

Alfred Schmidt and Kunibert G. Siebert

Design of Adaptive Finite
Element Software:
The Finite Element Toolbox
ALBERTA

SPIN Springer's internal project number, if known

– Monograph –

October 18, 2004

Springer

Berlin Heidelberg New York
Hong Kong London
Milan Paris Tokyo

Preface

During the last years, scientific computing has become an important research branch located between applied mathematics and applied sciences and engineering. Nowadays, in numerical mathematics not only simple model problems are treated, but modern and well-founded mathematical algorithms are applied to solve complex problems of real life applications. Such applications are demanding for computational realization and need suitable and robust tools for a flexible and efficient implementation. Modularity and abstract concepts allow for an easy transfer of methods to different applications.

Inspired by and parallel to the investigation of real life applications, numerical mathematics has built and improved many modern algorithms which are now standard tools in scientific computing. Examples are adaptive methods, higher order discretizations, fast linear and non-linear iterative solvers, multi-level algorithms, etc. These mathematical tools are able to reduce computing times tremendously and for many applications a simulation can only be realized in a reasonable time frame using such highly efficient algorithms.

A very flexible software is needed when working in both fields of scientific computing and numerical mathematics. We developed the toolbox **ALBERTA**¹ for meeting these requirements. Our intention in the design of **ALBERTA** is threefold: First, it is a toolbox for fast and flexible implementation of efficient software for real life applications, based on the modern algorithms mentioned above. Secondly, in an interplay with mathematical analysis, **ALBERTA** is an environment for improving existent, or developing new numerical methods. And finally, it allows the direct integration of such new or improved methods in existing simulation software.

Before having **ALBERTA**, we worked with a variety of solvers, each designed for the solution of one single application. Most of them were based on data structures specifically designed for one single application. A combination of different solvers or exchanging modules between programs was hard to do.

¹The original name of the toolbox was **ALBERT**. Due to copyright reasons, we had to rename it and we have chosen **ALBERTA**.

Facing these problems, we wanted to develop a general adaptive finite element environment, open for implementing a large class of applications, where an exchange of modules and a coupling of different solvers is easy to realize.

Such a toolbox has to be based on a solid concept which is still open for extensions as science develops. Such a solid concept can be derived from a mathematical abstraction of problem classes, numerical methods, and solvers. Our mathematical view of numerical algorithms, especially finite element methods, is based on our education and scientific research in the departments for applied mathematics at the universities of Bonn and Freiburg. This view point has greatly inspired the abstract concepts of ALBERTA as well as their practical realization, reflected in the main data structures. The robustness and flexible extensibility of our concept was approved in various applications from physics and engineering, like computational fluid dynamics, structural mechanics, industrial crystal growth, etc. as well as by the validation of new mathematical methods.

ALBERTA is a library with data structures and functions for adaptive finite element simulations in one, two, and three space dimension, written in the programming language ANSI-C. Shortly after finishing the implementation of the first version of ALBERTA and using it for first scientific applications, we confronted students with it in a course about finite element methods. The idea was to work on more interesting projects in the course and providing a strong foundation for an upcoming diploma thesis. Using ALBERTA in education then required a documentation of data structures and functions. The numerical course tutorials were the basis for a description of the background and concepts of adaptive finite elements.

The combination of the abstract and concrete description resulted in a manual for ALBERTA and made it possible that it is now used world wide in universities and research centers. The interest from other scientists motivated a further polishing of the manual as well as the toolbox itself, and resulted in this book.

These notes are organized as follows: In Chapter 1 we describe the concepts of adaptive finite element methods and its ingredients like the domain discretization, finite element basis functions and degrees of freedom, numerical integration via quadrature formulas for the assemblage of discrete systems, and adaptive algorithms.

The second chapter is a tutorial for using ALBERTA without giving much details about data structures and functions. The implementation of three model problems is presented and explained. We start with the easy and straight forward implementation of the Poisson problem to learn about the basics of ALBERTA. The examples with the implementation of a nonlinear reaction-diffusion problem and the time dependent heat equation are more involved and show the tools of ALBERTA for attacking more complex problems. The chapter is closed with a short introduction to the installation of the

ALBERTA distribution enclosed to this book in a UNIX/Linux environment. Visit the ALBERTA web site

<http://www.alberta-fem.de/>

for updates, more information, FAQ, contributions, pictures from different projects, etc.

The realization of data structures and functions in ALBERTA is based on the abstract concepts presented in Chapter 1. A detailed description of all data structures and functions of ALBERTA is given in Chapter 3. The book closes with separate lists of all data types, symbolic constants, functions, and macros.

The cover picture of this book shows the ALBERTA logo, combined with a locally refined cogwheel mesh [17], and the norm of the velocity from a calculation of edge tones in a flute [4].

Starting first as a two-men-project, ALBERTA is evolving and now there are more people maintaining and extending it. We are grateful for a lot of substantial contributions coming from: Michael Fried, who was the first brave man besides us to use ALBERT, Claus-Justus Heine, Daniel Köster, and Oliver Kriessl. Daniel and Claus in particular set up the GNU configure tools for an easy, platform-independent installation of the software.

We are indebted to the authors of the gltools, especially Jürgen Fuhrmann, and also to the developers of GRAPE, especially Bernard Haasdonk, Robert Klöfkorn, Mario Ohlberger, and Martin Rumpf.

We want to thank the Department of Mathematics at the University of Maryland (USA), in particular Ricardo H. Nochetto, where part of the documentation was written during a visit of the second author. We appreciate the invitation of the Isaac Newton Institute in Cambridge (UK) where we could meet and work intensively on the revision of the manual for three weeks.

We thank our friends, distributed all over the world, who have pointed out a lot of typos in the manual and suggested several improvements for ALBERTA.

Last but not least, ALBERTA would not have come into being without the stimulating atmosphere in the group in Freiburg, which was the perfect environment for working on this project. We want to express our gratitude to all former colleagues, especially Gerhard Dziuk.

Bremen and Augsburg, October 2004

Alfred Schmidt and Kunibert G. Siebert

Contents

Introduction	1
1 Concepts and abstract algorithms	9
1.1 Mesh refinement and coarsening	9
1.1.1 Refinement algorithms for simplicial meshes	12
1.1.2 Coarsening algorithm for simplicial meshes	18
1.1.3 Operations during refinement and coarsening	20
1.2 The hierarchical mesh	22
1.3 Degrees of freedom	24
1.4 Finite element spaces and finite element discretization	25
1.4.1 Barycentric coordinates	26
1.4.2 Finite element spaces	29
1.4.3 Evaluation of finite element functions	29
1.4.4 Interpolation and restriction during refinement and coarsening	32
1.4.5 Discretization of 2nd order problems	35
1.4.6 Numerical quadrature	38
1.4.7 Finite element discretization of 2nd order problems	39
1.5 Adaptive Methods	42
1.5.1 Adaptive method for stationary problems	42
1.5.2 Mesh refinement strategies	43
1.5.3 Coarsening strategies	47
1.5.4 Adaptive methods for time dependent problems	49
2 Implementation of model problems	55
2.1 Poisson equation	56
2.1.1 Include file and global variables	56
2.1.2 The main program for the Poisson equation	57
2.1.3 The parameter file for the Poisson equation	59
2.1.4 Initialization of the finite element space	60
2.1.5 Functions for leaf data	60

2.1.6	Data of the differential equation	62
2.1.7	The assemblage of the discrete system	63
2.1.8	The solution of the discrete system	65
2.1.9	Error estimation	66
2.2	Nonlinear reaction–diffusion equation	68
2.2.1	Program organization and header file	69
2.2.2	Global variables	71
2.2.3	The main program for the nonlinear reaction–diffusion equation	71
2.2.4	Initialization of the finite element space and leaf data	72
2.2.5	The build routine	72
2.2.6	The solve routine	73
2.2.7	The estimator for the nonlinear problem	73
2.2.8	Initialization of problem dependent data	75
2.2.9	The parameter file for the nonlinear reaction–diffusion equation	79
2.2.10	Implementation of the nonlinear solver	80
2.3	Heat equation	93
2.3.1	Global variables	93
2.3.2	The main program for the heat equation	94
2.3.3	The parameter file for the heat equation	96
2.3.4	Functions for leaf data	98
2.3.5	Data of the differential equation	99
2.3.6	Time discretization	100
2.3.7	Initial data interpolation	100
2.3.8	The assemblage of the discrete system	101
2.3.9	Error estimation	105
2.3.10	Time steps	107
2.4	Installation of ALBERTA and file organization	111
2.4.1	Installation	111
2.4.2	File organization	111
3	Data structures and implementation	113
3.1	Basic types, utilities, and parameter handling	113
3.1.1	Basic types	113
3.1.2	Message macros	114
3.1.3	Memory allocation and deallocation	118
3.1.4	Parameters and parameter files	122
3.1.5	Parameters used by the utilities	127
3.1.6	Generating filenames for meshes and finite element data	127
3.2	Data structures for the hierarchical mesh	128
3.2.1	Constants describing the dimension of the mesh	128
3.2.2	Constants describing the elements of the mesh	129
3.2.3	Neighbour information	129

3.2.4	Element indices	130
3.2.5	The BOUNDARY data structure	130
3.2.6	The local indexing on elements	132
3.2.7	The MACRO_EL data structure	132
3.2.8	The EL data structure	134
3.2.9	The EL_INFO data structure	135
3.2.10	The NEIGH , OPP_VERTEX and EL_TYPE macros	137
3.2.11	The INDEX macro	137
3.2.12	The LEAF_DATA_INFO data structure	138
3.2.13	The RC_LIST_EL data structure	140
3.2.14	The MESH data structure	141
3.2.15	Initialization of meshes	143
3.2.16	Reading macro triangulations	144
3.2.17	Writing macro triangulations	150
3.2.18	Import and export of macro triangulations from/to other formats	151
3.2.19	Mesh traversal routines	153
3.3	Administration of degrees of freedom	161
3.3.1	The DOF_ADMIN data structure	162
3.3.2	Vectors indexed by DOFs: The DOF*_VEC data structures	164
3.3.3	Interpolation and restriction of DOF vectors during mesh refinement and coarsening	167
3.3.4	The DOF_MATRIX data structure	168
3.3.5	Access to global DOFs: Macros for iterations using DOF indices	170
3.3.6	Access to local DOFs on elements	171
3.3.7	BLAS routines for DOF vectors and matrices	173
3.3.8	Reading and writing of meshes and vectors	173
3.4	The refinement and coarsening implementation	176
3.4.1	The refinement routines	176
3.4.2	The coarsening routines	182
3.5	Implementation of basis functions	183
3.5.1	Data structures for basis functions	184
3.5.2	Lagrange finite elements	190
3.5.3	Piecewise constant finite elements	191
3.5.4	Piecewise linear finite elements	191
3.5.5	Piecewise quadratic finite elements	195
3.5.6	Piecewise cubic finite elements	200
3.5.7	Piecewise quartic finite elements	204
3.5.8	Access to Lagrange elements	206
3.6	Implementation of finite element spaces	206
3.6.1	The finite element space data structure	206
3.6.2	Access to finite element spaces	207
3.7	Routines for barycentric coordinates	208

3.8	Data structures for numerical quadrature	210
3.8.1	The QUAD data structure	210
3.8.2	The QUAD_FAST data structure	212
3.8.3	Integration over sub-simplices (edges/faces)	215
3.9	Functions for the evaluation of finite elements	216
3.10	Calculation of norms for finite element functions.....	221
3.11	Calculation of errors of finite element approximations	222
3.12	Tools for the assemblage of linear systems	224
3.12.1	Assembling matrices and right hand sides	224
3.12.2	Data structures and function for matrix assemblage ...	227
3.12.3	Data structures for storing pre-computed integrals of basis functions	236
3.12.4	Data structures and functions for vector update	243
3.12.5	Dirichlet boundary conditions	247
3.12.6	Interpolation into finite element spaces	248
3.13	Data structures and procedures for adaptive methods	249
3.13.1	ALBERTA adaptive method for stationary problems ...	249
3.13.2	Standard ALBERTA marking routine	255
3.13.3	ALBERTA adaptive method for time dependent problems	256
3.13.4	Initialization of data structures for adaptive methods ..	260
3.14	Implementation of error estimators	263
3.14.1	Error estimator for elliptic problems	263
3.14.2	Error estimator for parabolic problems	266
3.15	Solver for linear and nonlinear systems	268
3.15.1	General linear solvers	268
3.15.2	Linear solvers for DOF matrices and vectors	272
3.15.3	Access of functions for matrix-vector multiplication....	274
3.15.4	Access of functions for preconditioning	275
3.15.5	Multigrid solvers	277
3.15.6	Nonlinear solvers	282
3.16	Graphics output	285
3.16.1	One and two dimensional graphics subroutines	286
3.16.2	gltools interface	290
3.16.3	GRAPE interface	293
	References	295
	Index	301
	Data types, symbolic constants, functions, and macros	311
	Data types	311
	Symbolic constants	311
	Functions	312
	Macros	315

Introduction

Finite element methods provide a widely used tool for the solution of problems with an underlying variational structure. Modern numerical analysis and implementations for finite elements provide more and more tools for the efficient solution of large-scale applications. Efficiency can be increased by using local mesh adaption, by using higher order elements, where applicable, and by fast solvers.

Adaptive procedures for the numerical solution of partial differential equations started in the late 70's and are now standard tools in science and engineering. Adaptive finite element methods are a meaningful approach for handling multi scale phenomena and making realistic computations feasible, specially in 3d.

There exists a vast variety of books about finite elements. Here, we only want to mention the books by Ciarlet [25], and Brenner and Scott [23] as the most prominent ones. The book by Brenner and Scott also contains an introduction to multi-level methods.

The situation is completely different for books about adaptive finite elements. Only few books can be found with introductory material about the mathematics of adaptive finite element methods, like the books by Verfürth [73], and Ainsworth and Oden [2]. Material about more practical issues like adaptive techniques and refinement procedures can for example be found in [3, 5, 8, 44, 46].

Another basic ingredient for an adaptive finite element method is the a posteriori error estimator which is main object of interest in the analysis of adaptive methods. While a general theory exists for these estimators in the case of linear and mildly nonlinear problems [10, 73], highly nonlinear problems usually still need a special treatment, see [24, 33, 54, 55, 69] for instance. There exist a lot of different approaches to (and a large number of articles about) the derivation of error estimates, by residual techniques, dual techniques, solution of local problems, hierarchical approaches, etc., a fairly incomplete list of references is [1, 3, 7, 13, 21, 36, 52, 72].

Although adaptive finite element methods in practice construct a sequence of discrete solutions which converge to the true solution, this convergence could only be proved recently for linear elliptic problem [50, 51, 52] and for the nonlinear Laplacian [70], based on the fundamental paper [31]. For a modification of the convergent algorithm in [50], quasi-optimality of the adaptive method was proved in [16] and [67].

During the last years there has been a great progress in designing finite element software. It is not possible to mention all freely available packages. Examples are [5, 11, 12, 49, 62], and an continuously updated list of other available finite element codes and resources can for instance be found at

http://www.engr.usask.ca/~macphed/finite/fe_resources/.

Adaptive finite element methods and basic concepts of ALBERTA

Finite element methods calculate approximations to the true solution in some finite dimensional function space. This space is built from *local function spaces*, usually polynomials of low order, on elements of a partitioning of the domain (the *mesh*). An adaptive method adjusts this mesh (or the local function space, or both) to the solution of the problem. This adaptation is based on information extracted from *a posteriori error estimators*.

The basic iteration of an adaptive finite element code for a stationary problem is

- assemble and solve the discrete system;
- calculate the error estimate;
- adapt the mesh, when needed.

For time dependent problems, such an iteration is used in each time step, and the step size of a time discretization may be subject to adaptivity, too.

The core part of every finite element program is the problem dependent assembly and solution of the discretized problem. This holds for programs that solve the discrete problem on a fixed mesh as well as for adaptive methods that automatically adjust the underlying mesh to the actual problem and solution. In the adaptive iteration, the assemblage and solution of a discrete system is necessary after each mesh change. Additionally, this step is usually the most time consuming part of that iteration.

A general finite element toolbox must provide flexibility in problems and finite element spaces while on the other hand this core part can be performed efficiently. Data structures are needed which allow an easy and efficient implementation of the problem dependent parts and also allow to use adaptive methods, mesh modification algorithms, and fast solvers for linear and nonlinear discrete problems by calling library routines. On one hand, large flexibility is needed in order to choose various kinds of finite element spaces, with higher order elements or combinations of different spaces for mixed methods or systems. On the other hand, the solution of the resulting discrete systems may

profit enormously from a simple vector-oriented storage of coefficient vectors and matrices. This also allows the use of optimized solver and BLAS libraries. Additionally, multilevel preconditioners and solvers may profit from hierarchy information, leading to highly efficient solvers for the linear (sub-) problems.

ALBERTA [59, 60, 62] provides all those tools mentioned above for the efficient implementation and adaptive solution of general nonlinear problems in two and three space dimensions. The design of the ALBERTA data structures allows a dimension independent implementation of problem dependent parts. The mesh adaptation is done by local refinement and coarsening of mesh elements, while the same local function space is used on all mesh elements.

Starting point for the design of ALBERTA data structures is the abstract concept of a finite element space defined (similar to the definition of a single finite element by Ciarlet [25]) as a triple consisting of

- a collection of *mesh elements*;
- a set of local *basis functions* on a single element, usually a restriction of global basis functions to a single element;
- a connection of local and global basis functions giving global *degrees of freedom* for a finite element function.

This directly leads to the definition of three main groups of data structures:

- data structures for geometric information storing the underlying mesh together with element coordinates, boundary type and geometry, etc.;
- data structures for finite element information providing values of local basis functions and their derivatives;
- data structures for algebraic information linking geometric data and finite element data.

Using these data structures, the finite element toolbox ALBERTA provides the whole abstract framework like finite element spaces and adaptive strategies, together with hierarchical meshes, routines for mesh adaptation, and the complete administration of finite element spaces and the corresponding degrees of freedom (DOFs) during mesh modifications. The underlying data structures allow a flexible handling of such information. Furthermore, tools for numerical quadrature, matrix and load vector assembly as well as solvers for (linear) problems, like conjugate gradient methods, are available.

A specific problem can be implemented and solved by providing just some problem dependent routines for evaluation of the (linearized) differential operator, data, nonlinear solver, and (local) error estimators, using all the tools above mentioned from a library.

Both geometric and finite element information strongly depend on the space dimension. Thus, mesh modification algorithms and basis functions are implemented for one (1d), two (2d), and three (3d) dimensions separately and are provided by the toolbox. Everything besides that can be formulated in such a way that the dimension only enters as a parameter (like size of local coordinate vectors, e.g.). For usual finite element applications this results in

a dimension independent programming, where all dimension dependent parts are hidden in a library. This allows a dimension independent programming of applications to the greatest possible extent.

The remaining parts of the introduction give a short overview over the main concepts, details are then given in Chapter 1.

The hierarchical mesh

The underlying mesh is a conforming triangulation of the computational domain into simplices, i.e. intervals (1d), triangles (2d), or tetrahedra (3d). The simplicial mesh is generated by refinement of a given initial triangulation. Refined parts of the mesh can be de-refined, but elements of the initial triangulation (*macro elements*) must not be coarsened. The refinement and coarsening routines construct a sequence of nested meshes with a hierarchical structure. In ALBERTA, the recursive refinement by bisection is implemented, using the notation of Kossaczky [44].

During refinement, new degrees of freedom are created. A single degree of freedom is shared by all elements which belong to the support of the corresponding finite element basis function (compare next paragraph). The mesh refinement routines must create a new DOF only once and give access to this DOF from all elements sharing it. Similarly, DOFs are handled during coarsening. This is done in cooperation with the DOF administration tool, see below.

The bisectioning refinement of elements leads naturally to nested meshes with the hierarchical structure of binary trees, one tree for every element of the initial triangulation. Every interior node of that tree has two pointers to the two children; the leaf elements are part of the actual triangulation, which is used to define the finite element space(s). The whole triangulation is a list of given macro elements together with the associated binary trees. The hierarchical structure allows the generation of most information by the hierarchy, which reduces the amount of data to be stored. Some information is stored on the (leaf) elements explicitly, other information is located at the macro elements and is transferred to the leaf elements while traversing through the binary tree. Element information about vertex coordinates, domain boundaries, and element adjacency can be computed easily and very fast from the hierarchy, when needed. Data stored explicitly at tree elements can be reduced to pointers to the two possible children and information about local DOFs (for leaf elements). Furthermore, the hierarchical mesh structure directly leads to multilevel information which can be used by multilevel preconditioners and solvers.

Access to mesh elements is available solely via routines which traverse the hierarchical trees; no direct access is possible. The traversal routines can give access to all tree elements, only to leaf elements, or to all elements which belong to a single hierarchy level (for a multilevel application, e.g.). In order to perform operations on visited elements, the traversal routines call a subroutine

which is given to them as a parameter. Only such element information which is needed by the current operation is generated during the tree traversal.

Finite elements

The values of a finite element function or the values of its derivatives are uniquely defined by the values of its DOFs and the values of the basis functions or the derivatives of the basis functions connected with these DOFs. We follow the concept of finite elements which are given on a single element S in local coordinates: Finite element functions on an element S are defined by a finite dimensional function space $\bar{\mathbb{P}}$ on a reference element \bar{S} and the (one to one) mapping $\lambda^S : \bar{S} \rightarrow S$ from the reference element \bar{S} to the element S . In this situation the non vanishing basis functions on an arbitrary element are given by the set of basis functions of $\bar{\mathbb{P}}$ in local coordinates λ^S . Also, derivatives are given by the derivatives of basis functions on $\bar{\mathbb{P}}$ and derivatives of λ^S .

Each local basis function on S is uniquely connected to a global degree of freedom, which can be accessed from S via the DOF administration tool. ALBERTA supports basis functions connected with DOFs, which are located at vertices of elements, at edges, at faces (in 3d), or in the interior of elements. DOFs at a vertex are shared by all elements which meet at this vertex, DOFs at an edge or face are shared by all elements which contain this edge or face, and DOFs inside an element are not shared with any other element. The support of the basis function connected with a DOF is the patch of all elements sharing this DOF.

For a very general approach, we only need a vector of the basis functions (and its derivatives) on \bar{S} and a function for the communication with the DOF administration tool in order to access the degrees of freedom connected to local basis functions. By such information every finite element function (and its derivatives) is uniquely described on every element of the mesh.

During mesh modifications, finite element functions must be transformed to the new finite element space. For example, a discrete solution on the old mesh yields a good initial guess for an iterative solver and a smaller number of iterations for a solution of the discrete problem on the new mesh. Usually, these transformations can be realized by a sequence of local operations. Local interpolations and restrictions during refinement and coarsening of elements depend on the function space $\bar{\mathbb{P}}$ and the refinement of \bar{S} only. Thus, the subroutine for interpolation during an atomic mesh refinement is the efficient implementation of the representation of coarse grid functions by fine grid functions on \bar{S} and its refinement. A restriction during coarsening is implemented using similar information.

Lagrange finite element spaces up to order four are currently implemented in one, two, and three dimensions. This includes the communication with the DOF administration as well as the interpolation and restriction routines.

Degrees of freedom

Degrees of freedom (DOFs) connect finite element data with geometric information of a triangulation. For general applications, it is necessary to handle several different sets of degrees of freedom on the same triangulation. For example, in mixed finite element methods for the Navier-Stokes problem, different polynomial degrees are used for discrete velocity and pressure functions.

During adaptive refinement and coarsening of a triangulation, not only elements of the mesh are created and deleted, but also degrees of freedom together with them. The geometry is handled dynamically in a hierarchical binary tree structure, using pointers from parent elements to their children. For data corresponding to DOFs, which are usually involved with matrix-vector operations, simpler storage and access methods are more efficient. For that reason every DOF is realized just as an integer index, which can easily be used to access data from a vector or to build matrices that operate on vectors of DOF data. This results in a very efficient access during matrix/vector operations and in the possibility to use libraries for the solution of linear systems with a sparse system matrix ([29], e.g.).

