to single grid elements. In contrast to the previous sections we now require that there is a finite element grid for the domain Ω . For simplicity we will assume that all bases consist of functions that are defined piecewise with respect to this grid, but it is actually sufficient to require that the restrictions of all basis functions to elements of the grid can be constructed cheaply.

Consider the restrictions of all basis functions $\lambda \in \Lambda$ of a given tree to a single fixed grid element e . Of these restricted functions, we discard all those that are constant zero functions on e . All others form the *local basis* on e

$$\Lambda|_e := \{\lambda|_e : \lambda \in \Lambda, \text{int}(\text{supp } \lambda) \cap e \neq \emptyset\}.$$

The local basis forms a tree that is isomorphic to the original function space basis tree, with each global function space basis Λ replaced by its local counterpart $\Lambda|_e$.

For a given index map ι of Λ , this natural isomorphism from the global to the local tree naturally induces a localized version of ι given by

$$\iota|_e : \Lambda|_e \rightarrow \mathcal{I}, \quad \iota|_e(\lambda_e) := \iota(\lambda).$$

This is the map that associates shape functions on a given grid element e to the multi-indices of the corresponding global basis functions. Note that the map $\iota|_e$ itself is not an index map in the sense of Definition 5 since $\iota|_e(\Lambda|_e)$ is only a subset of the index tree $\iota(\Lambda)$, and not always an index tree itself.

In order to index the basis functions in $\Lambda|_e$ efficiently we introduce an additional local index map

$$\iota_{\Lambda|_e}^{\text{local}} : \Lambda|_e \rightarrow \mathcal{N},$$

such that $\iota_{\Lambda|_e}^{\text{local}}(\Lambda|_e)$ is an index tree. The index $\iota_{\Lambda|_e}^{\text{local}}(\lambda|_e)$ is called the *local index* of λ (with respect to e). To distinguish it from the indices generated by ι we call $\iota(\lambda)$ the *global index* of λ . The local index is typically used to address the element stiffness matrix. In principle, this indexing can use another non-flat index tree, which does not have to coincide with the index tree for the global basis. This means that the local index of a shape function can again be a multi-index, but the types, lengths and orderings can be completely unrelated to the corresponding global indices. This would allow to use nested types for element stiffness matrices and load vectors. As explained in Chapter 2, the **dune-functions implementation** is fairly restrictive here, and only allows flat local indices, i.e., $\iota_{\Lambda|_e}^{\text{local}}(\Lambda|_e) \subset \mathbb{N}_0$.

In addition, we introduce for each leaf local basis $\hat{\Lambda}|_e$ of the full local basis tree another local index map

$$\iota_{\hat{\Lambda}|_e}^{\text{leaf-local}} : \hat{\Lambda}|_e \rightarrow \mathbb{N}_0.$$

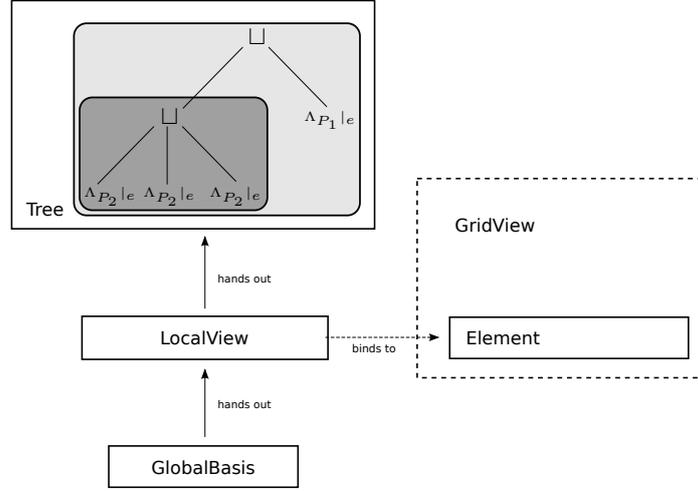


Figure 7: Overview of the classes making up the interface to finite element space bases

As there is no hierarchical structure involved, this index is simply a natural number. The index $\iota_{\hat{\Lambda}|e}^{\text{leaf-local}}(\lambda|e)$ is called the *leaf-local index* of λ (with respect to e).

In an actual programming interface one typically accesses basis functions by indices directly. We will later see that in `dune-functions` the leaf-local index is the shape function index of the `dune-localfunctions` module. Hence the `dune-functions` API needs to implement the map

$$\iota_e^{\text{leaf} \rightarrow \text{local}} := \iota_{\Lambda|e}^{\text{local}} \circ (\iota_{\hat{\Lambda}|e}^{\text{leaf-local}})^{-1}$$

mapping leaf-local indices to local indices and

$$\iota_e^{\text{local} \rightarrow \text{global}} := \iota|e \circ (\iota_{\Lambda|e}^{\text{local}})^{-1}$$

mapping local indices to global multi-indices.

2 Programmer interface for function space bases

The design of the `dune-functions` interface for bases of function spaces follows the ideas of the previous section. The main interface concept are global basis objects that represent trees of function space bases. These trees can be localized to individual elements of the grid. Such a localization provides access to the (tree of) shape functions there, together with the two shape-function index maps $\iota_e^{\text{leaf} \rightarrow \text{local}}$ and $\iota_e^{\text{local} \rightarrow \text{global}}$. The structure of the interface is visualized in Figure 7.

Notation

The `dune-functions` module is implemented in the C++ programming language. All header include paths start with `dune/functions/`, and all code is contained in the namespace `Dune::Functions`. Internally, `dune-functions` depends on the `dune-typetree` module, which implements abstract compile-time tree data structures. The global basis interface described below is not enforced by deriving from specific base classes. Instead, `dune-functions` is based on C++-style duck-typing [9], i.e., any C++ type providing the required interface is a valid implementation of that interface.

Throughout this text we will introduce the programmer interfaces by presenting the interface declaration, explaining its meaning, and giving examples of its usage. In order to distinguish interface declarations from code examples they are formatted differently. Furthermore, implementation defined types and arguments are highlighted. The following shows an example of an interface declaration:

```
// Declaration of type T referring to an implementation-defined type
using T = <implementation defined>;

// Declaration of method foo
T foo(int);

// Declaration of class Bar with implementation-defined constructor arguments
class Bar {
public:
    Bar(<args>);
};
```

In contrast, an example for using this interface would be formatted like this:

```
// Call foo() and store result
T t = foo(1);

// Construct an object of type Bar
auto bar = Bar(t, <more args>);
```

2.1 The interface for a global function space basis

We start by describing the user interface for global bases. Since we are discussing duck-typing interfaces, all class names used below are generic. A tree of global bases is implemented by one class which, in the following, we will call `GlobalBasis`, and which can have an arbitrary number of template parameters. All types and methods listed in the following interface declaration shall be public members of the generic implementation class `GlobalBasis`.

As each basis implementation may require its own specific data for construction, we do not enforce a precise set of constructors.

```
GlobalBasis(<implementation defined>);
```

Each `GlobalBasis` may have one or several constructors with implementation-dependent lists of arguments.

The main feature of a `GlobalBasis` is to give access to basis functions and their indices. Most of this access happens through the localization of the basis to single grid elements (Section 1.5). In the programmer interface, this localization is called `LocalView`. Objects of type `LocalView` are obtained from `GlobalBasis` objects through the method

```
using LocalView = <implementation defined>;
LocalView localView() const;
```

The precise return type of the `localView` method is implementation-dependent. Objects created by the method have undefined state, and need to be attached to individual grid elements in a process called *binding*. The details are explained in Section 2.2.

Several methods of a `GlobalBasis` provide information on the sizes of the bases contained in the tree. The total number of basis functions of the global basis is exported via the method

```
using size_type = <implementation defined>;
size_type dimension() const;
```

This method can be used to allocate vector containers if flat multi-indices are used. However, the information provided by the `dimension` method is generally not sufficient to allocate hierarchical containers to be accessed by more general multi-indices. Therefore, the basis provides more structural information of those multi-indices via the method

```
using SizePrefix = ReservedVector<implementation defined>;
size_type size(const SizePrefix& prefix) const;
```

The parameter `prefix` is a multi-index itself. If \mathcal{I} is the set of all global multi-indices of the basis and `prefix` is a prefix for this set, then `size(prefix)` returns the size $\deg_{\mathcal{I}}^+[\text{prefix}]$ of \mathcal{I} relative to `prefix` defined in (2), i.e., the number of direct children of the node `prefix` in the index tree. If `prefix` is not a prefix for \mathcal{I} the result is undefined. If `prefix` $\in \mathcal{I}$, i.e., the prefix is itself one of the multi-indices then the result is zero. The type `SizePrefix` is always a container of type `ReservedVector` (from the `dune-common` module). More details are given in Section 2.4. Like all other types used in the `GlobalBasis` interface, it is expected to be exported by the implementation class. For convenience there is also the method

```
size_type size() const;
```

returning the same value as `size({})`, i.e., the number of children of the root of the index tree. For a scalar basis, this is again the overall number of basis functions.

Finally, each `GlobalBasis` provides access to the grid view it is defined on by the method

```
const GridView& gridView() const;
```

The corresponding type is exported as `GridView`. If the grid view was modified (e.g., by local grid refinement), the result of calling any method of the basis is undefined until the basis has been explicitly updated. For this, call the method

```
void update(const GridView & gv);
```

which tells the basis to adapt its local state to the new grid view.

2.2 The user interface for a localized basis

The localization of a function space basis to a single grid element is represented by an interface called `LocalView`. Objects of type `LocalView` are returned by the method `GlobalBasis::localView()`, and there is no way to construct such objects directly. All types and methods listed in the following interface declaration are public members of the generic class `LocalView`.

A freshly constructed `LocalView` object is not completely initialized yet. To truly have the object represent the basis localization on a particular element, it must be *bound* to that element. This is achieved by calling

```
using GridView = typename GlobalBasis::GridView;
using Element = typename GridView::template Codim<0>::Entity;
void bind(const Element& e);
```

Once this method has been called, the `LocalView` object is fully set up and can be used. The call may incorporate expensive computations needed to precompute the local basis functions and their global indices. The local view can be bound to another element at any time by calling the `bind` method again. To set the local view back to the unbound state again, call the method

```
void unbind();
```

The local view will store a copy of the element it is bound to, which is accessible via

```
const Element& element() const;
```

A bound `LocalView` object provides information about the size of the local basis at the current element. The total number of basis functions associated to the local view at the current element is returned by

```
using size_type = typename GlobalBasis::size_type;
size_type size() const;
```

In the language of Chapter 1, this method computes the number $|\Lambda|_e$.