Using this realization of DOFs two major problems arise:

- During refinement of the mesh, new DOFs are added, and additional indices are needed. The total range of used indices has to be enlarged. At the same time, all vectors and matrices that use these DOF indices have to be adjusted in size, too.
- During coarsening of the mesh, DOFs are deleted. In general, the deleted DOF is not the one which corresponds to the largest integer index. Holes with unused indices appear in the total range of used indices and one has to keep track of all used and unused indices.

These problems are solved by a general DOF administration tool. During refinement, it enlarges the ranges of indices, if no unused indices produced by a previous coarsening are available. During coarsening, a book-keeping about used and unused indices is done. In order to reestablish a contiguous range of used indices, a compression of DOFs can be performed; all DOFs are renumbered such that all unused indices are shifted to the end of the index range, thus removing holes of unused indices. Additionally, all vectors and matrices connected to these DOFs are adjusted correspondingly. After this process, vectors do not contain holes anymore and standard operations like BLAS1 routines can be applied and yield optimal performance.

In many cases, information stored in DOF vectors has to be adjusted to the new distribution of DOFs during mesh refinement and coarsening. Each DOF vector can provide pointers to subroutines that implements these operations on data (which usually strongly depend on the corresponding finite element basis). Providing such a pointer, a DOF vector will automatically be transformed during mesh modifications.

All tasks of the DOF administration are performed automatically during refinement and coarsening for every kind and combination of finite elements defined on the mesh.

Adaptive solution of the discrete problem

The aim of adaptive methods is the generation of a mesh which is adapted to the problem such that a given criterion, like a tolerance for the estimated error between exact and discrete solution, is fulfilled by the finite element solution on this mesh. An optimal mesh should be as coarse as possible while meeting the criterion, in order to save computing time and memory requirements. For time dependent problems, such an adaptive method may include mesh changes in each time step and control of time step sizes. The philosophy implemented in ALBERTA is to change meshes successively by local refinement or coarsening, based on error estimators or error indicators, which are computed a posteriori from the discrete solution and given data on the current mesh.

Several adaptive strategies are proposed in the literature, that give criteria which mesh elements should be marked for refinement. All strategies are based on the idea of an equidistribution of the local error to all mesh elements. Babuška and Rheinboldt [3] motivate that for stationary problems a mesh is almost optimal when the local errors are approximately equal for all elements. So, elements where the error indicator is large will be marked for refinement, while elements with a small estimated indicator are left unchanged or are marked for coarsening. In time dependent problems, the mesh is adapted to the solution in every time step using a posteriori information like in the stationary case. As a first mesh for the new time step we use the adaptive mesh from the previous time step. Usually, only few iterations of the adaptive procedure are then needed for the adaptation of the mesh for the new time step. This may be accompanied by an adaptive control of time step sizes.

Given pointers to the problem dependent routines for assembling and solution of the discrete problems, as well as an error estimator/indicator, the adaptive method for finding a solution on a quasi-optimal mesh can be performed as a black-box algorithm. The problem dependent routines are used for the calculation of discrete solutions on the current mesh and (local) error estimates. Here, the problem dependent routines heavily make use of library tools for assembling system matrices and right hand sides for an arbitrary finite element space, as well as tools for the solution of linear or nonlinear discrete problems. On the other hand, any specialized algorithm may be added if needed. The marking of mesh elements is based on general refinement and coarsening strategies relying on the local error indicators. During the following mesh modification step, DOF vectors are transformed automatically to the new finite element spaces as described in the previous paragraphs.

Dimension independent program development

Using black-box algorithms, the abstract definition of basis functions, quadrature formulas and the DOF administration tool, only few parts of the finite element code depend on the dimension. Usually, all dimension dependent parts are hidden in the library. Hence, program development can be done in 1d or 2d, where execution is usually much faster and debugging is much easier (because of simple 1d and 2d visualization, e.g., which is much more involved in 3d). With no (or maybe few) additional changes, the program will then also work in 3d. This approach leads to a tremendous reduction of program development time for 3d problems.

Notations. For a differentiable function $f: \Omega \rightarrow \mathbb{R}$ on a domain $\Omega \subset \mathbb{R}^d$, $d = 1, 2, 3$, we set

$$\nabla f(x) = (f_{,x_1}(x), \dots, f_{,x_d}(x)) = \left(\frac{\partial}{\partial x_1} f(x), \dots, \frac{\partial}{\partial x_d} f(x) \right)$$

and

$$D^2 f(x) = (f_{,x_k x_l})_{k,l=1,\dots,d} = \left(\frac{\partial^2}{\partial x_k \partial x_l} f(x) \right)_{k,l=1,\dots,d}.$$

For a vector valued, differentiable function $f = (f_1, \dots, f_n): \Omega \rightarrow \mathbb{R}^n$ we write

$$\nabla f(x) = (f_{i,x_1}(x), \dots, f_{i,x_d}(x))_{i=1,\dots,n} = \left(\frac{\partial}{\partial x_1} f_i(x), \dots, \frac{\partial}{\partial x_d} f_i(x) \right)_{i=1,\dots,n}$$

and

$$D^2 f(x) = (f_{i,x_k x_l})_{\substack{i=1,\dots,n \\ k,l=1,\dots,d}} = \left(\frac{\partial^2}{\partial x_k \partial x_l} f_i(x) \right)_{\substack{i=1,\dots,n \\ k,l=1,\dots,d}}.$$

By $L^p(\Omega)$, $1 \leq p \leq \infty$, we denote the usual Lebesgue spaces with norms

$$\|f\|_{L^p(\Omega)} = \left(\int_{\Omega} |f(x)|^p dx \right)^{1/p} \quad \text{for } p < \infty$$

and

$$\|f\|_{L^\infty(\Omega)} = \operatorname{ess\,sup}_{x \in \Omega} |f(x)|.$$

The Sobolev space of functions $u \in L^2(\Omega)$ with weak derivatives $\nabla u \in L^2(\Omega)$ is denoted by $H^1(\Omega)$ with semi norm

$$|u|_{H^1(\Omega)} = \left(\int_{\Omega} |\nabla u(x)|^2 dx \right)^{1/2}$$

and norm

$$\|u\|_{H^1(\Omega)} = \left(\|u\|_{L^2(\Omega)}^2 + |u|_{H^1(\Omega)}^2 \right)^{1/2}.$$

Concepts and abstract algorithms

1.1 Mesh refinement and coarsening

In this section, we describe the basic algorithms for the local refinement and coarsening of simplicial meshes in two and three dimensions. In 1d the grid is built from intervals, in 2d from triangles, and in 3d from tetrahedra. We restrict ourselves here to simplicial meshes, for several reasons:

1. A simplex is one of the most simple geometric types and complex domains may be approximated by a set of simplices quite easily.
2. Simplicial meshes allow local refinement (see Fig. 1.1) without the need of non-conforming meshes (hanging nodes), parametric elements, or mixture of element types (which is the case for quadrilateral meshes, e.g., see Fig. 1.2).
3. Polynomials of any degree are easily represented on a simplex using local (barycentric) coordinates.




Fig. 1.1. Global and local refinement of a triangular mesh.

First of all we start with the definition of a simplex, parametric simplex and triangulation:

Definition 1.1 (Simplex).

- a) Let $a_0, \dots, a_d \in \mathbb{R}^n$ be given such that $a_1 - a_0, \dots, a_d - a_0$ are linear independent vectors in \mathbb{R}^n . The convex set




Fig. 1.2. Local refinements of a rectangular mesh: with hanging nodes, conforming closure using bisected rectangles, and conforming closure using triangles. Using a conforming closure with rectangles, a local refinement has always global effects up to the boundary.

$$S = \text{conv hull}\{a_0, \dots, a_d\}$$

is called a d -simplex in \mathbb{R}^n . For $k < d$ let

$$S' = \text{conv hull}\{a'_0, \dots, a'_k\} \subset \partial S$$

be a k -simplex with $a'_0, \dots, a'_k \in \{a_0, \dots, a_d\}$. Then S' is called a k -sub-simplex of S . A 0-sub-simplex is called vertex, a 1-sub-simplex edge and a 2-sub-simplex face.

b) The standard simplex in \mathbb{R}^d is defined by

$$\hat{S} = \text{conv hull}\{\hat{a}_0 = 0, \hat{a}_1 = e_1, \dots, \hat{a}_d = e_d\},$$

where e_i are the unit vectors in \mathbb{R}^d .

c) Let $F_S: \hat{S} \rightarrow S \subset \mathbb{R}^n$ be an invertible, differentiable mapping. Then S is called a parametric d -simplex in \mathbb{R}^n . The k -sub-simplices S' of S are given by the images of the k -sub-simplices \hat{S}' of \hat{S} . Thus, the vertices a_0, \dots, a_d of S are the points $F_S(\hat{a}_0), \dots, F_S(\hat{a}_d)$.

d) For a d -simplex S , we define

$$h_S := \text{diam}(S) \quad \text{and} \quad \rho_S := \sup\{2r; B_r \subset S \text{ is a } d\text{-ball of radius } r\},$$

the diameter and inball-diameter of S .

Remark 1.2. Every d -simplex S in \mathbb{R}^n is a parametric simplex. Defining the matrix $A_S \in \mathbb{R}^{n \times d}$ by

$$A_S = \begin{bmatrix} \vdots & & \vdots \\ a_1 - a_0 & \cdots & a_d - a_0 \\ \vdots & & \vdots \end{bmatrix},$$

the parameterization $F_S: \hat{S} \rightarrow S$ is given by

$$F_S(\hat{x}) = A_S \hat{x} + a_0. \tag{1.1}$$

Since F_S is affine linear it is differentiable. It is easy to check that $F_S: \hat{S} \rightarrow S$ is invertible and that $F_S(\hat{a}_i) = a_i$, $i = 0, \dots, d$ holds.

Definition 1.3 (Triangulation).

a) Let \mathcal{S} be a set of (parametric) d -simplices and define

$$\Omega = \text{interior} \bigcup_{S \in \mathcal{S}} S \subset \mathbb{R}^n.$$

We call \mathcal{S} a conforming triangulation of Ω , iff for two simplices $S_1, S_2 \in \mathcal{S}$ with $S_1 \neq S_2$ the intersection $S_1 \cap S_2$ is either empty or a complete k -sub-simplex of both S_1 and S_2 for some $0 \leq k < d$.

b) Let \mathcal{S}_k , $k \geq 0$, be a sequence of conforming triangulations. This sequence is called (shape) regular, iff

$$\sup_{k \in \mathbb{N}_0} \max_{S \in \mathcal{S}_k} \max_{\hat{x} \in \hat{S}} \text{cond}(DF_S^t(\hat{x}) \cdot DF_S(\hat{x})) < \infty \quad (1.2)$$

holds, where DF_S is the Jacobian of F_S and $\text{cond}(A) = \|A\| \|A^{-1}\|$ denotes the condition number.

Remark 1.4. For a sequence \mathcal{S}_k , $k \geq 0$, of non-parametric triangulations the regularity condition (1.2) is equivalent to the condition

$$\sup_{k \in \mathbb{N}_0} \max_{S \in \mathcal{S}_k} \frac{h_S}{\rho_S} < \infty.$$

In order to construct a sequence of triangulations, we consider the following situation: An initial (coarse) triangulation \mathcal{S}_0 of the domain is given. We call it *macro triangulation*. It may be generated by hand or by some mesh generation algorithm ([63, 65]).

Some (or all) of the simplices are marked for refinement, depending on some error estimator or indicator. The marked simplices are then refined, i.e. they are cut into smaller ones. After several refinements, some other simplices may be marked for coarsening. Coarsening tries to unite several simplices marked for coarsening into a bigger simplex. A successive refinement and coarsening will produce a sequence of triangulations $\mathcal{S}_0, \mathcal{S}_1, \dots$. The refinement of single simplices that we describe in the next section produces for every simplex of the macro triangulation only a finite and small number of similarity classes for the resulting elements. The coarsening is more or less the inverse process of refinement. This leads to a finite number of similarity classes for all simplices in the whole sequence of triangulations.

The refinement of non-parametric and parametric simplices is the same topological operation and can be performed in the same way. The actual children's shape of parametric elements additionally involves the children's parameterization. In the following we describe the refinement and coarsening for triangulations consisting of non-parametric elements. The refinement of parametric triangulations can be done in the same way, additionally using given parameterizations. Regularity for the constructed sequence can be obtained

with special properties of the parameterizations for parametric elements and the finite number of similarity classes for simplices.

Marking criteria and marking strategies for refinement and coarsening are subject of Section 1.5.

1.1.1 Refinement algorithms for simplicial meshes

For simplicial elements, several refinement algorithms are widely used. The discussion about and description of these algorithms mainly centers around refinement in 2d and 3d since refinement in 1d is straight forward.

One example is regular refinement (“red refinement”), which divides every triangle into four similar triangles, see Fig. 1.3. The corresponding refinement algorithm in three dimensions cuts every tetrahedron into eight tetrahedra, and only a small number of similarity classes occur during successive refinements, see [14, 15]. Unfortunately, hanging nodes arise during local regular refinement. To remove them and create a conforming mesh, in two dimensions some triangles have to be bisected (“green closure”). In three dimensions, several types of irregular refinement are needed for the green closure. This creates more similarity classes, even in two dimensions. Additionally, these bisected elements have to be removed before a further refinement of the mesh, in order to keep the triangulations shape regular.




Fig. 1.3. Global and local regular refinement of triangles and conforming closure by bisection.

Another possibility is to use bisection of simplices only. For every element (triangle or tetrahedron) one of its edges is marked as the *refinement edge*, and the element is refined into two elements by cutting this edge at its midpoint. There are several possibilities to choose such a refinement edge for a simplex, one example is to use the longest edge; Mitchell [48] compared different approaches. We focus on an algorithm where the choice of refinement edges on the macro triangulation prescribes the refinement edges for all simplices that are created during mesh refinement. This makes sure that shape regularity of the triangulations is conserved.

In two dimensions we use the *newest vertex* bisection (in Mitchell’s notation) and in three dimensions the bisection procedure of Kossaczky described in [44]. We use the convention, that all vertices of an element are given fixed *local indices*. Valid indices are 0, 1, for vertices of an interval, 0, 1, and 2 for

vertices of a triangle, and 0, 1, 2, and 3 for vertices of a tetrahedron. Now, the refinement edge for an element is fixed to be the edge between the vertices with local indices 0 and 1. Here we use the convention that in 1d the element itself is called “refinement edge”.

During refinement, the new vertex numbers, and thereby the refinement edges, for the newly created child simplices are prescribed by the refinement algorithm. For both children elements, the index of the newly generated vertex at the midpoint of this edge has the highest local index (2 resp. 3 for triangles and tetrahedra). These numbers are shown in Fig. 1.4 for 1d and 2d, and in Fig. 1.5 for 3d. In 1d and 2d this numbering is the same for all refinement levels. In 3d, one has to make some special arrangements: the numbering of the second child’s vertices does depend on the *type* of the element. There exist three different element types 0, 1, and 2. The type of the elements on the macro triangulation can be prescribed (usually type 0 tetrahedron). The type of the refined tetrahedra is recursively given by the definition that the type of a child element is $((\text{parent's type} + 1) \bmod 3)$. In Fig. 1.5 we used the following convention: for the index set $\{1, 2, 2\}$ on `child[1]` of a tetrahedron of type 0 we use the index 1 and for a tetrahedron of type 1 and 2 the index 2. Fig. 1.6 shows successive refinements of a type 0 tetrahedron, producing tetrahedra of types 1, 2, and 0 again.




Fig. 1.4. Numbering of nodes on parent and children for intervals and triangles.




Fig. 1.5. Numbering of nodes on parent and children for tetrahedra.

By the above algorithm the refinements of simplices are totally determined by the local vertex numbering on the macro triangulation, plus a prescribed type for every macro element in three dimensions. Furthermore, a successive refinement of every macro element only produces a small number of similarity




Fig. 1.6. Successive refinements of a type 0 tetrahedron.

classes. In case of the “generic” triangulation of a (unit) square in 2d and cube in 3d into two triangles resp. six tetrahedra (see Fig. 1.7 for a single triangle and tetrahedron from such a triangulation – all other elements are generated by rotation and reflection), the numbering and the definition of the refinement edge during refinement of the elements guarantee that always the longest edge will be the refinement edge and will be bisected, see Fig. 1.8.

The refinement of a given triangulation now uses the bisection of *single* elements and can be performed either *iteratively* or *recursively*. In 1d, bisection only involves the element which is subject to refinement and thus is a completely local operation. Both variants of refining a given triangulation are the same. In 2d and 3d, bisection of a single element usually involves other elements, resulting in two different algorithms.

For tetrahedra, the first description of such a refinement procedure was given by Bänsch using the iterative variant [8]. It abandons the requirement of one to one inter-element adjacencies during the refinement process and thus needs the intermediate handling of hanging nodes. Two recursive algorithms, which do not create such hanging nodes and are therefore easier to implement, are published by Kossaczky [44] and Maubach [46]. For a special class of

macro triangulations, they result in exactly the same tetrahedral meshes as the iterative algorithm.

In order to keep the mesh conforming during refinement, the bisection of an edge is allowed only when such an edge is the refinement edge for *all* elements which share this edge. Bisection of an edge and thus of all elements around the edge is the *atomic refinement operation*, and no other refinement operation is allowed. See Figs. 1.9 and 1.10 for the two and three-dimensional situations.




Fig. 1.7. Generic elements in two and three dimensions.




Fig. 1.8. Refined generic elements in two and three dimensions.




Fig. 1.9. Atomic refinement operation in two dimensions. The common edge is the refinement edge for both triangles.

If an element has to be refined, we have to collect all elements at its refinement edge. In two dimensions this is either the neighbour opposite this edge or there is no other element in the case that the refinement edge belongs




Fig. 1.10. Atomic refinement operation in three dimensions. The common edge is the refinement edge for all tetrahedra sharing this edge.

to the boundary. In three dimensions we have to loop around the edge and collect all neighbours at this edge. If for all collected neighbours the common edge is the refinement edge too, we can refine the whole patch at the same time by inserting one new vertex in the midpoint of the common refinement edge and bisecting every element of the patch. The resulting triangulation then is a conforming one.

But sometimes the refinement edge of a neighbour is not the common edge. Such a neighbour is *not compatibly divisible* and we have to perform first the atomic refinement operation at the neighbour's refinement edge. In 2d the child of such a neighbour at the common edge is then compatibly divisible; in 3d such a neighbour has to be bisected at most three times and the resulting tetrahedron at the common edge is then compatibly divisible. The recursive refinement algorithm now reads

Algorithm 1.5 (Recursive refinement of one simplex).

```

subroutine recursive_refine( $S$ ,  $S$ )
do
   $\mathcal{A} := \{S' \in S; S' \text{ is not compatibly divisible with } S\}$ 
  for all  $S' \in \mathcal{A}$  do
    recursive_refine( $S'$ ,  $S$ );
  end for
   $\mathcal{A} := \{S' \in S; S' \text{ is not compatibly divisible with } S\}$ 
until  $\mathcal{A} = \emptyset$ 

 $\mathcal{A} := \{S' \in S; S' \text{ is element at the refinement edge of } S\}$ 
for all  $S' \in \mathcal{A}$ 
  bisect  $S'$  into  $S'_0$  and  $S'_1$ 
   $S := S \setminus \{S'\} \cup \{S'_0, S'_1\}$ 
end for

```

In Fig. 1.11 we show a two-dimensional situation where recursion is needed. For all triangles, the longest edge is the refinement edge. Let us assume that triangles A and B are marked for refinement. Triangle A can be refined at once, as its refinement edge is a boundary edge. For refinement of triangle B, we have to recursively refine triangles C and D. Again, triangle

D can be directly refined, so recursion terminates there. This is shown in the second part of the figure. Back in triangle C, this can now be refined together with its neighbour. After this, also triangle B can be refined together with its neighbour.




Fig. 1.11. Recursive refinement in two dimensions. Triangles A and B are initially marked for refinement.

The refinement of a given triangulation \mathcal{S} where some or all elements are marked for refinement is then performed by

Algorithm 1.6 (Recursive refinement algorithm).

```

subroutine refine( $\mathcal{S}$ )
  for all  $S \in \mathcal{S}$  do
    if  $S$  is marked for refinement
      recursive_refine( $S, \mathcal{S}$ )
    end if
  end for

```

Since we use recursion, we have to guarantee that recursions terminates. Kossaczky [44] and Mitchell [48] proved

Theorem 1.7 (Termination and Shape Regularity). *The recursive refinement algorithm using bisection of single elements fulfills*

1. *The recursion terminates if the macro triangulation satisfies certain criteria.*
2. *We obtain shape regularity for all elements at all levels.*

Remark 1.8.

1. A first observation is, that simplices initially not marked for refinement are bisected, enforced by the refinement of a marked simplex. This is a necessity to obtain a conforming triangulation, also for the regular refinement.
2. It is possible to mark an element for more than one bisection. The natural choice is to mark a d -simplex S for d bisections. After d refinement steps all original edges of S are bisected. A simplex S is refined k times by refining the children S_1 and S_2 $k - 1$ times right after the refinement of S .

3. The recursion does not terminate for an arbitrary choice of refinement edges on the macro triangulation. In two dimensions, such a situation is shown in Fig. 1.12. The selected refinement edges of the triangles are shown by dashed lines. One can easily see, that there are no patches for the atomic refinement operation. This triangulation can only be refined if other choices of refinement edges are made, or by a non-recursive algorithm.




Fig. 1.12. A macro triangulation where recursion does not stop.

4. In two dimensions, for every macro triangulation it is possible to choose the refinement edges in such a way that the recursion terminates (selecting the ‘longest edge’). In three dimensions the situation is more complicated. But there is a maybe refined grid such that refinement edges can be chosen in such a way that recursion terminates [44].

1.1.2 Coarsening algorithm for simplicial meshes

The coarsening algorithm is more or less the inverse of the refinement algorithm. The basic idea is to collect all those elements that were created during the refinement at same time, i.e. the parents of these elements build a compatible refinement patch. The elements must only be coarsened if *all* involved elements are marked for coarsening and are of finest level locally, i.e. no element is refined further. The actual coarsening again can be performed in an *atomic coarsening operation* without the handling of hanging nodes. Information is passed from all elements onto the parents and the whole patch is coarsened at the same time by removing the vertex in the parent’s common refinement edge (see Figs. 1.13 and 1.14 for the atomic coarsening operation in 2d and 3d). This coarsening operation is completely local in 1d.

During refinement, the bisection of an element can enforce the refinement of an unmarked element in order to keep the mesh conforming. During coarsening, an element must only be coarsened if all elements involved in this operation are marked for coarsening. This is the main difference between refinement and coarsening. In an adaptive method this guarantees that elements with a large local error indicator marked for refinement are refined and no element is coarsened where the local error indicator is not small enough (compare Section 1.5.3).




Fig. 1.13. Atomic coarsening operation in two dimensions.




Fig. 1.14. Atomic coarsening operation in three dimensions.

Since the coarsening process is the inverse of the refinement, refinement edges on parent elements are again at their original position. Thus, further refinement is possible with a terminating recursion and shape regularity for all resulting elements.

Algorithm 1.9 (Local coarsening).

```

subroutine coarsen_element( $S, S$ )
   $\mathcal{A} := \{S' \in S; S' \text{ must not be coarsened with } S\}$ 
  if  $\mathcal{A} = \emptyset$ 
    for all child pairs  $S'_0, S'_1$  at common coarsening edge do
      coarse  $S'_0$  and  $S'_1$  into the parent  $S'$ 
       $S := S \setminus \{S'_0, S'_1\} \cup \{S'\}$ 
    end for
  end if

```

The following routine coarsens as many elements as possible of a given triangulation S :

Algorithm 1.10 (Coarsening algorithm).

```

subroutine coarsen( $S$ )
  for all  $S \in S$  do
    if  $S$  is marked for coarsening
      coarsen_element( $S, S$ )
    end if
  end for

```

Remark 1.11. Also in the coarsening procedure an element can be marked for several coarsening steps. Usually, the coarsening markers for all patch

elements are cleared if a patch must not be coarsened. If the patch must not be coarsened because one patch element is not of locally finest level but may coarsened more than once, elements stay marked for coarsening. A coarsening of the finer elements can result in a patch which may then be coarsened.

1.1.3 Operations during refinement and coarsening

The refinement and coarsening of elements can be split into four major steps, which are now described in detail.

Topological refinement and coarsening

The actual bisection of an element is performed as follows: the simplex is cut into two children by inserting a new vertex at the refinement edge. All objects like this new vertex, or a new edge (in 2d and 3d), or face (in 3d) have only to be created once on the refinement patch. For example, all elements *share* the new vertex and two children triangles share a common edge. The refinement edge is divided into two smaller ones which have to be adjusted to the respective children. In 3d all faces inside the patch are bisected into two smaller ones and this creates an additional edge for each face. All these objects can be shared by several elements and have to be assigned to them. If neighbour information is stored, one has to update such information for elements inside the patch as well as for neighbours at the patch's boundary.

In the coarsening process the vertex which is shared by all elements is removed, edges and faces are rejoined and assigned to the respective parent simplices. Neighbour information has to be reinstalled inside the patch and with patch neighbours.

Administration of degrees of freedoms

Single DOFs can be assigned to a vertex, edge, or face and such a DOF is shared by all simplices meeting at the vertex, edge, or face respectively. Finally, there may be DOFs on the element itself, which are not shared with any other simplex. At each object there may be a single DOF or several DOFs, even for several finite element spaces.

During refinement new DOFs are created. For each newly created object (vertex, edge, face, center) we have to create the exact amount of DOFs, if DOFs are assigned to the object. For example we have to create vertex DOFs at the midpoint of the refinement edge, if DOFs are assigned to a vertex. Again, DOFs must only be created once for each object and have to be assigned to all simplices having this object in common.

Additionally, all vectors and matrices using these DOFs have automatically to be adjusted in size.

Transfer of geometric data

Information about the children's/parent's shape has to be transformed. During refinement, for a simplex we only have to calculate the coordinates of the midpoint of the refinement edge, coordinates of the other vertices stay the same and can be handed from parent to children. If the refinement edge belongs to a curved boundary, the coordinates of the new vertex are calculated by projecting this midpoint onto the curved boundary. During coarsening, no calculations have to be done. The $d + 1$ vertices of the two children which are *not* removed are the vertices of the parent.

For the shape of parametric elements, usually more information has to be calculated. Such information can be stored in a DOF-vector, e.g., and may need DOFs on parent and children. Thus, information has to be assembled after installing the DOFs on the children and before deleting DOFs on the parent during refinement; during coarsening, first DOFs on the parent have to be installed, then information can be assembled, and finally the children's DOFs are removed.

Transformation of finite element information

Using iterative solvers for the (non-) linear systems, a good initial guess is needed. Usually, the discrete solution from the old grid, interpolated into the finite element space on the new grid, is a good initial guess. For piecewise linear finite elements we only have to compute the value at the newly created node at the midpoint of the refinement edge and this value is the mean value of the values at the vertices of the refinement edge:

$$u_h(\text{midpoint}) = \frac{1}{2}(u_h(\text{vertex } 0) + u_h(\text{vertex } 1)).$$

For linear elements an interpolation during coarsening is trivial since the values at the vertices of the parents stay the same.

For higher order elements more DOFs are involved, but only DOFs belonging to the refinement/coarsening patch. The interpolation strongly depends on the local basis functions and it is described in detail in Section 1.4.4.

Usually during coarsening information is lost (for example, we lose information about the value of a linear finite element function at the coarsening edge's midpoint). But linear functionals applied to basis functions that are calculated on the fine grid and stored in some coefficient vector can be transformed during coarsening without loss of information, if the finite element spaces are nested. This is also described in detail in Section 1.4.4. One application of this procedure is a time discretization, where L^2 scalar products of the new basis functions with the solution u_h^{old} from the last time step appear on the right hand side of the discrete problem.

Since DOFs can be shared by several elements, these operations are done on the whole refinement/coarsening patch. This avoids that coefficients of the

interpolant are calculated more than once for a shared DOF. During the restriction of a linear functional we have to add contribution(s) from one/several DOF(s) to some other DOF(s). Performing this operation on the whole patch makes it easy to guarantee that the contribution of a shared DOF is only added once.

1.2 The hierarchical mesh

There are basically two kinds of storing a finite element grid. One possibility is to store only the elements of the triangulation in a vector or a linked list. All information about elements is located at the elements. In this situation there is no direct information of a hierarchical structure needed, for example, for multigrid methods. Such information has to be generated and stored separately. During mesh refinement, new elements are added (at the end) to the vector or list of elements. During mesh coarsening, elements are removed. In case of an element vector, ‘holes’ may appear in the vector that contain no longer a valid element. One has to take care of them, or remove them by compressing the vector.

ALBERTA uses the second way of storing the mesh. It keeps information about the whole hierarchy of grids starting with the macro triangulation up to the actual one. Storing information about the whole hierarchical structure will need an additional amount of computer memory. On the other hand, we can save computer memory because such information which can be produced by the hierarchical structure does not have to be stored explicitly on each element.

The simplicial grid is generated by refinement of a given macro triangulation. Refined parts of the grid can be de-refined, but we can not coarsen elements of the macro triangulation. The refinement and coarsening routines, described in Section 1.1, construct a sequence of nested grids with a hierarchical structure. Every refined simplex is refined into two children. Elements that may be coarsened were created by refining the parent into these two elements and are now just coarsened back into this parent (compare Sections 1.1.1, 1.1.2).

Using this structure of the refinement/coarsening routines, every element of the macro triangulation is the root of a binary tree: every interior node of that tree has two pointers to the two children; the leaf elements are part of the actual triangulation, which is used to define the finite element space. The whole triangulation is a list of given macro elements together with the associated binary trees, compare Fig. 1.15.

Some information is stored on the (leaf) elements explicitly, other information is located at the macro elements and is transferred to the leaf elements while traversing through the binary tree. For instance, information about DOFs has to be stored explicitly for all (leaf) elements whereas geometric information can be produced using the hierarchical structure.

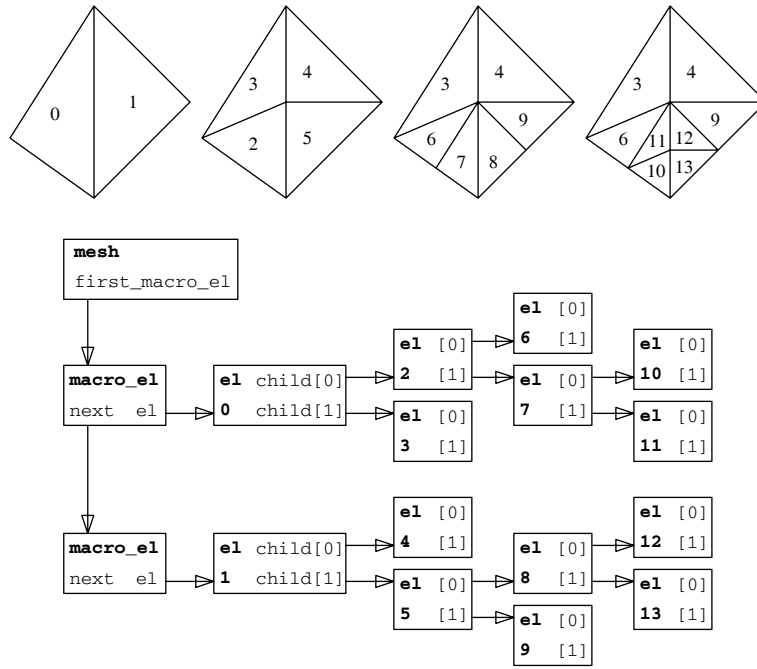


Fig. 1.15. Sample mesh refinement and corresponding element trees

Operations on elements can only be performed by using the mesh traversal routines described in Section 3.2.19. These routines need as arguments a flag which indicates which information should be present on the elements, which elements should be called (interior or leaf elements), and a pointer to a function which performs the operation on a single element. The traversal routines always start on the first macro element and go to the indicated elements of the binary tree at this macro element. This is done in a recursive way by first traversing through the subtree of the first child and then by traversing through the subtree of the second child. This recursion terminates if a leaf element is reached. After calling all elements of this tree we go to the next macro element, traverse through the binary tree located there, and so on until the end of the list of macro elements.

All information that should be available for mesh elements is stored explicitly for elements of the macro triangulation. Thus, all information is present on the macro level and is transferred to the other tree elements by transforming requested data from one element to its children. This can be done by simple calculations using the hierarchic structure induced by the refinement algorithm, compare Section 1.1.1.

As mentioned above, geometric data like coordinates of the element's vertices can be efficiently computed using the hierarchical structure (in the case of non-parametric elements and polyhedral boundary). Going from parent to

child only the coordinates of one vertex changes and the new ones are simply the mean value of the coordinates of two vertices at the refinement edge of the parent. The other vertex coordinates stay the same. Another example of such information is information about adjacent elements. Using adjacency information of the macro elements we can compute requested information for all elements of the mesh.

User data on leaf elements. Many finite element applications need special information on each element of the actual triangulation, i.e. the leaf elements of the hierarchical mesh. In adaptive codes this may be, for example, error indicators produced by an error estimator. Such information has only to be available on leaf elements and not for elements inside the binary tree.

The fact that leaf elements do not have children, and thus the pointers to such children in leaf element's data structures are not used, can be exploited by enabling access to special data via these pointers. So, special pointers for such data do not have to be included in an element data structure. Details about such an implementation are given in Section 3.2.12.

1.3 Degrees of freedom

Degrees of freedom (DOFs) connect finite element data with geometric information of a triangulation. Each finite element function is uniquely determined by the values (coefficients) of all its degrees of freedom.

For example, a continuous and piecewise linear finite element function can be described by the values of this function at all vertices of the triangulation. They build this function's degrees of freedom. A piecewise constant function is determined by its value in each element. In ALBERTA, every abstract DOF is realized as an integer index into vectors, which corresponds to the global index in a vector of coefficients.

For the definition of general finite element spaces DOFs located at vertices of elements, at edges (in 2d and 3d), at faces (in 3d), or in the interior of elements are needed. DOFs at a vertex are shared by all elements which meet at this vertex, DOFs at an edge or face are shared by all elements which contain this edge or face, and DOFs inside an element are not shared with any other element. The location of a DOF and the sharing between elements corresponds directly to the support of basis functions that are connected to them, see Fig. 1.16.

When DOFs and basis functions are used in a hierarchical manner, then the above applies only to a single hierarchical level. Due to the hierarchies, the supports of basis functions which belong to different levels do overlap.

For general applications, it may be necessary to handle several different sets of degrees of freedom on the same triangulation. For example, in mixed finite element methods for the Navier–Stokes problem, different polynomial degrees are used for discrete velocity and pressure functions. In Fig. 1.17, three




Fig. 1.16. Support of basis functions connected with a DOF at a vertex, edge, face (only in 3d), and the interior.

examples of DOF distributions for continuous finite elements in 2d are shown: piecewise quadratic finite elements \square (left), piecewise linear $+$ and piecewise quadratic \square finite elements (middle, Taylor–Hood element for Navier–Stokes: linear pressure and quadratic velocity), piecewise cubic $+$ and piecewise quartic \square finite elements (right, Taylor–Hood element for Navier–Stokes: quartic velocity and linear pressure).




Fig. 1.17. Examples of DOF distributions in 2d.

Additionally, different finite element spaces may use the same set of degrees of freedom, if appropriate. For example, higher order elements with Lagrange type basis or a hierarchical type basis can share the same abstract set of DOFs.

The DOFs are directly connected to the mesh and its elements, by the connection between local (on each element) and global degrees of freedom. On the other hand, an application uses DOFs only in connection with finite element spaces and basis functions. Thus, while the administration of DOFs is handled by the mesh, definition and access of DOFs is mainly done via finite element spaces.

1.4 Finite element spaces and finite element discretization

In the sequel we assume that $\Omega \subset \mathbb{R}^d$ is a bounded domain triangulated by \mathcal{S} , i.e.

$$\bar{\Omega} = \bigcup_{S \in \mathcal{S}} S.$$

The following considerations are also valid for a triangulation of an immersed surface (with $n > d$). In this situation one has to exchange derivatives (those with respect to x) by *tangential* derivatives (tangential to the actual element, derivatives are always taken element-wise) and the determinant of the parameterization's Jacobian has to be replaced by Gram's determinant of the parameterization. But for the sake of clearness and simplicity we restrict our considerations to the case $n = d$.

The values of a finite element function or the values of its derivatives are uniquely defined by the values of its DOFs and the values of the basis functions or the derivatives of the basis functions connected with these DOFs. Usually, evaluation of those values is performed element-wise. On a single element the value of a finite element function at a point x in this element is determined by the DOFs associated with this specific element and the values of the non vanishing basis functions at this point.

We follow the concept of finite elements which are given on a single element S in local coordinates. We distinguish two classes of finite elements:

Finite element functions on an element S defined by a finite dimensional function space $\bar{\mathbb{P}}$ on a *reference element* \bar{S} and the (one to one) mapping λ^S from the reference element \bar{S} to S . For this class the dependency on the actual element S is fully described by the mapping λ^S . For example, all Lagrange finite elements belong to this class.

Secondly, finite element functions depending on the actual element S . Hence, the basis functions are not fully described by $\bar{\mathbb{P}}$ and the one to one mapping λ^S . But using an initialization of the actual element (which defines a finite dimensional function space \mathbb{P} with information about the actual element), we can implement this class in the same way as the first one. This class is needed for Hermite finite elements which are not affine equivalent to the reference element. Examples in 2d are the *Hsieh-Clough-Tocher* or *HCT element* or the *Argyris element* where only the normal derivative at the midpoint of edges are used in the definition of finite element functions; both elements lead to functions which are globally $C^1(\bar{\Omega})$. The concrete implementation for this class in ALBERTA is future work.

All in all, for a very general situation, we only need a vector of basis functions (and its derivatives) on \bar{S} and a function which connects each of these basis functions with its degree of freedom on the element. For the second class, we additionally need an initialization routine for the actual element. By such information, every finite element function is uniquely described on every element of the grid.

1.4.1 Barycentric coordinates

For describing finite elements on simplicial grids, it is very convenient to use $d + 1$ barycentric coordinates as a local coordinate system on an element of

the triangulation. Using $d + 1$ local coordinates, the *reference simplex* \bar{S} is a subset of a hyper surface in \mathbb{R}^{d+1} :

$$\bar{S} := \{(\lambda_0, \dots, \lambda_d) \in \mathbb{R}^{d+1}; \lambda_k \geq 0, \sum_{k=0}^d \lambda_k = 1\}$$

On the other hand, for numerical integration on an element it is much more convenient to use the standard element $\hat{S} \in \mathbb{R}^d$ defined in Section 1.1 as

$$\hat{S} = \text{conv hull} \{\hat{a}_0 = 0, \hat{a}_1 = e_1, \dots, \hat{a}_d = e_d\}$$

where e_i are the unit vectors in \mathbb{R}^d ; using \hat{S} for the numerical integration, we only have to compute the determinant of the parameterization's Jacobian and not Gram's determinant.

The relation between a given simplex S , the reference simplex \bar{S} , and the standard simplex \hat{S} is now described in detail.

Let S be an element of the triangulation with vertices $\{a_0, \dots, a_d\}$; let $F_S: \hat{S} \rightarrow S$ be the diffeomorphic parameterization of S over \hat{S} with regular Jacobian DF_S , such that

$$F_S(\hat{a}_k) = a_k, \quad k = 0, \dots, d$$

holds. For a point $x \in S$ we set

$$\hat{x} = F_S^{-1}(x) \in \hat{S}.$$

For a simplex S the easiest choice for F_S is the unique affine mapping (1.1) defined on page 10. For an affine mapping, DF_S is constant. In the following, we assume that the parameterization F_S of a simplex S is affine.

For a simplex S the barycentric coordinates

$$\lambda^S(x) = (\lambda_0^S, \dots, \lambda_d^S)(x) \in \mathbb{R}^{d+1}$$

of some point $x \in \mathbb{R}^d$ are (uniquely) determined by the $(d + 1)$ equations

$$\sum_{k=0}^d \lambda_k^S(x) a_k = x \quad \text{and} \quad \sum_{k=0}^d \lambda_k^S(x) = 1.$$

The following relation holds:

$$x \in S \quad \text{iff} \quad \lambda_k^S(x) \in [0, 1] \text{ for all } k = 0, \dots, d \quad \text{iff} \quad \lambda^S \in \bar{S}.$$

On the other hand, each $\lambda \in \bar{S}$ defines a unique point $x^S \in S$ by

$$x^S(\lambda) = \sum_{k=0}^d \lambda_k a_k.$$

Thus, $x^S: \bar{S} \rightarrow S$ is an invertible mapping with inverse $\lambda^S: S \rightarrow \bar{S}$. The barycentric coordinates of x on S are the same as those of \hat{x} on \hat{S} , i.e. $\lambda^S(x) = \lambda^{\hat{S}}(\hat{x})$.

In the general situation, when F_S may not be affine, i.e. we have a parametric element, the barycentric coordinates λ^S are given by the inverse of the parameterization F_S and the barycentric coordinates on \hat{S} :

$$\lambda^S(x) = \lambda^{\hat{S}}(\hat{x}) = \lambda^{\hat{S}}(F_S^{-1}(x))$$

and the *world coordinates* of a point $x^S \in S$ with barycentric coordinates λ are given by

$$x^S(\lambda) = F_S \left(\sum_{k=0}^d \lambda_k \hat{a}_k \right) = F_S(x^{\hat{S}}(\lambda))$$

(see also Fig. 1.18).




Fig. 1.18. Definition of $\lambda^S: S \rightarrow \bar{S}$ via F_S^{-1} and $\lambda^{\hat{S}}$, and $x^S: \bar{S} \rightarrow S$ via $x^{\hat{S}}$ and F_S

Every function $f: S \rightarrow V$ defines (uniquely) two functions

$$\bar{f}: \bar{S} \rightarrow V \quad \text{and} \quad \hat{f}: \hat{S} \rightarrow V$$

$$\lambda \mapsto f(x^S(\lambda)) \quad \quad \hat{x} \mapsto f(F_S(\hat{x})).$$

Accordingly, $\hat{f}: \hat{S} \rightarrow V$ defines two functions $f: S \rightarrow V$ and $\bar{f}: \bar{S} \rightarrow V$, and $\bar{f}: \bar{S} \rightarrow V$ defines $f: S \rightarrow V$ and $\hat{f}: \hat{S} \rightarrow V$.

Assuming that a function space $\bar{\mathbb{P}} \subset C^0(\bar{S})$ is given, it uniquely defines function spaces $\hat{\mathbb{P}}$ and \mathbb{P}_S by

$$\hat{\mathbb{P}} = \left\{ \hat{\varphi} \in C^0(\hat{S}); \bar{\varphi} \in \bar{\mathbb{P}} \right\} \quad \text{and} \quad \mathbb{P}_S = \left\{ \varphi \in C^0(S); \bar{\varphi} \in \bar{\mathbb{P}} \right\}. \quad (1.3)$$

We can also assume that the function space $\hat{\mathbb{P}}$ is given and this space uniquely defines $\bar{\mathbb{P}}$ and \mathbb{P}_S in the same manner. In ALBERTA, we use the function space $\bar{\mathbb{P}}$ on \bar{S} ; the implementation of a basis $\{\bar{\varphi}^1, \dots, \bar{\varphi}^m\}$ of $\bar{\mathbb{P}}$ is much simpler

than the implementation of a basis $\{\hat{\varphi}^1, \dots, \hat{\varphi}^m\}$ of $\hat{\mathbb{P}}$ as we are able to use symmetry properties of the barycentric coordinates λ .

In the following we shall often drop the superscript S of λ^S and x^S . The mappings $\lambda(x) = \lambda^S(x)$ and $x(\lambda) = x^S(\lambda)$ are always defined with respect to the actual element $S \in \mathcal{S}$.

1.4.2 Finite element spaces

ALBERTA supports scalar and vector valued finite element functions. The basis functions are always real valued; thus, the coefficients of a finite element function in a representation by the basis functions are either scalars or vectors of length n .

For a given function space $\bar{\mathbb{P}}$ and some given real valued function space C on Ω , a finite element space X_h is defined by

$$X_h = X_h(\mathcal{S}, \bar{\mathbb{P}}, C) = \{\varphi \in C; \varphi|_S \in \mathbb{P}_S \text{ for all } S \in \mathcal{S}\}$$

for scalar finite elements. For vector valued finite elements, it is given by

$$X_h^n = X_h^n(\mathcal{S}, \bar{\mathbb{P}}, C) = \{\varphi \in C^n; \varphi_i|_S \in \mathbb{P}_S \text{ for all } i = 1, \dots, n, S \in \mathcal{S}\}.$$

The spaces \mathbb{P}_S are defined by $\bar{\mathbb{P}}$ via (1.3).

For conforming finite element discretizations, C is the continuous space X (for 2nd order problems, $C = X = H^1(\Omega)$). For non-conforming finite element discretizations, C may control the *non conformity* of the ansatz space which has to be controlled in order to obtain optimal error estimates (for example, in the discretization of the incompressible Navier–Stokes equation by the non-conforming Crouzeix–Raviart element, the finite element functions are continuous only in the midpoints of the edges).

The abstract definition of finite element spaces as a triple (mesh, local basis functions, and DOFs) now matches the mathematical definition as $X_h = X_h(\mathcal{S}, \bar{\mathbb{P}}, C)$ in the following way: The mesh is the underlying triangulation \mathcal{S} , the local basis functions are given by the function space $\bar{\mathbb{P}}$, and together with the relation between global and local degrees of freedom every finite element function satisfies the global continuity constraint C . This relation between global and local DOFs is now described in detail.

1.4.3 Evaluation of finite element functions

Let $\{\bar{\varphi}^1, \dots, \bar{\varphi}^m\}$ be a basis of $\bar{\mathbb{P}}$ and let $\{\varphi_1, \dots, \varphi_N\}$ be a basis of X_h , $N = \dim X_h$, such that for every $S \in \mathcal{S}$ and for all basis functions φ_j which do not vanish on S

$$\varphi_j|_S(x(\lambda)) = \bar{\varphi}^i(\lambda) \quad \text{for all } \lambda \in \bar{S}$$

holds with some $i \in \{1, \dots, m\}$ depending on j and S . Thus, the numbering of the basis functions in $\bar{\mathbb{P}}$ and the mapping x^S induces a local numbering of the

non vanishing basis functions on S . Denoting by J_S the index set of all non vanishing basis functions on S , the mapping $i_S : J_S \rightarrow \{1, \dots, m\}$ is one to one and uniquely connects the degrees of freedom of a finite element function on S with the local numbering of basis functions. If $j_S : \{1, \dots, m\} \rightarrow J_S$ denotes the inverse mapping of i_S , the connection between global and local basis functions is uniquely determined on each element S by

$$\varphi_j(x(\lambda)) = \bar{\varphi}^{i_S(j)}(\lambda), \quad \text{for all } \lambda \in \bar{S}, j \in J_S, \quad (1.4a)$$

$$\varphi_{j_S(i)}(x(\lambda)) = \bar{\varphi}^i(\lambda), \quad \text{for all } \lambda \in \bar{S}, i \in \{1, \dots, m\}. \quad (1.4b)$$

For $u_h \in X_h$ denote by (u_1, \dots, u_N) the *global coefficient vector* of the basis $\{\varphi_j\}$ with $u_j \in \mathbb{R}$ for scalar finite elements, $u_j \in \mathbb{R}^n$ for vector valued finite elements, i.e.

$$u_h(x) = \sum_{j=1}^N u_j \varphi_j(x) \quad \text{for all } x \in \Omega,$$

and the *local coefficient vector*

$$(u_S^1, \dots, u_S^m) = (u_{j_S(1)}, \dots, u_{j_S(m)})$$

of u_h on S with respect to the local numbering of the non vanishing basis functions (local numbering is denoted by a superscript index and global numbering is denoted by a subscript index). Using the local coefficient vector and the local basis functions we obtain the *local representation* of u_h on S :

$$u_h(x) = \sum_{i=1}^m u_S^i \bar{\varphi}^i(\lambda(x)) \quad \text{for all } x \in S.$$

In finite element codes, finite element functions are *not* evaluated at *world coordinates* x as in (1.4.3), but they are evaluated on a single element S at barycentric coordinates λ on S giving the value at the world coordinates $x(\lambda)$:

$$u_h(x(\lambda)) = \sum_{i=1}^m u_S^i \bar{\varphi}^i(\lambda) \quad \text{for all } \lambda \in \bar{S}.$$

The mapping j_S , which allows the access of the local coefficient vector from the global one, is the relation between local DOFs and global DOFs. Global DOFs for a finite element space are stored on the mesh elements at positions which are known to the DOF administration of the finite element space. Thus, the corresponding DOF administration will provide information for the implementation of the function j_S and therefore information for reading/writing local coefficients from/to global coefficient vectors (compare Section 3.5.1).