To allow preallocation of buffers for local functions, the method

```
size_type maxSize() const;
```

returns the maximum return value of the `size` method for all elements in the grid view associated to the global basis, i.e., it computes $\max_e |\Lambda|_e$. As this information

does not depend on a particular element, the method `maxSize` can even be called in unbound state.

As an example, suppose that `basis` is an object of type `Dune::Functions::TaylorHoodBasis`, which implements the Taylor–Hood basis that has been used for examples in the previous section. The following code loops over all elements of the grid view and prints the numbers of degrees of freedom per element:

```
auto localView = basis.localView();

for (auto&& element : elements(basis.gridView()))
{
    localView.bind(element);
    std::cout << "Element_ with_" << localView.size()
    << "_degrees_of_freedom" << std::endl;
}
```

Access to the actual local basis functions is provided by the method

```
using Tree = <implementation defined>;
const Tree& tree() const;
```

This encapsulates the set $\Lambda|_e$ of basis functions localized to the element e , organized in the tree of function space bases. While the tree itself can be queried in unbound state, the local view must be bound in order to use most of the trees methods. A detailed discussion of the interface of the tree object is given below.

For any of the local basis functions in the local tree accessible by `tree()` the global multi-index is provided by the method

```
using MultiIndex = <implementation defined>;
MultiIndex index(size_type i) const;
```

The argument for this method is the local index of the basis function within the tree as returned by `node.localIndex(k)`; here `node` is a leaf node of the tree provided by `tree()`, and `k` is the number of the shape function within the corresponding local finite element (see below). Hence the `index` method implements the map $\iota_e^{\text{local} \rightarrow \text{global}}$ introduced in Section 1.5, which maps local indices to global multi-indices. Accessing the same global index multiple times is expected to be cheap, because implementations are supposed to pre-compute and cache indices during `bind(Element)`. The result of calling `index(size_type)` in unbound state is undefined.

Extending the previous example a little, the following loop prints the global indices for each degree of freedom of each element.

```
auto localView = basis.localView();

for (auto&& element : elements(basis.gridView()))
{
    localView.bind(element);
    for (std::size_t i=0; i<localView.size(); i++)
        std::cout << localView.index(i) << std::endl;
}
```

When this code is run for a Taylor–Hood basis on a two-dimensional triangle grid, it will print 15 multi-indices per element, because a Taylor–Hood element has 12 velocity degrees of freedom and 3 pressure degrees of freedom per triangle.

Finally, the global basis of type `GlobalBasis` is known by the `LocalView` object, and exported by the `globalBasis` method.

```
using GlobalBasis = <implementation defined>;
const GlobalBasis& globalBasis() const;
```

Therefore, code that is given only a `LocalView` object can retrieve the global basis from it, and the grid view from there.

2.3 The user interface of the tree of local bases

The local view provides access to local basis functions of an element by exporting a `Tree` object, which keeps the local basis functions in its leaves. The tree structure is encoded in the type of the `Tree` object, using the infrastructure of the `dune-typetree` module.

The object returned by the `LocalView::tree` method is not actually a tree, but rather (a `const` reference to) the root node of the tree. To navigate within this tree, any non-leaf node allows to access its children using the two methods

```
template<class... ChildIndices>
auto child(ChildIndices... childIndices);

template<class ChildTreePath>
auto child(ChildTreePath childTreePath);
```

The arguments to these methods are the paths from the current node to the desired descendants. Notice that both methods only provide access to strict descendants, while accessing the node itself using an empty path is not supported.

For the first method, the path is passed as a sequence of indices. Indices referring to children of a power node can be passed as integer values, typically of type `std::size_t`. Indices referring to children of a composite node have to be passed statically as objects of type `Dune::index_constant<i>`.⁵ For convenience global constants `_0`, `_1`, ... of this type are implemented in the `Dune::Indices` namespace. Continuing the example of the previous section, if `localView` is a local view of the `TaylorHoodBasis` localized to a particular grid element, then the leaf node for the second velocity component can be obtained by

```
using namespace Indices; // Import namespace with index constants _0, _1, _2, etc
const auto& node = localView.tree().child(_0, 1);
```

Note how the index constant `Dune::Indices::_0` is used to address the velocity node, because the tree root is a composite node whose child nodes are of different type. Within the velocity subtree, all three children are identical, and the second one can be accessed by the run-time integer `1`.

⁵ ... which is a shortcut for `std::integral_constant<std::size_t, i>`.

The second `child` method allows to pass the tree path in a dedicated container. Such a container needs to handle sequences of static and run-time values. The `dune-functions` module uses `Dune::TypeTree::HybridTreePath` for this, which we describe in detail in Section 2.4. Using a `HybridTreePath` object, the example looks as follows:

```
auto treePath = Dune::TypeTree::treePath(_0, 1);
const auto& node = localView.tree().child(treePath);
```

At each of its leaf nodes, the localized basis function tree provides the set of all corresponding shape functions. The method for this is

```
using FiniteElement = <implementation defined>;
const FiniteElement& finiteElement() const;
```

The object returned by this method is a `LocalFiniteElement` as specified in the `dune-localfunctions` module. As such, it provides access to shape function values and derivatives, to the evaluation of degrees of freedom in the sense of [4], and to the assignment of local degrees of freedom to element faces. The numbering used by `dune-localfunctions` for the shape functions coincides with the leaf-local indices defined in Section 1.5. For example, if `node` is a leaf node in the localized Taylor–Hood tree, the following code prints all shape function values of the leaf shape function set at the point $(0, 0, 0)$ in local coordinates of the appropriate reference element:

```
const auto& localBasis = node.finiteElement().localBasis();
std::vector<double> values;
localBasis.evaluate({0,0,0}, values);
for (auto v : values)
    std::cout << v << std::endl;
```

To obtain the entries of the element stiffness matrix that corresponds to a given shape function from a given leaf node, the local index needs to be computed from the leaf-local index of that shape function. For each local basis function the method

```
size_type localIndex(size_type i) const;
```

returns the local index within all local basis functions of the current element associated to the full local tree. The argument to this method is the index of the local basis functions within the leaf. In other words, the `localIndex` method implements the map $\iota_e^{\text{leaf} \rightarrow \text{local}}$ introduced in Section 1.5. The return value is *not* a multi-index. While in principle all basis functions of the local subtree could be indexed using general multi-indices, the `dune-functions` module only supports flat indices here to keep the implementation simple.

While `LocalFiniteElement` objects are only available at leaf nodes, the following methods work at every node in the tree again. Calling

```
using size_type = <implementation defined>;
size_type size() const;
```

returns the total number of local basis functions within the subtree rooted at the present node. In particular, calling this method for the tree root yields the size of the element stiffness matrix.

Finally, all nodes provide access to the element which they are bound to via the method

```
using Element = <implementation defined>;
const Element& element() const;
```

2.4 Multi-indices

Multi-indices appear in several places in `dune-functions`. They are used as global indices to identify individual basis functions of a function space basis, and for indexing inner nodes of basis and index trees as well. From an implementation point of view, basis and index trees differ considerably. Only the localized basis tree explicitly appears in the programmer interface, whereas index trees appear only implicitly in the form of sets of indices with the appropriate structure (Definition 3). These differences require separate multi-index implementations for the different types of trees. We discuss implementations for both types of trees in turn.

2.4.1 Multi-index implementations for basis trees

The tree of localized basis functions is the only tree that explicitly appears in the `dune-functions` programmer interface. The tree structure is encoded as C++ type information using the tools from the `dune-typetree` module. Navigation in this tree requires to manipulate paths from the root to particular nodes. In principle, such a path is a sequence of integers.

To understand the implementation, remember that non-leaf tree nodes can be of two types, *power* and *composite* (Section 1.4). Since composite nodes have children of different types, it is not possible to access those children using a dynamic run-time index. Instead the child index in a composite node has to be encoded in a static way. For such situations, `dune-common` offers the type

```
template <std::size_t i> Dune::index_constant<i>;
```

which turns the number `i` into a type, and by this makes it accessible to compile-time expressions. On the other hand, all children of a power node have the same C++ type, and can be accessed using a dynamic index of type `std::size_t`.

In a typical tree, composite and power nodes appear together. It is therefore necessary to have a container that can store both compile-time and run-time integers. This is achieved by the class

```
template <class... I>
class HybridTreePath;
```

from the `dune-typetree` module.

Conceptually, a `HybridTreePath` is a fixed-size container, where each entry can be of different type. The types of the individual entries are passed as template parameters. If the type used for an entry is `std::size_t`, then this entry will have a dynamic value. If, on the other hand, the type is `Dune::index_constant<i>`, then its value is static, and can be used for compile-time decisions.

An object of type `HybridTreePath` can be used to access the nodes of a localized basis tree if dynamic tree path entries only appear as child indices for power nodes in the tree while all other entries are static. For example, to access the leaf nodes corresponding to the velocity components in the Taylor–Hood ansatz tree depicted in Figure 2 one would use multi-indices of the type

```
HybridTreePath<Dune::index_constant<0>, std::size_t>;
```

whereas the multi-index for the pressure leaf node would use the type

```
HybridTreePath<Dune::index_constant<1>>;
```

To construct objects of these types, call

```
using namespace Dune::Indices;
HybridTreePath<Dune::index_constant<0>, std::size_t> i00(_0,0);
HybridTreePath<Dune::index_constant<0>, std::size_t> i01(_0,1);
HybridTreePath<Dune::index_constant<0>, std::size_t> i02(_0,2);
```

for the velocity leaf nodes, and

```
HybridTreePath<Dune::index_constant<1>> i1(_1);
```

for the pressure node. The constants `_0`, `_1`, and `_2` are predefined in the namespace `Dune::Indices`.

This way of construction is overly verbose, because static indices have to be provided both as template and as constructor arguments. To simplify the construction of such objects, the `dune-typetree` module provides the helper function

```
template <class... I>
auto TypeTree::treePath(I... i);
```

which creates a `HybridTreePath` with the given entries. As this is a free method rather than a constructor, the entries have to be given only once, and their types are inferred. The multi-indices of the previous example can be constructed using

```
auto i00 = TypeTree::treePath(_0, 0);
auto i01 = TypeTree::treePath(_0, 1);
auto i02 = TypeTree::treePath(_0, 2);
auto i1 = TypeTree::treePath(_1);
```

which is much shorter. To access entries of a `HybridTreePath` object, the object has the method

```
template<std::size_t i>
constexpr decltype(auto) operator[](const Dune::index_constant<i>&) const;
```

Depending on the template parameter `i`, the return type is either `std::size_t` or `Dune::index_constant<i>`. Note that this construction is necessary although the function is already `constexpr`: Since `HybridTreePath` objects are intended to select children of different type in run-time contexts, they have to encode compile time index values into the run-time index objects. The latter is only possible by making the types of those index objects dependent on their compile time values.