Evaluating derivatives of finite element functions

Derivatives of finite element functions are also evaluated on single elements S in barycentric coordinates. In the calculation of derivatives in terms of the basis functions $\bar{\varphi}^i$, the Jacobian $\Lambda = \Lambda_S \in \mathbb{R}^{(\text{DIM}+1) \times \text{DIM_OF_WORLD}}$ of the barycentric coordinates on S is involved (we consider here only the case $\text{DIM} = \text{DIM_OF_WORLD} = d$):

$$\Lambda(x) = \begin{pmatrix} \lambda_{0,x_1}(x) & \lambda_{0,x_2}(x) & \cdots & \lambda_{0,x_d}(x) \\ \vdots & \vdots & & \vdots \\ \lambda_{d,x_1}(x) & \lambda_{d,x_2}(x) & \cdots & \lambda_{d,x_d}(x) \end{pmatrix} = \begin{pmatrix} -\nabla \lambda_0(x) - \\ \vdots \\ -\nabla \lambda_d(x) - \end{pmatrix}.$$

Now, using the chain rule we obtain for every function $\varphi \in \mathbb{P}_S$ and $x \in S$

$$\nabla \varphi(x) = \nabla (\bar{\varphi}(\lambda(x))) = \sum_{k=0}^d \bar{\varphi}_{,\lambda_k}(\lambda(x)) \nabla \lambda_k(x) = \Lambda^t(x) \nabla_\lambda \bar{\varphi}(\lambda(x))$$

and

$$D^2 \varphi(x) = \Lambda^t(x) D_\lambda^2 \bar{\varphi}(\lambda(x)) \Lambda(x) + \sum_{k=0}^d D^2 \lambda_k(x) \bar{\varphi}_{,\lambda_k}(\lambda(x)).$$

For a simplex S with an affine parameterization F_S , $\nabla \lambda_k$ is constant on S and we get

$$\nabla \varphi(x) = \Lambda^t \nabla_\lambda \bar{\varphi}(\lambda(x)) \quad \text{and} \quad D^2 \varphi(x) = \Lambda^t D_\lambda^2 \bar{\varphi}(\lambda(x)) \Lambda.$$

Using these equations, we immediately obtain

$$\nabla u_h(x) = \Lambda^t(x) \sum_{i=1}^m u_S^i \nabla_\lambda \bar{\varphi}^i(\lambda(x))$$

and

$$D^2 u_h(x) = \Lambda^t(x) \sum_{i=1}^m u_S^i D_\lambda^2 \bar{\varphi}^i(\lambda(x)) \Lambda(x) + \sum_{k=0}^d D^2 \lambda_k(x) \sum_{i=1}^m u_S^i \bar{\varphi}_{,\lambda_k}^i(\lambda(x)).$$

Since the evaluation is actually done in barycentric coordinates, this turns for $\lambda \in \bar{S}$ on S into

$$\nabla u_h(x(\lambda)) = \Lambda^t(x(\lambda)) \sum_{i=1}^m u_S^i \nabla_\lambda \bar{\varphi}^i(\lambda)$$

and

$$\begin{aligned} D^2 u_h(x(\lambda)) &= \Lambda^t(x(\lambda)) \sum_{i=1}^m u_S^i D_\lambda^2 \bar{\varphi}^i(\lambda) \Lambda(x(\lambda)) \\ &\quad + \sum_{k=0}^d D^2 \lambda_k(x(\lambda)) \sum_{i=1}^m u_S^i \bar{\varphi}_{,\lambda_k}^i(\lambda). \end{aligned}$$

Once the values of the basis functions, its derivatives, and the local coefficient vector (u_S^1, \dots, u_S^m) are known, the evaluation of u_h and its derivatives does only depend on Λ and can be performed by some general evaluation routine (compare Section 3.9).

1.4.4 Interpolation and restriction during refinement and coarsening

We assume the following situation: Let S be a (non-parametric) simplex with children S_0 and S_1 generated by the bisection of S (compare Algorithm 1.5). Let X_S , X_{S_0, S_1} be the finite element spaces restricted to S and $S_0 \cup S_1$ respectively.

Throughout this section we denote by $\{\varphi^i\}_{i=1, \dots, m}$ the basis of the coarse grid space X_S and by $\{\psi^j\}_{j=1, \dots, k}$ the basis functions of $X_{S_0 \cup S_1}$. For a function $u_h \in X_S$ we denote by $\mathbf{u}_\varphi = (u_\varphi^1, \dots, u_\varphi^m)^t$ the coefficient vector with respect to the basis $\{\varphi^i\}$ and for a function $v_h \in X_{S_0 \cup S_1}$ by $\mathbf{v}_\psi = (v_\psi^1, \dots, v_\psi^k)^t$ the coefficient vector with respect to $\{\psi^j\}$.

We now derive equations for the transformation of local coefficient vectors for finite element functions that are interpolated during refinement and coarsening, and for vectors storing values of a linear functional applied to the basis functions on the fine grid which are restricted to the coarse functions during coarsening.

Let

$$\mathbb{I}_{S_0, S_1}^S : X_S \rightarrow X_{S_0 \cup S_1}$$

be an interpolation operator. For nested finite element spaces, i.e. $X_S \subset X_{S_0 \cup S_1}$, every coarse grid function $u_h \in X_S$ belongs also to $X_{S_0 \cup S_1}$, so the natural choice is $\mathbb{I}_{S_0, S_1}^S = id$ on X_S (for example, Lagrange finite elements are nested). The interpolants $\mathbb{I}_{S_0, S_1}^S \varphi^i$ can be written in terms of the fine grid basis functions

$$\mathbb{I}_{S_0, S_1}^S \varphi^i = \sum_{j=1}^k a_{ij} \psi^j$$

defining the $(m \times k)$ -matrix

$$A = (a_{ij})_{\substack{i=1, \dots, m \\ j=1, \dots, k}}. \quad (1.5)$$

This matrix A is involved in the interpolation during refinement and the transformation of a linear functional during coarsening.

For the interpolation of functions during coarsening, we need an interpolation operator $\mathbb{I}_S^{S_0, S_1} : X_{S_0 \cup S_1} \rightarrow X_S$. The interpolants $\mathbb{I}_S^{S_0, S_1} \psi^j$ of the fine grid functions can now be represented by the coarse grid basis

$$\mathbb{I}_S^{S_0, S_1} \psi^j = \sum_{i=1}^m b_{ij} \varphi^i$$

defining the $(m \times k)$ -matrix

$$B = (b_{ij})_{\substack{i=1,\dots,m \\ j=1,\dots,k}}. \quad (1.6)$$

This matrix B is used for the interpolation during coarsening.

Both matrices do only depend on the set of local basis functions on parent and children. Thus, they depend on the reference element \bar{S} and one single bisection of the reference element into \bar{S}_0, \bar{S}_1 . The matrices do depend on the local numbering of the basis functions on the children with respect to the parent. Thus, in 3d the matrices depend on the element type of S also (for an element of type 0 the numbering of basis functions on \bar{S}_1 differs from the numbering on \bar{S}_1 for an element of type 1, 2). But all matrices can be calculated by the local set of basis functions on the reference element.

DOFs can be shared by several elements, compare Section 1.3. Every DOF is connected to a basis function which has a support on all elements sharing this DOF. Each DOF refers to one coefficient of a finite element function, and this coefficient has to be calculated only once during interpolation. During the restriction of a linear functional, contributions from several basis functions are added to the coefficient of another basis function. Here we have to control that for two DOFs, both shared by several elements, the contribution of the basis function at one DOF is only added once to the other DOF and vice versa. This can only be done by performing the interpolation and restriction on the whole refinement/coarsening patch at the same time.

Interpolation during refinement

Let $\mathbf{u}_\varphi = (u_\varphi^1, \dots, u_\varphi^m)^t$ be the coefficient vector of a finite element function $u_h \in X_S$ with respect to $\{\varphi^i\}$, and let $\mathbf{u}_\psi = (u_\psi^1, \dots, u_\psi^k)^t$ the coefficient vector of $\mathbb{I}_{S_0, S_1}^S u_h$ with respect to $\{\psi^j\}$. Using matrix A defined in (1.5) we conclude

$$\mathbb{I}_{S_0, S_1}^S u_h = \sum_{i=1}^m u_\varphi^i \mathbb{I}_{S_0, S_1}^S \varphi^i = \sum_{i=1}^m u_\varphi^i \sum_{j=1}^k a_{ij} \psi^j = \sum_{j=1}^k (A^t \mathbf{u}_\varphi)_j \psi^j,$$

or equivalently

$$\mathbf{u}_\psi = A^t \mathbf{u}_\varphi.$$

A subroutine which interpolates a finite element function during refinement is an efficient implementation of this matrix-vector multiplication.

Restriction during coarsening

In an (Euler, e.g.) discretization of a time dependent problem, the term $(u_h^{\text{old}}, \varphi_i)_{L^2(\Omega)}$ appears on the right hand side of the discrete system, where

u_h^{old} is the solution from the last time step. Such an expression can be calculated exactly, if the grid does not change from one time step to another. Assuming that the finite element spaces are nested, it is also possible to calculate this expression exactly when the grid was refined, since u_h^{old} belongs to the fine grid finite element space also. Usually, during coarsening information is lost, since we can not represent u_h^{old} exactly on a coarser grid. But we can calculate $(u_h^{\text{old}}, \psi_i)_{L^2(\Omega)}$ exactly on the fine grid; using the representation of the coarse grid basis functions φ_i by the fine grid functions ψ_i , we can transform data during coarsening such that $(u_h^{\text{old}}, \varphi_i)_{L^2(\Omega)}$ is calculated exactly for the coarse grid functions too.

More general, assume that the finite element spaces are nested and that we can evaluate a linear functional f exactly for all basis functions of the fine grid. Knowing the values $\mathbf{f}_\psi = (\langle f, \psi^1 \rangle, \dots, \langle f, \psi^k \rangle)^t$ for the fine grid functions, we obtain with matrix A from (1.5) for the values $\mathbf{f}_\varphi = (\langle f, \varphi^1 \rangle, \dots, \langle f, \varphi^m \rangle)^t$ on the coarse grid

$$\mathbf{f}_\varphi = A\mathbf{f}_\psi$$

since

$$\langle f, \varphi^i \rangle = \langle f, \sum_{j=1}^k a_{ij} \psi^j \rangle = \sum_{j=1}^k a_{ij} \langle f, \psi^j \rangle$$

holds (here we used the fact, that $\mathbb{I}_{S_0, S_1}^S = id$ on X_S since the spaces are nested).

Thus, given a functional f which we can evaluate exactly for all basis functions of a grid \tilde{S} and its refinements, we can also calculate $\langle f, \varphi^i \rangle$ exactly for all basis functions φ^i of a grid S obtained by refinement and coarsening of \tilde{S} in the following way: First refine all elements of the grid that have to be refined; calculate $\langle f, \varphi \rangle$ for all basis functions φ of this intermediate grid; in the last step coarsen all elements that may be coarsened and restrict this vector during each coarsening step as described above.

In ALBERTA the assemblage of the discrete system inside the adaptive method can be split into three steps: one initialization step before refinement, the second step between refinement and coarsening, and the last, and usually most important, step after coarsening, when all grid modifications are completed, see Section 3.13.1. The second assemblage step can be used for an exact computation of a functional $\langle f, \varphi \rangle$ as described above.

Interpolation during coarsening

Finally, we can interpolate a finite element function during coarsening. The matrix for transforming the coefficient vector $\mathbf{u}_\psi = (u_\psi^1, \dots, u_\psi^k)^t$ of a fine grid function u_h to the coefficient vector $\mathbf{u}_\varphi = (u_\varphi^1, \dots, u_\varphi^m)^t$ of the interpolant on the coarse grid is given by matrix B defined in (1.6):

$$\begin{aligned}
\mathbb{I}_S^{S_0, S_1} u_h &= \mathbb{I}_S^{S_0, S_1} \sum_{j=1}^k u_\psi^j \psi^j = \sum_{j=1}^k u_\psi^j \mathbb{I}_S^{S_0, S_1} \psi^j \\
&= \sum_{j=1}^k u_\psi^j \sum_{i=1}^m b_{ij} \varphi^i = \sum_{i=1}^m \left(\sum_{j=1}^k b_{ij} u_\psi^j \right) \varphi^i.
\end{aligned}$$

Thus we have the following equation for the coefficient vector of the interpolant of u_h :

$$\mathbf{u}_\varphi = B \mathbf{u}_\psi.$$

In contrast to the interpolation during refinement and the above described transformation of a linear functional, information is lost during an interpolation to the coarser grid.

Example 1.12 (Lagrange elements). Lagrange finite elements are connected to Lagrange nodes x^i . For linear elements, these nodes are the vertices of the triangulation, and for quadratic elements, the vertices and the edge-midpoints. The Lagrange basis functions $\{\varphi^i\}$ satisfy

$$\varphi_i(x_j) = \delta_{ij} \quad \text{for } i, j = 1, \dots, \dim X_h.$$

Consider the situation of a simplex S with children S_0, S_1 . Let $\{\varphi^i\}_{i=1, \dots, m}$ be the Lagrange basis functions of X_S with Lagrange nodes $\{x_\varphi^i\}_{i=1, \dots, m}$ on S and $\{\psi^j\}_{j=1, \dots, k}$ be the Lagrange basis functions of $X_{S_0 \cup S_1}$ with Lagrange nodes $\{x_\psi^j\}_{j=1, \dots, k}$ on $S_0 \cup S_1$. The Matrix A is then given by

$$a_{ij} = \varphi^i(x_\psi^j), \quad i = 1, \dots, m, \quad j = 1, \dots, k$$

and matrix B is given by

$$b_{ij} = \psi^j(x_\varphi^i), \quad i = 1, \dots, m, \quad j = 1, \dots, k.$$

1.4.5 Discretization of 2nd order problems

In this section we describe the assembling of the discrete system in detail. We consider the following second order differential equation in divergence form:

$$Lu := -\nabla \cdot A \nabla u + b \cdot \nabla u + cu = f \quad \text{in } \Omega, \quad (1.7a)$$

$$u = g \quad \text{on } \Gamma_D, \quad (1.7b)$$

$$\nu_\Omega \cdot A \nabla u = 0 \quad \text{on } \Gamma_N, \quad (1.7c)$$

where $A \in L^\infty(\Omega; \mathbb{R}^{d \times d})$, $b \in L^\infty(\Omega; \mathbb{R}^d)$, $c \in L^\infty(\Omega)$, and $f \in L^2(\Omega)$. By $\Gamma_D \subset \partial\Omega$ (with $|\Gamma_D| \neq 0$) we denote the Dirichlet boundary and we assume that the Dirichlet boundary values $g: \Gamma_D \rightarrow \mathbb{R}$ have an extension to some function $g \in H^1(\Omega)$.

By $\Gamma_N = \partial\Omega \setminus \Gamma_D$ we denote the Neumann boundary, and by ν_Ω we denote the outer unit normal vector on $\partial\Omega$. The boundary condition (1.7c) is a so called *natural* Neumann condition.

Equations (1.7) describe not only a simple model problem. The same kind of equations result from a linearization of nonlinear elliptic problems (for example by a Newton method) as well as from a time discretization scheme for (non-) linear parabolic problems.

Setting

$$X = H^1(\Omega) \quad \text{and} \quad \mathring{X} = \{v \in H^1(\Omega); v = 0 \text{ on } \Gamma_D\}$$

this equation has the following weak formulation: We are looking for a solution $u \in X$, such that $u \in g + \mathring{X}$ and

$$\int_{\Omega} \nabla \varphi \cdot A \nabla u + \varphi b \cdot \nabla u + c \varphi u \, dx = \int_{\Omega} f \varphi \, dx \quad (1.8)$$

for all $\varphi \in \mathring{X}$

Denoting by \mathring{X}^* the dual space of \mathring{X} we identify the differential operator L with the linear operator $L \in \mathcal{L}(X, \mathring{X}^*)$ defined by

$$\langle Lv, \varphi \rangle_{\mathring{X}^* \times \mathring{X}} := \int_{\Omega} \nabla \varphi \cdot A \nabla v + \int_{\Omega} \varphi b \cdot \nabla v + \int_{\Omega} c \varphi v \quad \text{for all } v, \varphi \in \mathring{X}$$

and the right hand side f with the linear functional $f \in \mathring{X}^*$ defined by

$$\langle F, \varphi \rangle_{\mathring{X}^* \times \mathring{X}} := \int_{\Omega} f \varphi \quad \text{for all } \varphi \in \mathring{X}.$$

With these identifications we use the following reformulation of (1.8): Find $u \in X$ such that

$$u \in g + \mathring{X} : \quad Lu = f \quad \text{in } \mathring{X}^* \quad (1.9)$$

holds.

Suitable assumptions on the coefficients imply that L is elliptic, i.e. there is a constant $C = C_{A,b,c,\Omega}$ such that

$$\langle L\varphi, \varphi \rangle_{\mathring{X}^* \times \mathring{X}} \geq C \|\varphi\|_X^2 \quad \text{for all } \varphi \in \mathring{X}.$$

The existence of a unique solution $u \in X$ of (1.9) is then a direct consequence of the Lax–Milgram–Theorem.

We consider a finite dimensional subspace $X_h \subset X$ for the discretization of (1.9) with $N = \dim X_h$. We set $\mathring{X}_h = X_h \cap \mathring{X}$ with $\mathring{N} = \dim \mathring{X}_h$. Let $g_h \in X_h$ be an approximation of $g \in X$. A discrete solution of (1.9) is then given by: Find $u_h \in X_h$ such that

$$u_h \in g_h + \mathring{X}_h : \quad Lu_h = f \quad \text{in } \mathring{X}_h^*, \quad (1.10)$$

i.e.

$$u_h \in g_h + \mathring{X}_h : \quad \langle Lu_h, \varphi_h \rangle_{\mathring{X}_h^* \times \mathring{X}_h} = \langle f, \varphi_h \rangle_{\mathring{X}_h^* \times \mathring{X}_h} \quad \text{for all } \varphi_h \in \mathring{X}_h$$

holds. If L is elliptic, we have a unique discrete solution $u_h \in X_h$ of (1.10), again using the Lax–Milgram–Theorem.

Choose a basis $\{\varphi_1, \dots, \varphi_N\}$ of X_h such that $\{\varphi_1, \dots, \varphi_{\mathring{N}}\}$ is a basis of \mathring{X}_h . For a function $v_h \in X_h$ we denote by $\mathbf{v} = (v_1, \dots, v_N)$ the coefficient vector of v_h with respect to the basis $\{\varphi_1, \dots, \varphi_N\}$, i.e.

$$v_h = \sum_{j=1}^N v_j \varphi_j.$$

Using (1.10) with test functions φ_i , $i = 1, \dots, \mathring{N}$, we get the following N equations for the coefficient vector $\mathbf{u} = (u_1, \dots, u_N)$ of u_h :

$$\begin{aligned} \sum_{j=1}^N u_j \langle L\varphi_j, \varphi_i \rangle_{\mathring{X}_h^* \times \mathring{X}_h} &= \langle f, \varphi_i \rangle_{\mathring{X}_h^* \times \mathring{X}_h} & \text{for } i = 1, \dots, \mathring{N}, \\ u_i &= g_i & \text{for } i = \mathring{N} + 1, \dots, N. \end{aligned}$$

Defining the *system matrix* \mathbf{L} by

$$\mathbf{L} := \begin{bmatrix} \langle L\varphi_1, \varphi_1 \rangle & \dots & \langle L\varphi_{\mathring{N}}, \varphi_1 \rangle & \langle L\varphi_{\mathring{N}+1}, \varphi_1 \rangle & \dots & \langle L\varphi_N, \varphi_1 \rangle \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \langle L\varphi_1, \varphi_{\mathring{N}} \rangle & \dots & \langle L\varphi_{\mathring{N}}, \varphi_{\mathring{N}} \rangle & \langle L\varphi_{\mathring{N}+1}, \varphi_{\mathring{N}} \rangle & \dots & \langle L\varphi_N, \varphi_{\mathring{N}} \rangle \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 \\ 0 & \dots & 0 & 0 & 1 & \dots & 0 \\ \vdots & \ddots & 0 & 0 & 0 & \ddots & \vdots \\ 0 & \dots & 0 & 0 & 0 & \dots & 1 \end{bmatrix} \quad (1.11)$$

and the *right hand side vector* or *load vector* \mathbf{f} by

$$\mathbf{f} := \begin{bmatrix} \langle f, \varphi_1 \rangle \\ \vdots \\ \langle f, \varphi_{\mathring{N}} \rangle \\ g_{\mathring{N}+1} \\ \vdots \\ g_N \end{bmatrix}, \quad (1.12)$$

we can write the discrete equations as the linear $N \times N$ system

$$\mathbf{L} \mathbf{u} = \mathbf{f}, \quad (1.13)$$

which has to be assembled and solved numerically.

1.4.6 Numerical quadrature

For the assemblage of the system matrix and right hand side vector of the linear system (1.13), we have to compute integrals, for example

$$\int_{\Omega} f(x) \varphi_i(x) dx.$$

For general data A , b , c , and f , these integrals can not be calculated exactly. Quadrature formulas have to be used in order to calculate the integrals approximately. Numerical integration in finite element methods is done by looping over all grid elements and using a quadrature formula on each element.

Definition 1.13 (Numerical quadrature). A numerical quadrature \hat{Q} on \hat{S} is a set $\{(w_k, \lambda_k) \in \mathbb{R} \times \mathbb{R}^{d+1}; k = 0, \dots, n_Q - 1\}$ of weights w_k and quadrature points $\lambda_k \in \hat{S}$ (i.e. given in barycentric coordinates) such that

$$\int_{\hat{S}} f(\hat{x}) d\hat{x} \approx \hat{Q}(f) := \sum_{k=0}^{n_Q-1} w_k f(\hat{x}(\lambda_k)).$$

It is called exact of degree p for some $p \in \mathbb{N}$ if

$$\int_{\hat{S}} q(\hat{x}) d\hat{x} = \hat{Q}(q) \quad \text{for all } q \in \mathbb{P}_p(\hat{S}).$$

It is called stable if

$$w_k > 0 \quad \text{for all } k = 0, \dots, n_Q - 1.$$

Remark 1.14. A given numerical quadrature \hat{Q} on \hat{S} defines for each element S a numerical quadrature Q_S . Using the transformation rule we define Q_S on an element S which is parameterized by $F_S: \hat{S} \rightarrow S$ and a function $f: S \rightarrow \mathbb{R}$:

$$\begin{aligned} \int_S f(x) dx &\approx Q_S(f) := \hat{Q}((f \circ F_S) |\det DF_S|) \\ &= \sum_{k=0}^{n_Q-1} w_k f(x(\lambda_k)) |\det DF_S(\hat{x}(\lambda_k))|. \end{aligned}$$

For a simplex S this results in

$$Q_S(f) = d! |S| \sum_{k=0}^{n_Q-1} w_k f(x(\lambda_k)).$$

1.4.7 Finite element discretization of 2nd order problems

Let $\bar{\mathbb{P}}$ be a finite dimensional function space on \bar{S} with basis $\{\bar{\varphi}^1, \dots, \bar{\varphi}^m\}$. For a conforming finite element discretization of (1.8) we use the finite element space $X_h = X_h(\mathcal{S}, \bar{\mathbb{P}}, X)$. For this space \hat{X}_h is given by $X_h(\mathcal{S}, \bar{\mathbb{P}}, \hat{X})$.

By the relation (1.4a) for global and local basis functions, we obtain for the j th component of the right hand side vector \mathbf{f}

$$\begin{aligned} \langle f, \varphi_j \rangle &= \int_{\Omega} f(x) \varphi_j(x) dx = \sum_{S \in \mathcal{S}} \int_S f(x) \varphi_j(x) dx \\ &= \sum_{\substack{S \in \mathcal{S} \\ S \subset \text{supp}(\varphi_j)}} \int_S f(x) \bar{\varphi}^{i_S(j)}(\lambda(x)) dx \\ &= \sum_{\substack{S \in \mathcal{S} \\ S \subset \text{supp}(\varphi_j)}} \int_{\hat{S}} f(F_S(\hat{x})) \bar{\varphi}^{i_S(j)}(\lambda(\hat{x})) |\det DF_S(\hat{x})| d\hat{x}, \end{aligned}$$

where S is parameterized by $F_S: \hat{S} \rightarrow S$. The above sum is reduced to a sum over all $S \subset \text{supp}(\varphi_j)$ which are only few elements due to the small support of φ_j .

The right hand side vector can be assembled in the following way: First, the right hand side vector is initialized with zeros. For each element S of \mathcal{S} we calculate the *element load vector* $\mathbf{f}_S = (f_S^1, \dots, f_S^m)^t$, where

$$f_S^i = \int_{\hat{S}} f(F_S(\hat{x})) \bar{\varphi}^i(\lambda(\hat{x})) |\det DF_S(\hat{x})| d\hat{x}, \quad i = 1, \dots, m. \quad (1.14)$$

Denoting again by $j_S: \{1, \dots, m\} \rightarrow J_S$ the function which connects the local DOFs with the global DOFs (defined in (1.4b)), the values f_S^i are then added to the $j_S(i)$ th component of the right hand side vector \mathbf{f} , $i = 1, \dots, m$.

For general f , the integrals in (1.14) can not be calculated exactly and we have to use a quadrature formula for the approximation of the integrals (compare Section 1.4.6). Given a numerical quadrature \hat{Q} on \hat{S} we approximate

$$\begin{aligned} f_S^i &\approx \hat{Q}((f \circ F_S)(\bar{\varphi}^i \circ \lambda) |\det DF_S|) \\ &= \sum_{k=0}^{n_Q-1} w_k f(x(\lambda_k)) \bar{\varphi}^i(\lambda_k) |\det DF_S(\hat{x}(\lambda_k))|. \end{aligned} \quad (1.15)$$

For a simplex S this is simplified to

$$f_S^i \approx d! |S| \sum_{k=0}^{n_Q-1} w_k f(x(\lambda_k)) \bar{\varphi}^i(\lambda_k).$$

In ALBERTA, information about values of basis functions and its derivatives as well as information about the connection of global and local DOFs

(i.e. information about j_S) is stored in special data structures for local basis functions (compare Section 3.5). By such information, the element load vector can be assembled by a general routine if a function for the evaluation of the right hand side is supplied. For parametric elements, a function for evaluating $|\det DF_S|$ is additionally needed. The assemblage into the global load vector \mathbf{f} can again be performed automatically, using information about the connection of global and local DOFs.

The calculation of the system matrix is also done element-wise. Additionally, we have to handle derivatives of basis functions. Looking first at the second order term we obtain by the chain rule (1.4.3) and the relation (1.4) for global and local basis functions

$$\begin{aligned} \int_S \nabla \varphi_i(x) \cdot A(x) \nabla \varphi_j(x) dx &= \int_S \nabla(\bar{\varphi}^{is(i)} \circ \lambda)(x) \cdot A(x) \nabla(\bar{\varphi}^{is(j)} \circ \lambda)(x) dx \\ &= \int_S \nabla_\lambda \bar{\varphi}^{is(i)}(\lambda(x)) \cdot (\Lambda(x) A(x) \Lambda^t(x)) \nabla_\lambda \bar{\varphi}^{is(j)}(\lambda(x)) dx, \end{aligned}$$

where Λ is the Jacobian of the barycentric coordinates λ on S . In the same manner we obtain for the first and zero order terms

$$\int_S \varphi_i(x) b(x) \cdot \nabla \varphi_j(x) dx = \int_S \bar{\varphi}^{is(i)}(\lambda(x)) (\Lambda(x) b(x)) \cdot \nabla_\lambda \bar{\varphi}^{is(j)}(\lambda(x)) dx$$

and

$$\int_S c(x) \varphi_i(x) \varphi_j(x) dx = \int_S c(x) \bar{\varphi}^{is(i)}(\lambda(x)) \bar{\varphi}^{is(j)}(\lambda(x)) dx.$$

Using on S the abbreviations

$$\begin{aligned} \bar{A}(\lambda) &:= (\bar{a}_{kl}(\lambda))_{k,l=0,\dots,d} := |\det DF_S(\hat{x}(\lambda))| \Lambda(x(\lambda)) A(x(\lambda)) \Lambda^t(x(\lambda)), \\ \bar{b}(\lambda) &:= (\bar{b}_l(\lambda))_{l=0,\dots,d} := |\det DF_S(\hat{x}(\lambda))| \Lambda(x(\lambda)) b(x(\lambda)), \quad \text{and} \\ \bar{c}(\lambda) &:= |\det DF_S(\hat{x}(\lambda))| c(x(\lambda)) \end{aligned}$$

and transforming the integrals to the reference simplex, we can write the *element matrix* $\mathbf{L}_S = (L_S^{ij})_{i,j=1,\dots,m}$ as

$$\begin{aligned} L_S^{ij} &= \int_{\hat{S}} \nabla_\lambda \bar{\varphi}^i(\lambda(\hat{x})) \cdot \bar{A}(\lambda(\hat{x})) \nabla_\lambda \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x} \\ &\quad + \int_{\hat{S}} \bar{\varphi}^i(\lambda(\hat{x})) \bar{b}(\lambda(\hat{x})) \cdot \nabla_\lambda \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x} \\ &\quad + \int_{\hat{S}} \bar{c}(\lambda(\hat{x})) \bar{\varphi}^i(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x}, \end{aligned}$$

or writing the matrix-vector and vector-vector products explicitly

$$\begin{aligned}
L_S^{ij} &= \sum_{k,l=0}^d \int_{\hat{S}} \bar{a}_{kl}(\lambda(\hat{x})) \bar{\varphi}_{,\lambda_k}^i(\lambda(\hat{x})) \bar{\varphi}_{,\lambda_l}^j(\lambda(\hat{x})) d\hat{x} \\
&\quad + \sum_{l=0}^d \int_{\hat{S}} \bar{b}_l(\lambda(\hat{x})) \bar{\varphi}^i(\lambda(\hat{x})) \bar{\varphi}_{,\lambda_l}^j(\lambda(\hat{x})) d\hat{x} \\
&\quad + \int_{\hat{S}} \bar{c}(\lambda(\hat{x})) \bar{\varphi}^i(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x},
\end{aligned}$$

$i, j = 1, \dots, m$. Using quadrature formulas \hat{Q}_2 , \hat{Q}_1 , and \hat{Q}_0 for the second, first and zero order term we approximate the element matrix

$$\begin{aligned}
L_S^{ij} &\approx \hat{Q}_2 \left(\sum_{k,l=0}^d (\bar{a}_{kl} \bar{\varphi}_{,\lambda_k}^i \bar{\varphi}_{,\lambda_l}^j) \circ \lambda \right) \\
&\quad + \hat{Q}_1 \left(\sum_{l=0}^d (\bar{b}_l \bar{\varphi}^i \bar{\varphi}_{,\lambda_l}^j) \circ \lambda \right) + \hat{Q}_0 \left((\bar{c} \bar{\varphi}^i \bar{\varphi}^j) \circ \lambda \right),
\end{aligned}$$

$i, j = 1, \dots, m$. Having access to functions for the evaluation of

$$\bar{a}_{kl}(\lambda_q), \quad \bar{b}_l(\lambda_q), \quad \bar{c}(\lambda_q)$$

at all quadrature points λ_q on S , L_S can be computed by some general routine. The assemblage into the system matrix can also be done automatically (compare the assemblage of the load vector).

Remark 1.15. Due to the small support of the global basis function, the system matrix is a sparse matrix, i.e. the maximal number of entries in all matrix rows is much smaller than the size of the matrix. Special data structures are needed for an efficient storage of sparse matrices and they are described in Section 3.3.4.

Remark 1.16. The calculation of $A(x(\lambda))$ usually involves the determinant of the parameterization's Jacobian $|\det DF_S(\hat{x}(\lambda))|$. Thus, a calculation of $|\det DF_S(\hat{x}(\lambda))| A(x(\lambda)) A(x(\lambda)) A^t(x(\lambda))$ may be much faster than the calculation of $A(x(\lambda)) A(x(\lambda)) A^t(x(\lambda))$ only; the same holds for the first order term.

Assuming that the coefficients A , b , and c are piecewise constant on a non-parametric triangulation, $\bar{A}(\lambda)$, $\bar{b}(\lambda)$, and $\bar{c}(\lambda)$ are constant on each simplex S and thus simplified to

$$\begin{aligned}
\bar{A}_S &= (\bar{a}_{kl})_{k,l=0,\dots,d} = d!|S| A|_S A^t, \\
\bar{b}_S &= (\bar{b}_l)_{l=0,\dots,d} = d!|S| A b|_S, \\
\bar{c}_S &= d!|S| c|_S.
\end{aligned}$$

For the approximation of the element matrix by quadrature we then obtain

$$\begin{aligned} L_S^{ij} \approx & \sum_{k,l=0}^d \bar{a}_{kl} \hat{Q}_2 \left((\bar{\varphi}_{,\lambda_k}^i \bar{\varphi}_{,\lambda_l}^j) \circ \lambda \right) \\ & + \sum_{l=0}^d \bar{b}_l \hat{Q}_1 \left((\bar{\varphi}^i \bar{\varphi}_{,\lambda_l}^j) \circ \lambda \right) + \bar{c}_S \hat{Q}_0 \left((\bar{\varphi}^i \bar{\varphi}^j) \circ \lambda \right) \end{aligned} \quad (1.16)$$

$i, j = 1, \dots, m$. Here, the numerical quadrature is only applied for approximating integrals of the basis functions on the standard simplex. These values can be computed only once, and can then be used on each simplex S . This will speed up the assembling of the system matrix. Additionally, for polynomial basis functions we can use quadrature formulas which integrate these integrals exactly.

As a result, using information about values of basis functions and their derivatives, and information about the connection of global and local DOFs, the linear system can be assembled automatically by some general routines. Only functions for the evaluation of given data have to be provided for special applications. The general assemble routines are described in Section 3.12.

1.5 Adaptive Methods

The aim of adaptive methods is the generation of a mesh which is adapted to the problem such that a given criterion, like a tolerance for the estimated error between exact and discrete solution, is fulfilled by the finite element solution on this mesh. An optimal mesh should be as coarse as possible while meeting the criterion, in order to save computing time and memory requirements. For time dependent problems, such an adaptive method may include mesh changes in each time step and control of time step sizes.

The philosophy implemented in ALBERTA is to change meshes successively by local refinement or coarsening, based on error estimators or error indicators, which are computed a posteriori from the discrete solution and given data on the current mesh.

1.5.1 Adaptive method for stationary problems

Let us assume that a triangulation \mathcal{S} of Ω , a finite element solution $u_h \in X_h$ to an elliptic problem, and an a posteriori error estimate

$$\|u - u_h\| \leq \eta(u_h) = \left(\sum_{S \in \mathcal{S}} \eta_S(u_h)^p \right)^{1/p}, \quad p \in [1, \infty) \quad (1.17)$$

on this mesh are given. If tol is a given allowed tolerance for the error, and $\eta(u_h) > tol$, the problem arises, where to refine the mesh in order to reduce

the error, while at the same time the number of unknowns should not become too large.

A global refinement of the mesh would lead to the best error reduction, but the amount of new unknowns might be much larger than needed to reduce the error below the given tolerance. Using local refinement, we hope to do much better.

The design of an “optimal” mesh, where the number of unknowns is as small as possible to keep the error below the tolerance, is an open problem and will probably be much too costly. Especially in the case of linear problems, the design of an optimal mesh will be much more expensive than the solution of the original problem, since the mesh optimization is a highly nonlinear problem. Usually, some heuristic arguments are then used in the algorithm. The aim is to produce a result that is “not too far” from an optimal mesh, but with a relatively small amount of additional work to generate it.

Several adaptive strategies are proposed in the literature, that give criteria which mesh elements should be marked for refinement. All strategies are based on the idea of an equidistribution of the local error to all mesh elements. Babuška and Rheinboldt [3] motivate that a mesh is almost optimal when the local errors are approximately equal for all elements. So, elements where the error indicator is large will be marked for refinement, while elements with a small error indicator are left unchanged.

The general outline of the adaptive algorithm for a stationary problem is the following. Starting from an initial triangulation \mathcal{S}_0 , we produce a sequence of triangulations \mathcal{S}_k , for $k = 1, 2, \dots$, until the estimated error is below the given tolerance:

Algorithm 1.17 (General adaptive refinement strategy).

```

Start with  $\mathcal{S}_0$  and error tolerance  $tol$ 
 $k := 0$ 
solve the discrete problem on  $\mathcal{S}_k$ 
compute global error estimate  $\eta$  and local indicators  $\eta_S$ 
while  $\eta > tol$  do
    mark elements for refinement (or coarsening)
    adapt mesh  $\mathcal{S}_k$ , producing  $\mathcal{S}_{k+1}$ 
     $k := k + 1$ 
    solve the discrete problem on  $\mathcal{S}_k$ 
    compute global error estimate  $\eta$  and local indicators  $\eta_S$ 
end while

```

1.5.2 Mesh refinement strategies

Since a discrete problem has to be solved in every iteration of this algorithm, the number of iterations should be as small as possible. Thus, the marking strategy should select not too few mesh elements for refinement in each cycle.

On the other hand, not much more elements should be selected than is needed to reduce the error below the given tolerance.

In the sequel, we describe several marking strategies that are commonly used in adaptive finite element methods.

The basic assumption for all marking strategies is the fact that the mesh is “optimal” when the local error is the same for all elements of the mesh. This optimality can be shown under some heuristic assumptions, see [3]. Since the true error is not known we try to equidistribute the local error indicators. This is motivated by the lower bound for error estimators of elliptic problems. This lower bound ensures that the local error is large if the local indicator is large and data of the problem is sufficiently resolved [2, 73]. As a consequence, elements with a large local error indicator should be refined, while elements with a very small local error indicator may be coarsened.

Global refinement. The simplest strategy is not really “adaptive” at all, at least not producing a locally refined mesh. It refines the mesh globally, until the given tolerance is reached.

If an a priori estimate for the error in terms of the maximal size of a mesh element h is known, where the error is bounded by a positive power of h , and if the error estimate tends to zero if the error gets smaller, then this strategy is guaranteed to produce a mesh and a discrete solution which meets the error tolerance.

But, in most cases, global refinement produces far too much mesh elements than are needed to meet the error tolerance.

Maximum strategy. Another very simple strategy is the maximum strategy. A threshold $\gamma \in (0, 1)$ is given, and all elements $S \in \mathcal{S}_k$ with

$$\eta_S > \gamma \max_{S' \in \mathcal{S}_k} \eta_{S'} \quad (1.18)$$

are marked for refinement. A small γ leads to more refinement and maybe non-optimal meshes, while a large γ leads to more cycles until the error tolerance is reached, but usually produces a mesh with less unknowns. Typically, a threshold value $\gamma = 0.5$ is used when the power p in (1.17) is $p = 2$ [72, 75].

Algorithm 1.18 (Maximum strategy).

```

Given parameter  $\gamma \in (0, 1)$ 
 $\eta_{\max} := \max(\eta_S, S \in \mathcal{S}_k)$ 
for all  $S$  in  $\mathcal{S}_k$  do
    if  $\eta_S > \gamma \eta_{\max}$  then mark  $S$  for refinement
end for

```

Equidistribution strategy. Let N_k be the number of mesh elements in \mathcal{S}_k . If we assume that the error indicators are equidistributed over all elements, i. e. $\eta_S = \eta_{S'}$ for all $S, S' \in \mathcal{S}_k$, then

$$\eta = \left(\sum_{S \in \mathcal{S}_k} \eta_S^p \right)^{1/p} = N_k^{1/p} \eta_S \stackrel{!}{=} \text{tol} \quad \text{and} \quad \eta_S = \frac{\text{tol}}{N_k^{1/p}}.$$

We can try to reach this equidistribution by refining all elements where it is violated because the error indicator is larger than $\text{tol}/N_k^{1/p}$. To make the procedure more robust, a parameter $\theta \in (0, 1)$, $\theta \approx 1$, is included in the method.

Algorithm 1.19 (Equidistribution strategy[36]).

```

Start with parameter  $\theta \in (0, 1)$ ,  $\theta \approx 1$ 
for all  $S$  in  $\mathcal{S}_k$  do
    if  $\eta_S > \theta \text{tol}/N_k^{1/p}$  then mark  $S$  for refinement
end for

```

If the error η is already near tol , then the choice $\theta = 1$ leads to the selection of only very few elements for refinement, which results in more iterations of the adaptive process. Thus, θ should be chosen smaller than 1, for example $\theta = 0.9$. Additionally, this accounts for the fact that the number of mesh elements increases, i. e. $N_{k+1} > N_k$, and thus the tolerance for local errors will be smaller after refinement.

Guaranteed error reduction strategy. Usually, it is not clear whether the adaptive refinement strategy Algorithm 1.17 using a marking strategy (other than global refinement) will converge and stop. Dörfler [31] describes a strategy with a guaranteed error reduction for the Poisson equation within a given tolerance.

We need the assumptions, that

- given data of the problem (like the right hand side) is sufficiently resolved by the initial mesh \mathcal{S}_0 with respect to the given tolerance (such that, for example, errors from the numerical quadrature are negligible),
- all edges of marked mesh elements are at least bisected by the refinement procedure (using regular refinement or two/three iterated bisections of triangles/tetrahedra, for example).

The idea is to refine a subset of the triangulation whose element errors sum up to a fixed amount of the total error η . Given a parameter $\theta_* \in (0, 1)$, the procedure is:

$$\text{Mark a set } \mathcal{A} \subseteq \mathcal{S}_k \text{ such that } \sum_{S \in \mathcal{A}} \eta_S^p \geq (1 - \theta_*)^p \eta^p.$$

It follows from the assumptions that the error will be reduced by at least a factor $\kappa < 1$ depending of θ_* and data of the problem. Selection of the set \mathcal{A} can be done in the following way. The maximum strategy threshold γ is reduced in small steps of size $\nu \in (0, 1)$, $\nu \ll 1$, until the maximum strategy marks a set which is large enough. This inner iteration is not costly in terms of CPU time as no computations are performed.

Algorithm 1.20 (Guaranteed error reduction strategy[31]).

Start with given parameters $\theta_* \in (0, 1)$, $\nu \in (0, 1)$

```

 $\eta_{\max} := \max(\eta_S, S \in \mathcal{S}_k)$ 
sum := 0
 $\gamma := 1$ 
while sum <  $(1 - \theta_*)^p \eta^p$  do
   $\gamma := \gamma - \nu$ 
  for all  $S$  in  $\mathcal{S}_k$  do
    if  $S$  is not marked
      if  $\eta_S > \gamma \eta_{\max}$ 
        mark  $S$  for refinement
        sum := sum +  $\eta_S^p$ 
      end if
    end if
  end for
end while

```

Using the above algorithm, Dörfler [30] describes a robust adaptive strategy also for the *nonlinear* Poisson equation $-\Delta u = f(u)$. It is based on a posteriori error estimates and a posteriori saturation criteria for the approximation of the nonlinearity.

Remark 1.21. Using this GERS strategy and an additional marking of elements due to data approximation, Morin, Nochetto, and Siebert [50, 51, 52] could remove the assumption that data is sufficiently resolved on \mathcal{S}_0 in order to prove convergence. The result is a simple and efficient adaptive finite element method for linear elliptic PDEs with a linear rate of convergence without any preliminary mesh adaptation.

Other refinement strategies. Babuška and Rheinboldt [3] describe an extrapolation strategy, which estimates the local error decay. Using this estimate, refinement of elements is done when the actual local error is larger than the biggest expected local error *after refinement*.

Jarusch [41] describes a strategy which generates quasi-optimal meshes. It is based on an optimization procedure involving the increase of a cost function during refinement and the profit while minimizing an energy functional.

For special applications, additional information may be generated by the error estimator and used by the adaptive strategy. This includes (anisotropic)

directional refinement of elements [43, 66], or the decision of local h - or p -enrichment of the finite element space [27, 61].

1.5.3 Coarsening strategies

Up to now we presented only refinement strategies. Practical experience indicates that for linear elliptic problems, no more is needed to generate a quasi-optimal mesh with nearly equidistributed local errors.

In time dependent problems, the regions where large local errors are produced can move in time. In stationary nonlinear problems, a bad resolution of the solution on coarse meshes may lead to some local refinement where it is not needed for the final solution, and the mesh could be coarsened again. Both situations result in the need to coarsen the mesh at some places in order to keep the number of unknowns small.

Coarsening of the mesh can produce additional errors. Assuming that these are bounded by an a posteriori estimate $\eta_{c,S}$, we can take this into account during the marking procedure.

Some of the refinement strategies described above can also be used to mark mesh elements for coarsening. Actually, elements will only be coarsened if all neighbour elements which are affected by the coarsening process are marked for coarsening, too. This makes sure that only elements where the error is small enough are coarsened, and motivates the coarsening algorithm in Section 1.1.2.

The main concept for coarsening is again the equidistribution of local errors mentioned above. Only elements with a very small local error estimate are marked for coarsening. On the other hand, such a coarsening tolerance should be small enough such that the local error *after coarsening* should not be larger than the tolerance used for refinement. If the error after coarsening gets larger than this value, the elements would be directly refined again in the next iteration (which may lead to a sequence of oscillating grid never meeting the desired criterion).

Usually, an upper bound μ for the mesh size power of the local error estimate is known, which can be used to determine the coarsening tolerance: if

$$\eta_S \leq ch_S^\mu,$$

then coarsening by undoing b bisections will enlarge the local error by a factor smaller than $2^{\mu b/\text{DIM}}$, such that the local coarsening tolerance tol_c should be smaller than

$$tol_c \leq \frac{tol_r}{2^{\mu b/\text{DIM}}},$$

where tol_r is the local refinement tolerance.

Maximum strategy. Given two parameters $\gamma > \gamma_c$, refine all elements S with

$$\eta_S^p > \gamma \max_{S' \in \mathcal{S}_k} \eta_{S'}^p$$

and mark all elements S with

$$\eta_S^p + \eta_{c,S}^p \leq \gamma_c \max_{S' \in \mathcal{S}_k} \eta_{S'}^p$$

for coarsening.

Equidistribution strategy. Equidistribution of the tolerated error tol leads to

$$\eta_S \approx \frac{tol}{N_k^{1/p}} \quad \text{for all } S \in \mathcal{S}.$$

If the local error at an element is considerably smaller than this mean value, we may coarsen the element without producing an error that is too large. All elements with

$$\eta_S > \theta \frac{tol}{N_k^{1/p}}$$

are marked for refinement, while all elements with

$$\eta_S + \eta_{c,S} \leq \theta_c \frac{tol}{N_k^{1/p}}$$

are marked for coarsening.

Guaranteed error reduction strategy. Similar to the refinement in Algorithm 1.20, Dörfler [32] describes a marking strategy for coarsening. Again, the idea is to coarsen a subset of the triangulation such that the additional error after coarsening is not larger than a fixed amount of the given tolerance tol . Given a parameter $\theta_c \in (0, 1)$, the procedure is:

$$\text{Mark a set } \mathcal{B} \subseteq \mathcal{S}_k \text{ such that } \sum_{S \in \mathcal{B}} \eta_S^p + \eta_{c,S}^p \leq \theta_c^p \eta^p.$$

The selection of the set \mathcal{B} can be done similar to Algorithm 1.20.

Remark 1.22. When local h - and p -enrichment and coarsening of the finite element space is used, then the threshold θ_c depends on the local degree of finite elements. Thus, local thresholds $\theta_{c,S}$ have to be used.