However, as objects of type `Dune::index_constant` can be implicitly converted to `std::size_t`, there is also

```
auto operator[] (std::size_t) const;
```

Hence, to get the first digit of a tree path it is possible to write

```
std::size_t a = myHybridTreePath[0];
```

For a tree path in the Taylor–Hood tree this will return 0 or 1 as expected. However, this return value is not usable in compile-time situations anymore.

These are just the more important methods of the `HybridTreePath` class. For a complete description see the online documentation of the `dune-typetree` module.

2.4.2 Multi-index implementations for index trees

Index trees are formed by the multi-indices that are used to label basis functions. Conceptually, there are two such trees in the `dune-functions` interface: the tree of global indices, and the tree of local indices. To keep the implementation simple, `dune-functions` only allows flat (i.e., single-digit) multi-indices for the local index tree. Therefore, only data types for global indices need to be discussed.

Unlike the tree paths of the previous section, global indices are run-time constructs. A single C++ type represents all such indices for a given basis, even if that basis has a non-trivial tree structure. The exact type is selected by the basis implementation, and can differ from basis to basis. It mainly depends on whether the index is uniform, i.e., whether all indices from the set have the same number of digits. Having purely dynamic multi-indices can be inconvenient when accessing containers such as `std::tuple` or `MultiTypeBlockVector` (from the `dune-istl` module). However, it has the advantage that standard run-time loops can be used to iterate over the indices.

Dynamic multi-indices are random-access containers holding entries of a fixed integer type. All implement a common interface, consisting of two member functions

```
std::size_t size() const;
auto operator[] (std::size_t) const;
```

The `size` method returns the number of digits of the multi-index, and `operator[]` allows to access each entry by its position. Since multi-indices are typically not changed by user code, both methods are `const`. The type used to represent the individual digits of multi-indices can be selected when instantiating `GlobalBasis` objects. The default type is `std::size_t`.

In the following we will give an overview of the types used to represent multi-indices in `dune-functions`. In the most general case, not all multi-indices for a given basis

have the same number of digits. As examples, consider columns 1, 2, 5, and 6 of Table 1, which give such numberings for the Taylor–Hood basis. In these cases, multi-indices are typically represented by the class

```
template <class T, int k>
class ReservedVector;
```

from the `dune-common` module, which is parameterized by the entry type `T` and a capacity `k`. It implements an STL-compatible random-access container with a dynamic size, which may not exceed `k` entries. In contrast to a fully dynamic vector implementation like `std::vector<T>`, the class `ReservedVector` stores its entries on the stack. This avoids dynamic memory management, and makes the implementation much more efficient. The global multi-indices typically have a small number of digits only with a known upper bound. Hence the overhead of always using a buffer of size `k` even for indices with less than `k` digits will typically be small.

However, many bases can be indexed by uniform index trees, i.e., sets of indices where all indices have the same number of digits. In that case, the capacity of a `ReservedVector` can be set to the correct length, and no buffer space is wasted. However, in addition to the buffer, each `ReservedVector` object has to store the container length, which is not needed when the index set is known to be uniform. `GlobalBasis` objects that implement uniform index sets can therefore opt to use a fixed-size container type like `std::array` instead of `ReservedVector`.

Finally, if the basis is indexed with a flat index, i.e., a multi-index with only a single digit, then using an array can be a bit cumbersome. Morally, flat multi-indices are simply natural numbers. However, if `i` is a `std::array` of length 1, using it to access the corresponding entry of a `std::vector` called `vec` has to be written as

```
auto value = vec[i[0]];
```

To allow the more intuitive syntax

```
auto value = vec[i];
```

`dune-functions` implements the `FlatMultiIndex` class for the case that the index of a basis tree is flat. Objects of type `FlatMultiIndex` behave like objects of type `std::array<T,1>`, but additionally, they allow to cast their content to `T&`. Therefore, objects of type `FlatMultiIndex` can be directly used like number types, and like multi-index types as well.

3 Constructing trees of function space bases

There are various ways to construct finite element bases in `dune-functions`. A set of standard bases is provided directly. These can then be combined to form trees. Conversely, subtrees can be extracted, and they act like complete bases in their own right.

3.1 Basis implementations provided by dune-functions

The `dune-functions` module contains a collection of standard finite element bases. These can be directly used in finite element simulation codes. At the time of writing there are:

- **LagrangeBasis**: Lagrange basis of order k , where k is a compile-time parameter. This implementation works on all kinds of conforming grids, including grids with more than one element type. At the time of writing, higher-order spaces are implemented only partially. Check the online class documentation for the current status.
- **LagrangeDGBasis**: Implements a k -th order Discontinuous-Galerkin (DG) basis with Lagrange shape functions. As a DG basis, it also works well on non-conforming grids. The polynomial order k is again a compile-time parameter.
- **RannacherTurekBasis**: An H^1 -nonconforming scalar basis, which adapts the idea of the Crouzeix–Raviart basis to cube grids [12].
- **BSplineBasis**: Implements a B-Spline basis on a structured, axis-aligned grid as described, e.g., in [5]. Arbitrary orders, dimensions, and knot vectors are supported, allowing, e.g., to work with C^1 elements for fourth-order differential equations.

Each `BSplineBasis` object implements a basis on a single patch, and the grid must correspond to this patch. For this to work, several restrictions apply for the grid. It must be structured and axis-aligned, and consist of (hyper-)cube elements only. Further, the element indices must be lexicographic and increase from the lower left to the upper right domain corner. The element spacing must match the knot spans. Unfortunately, not all these requirements can be checked for by the basis, so users have to be a bit careful. Using `YaspGrid` objects works well.

Unlike in standard finite element bases, in a B-spline basis the basis functions cannot be associated to grid entities such as vertices, edges, or elements. The `dune-localfunctions` programmer interface of a B-spline basis nevertheless mandates that a `LocalCoefficient` object must be available on each element, which assigns shape functions to faces of the reference element. For the `BSplineBasis`, the behavior of this object is undefined.

- **TaylorHoodBasis**: An implementation of a first-order Taylor–Hood basis. It exists mainly to serve as an example of how to directly implement a basis with a non-trivial tree. Generally, non-trivial product bases can be easily constructed in a generic way. This approach is described in Chapter 3.2 and it is the preferred way to construct a Taylor–Hood basis.

For all bases listed above, the shape functions provided by `tree.finiteElement()` are implemented in terms of coordinates of the reference element T_{ref} . That is, if a

grid element e is obtained by the transformation $\Phi_e : T_{\text{ref}} \rightarrow e$, then the implemented localized shape function representing the restriction of the basis function λ to the element e is given by $\hat{\lambda}|_e = \lambda \circ \Phi_e$. Finite elements that form non-affine families [4] may require additional transformations. This is the case for the following global bases implementations.

- **RaviartThomasBasis**: The standard Raviart–Thomas basis [2] for problems in $H(\text{div})$. Available for different orders and element types.
- **BrezziDouglasMariniBasis**: The standard Brezzi–Douglas–Marini basis, which is an alternative basis for $H(\text{div})$ -conforming problems [2].

Both bases require the Piola transformation to properly pull back the basis functions onto the reference element. This transformation is *not* performed by the `dune-functions` implementation, and is expected to happen in user code. For a detailed discussion of the template parameters and constructor arguments of the basis implementations listed above we refer to the online class documentation.

3.2 Combining bases into trees

The basis implementations of the previous section can be combined by multiplication to form new bases. This produces the tree structures described in Section 1. The multiplication code resides in the `BasisFactory` namespace, which is a nested namespace within `Dune::Functions::`. Therefore, the examples in this section need a

```
using namespace Dune::Functions::BasisFactory;
```

to compile.

The methods to combine bases into trees do not operate on the basis classes of the previous section directly. Rather, they combine so-called *pre-bases*, of which there is one for each basis. The reason for this is that it is technically challenging to combine the actual user-visible basis types in a tree hierarchy that itself again implements the interface of a hierarchical function space basis. Therefore, the multiplication operators are applied to pre-basis objects, and return pre-basis objects of the resulting tree. The pre-basis of the final basis tree can then be turned into an actual basis.

Since all pre-bases in the product pre-basis have to know some common information like, e.g., the grid view, doing this hierarchic construction manually is verbose and error prone. As a more user friendly and safer solution a global basis can be constructed by a call to

```
template <class GridView, class PreBasis>
auto makeBasis(const GridView& gridView, PreBasis&& preBasis);
```

The pre-basis argument encodes the hierarchic product. The actual basis is constructed automatically by the `makeBasis` function from the pre-bases in a consistent way. This

also determines a suitable multi-index type automatically, which otherwise would have to be done by the user. ⁶

In the simple-most case, the basis tree consists of a single leaf. This leaf is then, e.g., one of the basis implementations of the previous section. As a convention, for each global basis `FooBarBasis` there is a function `BasisFactory::fooBar()` (defined in the same header file as `FooBarBasis`), creating a suitable pre-basis object which stores all basis-specific information. That means that in particular you can write

```
auto raviartThomasBasis = makeBasis(gridView, raviartThomas<k>());
```

to obtain a Raviart–Thomas basis for the given grid view. This call to `makeBasis` is equivalent to constructing the basis directly:

```
RaviartThomasBasis<GridView,k> raviartThomasBasis(gridView);
```

Note that the `raviartThomas` function, just like the corresponding functions for other bases, does not need the grid view as parameter.

If `FooBarBasis` has template and/or constructor parameters, then by convention they are given in the same order as the template and method parameters of the `BasisFactory::fooBar()` function. As the only difference, the former has the grid view type and object prepended.

The pre-basis combining several bases in a product is called `CompositePreBasis`, defined in the header `dune/functions/functionspacebases/compositebasis.hh`. It implements a *composite* tree node as introduced in Definition 6. Analogously to the above description, a pre-basis for a tree with a composite root can be constructed using the global function

```
template <class... ChildPreBasis>
auto composite(ChildPreBasis&&... childPreBasis);
```

contained in the namespace `BasisFactory`. The method has an unspecified number of arguments, of unspecified type. The arguments are expected to be pre-basis objects themselves. They can either be plain pre-bases constructed by, e.g., `lagrange<1>()` or `raviartThomas<k>()`, or composite- or power pre-bases constructed by the `composite` or `power` function (see below), respectively.