Handling information loss during coarsening. Usually, some information is irreversibly destroyed during coarsening of parts of the mesh, compare Section 3.3.3. If the adaptive procedure iterates several times, it may occur that elements which were marked for coarsening in the beginning are not allowed to coarsen at the end. If the mesh was already coarsened, an error is produced which can not be reduced anymore.

One possibility to circumvent such problems is to delay the mesh coarsening until the final iteration of the adaptive procedure, allowing only refinements before. If the coarsening marking strategy is not too liberal (θ_c not too large), this should keep the error below the given bound.

Dörfler [32] proposes to keep all information until it is clear, after solving and by estimating the error on a (virtually) coarsened mesh, that the coarsening does not lead to an error which is too large.

1.5.4 Adaptive methods for time dependent problems

In time dependent problems, the mesh is adapted to the solution in every time step using a posteriori error estimators or indicators. This may be accompanied by an adaptive control of time step sizes, see below.

Bänsch [9] lists several different adaptive procedures (in space) for time dependent problems:

Explicit strategy: The current time step is solved once on the mesh from the previous time step, giving the solution u_h . Based on a posteriori estimates of u_h , the mesh is locally refined and coarsened. The problem is *not* solved again on the new mesh, and the solve–estimate–adapt process is *not* iterated.

This strategy is only usable when the solution is nearly stationary and does not change much in time, or when the time step size is very small. Usually, a given tolerance for the error can not be guaranteed with this strategy.

Semi-implicit strategy: The current time step is solved once on the mesh from the previous time step, giving an intermediate solution \tilde{u}_h . Based on a posteriori estimates of \tilde{u}_h , the mesh is locally refined and coarsened. This produces the final mesh for the current time step, where the discrete solution u_h is computed. The solve–estimate–adapt process is *not* iterated. This strategy works quite well, if the time steps are not too large, such that regions of refinement move too fast.

Implicit strategy A: In every time step starting from the previous time step's triangulation, a mesh is generated using local refinement and coarsening based on a posteriori estimates of a solution which is calculated on the current mesh. This solve–estimate–adapt process is iterated until the estimated error is below the given bound.

This guarantees that the estimated error is below the given bound. Together with an adaptive control of the time step size, this leads to global (in time) error bounds. If the time step size is not too large, the number of iterations of the solve–estimate–adapt process is usually very small.

Implicit strategy B: In every time step starting from the macro triangulation, a mesh is generated using local refinements based on a posteriori estimates of a solution which is calculated on the current (maybe quite coarse) mesh; no mesh coarsening is needed. This solve–estimate–adapt

process is iterated until the estimated error is below the given bound.

Like implicit strategy A, this guarantees error bounds. As the initial mesh for every time step is very coarse, the number of iterations of the solve–estimate–adapt process becomes quite large, and thus the algorithm might become expensive. On the other hand, a solution on a coarse grid is fast and can be used as a good initial guess for finer grids, which is usually better than using the solution from the old time step.

Implicit strategy B can also be used with anisotropically refined triangular meshes, see [37]. As coarsening of anisotropic meshes and changes of the anisotropy direction are still open problems, this implies that the implicit strategy A can not be used in this context.

The following algorithm implements one time step of the implicit strategy A. The adaptive algorithm ensures that the mesh refinement/coarsening is done at least once in each time step, even if the error estimate is below the limit. Nevertheless, the error might not be equally distributed over all elements; for some simplices the local error estimates might be bigger than allowed.

Algorithm 1.23 (Implicit strategy A).

```

Start with given parameters  $tol$  and time step size  $\tau$ ,
the solution  $u_n$  from the previous time step on grid  $\mathcal{S}_n$ 

 $\mathcal{S}_{n+1} := \mathcal{S}_n$ 
solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
compute error estimates on  $\mathcal{S}_{n+1}$ 
do
  mark elements for refinement or coarsening
  if elements are marked then
    adapt mesh  $\mathcal{S}_{n+1}$  producing a modified  $\mathcal{S}_{n+1}$ 
    solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
    compute error estimates on  $\mathcal{S}_{n+1}$ 
  end if
while  $\eta > tol$ 

```

Adaptive control of the time step size

A posteriori error estimates for parabolic problems usually consist of four different types of terms:

- terms estimating the initial error;
- terms estimating the error from discretization in space;
- terms estimating the error from mesh coarsening between time steps;
- terms estimating the error from discretization in time.

Thus, the total estimate can be split into parts

$$\eta_0, \eta_h, \eta_c, \text{ and } \eta_\tau$$

estimating these four different error parts. Usually, the error estimate can be written like

$$\|u(t_N) - u_N\| \leq \eta_0 + \max_{1 \leq n \leq N} \left(\eta_{\tau,n} + \left(\sum_{S \in \mathcal{S}_n} (\eta_{h,S}^p + \eta_{c,S}^p) \right)^{\frac{1}{p}} \right).$$

When a bound tol is given for the total error produced in each time step, the widely used strategy is to allow one fixed portion $\Gamma_0 tol$ to be produced by the discretization of initial data, a portion $\Gamma_h tol$ to be produced by the spatial discretization, and another portion $\Gamma_\tau tol$ of the error to be produced by the time discretization, with $\Gamma_0 + \Gamma_h + \Gamma_\tau \leq 1.0$. The adaptive procedure now tries to adjust time step sizes and meshes such that

$$\eta_0 \approx \Gamma_0 tol$$

and in every time step

$$\eta_\tau \approx \Gamma_\tau tol \quad \text{and} \quad \eta_h^p + \eta_c^p \approx (\Gamma_h tol)^p.$$

The adjustment of the time step size can be done via extrapolation techniques known from numerical methods for ordinary differential equations, or iteratively: The algorithm starts from the previous time step size τ_{old} or from an initial guess. A parameter $\delta_1 \in (0, 1)$ is used to reduce the step size until the estimate is below the given bound. If the error is smaller than the bound, the step size is enlarged by a factor $\delta_2 > 1$ (usually depending on the order of the time discretization). In this case, the actual time step is not recalculated, only the initial step size for the next time step is changed. Two additional parameters $\theta_1 \in (0, 1)$, $\theta_2 \in (0, \theta_1)$ are used to keep the algorithm robust, just like it is done in the equidistribution strategy for the mesh adaption. The algorithm starts from the previous time step size τ_{old} or from an initial guess.

If $\delta_1 \approx 1$, consecutive time steps may vary only slightly, but the number of iterations for getting the new accepted time step may increase. Again, as each iteration includes the solution of a discrete problem, this value should be chosen not too large. For a first order time discretization scheme, a common choice is $\delta_1 \approx 1/\sqrt{2}$.

Algorithm 1.24 (Time step size control).

Start with parameters $\delta_1 \in (0, 1)$, $\delta_2 > 1$, $\theta_1 \in (0, 1)$, $\theta_2 \in (0, \theta_1)$

$\tau := \tau_{old}$

Solve time step problem and estimate the error

while $\eta_\tau > \theta_1 \Gamma_\tau tol$ do

$\tau := \delta_1 \tau$

```

    Solve time step problem and estimate the error
  end while
  if  $\eta_\tau \leq \theta_2 \Gamma_\tau \text{tol}$  then
     $\tau := \delta_2 \tau$ 
  end if

```

The above algorithm controls only the time step size, but does not show the mesh adaption. There are several possibilities to combine both controls. An inclusion of the grid adaption in every iteration of Algorithm 1.24 can result in a large number of discrete problems to solve, especially if the time step size is reduced more than once. A better procedure is first to do the step size control with the old mesh, then adapt the mesh, and after this check the time error again. In combination with the implicit strategy A, this procedure leads to the following algorithm for one single time step

Algorithm 1.25 (Time and space adaptive algorithm).

```

Start with given parameter  $\text{tol}$ ,  $\delta_1 \in (0,1)$ ,  $\delta_2 > 1$ ,  $\theta_1 \in (0,1)$ ,
 $\theta_2 \in (0,\theta_1)$ , the solution  $u_n$  from the previous time step
on grid  $\mathcal{S}_n$  at time  $t_n$  with time step size  $\tau_n$ 

 $\mathcal{S}_{n+1} := \mathcal{S}_n$ 
 $\tau_{n+1} := \tau_n$ 
 $t_{n+1} := t_n + \tau_{n+1}$ 
solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
compute error estimates on  $\mathcal{S}_{n+1}$ 

while  $\eta_\tau > \theta_1 \Gamma_\tau \text{tol}$ 
   $\tau_{n+1} := \delta_1 \tau_{n+1}$ 
   $t_{n+1} := t_n + \tau_{n+1}$ 
  solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
  compute error estimates on  $\mathcal{S}_{n+1}$ 
end while

do
  mark elements for refinement or coarsening
  if elements are marked then
    adapt mesh  $\mathcal{S}_{n+1}$  producing a modified  $\mathcal{S}_{n+1}$ 
    solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
    compute estimates on  $\mathcal{S}_{n+1}$ 
  end if
  while  $\eta_\tau > \theta_1 \Gamma_\tau \text{tol}$ 
     $\tau_{n+1} := \delta_1 \tau_{n+1}$ 
     $t_{n+1} := t_n + \tau_{n+1}$ 
    solve the discrete problem for  $u_{n+1}$  on  $\mathcal{S}_{n+1}$  using data  $u_n$ 
    compute error estimates on  $\mathcal{S}_{n+1}$ 
  end while
while  $\eta_h > \text{tol}$ 

```

```

if  $\eta_\tau \leq \theta_2 \Gamma_\tau \text{tol}$  then
   $\tau_{n+1} := \delta_2 \tau_{n+1}$ 
end if

```

The adaptive a posteriori approach can be extended to the adaptive choice of the order of the time discretization: Bornemann [18, 19, 20] describes an adaptive variable order time discretization method, combined with implicit strategy B using the extrapolation marking strategy for the mesh adaption.

Implementation of model problems

In this chapter we describe the implementation of two stationary model problems (the linear Poisson equation and a nonlinear reaction–diffusion equation) and of one time dependent model problem (the heat equation). Here we give an overview how to set up an ALBERTA program for various applications. We do not go into detail when referring to ALBERTA data structures and functions. A detailed description can be found in Chapter 3. We start with the easy and straight forward implementation of the Poisson problem to learn about the basics of ALBERTA. The examples with the implementation of the nonlinear reaction-diffusion problem and the time dependent heat equation are more involved and show the tools of ALBERTA for attacking more complex problems. Removing all L^AT_EX descriptions of functions and variables results in the source code for the adaptive solvers.

During the installation of ALBERTA (described in Section 2.4) a subdirectory DEMO with sources and makefiles for these model problems is created. The corresponding ready-to-run programs can be found in the files `ellipt.c`, `heat.c`, and `nonlin.c`, `nlprob.c`, `nlsolve.c` in the subdirectory DEMO/src/Common/. Executable programs for different space dimensions can be generated in the subdirectories DEMO/src/1d/, DEMO/src/2d/, and DEMO/src/3d/ by calling `make ellipt`, `make nonlin`, and `make heat`. Graphics output for all problems is generated via a routine

```
void graphics(MESH *mesh, DOF_REAL_VEC *u_h, REAL (*get_est)(EL *));
```

which shows the geometry given by `mesh`, as well as finite element function values given by `u_h`, or local estimator values when parameter `get_est` is given, all in separate graphics windows. The source of this routine is DEMO/src/Common/graphics.c, which is self explaining and not described here in detail.

2.1 Poisson equation

In this section we describe a model implementation for the Poisson equation

$$\begin{aligned} -\Delta u &= f && \text{in } \Omega \subset \mathbb{R}^d, \\ u &= g && \text{on } \partial\Omega. \end{aligned}$$

This is the most simple elliptic problem (yet very important in many applications), but the program presents all major ingredients for general scalar stationary problems. Modifications needed for a nonlinear problem are presented in Section 2.2.




Fig. 2.1. Solution of the linear Poisson problem and corresponding mesh. The pictures were produced by GRAPE.

Data and parameters described below lead in 2d to the solution and mesh shown in Fig. 2.1. The implementation of the Poisson problem is split into several major steps which are now described in detail.

2.1.1 Include file and global variables

All ALBERTA source files must include the header file `alberta.h` with all ALBERTA type definitions, function prototypes and macro definitions:

```
#include <alberta.h>
```

For the linear scalar elliptic problem we use four global pointers to data structures holding the finite element space and components of the linear system of equations. These are used in different subroutines where such information cannot be passed via parameters.

```
static const FE_SPACE *fe_space;
static DOF_REAL_VEC   *u_h = nil;
static DOF_REAL_VEC   *f_h = nil;
static DOF_MATRIX     *matrix = nil;
```

fe_space: a pointer to the actually used finite element space; it is initialized by the function `init_dof_admin()` which is called by `GET_MESH()`, see Section 2.1.4;

u_h: a pointer to a DOF vector storing the coefficients of the discrete solution; it is initialized on the first call of `build()` which is called by `adapt_method_stat()`, see Section 2.1.7;

f_h: a pointer to a DOF vector storing the load vector; it is initialized on the first call of `build()`;

matrix: a pointer to a DOF matrix storing the system matrix; it is initialized on the first call of `build()`;

The data structure **FE_SPACE** is explained in Section 3.2.14, **DOF_REAL_VEC** in Section 3.3.2, and **DOF_MATRIX** in Section 3.3.4. Details about DOF administration **DOF_ADMIN** can be found in Section 3.3.1 and about the data structure **MESH** for a finite element mesh in Section 3.6.1.

2.1.2 The main program for the Poisson equation

The main program is very simple, it just includes the main steps needed to implement any stationary problem. Special problem-dependent aspects are hidden in other subroutines described below.

We first read a parameter file (indicating which data, algorithms, and solvers should be used; the file is described below in Section 2.1.3).

Then we initialize the mesh (the basic geometric data structure), and read the macro triangulation (including an initial global refinement, if necessary). The subdirectories **MACRO** in the **DEMO/src/*d** directories contain data for several sample macro triangulations. How to read and write macro triangulation files is explained in Section 3.2.16. The macro file name and the number of global refinements are given in the parameter file. Now, the domain's geometry is defined, and a finite element space is automatically generated via the `init_dof_admin()` routine described in Section 2.1.4 below. A call to `graphics()` displays the initial mesh.

The basic algorithmic data structure **ADAPT_STAT** introduced in Section 3.13.1 specifies the behaviour of the adaptive finite element method for stationary problems. A pre-initialized data structure is accessed by the function `get_adapt_stat()`; the most important members (`adapt->tolerance`, `adapt->strategy`, etc.) are automatically initialized with values from the parameter file; other members can be also initialized by adding similar lines for these members to the parameter file (compare Section 3.13.4). Eventually, function pointers for the problem dependent routines have to be set (`estimate`, `get_el_est`, `build`, `solve`). Since the assemblage is done in one step after all mesh modifications, only `adapt->build_after_coarsen` is used, no assemblage is done before refinement or before coarsening. These additional assemblage steps are possible and may be needed in a more general application, for details see Section 3.13.1.

The adaptive procedure is started by a call of `adapt_method_stat()`. This automatically solves the discrete problem, computes the error estimate, and refines the mesh until the given tolerance is met, or the maximal number of iterations is reached, compare Section 3.13.1. Finally, `WAIT REALLY` allows an inspection of the final solution by preventing a direct program exit with closure of the graphics windows.

```
int main(int argc, char **argv)
{
    FUNCNAME("main");
    MESH    *mesh;
    int      n_refine = 0;
    static ADAPT_STAT *adapt;
    char      filename[100];

    /*-----*/
    /*  first of all, init parameters of the init file          */
    /*-----*/

    init_parameters(0, "INIT/ellipt.dat");

    /*-----*/
    /*  get a mesh, and read the macro triangulation from file  */
    /*-----*/

    mesh = GET_MESH("ALBERTA mesh", init_dof_admin, init_leaf_data);
    GET_PARAMETER(1, "macro file name", "%s", filename);
    read_macro(mesh, filename, nil);
    GET_PARAMETER(1, "global refinements", "%d", &n_refine);
    global_refine(mesh, n_refine * DIM);
    graphics(mesh, nil, nil);

    /*-----*/
    /*  init adapt structure and start adaptive method          */
    /*-----*/

    adapt = get_adapt_stat("ellipt", "adapt", 2, nil);
    adapt->estimate = estimate;
    adapt->get_el_est = get_el_est;
    adapt->build_after_coarsen = build;
    adapt->solve = solve;

    adapt_method_stat(mesh, adapt);
    WAIT REALLY;
    return(0);
}
```

2.1.3 The parameter file for the Poisson equation

The following parameter file INIT/ellipt.dat is used for the `ellipt.c` program:

```
macro file name:      Macro/macro.amc
global refinements:  0
polynomial degree:    3

% graphic windows: solution, estimate, and mesh if size > 0
graphic windows:      300 300 300
% for graphics you can specify the range for the values of
% discrete solution for displaying: min max
% automatical scaling by display routine if min >= max
graphic range: 0.0 -1.0

solver:                2 % 1: BICGSTAB 2: CG 3: GMRES 4: ODIR 5: ORES
solver max iteration:  1000
solver restart:        10 % only used for GMRES
solver tolerance:      1.e-8
solver info:           2
solver precon:         2 % 0: no precon 1: diag precon
                        % 2: HB precon 3: BPX precon

error norm:            1 % 1: H1_NORM, 2: L2_NORM
estimator C0:          0.1 % constant of element residual
estimator C1:          0.1 % constant of jump residual
estimator C2:          0.0 % constant of coarsening estimate

adapt->strategy:       2 % 0: no adaption 1: GR 2: MS 3: ES 4: GERS
adapt->tolerance:       1.e-4
adapt->MS_gamma:        0.5
adapt->max_iteration:  20
adapt->info:            8

WAIT: 1
```

The file `Macro/macro.amc` storing data about the macro triangulation for $\Omega = (0, 1)^d$ can be found in Section 3.2.16 for 2d and 3d. The `polynomial degree` parameter selects the third order finite elements. By `graphic windows`, the number and sizes of graphics output windows are selected. This line is used by the `graphics()` routine. For 1d and 2d graphics, the range of function values might be specified (used for graph coloring and height). The solver for the linear system of equations is selected (here: the conjugate gradient solver), and corresponding parameters like preconditioner and tolerance. Parameters for the error estimator include values of different constants and selection of the error norm to be estimated (H^1 - or L^2 -norm, selection leads to multiplication with different powers of the local mesh size in the error indicators), see Section 3.14.1. An error tolerance and selection of a marking strategy with

corresponding parameters are main data given to the adaptive method. Finally, the `WAIT` parameter specifies whether the program should wait for user interaction at additional breakpoints, whenever a `WAIT` statement is executed as in the routine `graphics()` for instance.

The solution and corresponding mesh in 2d for the above parameters are shown in Fig. 2.1. As optimal parameter sets might differ for different space dimensions, separate parameter files exist in `1d/INIT/`, `2d/INIT/`, and `3d/INIT/`.

2.1.4 Initialization of the finite element space

During the initialization of the mesh by `GET_MESH()` in the main program, we have to specify all DOFs which we want to use during the simulation on the mesh. The initialization of the DOFs is implemented in the function `init_dof_admin()` which is called by `GET_MESH()`. For details we refer to Sections 3.2.15 and 3.6.2.

For the scalar elliptic problem we need one finite element space for the discretization. In this example, we use Lagrange elements and we initialize the degree of the elements via a parameter. The corresponding `fe_space` is accessed by `get_fe_space()` which automatically stores at the mesh information about the DOFs used by this finite element space.

It is possible to access several finite element spaces inside this function, for instance in a mixed finite element method, compare Section 3.6.2.

```
void init_dof_admin(MESH *mesh)
{
    FUNCNAME("init_dof_admin");
    int          degree = 1;
    const BAS_FCTS *lagrange;

    GET_PARAMETER(1, "polynomial degree", "%d", &degree);
    lagrange = get_lagrange(degree);
    TEST_EXIT(lagrange)("no lagrange BAS_FCTS\n");
    fe_space = get_fe_space(mesh, lagrange->name, nil, lagrange);
    return;
}
```

2.1.5 Functions for leaf data

As explained in Section 3.2.12, we can “hide” information which is only needed on a leaf element at the pointer to the second child. Such information, which we use here, is the local error indicator on an element. For this elliptic problem we need one `REAL` for storing this element indicator.

During mesh initialization by `GET_MESH()` in the main program, we have to give information about the size of leaf data to be stored and how to transform

leaf data from parent to children during refinement and vice versa during coarsening. The function `init_leaf_data()` initializes the leaf data used for this problem and is called by `GET_MESH()`. Here, leaf data is one structure `struct ellipt_leaf_data` and no transformation during mesh modifications is needed. The details of the `LEAF_DATA_INFO` data structure are stated in Section 3.2.12.

The error estimation is done by the library function `ellipt_est()`, see Section 3.14.1. For `ellipt_est()`, we need a function which gives read and write access to the local element error, and for the marking function of the adaptive procedure, we need a function which returns the local error indicator, see Section 3.13.1. The indicator is stored as the `REAL` member `estimate` of `struct ellipt_leaf_data` and the function `rw_el_est()` returns for each element a pointer to this member. The function `get_el_est()` returns the value stored at that member for each element.

```

struct ellipt_leaf_data
{
    REAL estimate;           /* one real for the estimate */
};

void init_leaf_data(LEAF_DATA_INFO *leaf_data_info)
{
    leaf_data_info->leaf_data_size = sizeof(struct ellipt_leaf_data);
    leaf_data_info->coarsen_leaf_data = nil; /* no transformation */
    leaf_data_info->refine_leaf_data = nil; /* no transformation */
    return;
}

static REAL *rw_el_est(EL *el)
{
    if (IS_LEAF_EL(el))
        return(&((struct ellipt_leaf_data *)LEAF_DATA(el))->estimate);
    else
        return(nil);
}

static REAL get_el_est(EL *el)
{
    if (IS_LEAF_EL(el))
        return(((struct ellipt_leaf_data *)LEAF_DATA(el))->estimate);
    else
        return(0.0);
}

```

2.1.6 Data of the differential equation

Data for the Poisson problem are the right hand side f and boundary values g . For test purposes it is convenient to have access to an exact solution of the problem. In this example we use the function

$$u(x) = e^{-10|x|^2}$$

as exact solution, resulting in

$$\nabla u(x) = -20 x e^{-10|x|^2}$$

and

$$f(x) = -\Delta u(x) = -(400|x|^2 - 20d) e^{-10|x|^2}.$$

Here, d denotes the space dimension, $\Omega \subset \mathbb{R}^d$. The functions `u` and `grd_u` are the implementation of u and ∇u and are optional (and usually not known for a general problem). The functions `g` and `f` are implementations of the boundary values and the right hand side and are not optional.

```
static REAL u(const REAL_D x)
{
    return(exp(-10.0*SCP_DOW(x,x)));
}

static const REAL *grd_u(const REAL_D x)
{
    static REAL_D grd;
    REAL          ux = exp(-10.0*SCP_DOW(x,x));
    int           n;

    for (n = 0; n < DIM_OF_WORLD; n++)
        grd[n] = -20.0*x[n]*ux;

    return(grd);
}

static REAL g(const REAL_D x)    /* boundary values, not optional */
{
    return(u(x));
}

static REAL f(const REAL_D x)    /* -Delta u, not optional      */
{
    REAL  r2 = SCP_DOW(x,x), ux = exp(-10.0*r2);
    return(-(400.0*r2 - 20.0*DIM)*ux);
}
```


2.1.7 The assemblage of the discrete system

For the assemblage of the discrete system we use the tools described in Sections 3.12.2, 3.12.4, and 3.12.5. For the matrix assemblage we have to provide an element-wise description of the differential operator. Following the description in Section 1.4.7 we provide the function `init_element()` for an initialization of the operator on an element and the function `LALt()` for the computation of $\det|DF_S|AAA^t$ on the actual element, where A is the Jacobian of the barycentric coordinates, DF_S the the Jacobian of the element parameterization, and A the matrix of the second order term. For $-\Delta$, we have $A = id$ and $\det|DF_S|AA^t$ is the description of differential operator since no lower order terms are involved.