As an example, to combine a Raviart–Thomas basis with a zero-order Lagrange basis (let’s say for solving the mixed formulation of the Poisson equation [3]), the appropriate call is

```
auto mixedBasis = makeBasis(
    gridView,
    composite(
        raviartThomas<0>()),
```

⁶The interface description is in fact slightly simplified: The user-provided arguments of `makeBasis` are not pre-bases themselves but pre-basis-factory objects that can construct the corresponding pre-bases. This mechanism allows to delay passing the shared information (e.g. the grid view) to the construction of the real pre-bases which is triggered by `makeBasis`. However, to simplify the presentation we will ignore the technical difference of a pre-basis and its pre-basis-factory in the following.

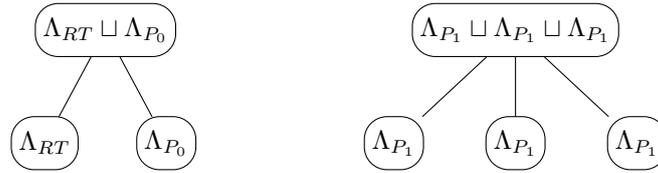


Figure 8: Example composite bases

```
    lagrange<0>()
  ));
```

Combining three copies of a first-order Lagrange basis for a displacement field in elasticity theory is done by

```
auto displacementBasis = makeBasis(
  gridView,
  composite(
    lagrange<1>(),
    lagrange<1>(),
    lagrange<1>()
  ));
```

The examples produce the trees shown in Figure 8.

The second example is not as elegant as it could be. First of all, it is inconvenient and unnecessarily wordy to list the same scalar Lagrange basis three times. Secondly, the required number may depend on a parameter. Finally, the implementation can benefit from the explicit knowledge that all children are equal. For these reasons, `dune-functions` offers a second way to combine bases: The `PowerPreBasis` to be constructed by the factory method `BasisFactory::power()`. The interface is again a single method

```
template<std::size_t k, class ChildPreBasis>
auto power(ChildPreBasis&& childPreBasis)
```

provided in the file `dune/functions/functionspacebases/powerbasis.hh`. It combines `k` copies of a subtree of type `ChildPreBasis` in a new tree. Therefore, the displacement vector field basis from above is more easily written as

```
auto displacementBasis = makeBasis(
  gridView,
  power<3>(
    lagrange<1>()
  ));
```

Since `composite` and `power` create pre-bases themselves, all these techniques can be combined. To obtain the `p`-th order Taylor–Hood basis, write

```
auto taylorHoodBasis = makeBasis(
```

```

gridView,
composite(
  power<dim>(
    lagrange<p+1>()),
  lagrange<p>()
));

```

The call to `power` produces the `dim`-component Lagrange basis of order `p+1` for the velocity, and the call to `composite` combines this with a `p`-th order Lagrange basis for the pressure. Note that this is the preferred way to construct a Taylor–Hood basis in contrast to

```

auto taylorHoodBasis1 = makeBasis(gridView, taylorHood());
and
auto taylorHoodBasis2 = TaylorHoodBasis<GridView>(gridView);

```

These variants mainly exist as an implementation example.

The previous discussion has left out the question of how the degrees of freedom in the combined tree are numbered. In Section 1.3 it was explained how the indices of the degrees of freedom form a separate tree by their multi-index structure, and how this tree is constructed from the basis tree by a set of strategies. These ideas are reflected in the design of the `dune-functions` programmer interface. First of all, each of the bases of Section 3.1 implements a numbering of its degrees of freedom, and generally these numberings cannot be changed. To select a degree of freedom numbering for a non-trivial basis, each call to `composite` or `power` can be augmented by an additional flag indicating an `IndexMergingStrategy`. The four implemented strategies are

- `BlockedLexicographic`
- `BlockedInterleaved`
- `FlatLexicographic`
- `FlatInterleaved`

and have been described in Section 1.4. For each strategy `FooBar` there is a function `BasisFactory::fooBar()` creating the flag in the header `functionspacebases/basistags.hh`. For example, a Taylor–Hood basis with the indexing listed in the second column (labeled BL(BI)) of Table 1 can be created using

```

auto taylorHoodBasis = makeBasis(
  gridView,
  composite(
    power<dim>(
      lagrange<p+1>(),
      blockedInterleaved()),
    lagrange<p>(),
    blockedLexicographic()
  ));

```

This will lead to multi-indices of length three and two for velocity and pressure degrees of freedom, respectively. The same ordering of basis functions with a uniform indexing scheme with multi-index length two (Column 4 labeled BL(FI) in Table 1) is obtained by

```
auto taylorHoodBasis = makeBasis(
  gridView,
  composite(
    power<dim>(
      lagrange<p+1>(),
      flatInterleaved(),
      lagrange<p>(),
      blockedLexicographic()
    )
  ));
```

Finally, a flat indexing scheme still preserving the same ordering (Column 8 labeled FL(FI) in Table 1) is obtained by

```
auto taylorHoodBasis = makeBasis(
  gridView,
  composite(
    power<dim>(
      lagrange<p+1>(),
      flatInterleaved(),
      lagrange<p>(),
      flatLexicographic()
    )
  ));
```

If no strategy is given, `composite` will use the `BlockedLexicographic` strategy, whereas `power` will use `BlockedInterleaved`.

4 Treating subtrees as separate bases

The previous section has shown how trees of bases can be combined to form bigger trees. It is also possible to extract subtrees from other trees and treat these subtrees as basis trees in their own right. The programmer interface for such subtree bases is called `SubspaceBasis`. It mostly coincides with the interface of a global basis, but additionally to the `GlobalBasis` interface the `SubspaceBasis` provides information about how the subtree is embedded into the global basis. More specifically, the method

```
const <implementation defined>& rootBasis() const
```

provides access to the root basis, and the method

```
using PrefixPath = TypeTree::HybridTreePath<implementation defined>;
const PrefixPath& prefixPath() const
```

returns the path of the subtree associated to the `SubspaceBasis` within the full tree. For convenience a global basis behaves like a trivial `SubspaceBasis`, i.e., it has the method `rootBasis` returning the basis itself, and `prefixPath` returning an empty

tree-path. Note that a `SubspaceBasis` differs from a full global basis because the global multi-indices are the same as the ones of the root basis, and thus they are in general neither consecutive nor zero-based. Instead, those multi-indices allow to access containers storing coefficients for the full root basis.

`SubspaceBasis` objects are created using a global factory function from the root basis and the path to the desired subtree. The path can either be passed as a single `HybridTreePath` object (see Section 2.4), or as a sequence of individual indices.

```
template<class RootBasis, class... PathIndices>
auto subspaceBasis(const RootBasis& rootBasis,
                  const TypeTree::HybridTreePath<PathIndices...& prefixPath);

template<class RootBasis, class... PathIndices>
auto subspaceBasis(const RootBasis& rootBasis, const PathIndices&... indices);
```

For example, suppose that `taylorHoodBasis` is any one of the implementations of the Taylor–Hood basis defined in Section 3.2. Then

```
auto velocityBasis = subspaceBasis(taylorHoodBasis, _0);
```

will extract the subtree of velocity degrees of freedom, and

```
auto pressureBasis = subspaceBasis(taylorHoodBasis, _1);
```

will extract the (trivial) subtree of pressure degrees of freedom. The possibly non-consecutive multi-indices of a `SubspaceBasis` are best illustrated by extracting a single velocity component

```
auto velocityZBasis = subspaceBasis(taylorHoodBasis, _0, 2);
```

For this example the following table shows the multi-indices of the `SubspaceBasis` extracted from the full basis, with columns representing the different index merging strategies also used in Table 1:

	BL(BL)	BL(BI)	BL(FL)	BL(FI)	FL(BL)	FL(BI)	FL(FL)	FL(FI)
$v_{x_2,0}$	(0, 2, 0)	(0, 0, 2)	(0, $2n_2 + 0$)	(0, 0 + 2)	(2, 0)	(0, 2)	($2n_2 + 0$)	(0 + 2)
$v_{x_2,1}$	(0, 2, 1)	(0, 1, 2)	(0, $2n_2 + 1$)	(0, 3 + 2)	(2, 1)	(1, 2)	($2n_2 + 1$)	(3 + 2)
$v_{x_2,2}$	(0, 2, 2)	(0, 2, 2)	(0, $2n_2 + 2$)	(0, 6 + 2)	(2, 2)	(2, 2)	($2n_2 + 2$)	(6 + 2)
$v_{x_2,3}$	(0, 2, 3)	(0, 3, 2)	(0, $2n_2 + 3$)	(0, 9 + 2)	(2, 3)	(3, 2)	($2n_2 + 3$)	(9 + 2)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

`SubspaceBasis` objects can be combined with coefficient vectors to represent vector- and scalar-valued discrete functions. The interface for this construction is discussed in the next section.

5 Combining global bases and coefficient vectors

Function space bases and coefficient vectors are combined to yield discrete functions, by the linear combination shown exemplarily in (3). Such discrete functions can then, e.g., be written to a file, or handed to some post-processing agent. Conversely, discrete

and non-discrete functions can be projected onto the span of a basis, which yields a corresponding coefficient vector. In `dune-functions`, this process is called *interpolation*, although it is not always an interpolation in the strict sense of the word.

5.1 Vector backends

In both cases, individual basis functions need to be associated with corresponding entries of a container data type that holds vector coefficients. While trivial in theory, in practice there is a gap here because the multi-index types used by `dune-functions` to label basis functions (Section 2.4.2) cannot be used to access entries of standard random-access containers.

The gap is bridged by a concept called *vector backends*. These are shim classes that abstract away implementation details of particular container classes, and make them addressable by multi-indices. The `dune-functions` module currently offers such a backend for containers from `dune-istl` and the C++ standard library, but others can be added easily. This makes it possible to combine `dune-functions` function space bases with basically any linear algebra implementation.

There are two parts to the vector backend concept: When interpreting a vector of given coefficients with respect to a basis, access is only required in a non-mutable way. In `dune-functions` this functionality is encoded in the `ConstVectorBackend` concept which solely requires direct access by `operator[]` using the multi-indices provided by the function space basis:

```
auto operator[](Basis::MultiIndex) const;
```

For interpolation of given functions a corresponding mutable access is needed as well. Furthermore, it must be possible to resize the vector to match the index tree generated by the basis. These two additional methods make up the `VectorBackend` concept:

```
auto operator[](Basis::MultiIndex);
void resize(const Basis&);
```

Note that the argument of the `resize` member function is not a number, but the basis itself. This is necessary because resizing nested containers requires information about the whole index tree.