For passing information about the Jacobian A of the barycentric coordinates and $\det|DF_S|$ from the function `init_element()` to the function `LALt()` we use the data structure `struct op_info` which stores the Jacobian and the determinant. The function `init_element()` calculates the Jacobian and the determinant by the library function `el_grd_lambda()` and the function `LALt()` uses these values in order to compute $\det|DF_S|AA^t$.

Pointers to these functions and to one structure `struct op_info` are members of a structure `OPERATOR_INFO` which is used for the initialization of a function for the automatic assemblage of the global system matrix. For more general equations with lower order terms, additional functions `Lb0`, `Lb1`, and/or `c` have to be defined at that point. The full description of the function `fill_matrix_info()` for general differential operators is given in Section 3.12.2. Currently, the functions `init_element()` and `LALt()` only work properly for `DIM_OF_WORLD == DIM`. The initialization of the `EL_MATRIX_INFO` structure is only done on the first call of the function `build()` which is called by `adapt_method_stat()` during the adaptive cycle (compare Section 3.13.1).

By calling `dof_compress()`, unused DOF indices are removed such that the valid DOF indices are consecutive in their range. This guarantees optimal performance of the BLAS1 routines used in the iterative solvers and `admin->size_used` is the dimension of the current finite element space. This dimension is printed for information.

On the first call, `build()` also allocates the DOF vectors `u_h` and `f_h`, and the DOF matrix `matrix`. The vector `u_h` additionally is initialized with zeros and the function pointers for an automatic interpolation during refinement and coarsening are adjusted to the predefined functions in `fe_space->bas_fcts`. The load vector `f_h` and the system matrix `matrix` are newly assembled on each call of `build()`. Thus, there is no need for interpolation during mesh modifications or initialization.

On each call of `build()` the matrix is assembled by first clearing the matrix using the function `clear_dof_matrix()` and then adding element contributions by `update_matrix()`. This function will call `init_element()` and `LALt()` on each element.

The load vector `f_h` is then initialized with zeros and the right hand side is added by `L2scp_fct_bas()`. Finally, Dirichlet boundary values are set for all Dirichlet DOFs in the load vector and the discrete solution `u_h` by `dirichlet_bound()`, compare Section 3.12.4 and 3.12.5.

```

struct op_info
{
    REAL_D  Lambda[DIM+1]; /* gradient of barycentric coordinates */
    REAL    det;           /* |det D F_S| */
};

static void init_element(const EL_INFO *el_info, const QUAD *quad[3],
                        void *ud)
{
    struct op_info *info = ud;

    info->det = el_grd_lambda(el_info, info->Lambda);
    return;
}

const REAL (*LALt(const EL_INFO *el_info, const QUAD *quad,
                  int iq, void *ud))[DIM+1]
{
    struct op_info *info = ud;
    int i, j, k;
    static REAL LALt[DIM+1][DIM+1];

    for (i = 0; i <= DIM; i++)
        for (j = i; j <= DIM; j++)
        {
            for (LALt[i][j] = k = 0; k < DIM_OF_WORLD; k++)
                LALt[i][j] += info->Lambda[i][k]*info->Lambda[j][k];
            LALt[i][j] *= info->det;
            LALt[j][i] = LALt[i][j];
        }

    return((const REAL (*)(DIM+1)) LALt);
}

static void build(MESH *mesh, U_CHAR flag)
{
    FUNCNAME("build");
    static const EL_MATRIX_INFO *matrix_info = nil;
    const QUAD *quad;

    dof_compress(mesh);
    MSG("%d DOFs for %s\n", fe_space->admin->size_used,
        fe_space->name);
}

```

```

if (!u_h)          /* access matrix and vector for linear system */
{
    matrix = get_dof_matrix("A", fe_space);
    f_h     = get_dof_real_vec("f_h", fe_space);
    u_h     = get_dof_real_vec("u_h", fe_space);
    u_h->refine_interpol = fe_space->bas_fcts->real_refine_inter;
    u_h->coarse_restrict = fe_space->bas_fcts->real_coarse_inter;
    dof_set(0.0, u_h);    /* initialize u_h ! */
}

if (!matrix_info) /* information for matrix assembling */
{
    OPERATOR_INFO o_info = {nil};

    o_info.row_fe_space = o_info.col_fe_space = fe_space;
    o_info.init_element = init_element;
    o_info.LALt         = LALt;
    o_info.LALt_pw_const = true;
    o_info.LALt_symmetric = true;
    o_info.use_get_bound = true;
    o_info.user_data = MEM_ALLOC(1, struct op_info);
    o_info.fill_flag = CALL_LEAF_EL|FILL_COORDS;

    matrix_info = fill_matrix_info(&o_info, nil);
}

clear_dof_matrix(matrix);          /* assembling of matrix */
update_matrix(matrix, matrix_info);

dof_set(0.0, f_h);                /* assembling of load vector */
quad = get_quadrature(DIM, 2*fe_space->bas_fcts->degree - 2);
L2scp_fct_bas(f, quad, f_h);

dirichlet_bound(g, f_h, u_h, nil); /* boundary values */
return;
}

```

2.1.8 The solution of the discrete system

The function `solve()` computes the solution of the resulting linear system. It is called by `adapt_method_stat()` (compare Section 3.13.1). The system matrix for the Poisson equation is positive definite and symmetric for non-Dirichlet DOFs. Thus, the solution of the resulting linear system is rather easy and we can use any preconditioned Krylov-space solver (`oem_solve_s()`), compare Section 3.15.2. On the first call of `solve()`, the parameters for the linear solver are initialized and stored in `static` variables. For the OEM solver

we have to initialize the `solver`, the tolerance `tol` for the residual, a maximal number of iterations `miter`, the level of information printed by the linear solver, and the use of a preconditioner by the parameter `icon`, which may be 0 (no preconditioning), 1 (diagonal preconditioning), 2 (hierarchical basis preconditioning), or 3 (BPX preconditioning). If GMRes is used, then the dimension of the Krylov-space for the minimizing procedure is needed, too.

After solving the discrete system, the discrete solution (and mesh) is displayed by calling `graphics()`.

```
static void solve(MESH *mesh)
{
    FUNCNAME("solve");
    static REAL      tol = 1.e-8;
    static int       miter = 1000, info = 2, icon = 1, restart = 0;
    static OEM_SOLVER solver = NoSolver;

    if (solver == NoSolver)
    {
        tol = 1.e-8;
        GET_PARAMETER(1, "solver", "%d", &solver);
        GET_PARAMETER(1, "solver tolerance", "%f", &tol);
        GET_PARAMETER(1, "solver precon", "%d", &icon);
        GET_PARAMETER(1, "solver max iteration", "%d", &miter);
        GET_PARAMETER(1, "solver info", "%d", &info);
        if (solver == GMRes)
            GET_PARAMETER(1, "solver restart", "%d", &restart);
    }
    oem_solve_s(matrix, f_h, u_h, solver, tol, icon, restart, miter,
                info);

    graphics(mesh, u_h, nil);
    return;
}
```

2.1.9 Error estimation

The last ingredient missing for the adaptive procedure is a function for an estimation of the error. For an elliptic problem with constant coefficients in the second order term this can be done by the library function `ellipt_est()` which implements the standard residual type error estimator and is described in Section 3.14.1. `ellipt_est()` needs a pointer to a function for writing the local error indicators (the function `rw_el_est()` described above in Section 2.1.5) and a function for the evaluation of the lower order terms of the element residuals at quadrature nodes. For the Poisson equation, this function has to return the negative value of the right hand side f at that node (which is implemented in `r()`). Since we only have to evaluate the right hand side

f , the init flag `r_flag` is zero. For an equation with lower order term involving the discrete solution or its derivative this flag has to be `INIT_UH` and/or `INIT_GRD_UH`, if needed by `r()`, compare Example 3.34.

The function `estimate()`, which is called by `adapt_method_stat()`, first initializes parameters for the error estimator, like the estimated norm and constants in front of the residuals. On each call the error estimate is computed by `ellipt_est()`. The degrees for quadrature formulas are chosen according to the degree of finite element basis functions. Additionally, as the exact solution for our test problem is known (defined by `u()` and `grd_u()`), the true error between discrete and exact solutions is calculated by the function `H1_err()` or `L2_err()`, and the ratio of the true and estimated errors is printed (which should be approximately constant). The experimental orders of convergence of the estimated and exact errors are calculated, which should both be, when using global refinement with DIM bisection refinements, `fe_space->bas_fcts->degree` for the H^1 norm and `fe_space->bas_fcts->degree+1` for the L^2 norm. Finally, the error indicators are displayed by calling `graphics()`.

```
static REAL r(const EL_INFO *el_info, const QUAD *quad, int iq,
              REAL uh_iq, const REAL_D grd_uh_iq)
{
    REAL_D      x;
    coord_to_world(el_info, quad->lambda[iq], x);
    return(-f(x));
}

#define EOC(e, eo) log(eo/MAX(e, 1.0e-15))/M_LN2

static REAL estimate(MESH *mesh, ADAPT_STAT *adapt)
{
    FUNCNAME("estimate");
    static int      degree, norm = -1;
    static REAL     C[3] = {1.0, 1.0, 0.0};
    static REAL     est, est_old = -1.0, err, err_old = -1.0;
    static FLAGS    r_flag = 0; /* INIT_UH|INIT_GRD_UH, if needed */
    REAL_DD         A = {{0.0}};
    int             n;
    const QUAD      *quad;

    for (n = 0; n < DIM_OF_WORLD; n++)
        A[n][n] = 1.0; /* set diag. of A; other elements are zero */

    if (norm < 0)
    {
        norm = H1_NORM;
        GET_PARAMETER(1, "error norm", "%d", &norm);
        GET_PARAMETER(1, "estimator CO", "%f", &C[0]);
    }
}
```

```

    GET_PARAMETER(1, "estimator C1", "%f", &C[1]);
    GET_PARAMETER(1, "estimator C2", "%f", &C[2]);
}
degree = 2*u_h->fe_space->bas_fcts->degree;
est = ellipt_est(u_h, adapt, rw_el_est, nil, degree, norm, C,
                (const REAL_D *) A, r, r_flag);

MSG("estimate   = %.8le", est);
if (est_old >= 0)
    print_msg(" EOC: %.2lf\n", EOC(est,est_old));
else
    print_msg("\n");
est_old = est;

quad = get_quadrature(DIM, degree);
if (norm == L2_NORM)
    err = L2_err(u, u_h, quad, 0, nil, nil);
else
    err = H1_err(grd_u, u_h, quad, 0, nil, nil);

MSG("||u-uh||%s = %.8le", norm == L2_NORM ? "L2" : "H1", err);
if (err_old >= 0)
    print_msg(" EOC: %.2lf\n", EOC(err,err_old));
else
    print_msg("\n");
err_old = err;
MSG("||u-uh||%s/estimate = %.2lf\n",
    norm == L2_NORM ? "L2" : "H1", err/MAX(est,1.e-15));

graphics(mesh, nil, get_el_est);
return(adapt->err_sum);
}

```

2.2 Nonlinear reaction–diffusion equation

In this section, we discuss the implementation of a stationary, nonlinear problem. Due to the nonlinearity, the computation of the discrete solution is more complex. The solver for the nonlinear reaction–diffusion equation and the solver for Poisson equation, described in Section 2.1, thus mainly differ in the routines `build()` and `solve()`.

Here we describe the solution by a Newton method, which involves the assemblage and solution of a linear system in each iteration. Hence, we do not split the assemble and solve routines in `build()` and `solve()` as in the solver for the Poisson equation (compare Sections 2.1.7 and 2.1.8), but only set Dirichlet boundary values for the initial guess in `build()` and solve the nonlinear equation (including the assemblage of linearized systems) in `solve()`. The

actual solution process is implemented by several subroutines in the separate file `nlsolve.c`, see Sections 2.2.5 and 2.2.6.

Additionally we describe a simple way to handle different problem data easily, see Sections 2.2.1 and 2.2.8.

We consider the following nonlinear reaction–diffusion equation:

$$-k\Delta u + \sigma u^4 = f + \sigma u_{ext}^4 \quad \text{in } \Omega \subset \mathbb{R}^d, \quad (2.1a)$$

$$u = g \quad \text{on } \partial\Omega. \quad (2.1b)$$

For $\Omega \subset \mathbb{R}^2$, this equation models the heat transport in a thin plate Ω which radiates heat and is heated by an external heat source f . Here, k is the constant heat conductivity, σ the Stefan–Boltzmann constant, g the temperature at the edges of the plate and u_{ext} the temperature of the surrounding space (absolute temperature in $^\circ K$).

The solver is applied to following data:

- For testing the solver we again use the ‘exponential peak’

$$u(x) = e^{-10|x|^2}, \quad x \in \Omega = (-1, 1)^d, \quad k = 1, \quad \sigma = 1, \quad u_{ext} = 0.$$

- In general (due to the nonlinearity), the problem is not uniquely solvable; depending on the initial guess for the nonlinear solver at least two discrete solutions can be obtained by using data

$$\Omega = (0, 1)^d, \quad k = 1, \quad \sigma = 1, \quad f \equiv 1, \quad g \equiv 0, \quad u_{ext} = 0.$$

and the interpolant of

$$u_0(x) = 4^d U_0 \prod_{i=1}^d x_i (1 - x_i) \quad \text{with } U_0 \in [-5.0, 1.0].$$

as initial guess for the discrete solution on the coarsest grid.

- The last application now addresses a physical problem in 2d with following data:

$$\Omega = (-1, 1)^2, \quad k = 2, \quad \sigma = 5.67\text{e-}8, \quad g \equiv 300, \quad u_{ext} = 273,$$

$$f(x) = \begin{cases} 150, & \text{if } x \in (-\frac{1}{2}, \frac{1}{2})^2, \\ 0, & \text{otherwise.} \end{cases}$$

2.2.1 Program organization and header file

The implementation is split into three source files:

nonlin.c: main program with all subroutines for the adaptive procedure; initializes DOFs, leaf data and problem dependent data in `main()` and the `solve()` routine calls the nonlinear solver;

nlprob.c: definition of problem dependent data;

nlsolve.c: implementation of the nonlinear solver.

Data structures used in all source files, and prototypes of functions are defined in the header file **nonlin.h**, which includes the **alberta.h** header file on the first line. This file is included by all three source files.

```
typedef struct prob_data PROB_DATA;
struct prob_data
{
    REAL          k, sigma;

    REAL          (*g)(const REAL_D x);
    REAL          (*f)(const REAL_D x);

    REAL          (*u0)(const REAL_D x);

    REAL          (*u)(const REAL_D x);
    const REAL    *(*grd_u)(const REAL_D x);
};

/*--- file nlprob.c -----*/
const PROB_DATA *init_problem(MESH *mesh);

/*--- file nlsolve.c -----*/
int nlsolve(DOF_REAL_VEC *, REAL, REAL, REAL (*)(const REAL_D));

/*--- file graphics.c -----*/
void graphics(MESH *, DOF_REAL_VEC *, REAL (*)(EL *));
```

The data structure **PROB_DATA** yields following information:

k: diffusion coefficient (constant heat conductivity);

sigma: reaction coefficient (Stefan–Boltzmann constant);

g: pointer to a function for evaluating boundary values;

f: pointer to a function for evaluating the right-hand side ($f + \sigma u_{ext}^4$);

u0: pointer to a function for evaluating an initial guess for the discrete solution on the macro triangulation, if not **nil**;

u: pointer to a function for evaluating the true solution, if not **nil** (only for test purpose);

grd_u: pointer to a function for evaluating the gradient of the true solution, if not **nil** (only for test purpose).

The function **init_problem()** initializes problem data, like boundary values, right hand side, etc. which is stored in a **PROB_DATA** structure and reads data of the macro triangulation for the actual problem. The function **nlsolve()** implements the nonlinear solver by a Newton method including the assemblage and solution of the linearized sub-problems and **graphics()**

is the routine for visualization already known from the solver for the Poisson equation, compare Section 2.1.

2.2.2 Global variables

In the main source file for the nonlinear solver `nonlin.c` we use the following global variables:

```
#include "nonlin.h"

static const FE_SPACE *fe_space;
static DOF_REAL_VEC *u_h = nil;
static const PROB_DATA *prob_data = nil;
```

As in the solver for the linear Poisson equation, we have a pointer to the used `fe_space` and the discrete solution `u_h`. In this file, we do not need a pointer to a `DOF_MATRIX` for storing the system matrix and a pointer to a `DOF_REAL_VEC` for storing the right hand side. The system matrix and right hand side are handled by the nonlinear solver `nlsolve()`, implemented in `nlsolve.c`. Data about the problem is handled via the `prob_data` pointer.

2.2.3 The main program for the nonlinear reaction–diffusion equation

The main program is very similar to the main program of the Poisson problem described in Section 2.1.2. A new feature is that besides a parameter initialization from the file `INIT/nonlin.dat`, parameters can also be defined and overwritten via additional arguments on the command line. For the definition of a parameter via the command line we need for each parameter a pair of arguments: the key (without terminating ‘:’) and the value.

```
nonlin "problem number" 1
```

will for instance overwrite the value of `problem number` in `nonlin.dat` with the value 1.

After processing command line arguments, the mesh with the used DOFs and leaf data is initialized, problem dependent data, including the macro triangulation, are initialized by `init_problem(mesh)` (see Section 2.2.8), the structure for the adaptive method is filled and finally the adaptive method is started.

```
int main(int argc, char **argv)
{
    FUNCNAME("main");
    MESH *mesh;
    ADAPT_STAT *adapt; int k;

    /*-----*/
    /* first of all, init parameters from init file and command line */
```

```

/*-----*/
    init_parameters(0, "INIT/nonlin.dat");
    for (k = 1; k+1 < argc; k += 2)
        ADD_PARAMETER(0, argv[k], argv[k+1]);

/*-----*/
/*  get a mesh with DOFs and leaf data */
/*-----*/
    mesh = GET_MESH("mesh", init_dof_admin, init_leaf_data);

/*-----*/
/*  init problem dependent data and read macro triangulation */
/*-----*/
    prob_data = init_problem(mesh);

/*-----*/
/*  init adapt struture and start adaptive method */
/*-----*/
    adapt = get_adapt_stat("nonlin", "adapt", 1, nil);
    adapt->estimate = estimate;
    adapt->get_el_est = get_el_est;
    adapt->build_after_coarsen = build;
    adapt->solve = solve;

    adapt_method_stat(mesh, adapt);

    WAIT REALLY;
    return(0);
}

```

2.2.4 Initialization of the finite element space and leaf data

The functions for initializing DOFs to be used (`init_dof_admin()`), leaf data (`init_leaf_data()`), and for accessing leaf data (`rw_el_est()`, `get_el_est()`) are exactly the same as in the solver for the linear Poisson equation, compare Sections 2.1.4 and 2.1.5.

2.2.5 The build routine

As mentioned above, inside the build routine we only access one vector for storing the discrete solution. On the coarsest grid, the discrete solution is initialized with zeros, or by interpolating the function `prob_data->u0`, which implements an initial guess for the discrete solution. On a refined grid we do not initialize the discrete solution again. Here, we use the discrete solution from the previous step, which is interpolated during mesh modifications, as an initial guess.

In each adaptive cycle, Dirichlet boundary values are set for the discrete solution. This ensures $u_0 \in g_h + \mathring{X}_h$ for the initial guess of the Newton method.

```
static void build(MESH *mesh, U_CHAR flag)
{
    FUNCNAME("build");

    dof_compress(mesh);
    MSG("%d DOFs for %s\n", fe_space->admin->size_used,
        fe_space->name);

    if (!u_h)          /* access and initialize discrete solution */
    {
        u_h = get_dof_real_vec("u_h", fe_space);
        u_h->refine_interpol = fe_space->bas_fcts->real_refine_inter;
        u_h->coarse_restrict = fe_space->bas_fcts->real_coarse_inter;
        if (prob_data->u0)
            interpol(prob_data->u0, u_h);
        else
            dof_set(0.0, u_h);
    }
    dirichlet_bound(prob_data->g, u_h, nil, nil);

    return;
}
```

2.2.6 The solve routine

The `solve()` routine solves the nonlinear equation by calling the function `nlsolve()` which is implemented in `nlsolve.c` and described below in Section 2.2.10. After solving the discrete problem, the new discrete solution is displayed via the `graphics()` routine.

```
static void solve(MESH *mesh)
{
    nlsolve(u_h, prob_data->k, prob_data->sigma, prob_data->f);
    graphics(mesh, u_h, nil);
    return;
}
```

2.2.7 The estimator for the nonlinear problem

In comparison to the Poisson program, the function `r()` which implements the lower order term in the element residual changes due to the term σu^4 in the differential operator, compare Section 3.14.1. The right hand side $f + \sigma u_{ext}^4$ is already implemented in the function `prob_data->f()`.