For the vector types implemented in the `dune-istl` and `dune-common` modules, such a backend can be obtained using

```
template<class SomeDuneISTLVector>
auto istlVectorBackend(SomeDuneISTLVector& x);

template<class SomeDuneISTLVector>
auto istlVectorBackend(const SomeDuneISTLVector& x);
```

Depending on the `const`-ness of the argument, the resulting object implements the `VectorBackend` or only the `ConstVectorBackend` interface. Even though these meth-

ods have the `istl` prefix in their names, they actually also work well for containers from the C++ standard library like `std::array` and `std::vector`.

Since a non-trivial product function space basis corresponds to functions with a non-scalar range, we additionally have to map the components of the spanned product function space to components of a function range type. If, for example, the functions from the power function space generated by the basis

```
auto basis = makeBasis(gridView, power<dim>(lagrange<1>()));
```

should be interpreted as vector fields, one would map the `dim` leaf nodes of this basis to the `dim` entries of a `FieldVector<double, dim>`. Whenever combining bases and range types `dune-function` uses a default mapping generalizing this idea to more complex nested bases: Assume that `y` is an object of the function range type. Then a leaf node with tree path `i0, ..., in` is associated to the entry `y[i0]... [in]`. In the exceptional case that the range type does not provide an `operator[]` it is directly used for all leaf nodes in the ansatz space. This last rule allows to interpolate a scalar function into all components of a basis at once. For additional flexibility, users can also provide custom mappings to be used instead of this one. However, we will not discuss the corresponding interface and rely on the implicitly used default implementation in the following.

5.2 Interpreting coefficient vectors as finite element functions

To combine a basis and a coefficient vector to a discrete function that can be evaluated point-wise in Ω , `dune-functions` provides the function

```
template<class Range, class B, class C>
auto makeDiscreteGlobalBasisFunction(const B& basis, const C& coefficients);
```

For given basis `basis` and coefficient vector `coefficients` this returns an object representing the corresponding finite element function. This object implements the `GridViewFunction` concept for the grid view the basis is defined on, described in [6] with the range type `Range`. The `basis` can either be a global basis or a `SubspaceBasis`. In the latter case the coefficient vector has to correspond to the full basis nevertheless, but only the coefficients associated with the subspace basis functions will be used.

For the type `C` used to represent the coefficient vector there are two choices. Either it implements the `ConstVectorBackend` concept, as for example all objects returned by the `istlVectorBackend` method do. If `C` does not implement this concept, then the code assumes that it is a `dune-istl`-style container and tries to wrap it with an `ISTLVectorBackend`. That way, `makeDiscreteGlobalBasisFunction` can be called with `dune-istl` or STL containers directly, but all others have to be wrapped in an appropriate backend explicitly.

Notice that the range type can in general not be determined automatically from the basis and coefficient type because there are multiple possible types to implement this. For example a scalar function could return `double` or `FieldVector<double, 1>`. Hence the range type `Range` has to be given explicitly by the user. The mapping from the

different leaf nodes of the basis to the entries of `Range` follows the procedure described in Section 5.1.

To give an example how `makeDiscreteGlobalBasisFunction` is used, we construct yet another instance of the Taylor–Hood basis

```
auto taylorHoodBasis = makeBasis(
    gridView,
    composite(
        power<dim>(
            lagrange<p+1>()),
        lagrange<p>()
    ));
```

By default, the merging strategies are `BlockedLexicographic` for the composite node, and `BlockedInterleaved` for the power node. The resulting indices are the ones from Column 2 of Table 1. An appropriate vector container type for this is

```
using Vectortype = TupleVector<
    BlockVector<FieldVector<double,dim> >,
    BlockVector<FieldVector<double,1> > >;
```

Let `x` be an object of this type. To obtain the corresponding velocity field as a discrete function, write

```
// Create SubspaceBasis for the velocity field
auto velocityBasis = subspaceBasis(taylorHoodBasis, _0);

// Fix a range type for the velocity field
using VelocityRange = FieldVector<double,dim>;

// Create a function for the velocity field only
// but using the vector x for the full taylorHoodBasis.
auto velocityFunction
    = makeDiscreteGlobalBasisFunction<VelocityRange>(velocityBasis, x);
```

Notice that the `dim` leaf nodes of the function space tree spanned by `velocityBasis` are automatically mapped to the `dim` components of the `VelocityRange` type. The resulting function created in the last line implements the full `GridViewFunction` interface described in [6]. For example it can be directly passed to the `VTKWriter` class of the `dune-grid` module to write the velocity field as a VTK vector field. See the end of Section 6.3.1 for how this is done.

5.3 Interpolation

In various parts of a finite element or finite volume simulation code, given functions need to be interpolated into spaces spanned by a global basis. For example, initial iterates may be given in closed form, but need to be transferred to a finite element representation to be usable. Similarly, Dirichlet values given in closed form may need to be interpolated on the set of Dirichlet degrees of freedom. Depending on the finite element space, interpolation may take different forms. Nodal interpolation is the

natural choice for Lagrange elements, but for other spaces L^2 -projections or Hermite-type interpolation may be more appropriate.

The `dune-functions` module provides a set of methods for interpolation in the file `dune/functions/functionspacebases/interpolation.hh`. These methods are canonical in the sense that they use the `LocalInterpolation` functionality on each element for the interpolation. This is appropriate for a lot, but not all finite element spaces. For example, no reasonable local interpolation can be defined for B-spline bases, and therefore the standard interpolation functionality cannot be used with the `BSplineBasis` class. This approach also fails for non-affine finite elements, because `LocalInterpolation` is not applying non-standard transformations to the reference element.

The interpolation functionality is implemented in two global functions. The first deals with the simple case of a given function and basis, where the function is to be projected onto the span of the basis, yielding a coefficient vector describing the result.

```
template <class Basis, class C, class F>
void interpolate(const Basis& basis, C&& coefficients, const F& f);
```

Note that this will only work if the range type of `f` and the global basis `basis` are compatible. `dune-functions` implements a compatibility layer that allows to use different vector (or matrix) types from the dune core modules and scalar types like, e.g. `double` for the range of `f` as long as the number of scalar entries of this range type is the same as the dimension of the range space of the function space spanned by the basis. This also implies the assumption that the coefficients for individual basis functions are scalar. The type of the coefficient vector `coefficients` either has to implement the `VectorBackend` concept or to be wrappable by the `istlVectorBackend`. For example, consider the function

$$f_1 : \mathbb{R}^2 \rightarrow \mathbb{R} \quad f_1 = \exp(-\|x\|^2)$$

implemented as

```
47 auto f1 = [](const FieldVector<double,2>& x)
48 {
49     return exp(-1.0*x.two_norm2());
50 };
```

Additionally, consider a scalar second-order Lagrange space

```
54 Functions::LagrangeBasis<GridView,2> p2basis(gridView);
```

and an empty coefficient vector `x1`, not necessarily of correct size:

```
58 std::vector<double> x1;
```

Then, the single line

```
62 interpolate(p2basis, x1, f1);
```

will fill `x1` with the nodal values of the function `f1`.

This interpolation works equally well for non-trivial basis trees and subtrees obtained by the `subspaceBasis` function, provided that the coefficient vector matches the basis and that the function range can be mapped to the product space associated to the basis. Suppose there is a Taylor–Hood basis for a two-dimensional grid that uses flat multi-indices to label its degrees of freedom:

```

67     using namespace Functions::BasisBuilder;
68
69     auto taylorHoodBasis = makeBasis(
70         gridView,
71         composite(
72             power<dim>(
73                 lagrange<2>(),
74                 flatLexicographic()),
75             lagrange<1>(),
76             flatLexicographic()
77         ));

```

A suitable coefficient vector for such a basis is for example

```

81     BlockVector<FieldVector<double,1>> x2;

```

In this situation, interpolation of `f1` into the pressure components of the Taylor–Hood basis can be achieved by

```

86     using namespace Indices;
87     interpolate(subspaceBasis(taylorHoodBasis, _1), x2, f1);

```

Similarly we can interpolate a given vector field `f2` into the non-trivial subtree representing the velocity using

```

91     auto f2 = [](const FieldVector<double,2>& x) {
92         return x;
93     };
94     interpolate(subspaceBasis(taylorHoodBasis, _0), x2, f2);

```

It is even possible to interpolate into the full `taylorHoodBasis` if the range type of the provided function has the same nesting structure as the basis.

In some situations it is also desirable to interpolate only on a part of the domain. Algebraically, the interpolation is then performed as before, but only a subset of all coefficients are written. The most frequent use-case is the interpolation of Dirichlet data onto the algebraic degrees of freedom on the Dirichlet boundary. All other degrees of freedom must not be touched, as they contain, e.g., a suitable initial iterate obtained by some other means.

To support this kind of interpolation, a variant of the `interpolate` method allows to explicitly mark a subset of coefficient vector entries to be written.

```

template <class B, class C, class F, class BV>
void interpolate(const B& basis,
               C&& coefficients,
               const F& f,
               const BV& bitVector);

```

Conceptually, the additional `bitVector` argument must be a container of booleans having the same nesting structure as `coefficients`. Its entries are treated as boolean values indicating if the corresponding entry of `coefficients` should be written. For example, for flat global indices `std::vector<bool>` and `std::vector<char>` work nicely. The class `BitSetVector<N>` (from the `dune-common` module) can be used as a space-optimized alternative to `std::vector<std::bitset<N>>`. For example, to interpolate the `f2` function defined above only into the boundary velocity degrees of freedom, first set up a suitable bit-vector:

```

98   BlockVector<FieldVector<char,1>> isBoundary;
99   auto isBoundaryBackend = Functions::istlVectorBackend(isBoundary);
100  isBoundaryBackend.resize(taylorHoodBasis);
101  isBoundary = false;
102  forEachBoundaryDOF(subspaceBasis(taylorHoodBasis, _0),
103    [&] (auto&& index) {
104      isBoundaryBackend[index] = true;
105    });

```

The actual interpolation is then a single line:

```

109  interpolate(subspaceBasis(taylorHoodBasis, _1), x2, f2, isBoundary);

```

See Chapter 6.3.1 for a more involved example. The `forEachBoundaryDOF` method that loops over all boundary degrees of freedom is defined in the file `dune/functions/functionspacebases/boundarydofs.hh`.

6 Example: Solving the Stokes equation with dune-functions

We close the paper by showing a complete example program that demonstrates a lot of the techniques presented so far. The example program will solve the stationary Stokes problem using Taylor–Hood finite elements. The example is contained in a single file, which comes as part of the `dune-functions` source tree, in `dune-functions/examples/stokes-taylorhood.cc`. If you read this document in electronic form, the file can also be accessed by clicking on the icon in the margin.



6.1 The Stokes equation

The Stokes equation models a viscous incompressible fluid in a d -dimensional domain Ω . There are two unknowns in this problem: a stationary fluid velocity field $\mathbf{u} : \Omega \rightarrow \mathbb{R}^d$, and the fluid pressure $p : \Omega \rightarrow \mathbb{R}$. Together, they have to solve the boundary value problem

$$\begin{aligned}
 -\Delta \mathbf{u} - \nabla p &= 0 && \text{in } \Omega, \\
 \operatorname{div} \mathbf{u} &= 0 && \text{in } \Omega, \\
 \mathbf{u} &= \mathbf{g} && \text{on } \partial\Omega,
 \end{aligned}$$

where we have omitted the physical parameters. The boundary value problem only determines the pressure p up to a constant function. The pressure is therefore usually normalized such that $\int_{\Omega} p \, dx = 0$.

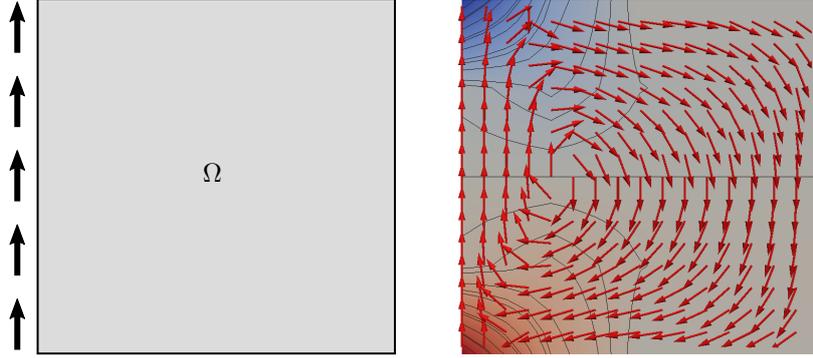


Figure 9: Driven cavity. Left: domain and boundary conditions; right: simulation result. The arrows show the *normalized* velocity.

Due to the constraint $\operatorname{div} \mathbf{u} = 0$, the corresponding weak form of the equation is a saddle-point problem. Introduce the spaces

$$\mathbf{H}_{\mathbf{g}}^1(\Omega) := \{\mathbf{v} \in H^1(\Omega, \mathbb{R}^d) : \operatorname{tr} \mathbf{v} = \mathbf{g}\},$$

$$L_{2,0}(\Omega) := \left\{ q \in L_2(\Omega) : \int_{\Omega} q \, dx = 0 \right\},$$

and the bilinear forms

$$a(\mathbf{u}, \mathbf{v}) := \int_{\Omega} \nabla \mathbf{u} \nabla \mathbf{v} \, dx, \quad \text{and} \quad b(\mathbf{v}, q) := \int_{\Omega} \operatorname{div} \mathbf{v} \cdot q \, dx.$$

Then the weak form of the Stokes equation is: Find $(\mathbf{u}, p) \in \mathbf{H}_{\mathbf{g}}^1(\Omega) \times L_{2,0}(\Omega)$ such that

$$\begin{aligned} a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) &= 0 & \text{for all } \mathbf{v} \in \mathbf{H}_0^1(\Omega) \\ b(\mathbf{u}, q) &= 0 & \text{for all } q \in L_{2,0}(\Omega). \end{aligned}$$

If \mathbf{g} is sufficiently smooth, this variational problem has a unique solution. The Taylor–Hood element is the standard way to discretize this saddle point problem [3], and will be used in the following implementation.

6.2 The driven-cavity benchmark

For our example we choose to simulate a two-dimensional driven-cavity. This is a standard benchmark for the Stokes problem in the literature [13]. Let Ω be the unit square $[0, 1]^2$, and set the Dirichlet boundary conditions for the velocity \mathbf{u} to

$$\mathbf{u}(x) = \mathbf{g}(x) = \begin{cases} (0, 1)^T & \text{if } x \in \{0\} \times [0, 1] \\ (0, 0)^T & \text{elsewhere on } \partial\Omega. \end{cases}$$

The interpretation of this is a fluid container that is closed on all but one side. While the fluid remains motionless on the closed sides, an external agent drives a constant upward motion on the left vertical side. The domain and boundary conditions are depicted in Figure 9, left. The corresponding solution is shown on the right side of the same figure. The velocity forms a vortex, while the pressure forms extrema in the two left corners.

6.3 Implementation

The implementation consists of an assembler for the Stokes problem and a `main` method. Both will be discussed in the following.

6.3.1 The main method

The `main` method sets up the algebraic Stokes problem, calls a linear solver, and writes the result to a VTK file. It begins by setting up MPI and the grid. We choose to discretize the domain using a structured 4×4 quadrilateral grid, which we get by using the `YaspGrid` grid implementation from the `dune-grid` module. Note that there is the line

```
40 using namespace Dune;
```

at the top of the file, so this namespace is imported completely. Additionally, everything in the `dune-functions` module is contained in the namespace `Functions`. This namespace is not imported; instead, the prefix `Functions::` is always given explicitly.

```
299 int main (int argc, char *argv[]) try
300 {
301     // Set up MPI, if available
302     MPIHelper::instance(argc, argv);

310     const int dim = 2;
311     using GridType = YaspGrid<dim>;
312     FieldVector<double,dim> upperRight = {1, 1};
313     std::array<int,dim> elements = {{4, 4}};
314     GridType grid(upperRight,elements);
315
316     using GridView = typename GridType::LeafGridView;
317     GridView gridView = grid.leafGridView();
```

The `gridView` object is the non-hierarchical finite element grid that we will use for the computation. On this grid view, we then set up the function space basis for the Taylor–Hood element. This is as simple as

```
326     using namespace Functions::BasisFactory;
327
328     constexpr std::size_t p = 1; // pressure order for Taylor–Hood
329
330     auto taylorHoodBasis = makeBasis(
331         gridView,
```

```

332     composite(
333         power<dim>(
334             lagrange<p+1>(),
335             blockedInterleaved()),
336         lagrange<p>()
337     );

```

This way of constructing a Taylor–Hood basis from instances of Lagrange bases has been discussed in Section 3.2. The indexing strategies used here are `BlockedInterleaved` for the velocity subtree and `BlockedLexicographic` (the default) for the root. This results in the indexing scheme spelled out in Column 2 of Table 1.

Before being able to assemble the stiffness matrix of the Stokes system we need to pick suitable data structures for the linear algebra. These data structures should have a blocking structure that matches the multi-indices used by the Taylor–Hood basis we just constructed. More concretely, the appropriate vector container will be a pair of vectors, where the first one has entries in \mathbb{R}^d for the velocity and the second one has entries in \mathbb{R} for the pressure degrees of freedom. Analogously, the matrix must consist of 2×2 large sparse matrices where the $(0,0)$ -block has entries in $\mathbb{R}^{d \times d}$, the $(0,1)$ -block has entries in $\mathbb{R}^{d \times 1}$, the $(1,0)$ -block has entries in $\mathbb{R}^{1 \times d}$, and the $(1,1)$ -block has entries in \mathbb{R} (as on the right side of Figure 6). The following code sets up such vector and matrix types for this. It uses the nesting machinery from `dune-istl`, but data types from other linear algebra libraries could be used as well.

```

361     using VelocityVector = BlockVector<FieldVector<double,dim>>;
362     using PressureVector = BlockVector<FieldVector<double,1>>;
363     using VectorType = MultiTypeBlockVector<VelocityVector, PressureVector>;
364
365     using Matrix00 = BCRSMMatrix<FieldMatrix<double,dim,dim>>;
366     using Matrix01 = BCRSMMatrix<FieldMatrix<double,dim,1>>;
367     using Matrix10 = BCRSMMatrix<FieldMatrix<double,1,dim>>;
368     using Matrix11 = BCRSMMatrix<FieldMatrix<double,1,1>>;
369     using MatrixRow0 = MultiTypeBlockVector<Matrix00, Matrix01>;
370     using MatrixRow1 = MultiTypeBlockVector<Matrix10, Matrix11>;
371     using MatrixType = MultiTypeBlockMatrix<MatrixRow0, MatrixRow1>;

```

Note that `VectorType` and `MatrixType` are no classical containers, because the entries have non-uniform types. Rather, they are constructed similarly to `std::tuple` from the C++ standard library. However, it must be emphasized that the use of such advanced data structures is by no means mandatory. As detailed in Sections 1.4 and 3.2 it is trivial to make the Taylor–Hood basis use flat global indices, which work directly with standard container types like `std::vector`.

Now that we have chosen the C++ types for the matrix and vector data structures we can actually assemble the system. Assembling the right-hand-side vector `rhs` is easy, because, apart from the Dirichlet boundary data (which we will insert later), all its entries are zero. An all-zero vector of the correct type and size is set up by the following lines

```

384     VectorType rhs;

```

```

385
386     auto rhsBackend = Dune::Functions::istlVectorBackend(rhs);
387
388     rhsBackend.resize(taylorHoodBasis);
389     rhs = 0;

```

The object returned by `istlVectorBackend` connects the `dune-functions` basis with `dune-istl` linear algebra containers. In particular, it offers convenient resizing of an entire hierarchy of nested vectors from given function space basis trees. Line 389 fills the entire vector with zeros in one go, but observe that this is actually a `dune-istl` feature.

To obtain the stiffness matrix we first create an empty matrix object of the correct type. The actual assembly is factored out into a separate method.

```

393     MatrixType stiffnessMatrix;
394     assembleStokesMatrix(taylorHoodBasis, stiffnessMatrix);

```

As the matrix assembly is a central part of this example we explain it in detail below, after having covered the `main` method.

Suppose now that we have the correct stiffness matrix assembled in the object `stiffnessMatrix`. We still need to modify the linear system to include the Dirichlet boundary information. In a first step we need to determine all degrees of freedom with Dirichlet boundary conditions. To store this information we use a vector of flags which has the same structure as `VectorType` and is again initialized using the `ISTLVectorBackend`.

```

403     using VelocityBitVector = BlockVector<FieldVector<char,dim>>;
404     using PressureBitVector = BlockVector<FieldVector<char,1>>;
405     using BitVectorType
406         = MultiTypeBlockVector<VelocityBitVector, PressureBitVector>;
407
408     BitVectorType isBoundary;
409
410     auto isBoundaryBackend = Dune::Functions::istlVectorBackend(isBoundary);
411     isBoundaryBackend.resize(taylorHoodBasis);
412     isBoundary = false;

```

We now want to mark all the velocity degrees of freedom on the Dirichlet boundary. In the driven-cavity example, the entire boundary is Dirichlet boundary. For convenience, `dune-functions` provides the method

```

template<class Basis, class F>
void forEachBoundaryDOF(const Basis& basis, F&& f);

```

(in the file `dune/functions/functionspacebases/boundarydofs.hh`), which implements a loop over all degrees of freedom associated to entities located on the domain boundary. The algorithm will invoke the callback function `f` for each such degree of freedom, passing its global index as the callback argument. To mark the boundary degrees of freedom for the velocity subtree write:

```

416 using namespace Indices;
417 Functions::forEachBoundaryDOF(
418     Functions::subspaceBasis(taylorHoodBasis, _0),
419     [&] (auto&& index) {
420         isBoundaryBackend[index] = true;
421     });

```

The `forEachBoundaryDOF` algorithm only considers velocity degrees of freedom because we called it with the corresponding subspace basis. Nevertheless, the global indices handed out by `forEachBoundaryDOF` correspond to the full tree, and can therefore be used to access the `isBoundary` container via the `ISTLVectorBackend` (Section 4).

Now that we have determined the set of Dirichlet degrees of freedom, we define a method implementing the actual Dirichlet values function `g`, and interpolate that into the right-hand-side vector `rhs`.

```

425 using Coordinate = GridView::Codim<0> ::Geometry::GlobalCoordinate;
426 using VelocityRange = FieldVector<double,dim>;
427 auto&& velocityDirichletData = [] (Coordinate x)
428 {
429     return VelocityRange{0.0, double(x[0] < 1e-8)};
430 };
431
432 Functions::interpolate(
433     Functions::subspaceBasis(taylorHoodBasis, _0), rhs,
434     velocityDirichletData,
435     isBoundary);

```

Observe how the `dune-functions` interface allows to interpolate C++11 lambdas, which makes the code very short and readable. Again the operation is constrained to the velocity degrees of freedom by passing the corresponding subspace basis only. The `isBoundary` vector given as the last argument restricts the interpolation to only the boundary degrees of freedom which we marked before.

The stiffness matrix is modified in a more manual fashion. For each Dirichlet degree of freedom we need to fill the corresponding matrix row with zeros, and write a 1 on the diagonal. The following algorithm does this by looping over all grid elements, and for each element looping over all Dirichlet degrees of freedom. This is less efficient than simply looping over all matrix rows, but it allows to avoid implementing iterators for the nested sparse matrix data types.

```

444 auto localView = taylorHoodBasis.localView();
445 for(const auto& element : Dune::elements(taylorHoodBasis.gridView()))
446 {
447     localView.bind(element);
448     for (size_t i=0; i<localView.size(); ++i)
449     {
450         auto row = localView.index(i);
451         // If row corresponds to a boundary entry, modify
452         // it to be an identity matrix row
453         if (isBoundaryBackend[row])

```

```

454     for (size_t j=0; j<localView.size(); ++j)
455     {
456         auto col = localView.index(j);
457         matrixEntry(stiffnessMatrix, row, col) = (i==j) ? 1 : 0;
458     }
459 }
460 }

```

Access to the matrix entries needs a matrix analogue to the vector backends—a translation layer that converts multi-indices to the correct sequence of instructions required to access the matrix data structure. Such a backend is more challenging to write, as it requires handling row and column indices at the same time. At the time of writing `dune-functions` does not provide matrix backends. Instead, the example code uses a small helper method `matrixEntry` that is defined in the example file itself. It is much simpler than a generic matrix backend, because it is written directly for the matrix data type of the Stokes problem.

```

217 template<class Matrix, class MultiIndex>
218 decltype(auto) matrixEntry(
219     Matrix& matrix, const MultiIndex& row, const MultiIndex& col)
220 {
221     using namespace Indices;
222     if ((row[0]==0) and (col[0]==0))
223         return matrix[_0][_0][row[1]][col[1]][row[2]][col[2]];
224     if ((row[0]==0) and (col[0]==1))
225         return matrix[_0][_1][row[1]][col[1]][row[2]][0];
226     if ((row[0]==1) and (col[0]==0))
227         return matrix[_1][_0][row[1]][col[1]][0][col[2]];
228     return matrix[_1][_1][row[1]][col[1]][0][0];
229 }

```

Notice that the outer indices for the `MultiTypeBlockMatrix` have to be encoded statically, because different matrix entries have different types. The additional `[0][0]` in Line 228 is necessary because the entries of the lower-right matrix diagonal block are of type `FieldMatrix<double,1,1>` (see Line 368), and the `[0][0]` is needed to get the `double` from that. Similarly the entries of the $(0,1)$ - and $(1,0)$ -blocks of the matrix are interpreted as single-column and single-row matrices, respectively, such that corresponding `[0]` indices have to be inserted.

Finally, we can solve the linear system. Dedicated Stokes solvers frequently operate on some sort of Schur complement, and hence they need direct access to the submatrices [8]. This can be elegantly done using the nested matrix type used in this example. However, efficiently solving the Stokes system is an art, which we do not want to get into here. Instead, we use a GMRes solver, without any preconditioner at all. This is known to converge, albeit slowly. The advantage is that it can be written down in very few lines. The following code shows a typical way of using `dune-istl` to solve a linear system of equations, and is not particular to `dune-functions` at all.

```

467     // Start from the rhs vector; that way the Dirichlet entries are already correct
468     VectorType x = rhs;

```

```

469
470 // Technicality: turn the matrix into a linear operator
471 MatrixAdapter<MatrixType,VectorType,VectorType> stiffnessOperator(stiffnessMatrix);
472
473 // Fancy (but only) way to not have a preconditioner at all
474 Richardson<VectorType,VectorType> preconditioner(1.0);
475
476 // Construct the actual iterative solver
477 RestartedGMResSolver<VectorType> solver(
478     stiffnessOperator, // operator to invert
479     preconditioner,   // preconditioner for iteration
480     1e-10,            // desired residual reduction factor
481     500,              // number of iterations between restarts
482     500,              // maximum number of iterations
483     2);              // verbosity of the solver
484
485 // Object storing some statistics about the solving process
486 InverseOperatorResult statistics;
487
488 // Solve!
489 solver.apply(x, rhs, statistics);

```

Observe how the `RestartedGMResSolver` object is completely oblivious to the fact that the matrix has a nesting structure.

Once the iterative solver has terminated, the result is written to a VTK file. However, as the `VTKWriter` class from the `dune-grid` module expects discrete functions rather than coefficient vectors, we construct velocity and pressure discrete functions by combining the appropriate coefficient vectors and basis subtrees:

```

497 using VelocityRange = FieldVector<double,dim>;
498 using PressureRange = double;
499
500 auto velocityFunction
501     = Functions::makeDiscreteGlobalBasisFunction<VelocityRange>(
502         Functions::subspaceBasis(taylorHoodBasis, _0), x);
503 auto pressureFunction
504     = Functions::makeDiscreteGlobalBasisFunction<PressureRange>(
505         Functions::subspaceBasis(taylorHoodBasis, _1), x);

```

Then, we write the resulting velocity as a vector field, and the resulting pressure as a scalar field. We subsample the grid twice, because the `VTKWriter` class natively only displays piecewise linear functions.

```

513 SubsamplingVTKWriter<GridView> vtkWriter(
514     gridView,
515     refinementLevels(2));
516 vtkWriter.addVertexData(
517     velocityFunction,
518     VTK::FieldInfo("velocity", VTK::FieldInfo::Type::vector, dim));
519 vtkWriter.addVertexData(

```

```

520         pressureFunction,
521         VTK::FieldInfo("pressure", VTK::FieldInfo::Type::scalar, 1));
522     vtkWriter.write("stokes-taylorhood-result");

```

When run, this program produces a file called `stokes-taylorhood-result.vtu`. The file can be opened in `PARAVIEW`, and the outcome looks like the image on the right in Figure 9.

6.3.2 The global assembler

Now that we have covered the `main` method, we can turn to the assembler for the Stokes stiffness matrix. We begin with the global assembler, which is implemented in the method `assembleStokesMatrix` called in Line 394 of the `main` method. The global assembler sets up the matrix pattern, loops over all elements, and accumulates the element stiffness matrices in the global matrix. The signature of the method is

```

245 template <class Basis, class MatrixType>
246 void assembleStokesMatrix(const Basis& basis, MatrixType& matrix)

```

The only arguments it gets are the finite element basis and the matrix to fill. Observe that the Taylor–Hood basis is not hard-wired here, so we could call the method with a different basis. However, not surprisingly the local assembler for the Stokes problem makes relatively tight assumptions on the basis tree structure, so there is relatively little practical freedom here. Ideally, a global assembler should be fully generic, and all knowledge about the current spaces and differential operators should be confined to the local assembler. Real discretization frameworks like `dune-pdelab` do achieve this separation, but for our example here we are less strict, to avoid technicalities.

The first few lines of the `assembleStokesMatrix` method set up the matrix occupation pattern, and initialize all matrix entries with zero.

```

250     // Set matrix size and occupation pattern
251     setOccupationPattern(basis, matrix);
252
253     // Set all entries to zero
254     matrix = 0;

```

The method `setOccupationPattern` that constructs the matrix pattern is included in the example file itself. It is easy to understand for everyone who understands the rest of the assembly code, and we therefore omit a detailed description.

Next comes the actual element loop. We first request a `localView` object from the finite element basis:

```

259     auto localView = basis.localView();

```

After that starts the loop over the grid elements. For each element, we bind the `localView` object to the element. From now on all enquiries to the local view will implicitly refer to this element.

```

264     for (const auto& element : elements(basis.gridView()))
265     {

```

```

266 // Bind the local FE basis view to the current element
267 localView.bind(element);

```

We then create the element stiffness matrix, and call the separate method `getLocalMatrix` to fill it. For simplicity the code uses a dense matrix type even though it is known a priori that the stationary Stokes matrix does not contain entries in the pressure diagonal block. As local shape function indices are flat, a matrix data type without nesting is used for the element stiffness matrix:

```

273 Matrix<FieldMatrix<double,1,1> > elementMatrix;
274 getLocalMatrix(localView, elementMatrix);

```

The `getLocalMatrix` method is discussed in detail below. It gets only the `localView` object in addition to the `elementMatrix`. The former object contains all necessary information. After the call to `getLocalMatrix` the `elementMatrix` object contains the element stiffness matrix for the current element. The code loops over the entries of the element stiffness matrix and adds them onto the global matrix.

```

279 for (size_t i=0; i<elementMatrix.N(); i++)
280 {
281 // The global index of the i-th local degree of freedom of the element 'e'
282 auto row = localView.index(i);
283
284 for (size_t j=0; j<elementMatrix.M(); j++ )
285 {
286 // The global index of the j-th local degree of freedom of the element 'e'
287 auto col = localView.index(j);
288 matrixEntry(matrix, row, col) += elementMatrix[i][j];
289 }
290 }

```

The type returned in Lines 282 and 287 for the global row and column indices is a multi-index. It has length 3 for velocity degrees of freedom and length 2 for pressure degrees of freedom. Line 288 uses the helper function `matrixEntry` again to access the nested global stiffness matrix using those multi-indices.

The preceding loops write into all four of the matrix blocks, even though it is known that for the Stokes system the lower right block contains only zeros. A more optimized version of the code would leave out the lower right submatrix altogether.

6.3.3 The local assembler

It remains to investigate the method that assembles the element stiffness matrices. Its signature is

```

45 template <class LocalView>
46 void getLocalMatrix(
47     const LocalView& localView,
48     Matrix<FieldMatrix<double,1,1>>& elementMatrix)

```

It only receives the local view of the Taylor–Hood basis, expected to be bound to an element, and the empty element matrix. There is no explicit requirement that

the `LocalView` object be a local view of a Taylor–Hood basis, but the assumption is made implicitly in various parts of the local assembler. The first few lines of the `getLocalMatrix` method gather some information about the element the method is to work on. In particular, from the `localView` object it extracts the element itself, and the element’s dimension and geometry

```
53 using Element = typename LocalView::Element;
54 const Element element = localView.element();
55
56 const int dim = Element::dimension;
57 auto geometry = element.geometry();
```

Next, the element stiffness matrix is initialized. The `localView` object knows the total number of degrees of freedom of the element it is bound to, and since the matrix has only scalar entries this is the correct number of matrix rows and columns:

```
62 elementMatrix.setSize(localView.size(), localView.size());
63 elementMatrix = 0; // fills the entire matrix with zeros
```

Finally, we ask for the set of velocity and pressure shape functions:

```
68 using namespace Indices;
69 const auto& velocityLocalFiniteElement
70     = localView.tree().child(_0,0).finiteElement();
71 const auto& pressureLocalFiniteElement
72     = localView.tree().child(_1).finiteElement();
```

The two objects returned in Lines 69–72 are `LocalFiniteElements` in the `dune-localfunctions` sense of the word. These lines also show the tree structure of the Taylor–Hood basis in action: The expression

```
localView.tree().child(_0,0)
```

returns the first child of the first child of the root, i.e., the basis for the first component of the velocity field, and

```
localView.tree().child(_1)
```

is the basis for the pressure space. As the root of the tree combines two different bases, the static identifiers `_0` and `_1` from the `Dune::TypeTree::Indices` namespace are needed to specify its children. The inner node for the velocities combines d times the same basis, and hence the normal integer 0 can be used to address its first child. This particular implementation of the local Stokes assembler is actually “cheating”, because it exploits the knowledge that the same basis is used for all velocity components. Therefore, only the first leaf of the velocity subtree is acquired in Line 69, and then used for all components. Using separate local finite elements is wasteful because the same shape function values and gradients would be computed multiple times.

Next, the code constructs a suitable quadrature rule and loops over the quadrature points. The formula for the quadrature order combines information about the element type, the shape functions, and the differential operator. It computes the lowest order that will integrate the weak form of the Stokes equation exactly on a cube grid.

```

77   int order = 2*(dim*velocityLocalFiniteElement.localBasis().order()-1);
78   const auto& quad = QuadratureRules<double, dim>::rule(element.type(), order);
79
80   // Loop over all quadrature points
81   for (const auto& quadPoint : quad)
82   {

```

The quadrature loop starts like similar local assembler codes seen elsewhere. First, we get the inverse transposed Jacobian of the map from the reference element to the grid element, and the Jacobian determinant for the integral transformation formula:

```

85   // The transposed inverse Jacobian of the map from the
86   // reference element to the element
87   const auto jacobianInverseTransposed
88       = geometry.jacobianInverseTransposed(quadPoint.position());
89
90   // The multiplicative factor in the integral transformation formula
91   const auto integrationElement
92       = geometry.integrationElement(quadPoint.position());

```

With these preparations done, we can assemble the first part of the stiffness matrix, corresponding to the velocity–velocity coupling. For two d -valued velocity basis functions $\varphi_i^k = \mathbf{e}_k \varphi_i$ and $\varphi_j^l = \mathbf{e}_l \varphi_j$ we need to compute

$$a_e(\varphi_i^k, \varphi_j^l) := \int_e \nabla \varphi_i^k \nabla \varphi_j^l dx = \delta_{kl} \int_e \nabla \varphi_i \nabla \varphi_j dx$$

on the current element e , where φ_i and φ_j are the corresponding scalar basis functions. The code first computes the derivatives of the velocity shape functions at the current quadrature point, and then uses the matrix in `jacobianInverseTransposed` to transform the shape functions gradients to gradients of the actual basis functions defined on the grid element.

```

101   std::vector<FieldMatrix<double,1,dim> > referenceGradients;
102   velocityLocalFiniteElement.localBasis().evaluateJacobian(
103       quadPoint.position(),
104       referenceGradients);
105
106   // Compute the shape function gradients on the grid element
107   std::vector<FieldVector<double,dim> > gradients(referenceGradients.size());
108   for (size_t i=0; i<gradients.size(); i++)
109       jacobianInverseTransposed.mv(referenceGradients[i][0], gradients[i]);

```

With the velocity basis function gradients at hand we can assemble the velocity contribution to the stiffness matrix:

```

114   for (size_t i=0; i<velocityLocalFiniteElement.size(); i++)
115       for (size_t j=0; j<velocityLocalFiniteElement.size(); j++)
116           for (size_t k=0; k<dim; k++)
117               {
118                   size_t row = localView.tree().child(_0,k).localIndex(i);

```

```

119         size_t col = localView.tree().child(_0,k).localIndex(j);
120         elementMatrix[row][col] += ( gradients[i] * gradients[j] )
121                                     * quadPoint.weight() * integrationElement;
122     }

```

Noteworthy here are the Lines 118–119 which, for two given shape functions from the finite element basis tree, obtain the flat numbering used to index the element stiffness matrix. The expression `child(_0,k)` singles out the tree leaf for the k -th component of the velocity basis. The loop variables i and j run over the shape functions in this set, and

```
localView.tree().child(_0,k).localIndex(i);
```

returns the corresponding scalar index for this shape function in the set of *all* shape functions of the Taylor–Hood basis on this element. This is the *local* index $l_{\Lambda|_e}^{\text{local}}(\cdot)$ of Section 1.5. Line 121 then updates the corresponding (scalar) element matrix entry with the correctly weighted product of the two gradients $\nabla\varphi_i$ and $\nabla\varphi_j$.

Once this part is understood, computing the velocity–pressure coupling terms is easy. For a given velocity shape function φ_i^k and pressure shape function θ_j we need to compute

$$b_e(\varphi_i^k, \theta_j) := \int_e \operatorname{div} \varphi_i^k \cdot \theta_j \, dx = \int_e \sum_{l=1}^d \frac{\partial(\varphi_i^k)_l}{\partial x_l} \cdot \theta_j \, dx = \int_e \frac{\partial \varphi_i}{\partial x_k} \cdot \theta_j \, dx = \int_e (\nabla \varphi_i)_k \cdot \theta_j \, dx.$$

At this point in the code the value of $\nabla\varphi_i$ at the current quadrature point has been computed already, but value of θ_j is still unknown. The values for all i are evaluated by the following two lines:

```

131     std::vector<FieldVector<double,1> > pressureValues;
132     pressureLocalFiniteElement.localBasis().evaluateFunction(
133         quadPoint.position(),
134         pressureValues);

```

Then, the actual matrix assembly of the bilinear form $b_e(\cdot, \cdot)$ is

```

139     for (size_t i=0; i<velocityLocalFiniteElement.size(); i++)
140         for (size_t j=0; j<pressureLocalFiniteElement.size(); j++)
141             for (size_t k=0; k<dim; k++)
142                 {
143                     size_t vIndex = localView.tree().child(_0,k).localIndex(i);
144                     size_t pIndex = localView.tree().child(_1).localIndex(j);
145
146                     elementMatrix[vIndex][pIndex] +=
147                         gradients[i][k] * pressureValues[j]
148                         * quadPoint.weight() * integrationElement;
149                     elementMatrix[pIndex][vIndex] +=
150                         gradients[i][k] * pressureValues[j]
151                         * quadPoint.weight() * integrationElement;
152                 }

```

Line 143 computes the flat local index of φ_i^k again, and Line 144 computes the index for θ_j (remember that $_1$ denotes the pressure basis). Finally, Lines 146–151 then compute the integrand value $(\nabla\varphi_i)_k \cdot \theta_j$, and add the resulting terms to the matrix. This concludes the implementation of the local assembler for the Stokes problem.

References

- [1] Randolph E. Bank. Hierarchical bases and the finite element method. *Acta Numerica*, 5:1–43, 1996. doi: 10.1017/S0962492900002610.
- [2] Daniele Boffi, Franco Brezzi, and Michel Fortin. *Mixed Finite Element Methods and Applications*. Springer, 2013.
- [3] Dietrich Braess. *Finite Elemente*. Springer, 2013.
- [4] Philippe G. Ciarlet. *The Finite Element Method for Elliptic Problems*. North-Holland, 1978.
- [5] J. Austin Cottrell, Thomas J. R. Hughes, and Yuri Bazilevs. *Isogeometric Analysis*. Wiley, 2009.
- [6] Christian Engwer, Carsten Gräser, Steffen Müthing, and Oliver Sander. The interface for functions in the dune-functions module. *Archive of Numerical Software*, 5(1):95–105, 2017. doi: 10.11588/ans.2017.1.27683.
- [7] Jan S. Hesthaven and Tim Warburton. *Nodal Discontinuous Galerkin Methods*. Springer, 2008.
- [8] Volker John. *Finite Element Methods for Incompressible Flow Problems*. Springer, 2016.
- [9] Andrew Koenig and Barbara E. Moo. Templates and duck typing. *Dr. Dobbs's*, 2005. URL www.drdobbs.com/templates-and-duck-typing/184401971. online.
- [10] Nicolas Moës, John Dolbow, and Ted Belytschko. A finite element method for crack growth without remeshing. *International Journal for Numerical Methods in Engineering*, 46(1):131–150, 1999.
- [11] Steffen Müthing. *A Flexible Framework for Multi Physics and Multi Domain PDE Simulations*. PhD thesis, Universität Stuttgart, 2015.
- [12] Rolf Rannacher and Stefan Turek. Simple nonconforming quadrilateral Stokes element. *Numerical Methods for Partial Differential Equations*, 8:97–111, 1992.
- [13] Rob Schreiber and Herbert B. Keller. Driven cavity flows by efficient numerical techniques. *Journal of Computational Physics*, 49(2):310–333, 1983.