In the function `estimate()` we have to initialize the diagonal of **A** with the heat conductivity `prob_data->k` and for the function `r()` we need the values

of u_h at the quadrature node, thus `r_flag = INIT_UH` is set. The initialization of parameters for the estimator is the same as in Section 2.1.9. The true error can be computed only for the first application, where the true solution is known (`prob_data->u()` and `prob_data->grd_u()` are not `nil`). Finally, the error indicator is displayed by `graphics()`.

```
static REAL r(const EL_INFO *el_info, const QUAD *quad, int iq,
              REAL uh_iq, const REAL_D grd_uh_iq)
{
    REAL_D    x;
    REAL      uhx2 = SQR(uh_iq);

    coord_to_world(el_info, quad->lambda[iq], x);
    return(prob_data->sigma*uhx2*uhx2 - (*prob_data->f)(x));
}

#define EOC(e, eo) log(eo/MAX(e, 1.0e-15))/M_LN2

static REAL estimate(MESH *mesh, ADAPT_STAT *adapt)
{
    FUNCNAME("estimate");
    static int    degree, norm = -1;
    static REAL   C[3] = {1.0, 1.0, 0.0};
    static REAL   est, est_old = -1.0, err = -1.0, err_old = -1.0;
    static REAL   r_flag = INIT_UH;
    REAL_DD       A = {{0.0}};
    int           n;

    for (n = 0; n < DIM_OF_WORLD; n++)
        A[n][n] = prob_data->k; /* set diag.; other elements are 0 */

    if (norm < 0)
    {
        norm = H1_NORM;
        GET_PARAMETER(1, "error norm", "%d", &norm);
        GET_PARAMETER(1, "estimator C0", "%f", C);
        GET_PARAMETER(1, "estimator C1", "%f", C+1);
        GET_PARAMETER(1, "estimator C2", "%f", C+2);
    }
    degree = 2*u_h->fe_space->bas_fcts->degree;
    est = ellipt_est(u_h, adapt, rw_el_est, nil, degree, norm, C,
                    (const REAL_D *) A, r, r_flag);

    MSG("estimate   = %.8le", est);
    if (est_old >= 0)
        print_msg(" EOC: %.2lf\n", EOC(est, est_old));
    else
        print_msg("\n");
}
```

```

est_old = est;

if (norm == L2_NORM && prob_data->u)
    err = L2_err(prob_data->u, u_h, nil, 0, nil, nil);
else if (norm == H1_NORM && prob_data->grd_u)
    err = H1_err(prob_data->grd_u, u_h, nil, 0, nil, nil);

if (err >= 0)
{
    MSG("||u-uh||%s = %.8le", norm == L2_NORM ? "L2" : "H1", err);
    if (err_old >= 0)
        print_msg(", EOC: %.2lf\n", EOC(err,err_old));
    else
        print_msg("\n");
    err_old = err;
    MSG("||u-uh||%s/estimate = %.2lf\n",
        norm == L2_NORM ? "L2" : "H1", err/MAX(est,1.e-15));
}
graphics(mesh, nil, get_el_est);
return(adapt->err_sum);
}

```

2.2.8 Initialization of problem dependent data

The file `nlprob.c` contains all problem dependent data. On the first line, `nonlin.h` is included and then two variables for storing the values of the heat conductivity and the Stefan–Boltzmann constant are declared. These values are used by several functions:

```

#include "nonlin.h"
static REAL k = 1.0, sigma = 1.0;

```

The following functions are used in the first example for testing the non-linear solver (problem number: 0):

```

static REAL u_0(const REAL_D x)
{
    REAL x2 = SCP_DOW(x,x);
    return(exp(-10.0*x2));
}

static const REAL *grd_u_0(const REAL_D x)
{
    static REAL_D grd;
    REAL ux = exp(-10.0*SCP_DOW(x,x));
    int n;

    for (n = 0; n < DIM_OF_WORLD; n++)
        grd[n] = -20.0*x[n]*ux;
}

```

```

    return(grd);
}

static REAL f_0(const REAL_D x)
{
    REAL r2 = SCP_DOW(x,x), ux = exp(-10.0*r2), ux4 = ux*ux*ux*ux;
    return(sigma*ux4 - k*(400.0*r2 - 20.0*DIM)*ux);
}

```

For the computation of a stable and an unstable (but non-physical) solution, depending on the initial choice of the discrete solution, the following functions are used, which also use a global variable `U0`. Such an unstable solution in 3d is shown in Fig. 2.2. Data is given as follows (**problem number: 1**):

```

static REAL U0 = 0.0;

static REAL g_1(const REAL_D x)
{
    #if DIM_OF_WORLD == 1
        return(4.0*U0*x[0]*(1.0-x[0]));
    #endif
    #if DIM_OF_WORLD == 2
        return(16.0*U0*x[0]*(1.0-x[0])*x[1]*(1.0-x[1]));
    #endif
    #if DIM_OF_WORLD == 3
        return(64.0*U0*x[0]*(1.0-x[0])*x[1]*(1.0-x[1])*x[2]*(1.0-x[2]));
    #endif
}

static REAL f_1(const REAL_D x)
{
    return(1.0);
}

```

The last example needs functions for boundary data and right hand side and variables for the temperature at the edges, and σu_{ext}^4 . A solution to this problem is depicted in Fig. 2.3 and problem data is (**problem number: 2**):

```

static REAL g2 = 300.0, sigma_uext4 = 0.0;
static REAL g_2(const REAL_D x)
{
    return(g2);
}

```

xy-plane, x=0 y=0 z=0.5

xy-plane, x=0 y=0 z=0.5




Fig. 2.2. Graph of the unstable solution with corresponding mesh of the nonlinear reaction–diffusion problem in 3d on the clipping plane $z = 0.5$. The pictures were produced by the gltools.

```
static REAL f_2(const REAL_D x)
{
    if (x[0] >= -0.25 && x[0] <= 0.25 && x[1] >= -0.25 && x[1] <= 0.25)
        return(150.0 + sigma_uext4);
    else
        return(sigma_uext4);
}
```

Depending on the chosen problem via the parameter `problem number`, the function `init_problem()` initializes the entries of a `PROB_DATA` structure, adjusts the corresponding function pointers, reads the macro triangulation, performs some initial refinement and returns a pointer to the filled `PROB_DATA` structure. Information store in `PROB_DATA` is then used in the `build()` and the `nlsolve()` routines.

```
const PROB_DATA *init_problem(MESH *mesh)
{
    FUNCNAME("init_problem");
    static PROB_DATA prob_data = {0};
    int pn = 2, n_refine = 0;

    GET_PARAMETER(1, "problem number", "%d", &pn);
    switch (pn)
    {
    case 0: /*- problem with known true solution -----*/
        k = 1.0;
```

rotation: x=-48.08 y=0.00 z=1.77

rotation: x=-48.08 y=0.00 z=1.77



Fig. 2.3. Graph of the solution to the physical problem with corresponding mesh of the nonlinear reaction-diffusion problem in 2d. The pictures were produced by the gltools.

```

sigma = 1.0;

prob_data.g = u_0;
prob_data.f = f_0;

prob_data.u = u_0;
prob_data.grd_u = grd_u_0;

read_macro(mesh, "Macro/macro-big.amc", nil);
break;
case 1: /*- problem for computing a stable and an unstable sol. -*/
k = 1.0;
sigma = 1.0;

prob_data.g = g_1;
prob_data.f = f_1;

prob_data.u0 = g_1;
GET_PARAMETER(1, "U0", "%f", &U0);

read_macro(mesh, "Macro/macro.amc", nil);
break;
case 2: /*- physical problem -----*/
k = 2.0;
sigma = 5.67e-8;
sigma_uext4 = sigma*273*273*273*273;

prob_data.g = g_2;

```



```

    prob_data.f = f_2;
    read_macro(mesh, "Macro/macro-big.amc", nil);
    break;
default:
    ERROR_EXIT("no problem defined with problem no. %d\n", pn);
}
prob_data.k = k;
prob_data.sigma = sigma;

GET_PARAMETER(1, "global refinements", "%d", &n_refine);
global_refine(mesh, n_refine*DIM);

return(&prob_data);
}

```

2.2.9 The parameter file for the nonlinear reaction–diffusion equation

The following parameter file INIT/nonlin.dat is read by main().

```

problem number:      2
global refinements:  1
polynomial degree:   2

U0:      -5.0          % height of initial guess for Problem 1

% graphic windows: solution, estimate, and mesh if size > 0
graphic windows:      300 300 300
% for graphics you can specify the range for the values of
% discrete solution for displaying:  min max
% automatical scaling by display routine if min >= max
graphic range:  1.0 0.0

newton tolerance: 1.e-6      % tolerance for Newton
newton max. iter: 50         % maximal number of iterations of Newton
newton info:      6          % information level of Newton
newton restart:   10         % for step size control

linear solver max iteration: 1000
linear solver restart:      10 % only used for GMRES
linear solver tolerance:    1.e-8
linear solver info:         0
linear solver precon:       1 % 0: no precon 1: diag precon
                             % 2: HB precon 3: BPX precon

error norm:      1 % 1: H1_NORM, 2: L2_NORM
estimator C0:    0.1 % constant of element residual
estimator C1:    0.1 % constant of jump residual

```

```

estimator C2:          0.0 % constant of coarsening estimate

adapt->strategy:       2   % 0: no adaption 1: GR 2: MS 3: ES 4:GERS
adapt->tolerance:       1.e-1
adapt->MS_gamma:        0.5
adapt->max_iteration:   15
adapt->info:            4
WAIT: 1

```

Besides the parameters for the Newton solver and the height of the initial guess U_0 in Problem 1, the file is very similar to the parameter file `ellipt.dat` for the Poisson problem, compare Section 2.1.3. As mentioned above, additional parameters may be defined or overwritten by command line arguments, see Section 2.2.3.

2.2.10 Implementation of the nonlinear solver

In this section, we now describe the solution of the nonlinear problem which differs most from the solver for the Poisson equation. It is the last module missing for the adaptive solver. We use the abstract Newton methods of Section 3.15.6 for solving

$$u_h \in g_h + \mathring{X}_h : \quad F(u_h) = 0 \quad \text{in } \mathring{X}_h^*,$$

where $g_h \in X_h$ is an approximation to boundary data g . Using the classical Newton method, we start with an initial guess $u_0 \in g_h + \mathring{X}_h$, where Dirichlet boundary values are set in the `build()` routine (compare Section 2.2.5). For $m \geq 0$ we compute

$$d_m \in \mathring{X}_h : \quad DF(u_m)d_m = F(u_m) \quad \text{in } \mathring{X}_h^*$$

and set

$$u_{m+1} = u_m - d_m$$

until some suitable norm $\|d_m\|$ or $\|F(u_{m+1})\|$ is sufficiently small. Since the correction d_m satisfies $d_m \in \mathring{X}_h$, all Newton iterates u_m satisfy $u_m \in g_h + \mathring{X}_h$, $m \geq 0$. Newton methods with step size control solve similar defect equations and perform similar update steps, compare Section 3.15.6.

For $v \in g_h + \mathring{X}_h$ the functional $F(v) \in \mathring{X}_h^*$ of the nonlinear reaction–diffusion equation is defined by

$$\langle F(v), \varphi_j \rangle_{\mathring{X}_h^* \times \mathring{X}_h} = \int_{\Omega} k \nabla \varphi_j \nabla v + \sigma \varphi_j v^4 dx - \int_{\Omega} (f + u_{ext}^4) \varphi_j dx \quad (2.2)$$

for all $\varphi_j \in \mathring{X}_h$, and the Frechet derivative $DF(v)$ of F is given for all $\varphi_i, \varphi_j \in \mathring{X}_h$ by

$$\langle DF(v) \varphi_i, \varphi_j \rangle_{\mathring{X}_h^* \times \mathring{X}_h} = \int_{\Omega} k \nabla \varphi_j \nabla \varphi_i + 4\sigma v^3 \varphi_j \varphi_i dx. \quad (2.3)$$

The Newton solvers need a function for assembling the right hand side vector of the discrete system (2.2), and the system matrix of the linearized equation (2.3) for some given v in X_h . The system matrix is always symmetric. It is positive definite, if $v \geq 0$, and is then solved by the conjugate gradient method. For $v \not\geq 0$ BiCGStab is used. We choose the H^1 semi–norm as problem dependent norm $\|\cdot\|$.

Problem dependent data structures for assembling and solving

Similar to the assemblage of the system matrix for the Poisson problem, we define a data structure `struct op_info` in order to pass information to the routines which describe the differential operator. In the assembling of the linearized system around a given finite element function v we additionally need the diffusion coefficient k and reaction coefficient σ . In general, v is not constant on the elements, thus we have to compute the zero order term by numerical quadrature on each element. For this we need access to the used quadrature for this term, and a vector storing the values of v for all quadrature nodes.

```
struct op_info
{
  REAL_D  Lambda[DIM+1]; /* gradient of barycentric coordinates */
  REAL    det;           /* |det D F_S| */

  REAL    k, sigma;      /* diffusion and reaction coefficient */

  const QUAD_FAST *quad_fast; /* quad_fast for the zero order term */
  const REAL      *v_qp;      /* v at quadrature nodes of quad_fast */
};
```

The general Newton solvers pass data about the actual problem by void pointers to the problem dependent routines. Information that is used by these routines are collected in the data structure `NEWTON_DATA`

```
typedef struct newton_data NEWTON_DATA;
struct newton_data
{
  const FE_SPACE  *fe_space; /* used finite element space */

  REAL    k; /* diffusion coefficient */
  REAL    sigma; /* reaction coefficient */
  REAL    (*f)(const REAL_D); /* compute f + sigma u_ext^4 */

  DOF_MATRIX *DF; /* pointer to system matrix */

  /*--- parameters for the linear solver -----*/
  OEM_SOLVER solver; /* used solver: CG (v >= 0) else BiCGStab */
  REAL    tolerance;
  int     max_iter;
};
```

```

    int      icon;
    int      restart;
    int      info;
};

```

All entries of this structure besides `solver` are initialized in the function `nl_solve()`. The entry `solver` is set every time the linearized matrix is assembled.

The assembling routine

Denote by $\{\varphi_0, \dots, \varphi_{\hat{N}}\}$ the basis of \hat{X}_h , by $\{\varphi_0, \dots, \varphi_N\}$ the basis of X_h . Let \mathbf{A} be the stiffness matrix, i.e.

$$A_{ij} = \begin{cases} \int_{\Omega} k \nabla \varphi_j \nabla \varphi_i \, dx & i = 0, \dots, N, j = 0, \dots, \hat{N} \\ \delta_{ij} & i = 0, \dots, N, j = \hat{N} + 1, \dots, N, \end{cases}$$

and $\mathbf{M} = \mathbf{M}(v)$ the mass matrix, i.e.

$$M_{ij} = \begin{cases} \int_{\Omega} \sigma v^3 \varphi_j \varphi_i \, dx & i = 0, \dots, N, j = 0, \dots, \hat{N} \\ 0 & i = 0, \dots, N, j = \hat{N} + 1, \dots, N. \end{cases}$$

The system matrix \mathbf{L} , representing $DF(v)$, of the linearized equation is then given as

$$\mathbf{L} = \mathbf{A} + 4\mathbf{M}.$$

The right hand side vector \mathbf{F} , representing $F(v)$ is for all non-Dirichlet DOFs j given by

$$\begin{aligned} F_j &= \int_{\Omega} k \nabla v \nabla \varphi_j + \sigma v^4 \varphi_j \, dx - \int_{\Omega} (f + \sigma u_{ext}^4) \varphi_j \, dx \\ &= (\mathbf{A} \mathbf{v} + \mathbf{M} \mathbf{v})_j - \int_{\Omega} (f + \sigma u_{ext}^4) \varphi_j \, dx, \end{aligned} \quad (2.4)$$

where \mathbf{v} denotes the coefficient vector of v . Thus, we want to use information assembled into \mathbf{A} and \mathbf{M} for both system matrix and right hand side vector.

Unfortunately, this can not be done *after* assembling $\mathbf{A} + 4\mathbf{M}$ into the system matrix \mathbf{L} due to the different scaling of \mathbf{M} in the system matrix (factor 4) and right hand side (factor 1). Storing both matrices \mathbf{A} and \mathbf{M} is too costly, since matrices are the objects in finite element codes which need most memory.

The solution to this problem comes from the observation, that (2.4) holds also element-wise for the element contributions of the right hand side and element matrices \mathbf{A}_S and \mathbf{M}_S when replacing \mathbf{v} by the local coefficient vector \mathbf{v}_S . Hence, on elements S we compute the element contributions of \mathbf{A}_S and \mathbf{M}_S , add them to the system matrix, and use them and the local coefficient vector \mathbf{v}_S for adding the right hand side contribution to the load vector.

The resulting assembling routine is more complicated in comparison to the very simple routine used for the linear Poisson problem. On the other hand, using ALBERTA routines for the computation of element matrices, extracting local coefficient vectors, and boundary information, the routine is still rather easy to implement. The implementation does yet not depend on the actually used set of local basis functions.

The function `update()` which is now described in detail, can be seen as an example for the very flexible implementation of rather complex nonlinear and time dependent problems which often show the same structure (compare the implementation of the assembling routine for the time dependent heat equation, Section 2.3.8). It demonstrates the functionality and flexibility of the ALBERTA tools: the assemblage of complex problems is still quite easy, whereas the resulting code is quite efficient.

Similar to the linear Poisson solver, we provide a function `LALt()` for the second order term. Besides the additional scaling by the heat conductivity k , it is exactly the same as for the Poisson problem. For the nonlinear reaction–diffusion equation we also need a function `c()` for the zero order term. This term is assembled using element-wise quadrature and thus needs information about the function v used in the linearization at all quadrature nodes. Information for `LALt()` and `c()` is stored in the data structure `struct op_info`, see above. The members of this structure are initialized during mesh traversal in `update()`.

```
static const REAL (*LALt(const EL_INFO *el_info, const QUAD *quad,
                        int iq, void *ud))[DIM+1]
{
    struct op_info *info = ud;
    REAL          fac = info->k*info->det;
    int           i, j, k;
    static REAL   LALt[DIM+1][DIM+1];

    for (i = 0; i <= DIM; i++)
    {
        for (j = i; j <= DIM; j++)
        {
            for (LALt[i][j] = k = 0; k < DIM_OF_WORLD; k++)
                LALt[i][j] += info->Lambda[i][k]*info->Lambda[j][k];
            LALt[i][j] *= fac;
            LALt[j][i] = LALt[i][j];
        }
    }
    return((const REAL (*)(DIM+1)) LALt);
}

static REAL c(const EL_INFO *el_info, const QUAD *quad, int iq,
              void *ud)
```

```

{
    struct op_info *info = ud;
    REAL v3;

    TEST_EXIT(info->quad_fast->quad == quad)("quads differ\n");

    v3 = info->v_qp[iq]*info->v_qp[iq]*info->v_qp[iq];

    return(info->sigma*info->det*v3);
}

```

As mentioned above, we use a general Newton solver and a pointer to the `update()` routine is adjusted inside the function `nlsolve()` in the data structure for this solver. Such a solver does not have any information about the actual problem, nor information about the ALBERTA data structures for storing DOF vectors and matrices. This is also reflected in the arguments of `update()`:

```

static void update(void *ud, int dim, const REAL *v, int up_DF,
                  REAL *F);

```

Here, `dim` is the dimension of the discrete nonlinear problem, `v` is a *vector* storing the coefficients of the finite element function which is used for the linearization, `up_DF` is a flag indicating whether $DF(v)$ should be assembled or not. If `F` is not `nil`, then $F(v)$ should be assembled and stored in the *vector* `F`. Information about the ALBERTA finite element space, a pointer to a DOF matrix, etc. can be passed to `update()` by the `ud` pointer. The declaration

```

NEWTON_DATA    *data = ud;

```

converts the `void *` pointer `ud` into a pointer `data` to a structure `NEWTON_DATA` which gives access to all information, used for the assembling (see above). This structure is initialized in `nlsolve()` before starting the Newton method.

The `update()` routine is split into three main parts: an initialization of the assembling functions (only done on the first call), a conversion of the vectors that are arguments to the routine into DOF vectors, and finally the assembling.

Initialization of the assembling functions. The initialization of ALBERTA functions for the assembling is similar to the initialization in the `build()` routine of the linear Poisson equation (compare Section 2.1.7). There are minor differences:

1. In addition to the assemblage of the 2nd order term (see the function `LALt()`), we now have to assemble the zero order term too (see the function `c()`). The integration of the zero order term has to be done by using an element wise quadrature which needs the values of v^3 at all quadrature nodes. The two element matrices are computed separately. This makes it possible to use them for the system matrix and right hand side.

2. In the solver for the Poisson problem, we have filled an `OPERATOR_INFO` structure with information about the differential operator. This structure is an argument to `fill_matrix_info()` which returns a pointer to a structure `EL_MATRIX_INFO`. This pointer is used for the complete assemblage of the system matrix by some `ALBERTA` routine. A detailed description of this structures and the general assemblage routines for matrices can be found in Section 3.12.2. Here, we want to use only the function for computing the element matrices. Thus, we only need the entries `el_matrix_fct()` and `fill_info` of the `EL_MATRIX_INFO` structure, which are used to compute the element matrix (`fill_info` is the second argument to `el_matrix_fct()`). We initialize a function pointer `fill_a` with data pointer `a_info` for the computation of the element matrix A_S and a function pointer `fill_c` with data pointer `c_info` for the computation M_S .

All other information inside the `EL_MATRIX_INFO` structure is used for the automatic assembling of element matrices into the system matrix by `update_matrix()`. Such information can be ignored here, since this is now done in `update()`.

3. For the assembling of the element matrix into the system matrix and the element contribution of the right hand side into the load vector we need information about the number of local basis functions, `n_phi`, and how to access global DOFs from the elements, `get_dof()`. This function uses the DOF administration `admin` of the finite element space. We also need information about the boundary type of the local basis functions, `get_bound()`, and for the computation of the values of v at quadrature nodes, we have to extract the local coefficient vector from the global one, `get_v_loc()`. These functions and the number of local basis functions can be accessed via the `bas_fcts` inside the `data->fe_space` structure. The used `admin` is the `admin` structure in `data->fe_space`. For details about these functions we refer to Sections 3.5.1, 3.3.6, and 1.4.3.

Conversion of the vectors into DOF vectors. The input vector \mathbf{v} of `update()` is a vector storing the coefficients of the function used for the linearization. It is not a DOF vector, but `ALBERTA` routines for extracting a local coefficient vector need a DOF vector. Thus, we have to “convert” \mathbf{v} into some DOF vector `dof_v`. This is done by setting the members `fe_space`, `size`, and `vec` in the structure `dof_v` to the used finite element space, `data->fe_space`, the size of the vector, `dim`, and the vector, \mathbf{v} . In the assignment of the vector we have to use a cast to `(REAL *)` since \mathbf{v} is a `const REAL *` whereas the member `dof_v.vec` is `REAL *` only. This is necessary whenever a `const REAL *` vector has to be converted into a DOF vector. Nevertheless, values of such vectors like \mathbf{v} must not be changed!

After this initialization, all `ALBERTA` tools working on DOF vectors can be used. But this vector is *not* linked to the list of DOF vectors of

`fe_space->admin` and are *not* administrated by this `admin` (not enlarged during mesh modifications, e.g.)!

In the same way we have to convert `F` to a DOF vector `dof_F` if `F` is not `nil`.

The assemblage of the linearized system. If the system matrix has to be assembled, then the DOF matrix `data->DF` is cleared and we check which solver can be used for solving the linearized equation.

If the right hand side has to be assembled, then this vector is initialized with values

$$-\int_{\Omega} (f + \sigma u_{ext}^4) \varphi_j dx.$$

For the assemblage of the element contributions we use the non-recursive mesh traversal routines. On each element we access the local coefficient vector `v_loc`, the global DOFs `dof` and boundary types `bound` of the local basis functions.

Next, we initialize the Jacobian of the barycentric coordinates and compute the values of v at the quadrature node by `uh_at_qp()`. Hence v^3 can easily be calculated in `c()` at all quadrature nodes. Routines for evaluating finite element functions and their derivatives are described in detail in Section 3.9.

Now, all members of `struct op_info` are initialized, and we compute the element matrices \mathbf{A}_S by the function `fill_a()` and \mathbf{M}_S by the function `fill_c()`.

These contributions are added to the system matrix if `up_DF` is not zero. Finally, the right hand side contributions for all non Dirichlet DOFs are computed, and zero Dirichlet boundary values are set for Dirichlet DOFs, if `F` is not `nil`.

```
static void update(void *ud, int dim, const REAL *v, int up_DF,
                  REAL *F)
{
    FUNCNAME("update");
    static struct op_info *op_info = nil;
    static const REAL    **(*fill_a)(const EL_INFO *, void *) = nil;
    static void          *a_info = nil;
    static const REAL    **(*fill_c)(const EL_INFO *, void *) = nil;
    static void          *c_info = nil;

    static const DOF_ADMIN *admin = nil;
    static int             n_phi;
    static const REAL      *(*get_v_loc)(const EL *, const DOF_REAL_VEC *,
                                         REAL *);
    static const DOF        *(*get_dof)(const EL *, const DOF_ADMIN *,
                                         DOF *);
    static const S_CHAR     *(*get_bound)(const EL_INFO *, S_CHAR *);
```



```

NEWTON_DATA      *data = ud;
TRAVERSE_STACK  *stack = get_traverse_stack();
const EL_INFO    *el_info;
FLAGS           fill_flag;
DOF_REAL_VEC     dof_v = {nil, nil, "v"};
DOF_REAL_VEC     dof_F = {nil, nil, "F(v)"};

/*-----*/
/*  init functions for assembling DF(v) and F(v)  */
/*-----*/

if (admin != data->fe_space->admin)
{
    OPERATOR_INFO      o_info2 = {nil}, o_info0 = {nil};
    const EL_MATRIX_INFO *matrix_info;
    const BAS_FCTS      *bas_fcts = data->fe_space->bas_fcts;
    const QUAD          *quad = get_quadrature(DIM, 2*bas_fcts->degree-2);

    admin = data->fe_space->admin;
    n_phi   = bas_fcts->n_bas_fcts;
    get_dof  = bas_fcts->get_dof_indices;
    get_bound = bas_fcts->get_bound;
    get_v_loc = bas_fcts->get_real_vec;;

    if (!op_info) op_info = MEM_ALLOC(1, struct op_info);

    o_info2.row_fe_space = o_info2.col_fe_space = data->fe_space;

    o_info2.quad[2]      = quad;
    o_info2.LALt         = LALt;
    o_info2.LALt_pw_const = true;
    o_info2.LALt_symmetric = true;
    o_info2.user_data     = op_info;

    matrix_info = fill_matrix_info(&o_info2, nil);
    fill_a = matrix_info->el_matrix_fct;
    a_info = matrix_info->fill_info;

    o_info0.row_fe_space = o_info0.col_fe_space = data->fe_space;

    o_info0.quad[0]      = quad;
    o_info0.c            = c;
    o_info0.c_pw_const   = false;
    o_info0.user_data     = op_info;

    op_info->quad_fast = get_quad_fast(bas_fcts, quad, INIT_PHI);

    matrix_info = fill_matrix_info(&o_info0, nil);

```

```

    fill_c = matrix_info->el_matrix_fct;
    c_info = matrix_info->fill_info;
}

/*-----*/
/* make a DOF vector from input vector v_vec */
/*-----*/
dof_v.fe_space = data->fe_space;
dof_v.size     = dim;
dof_v.vec      = (REAL *) v;
/*-----*/
/* cast of v is needed; dof_v.vec isn't const REAL * */
/* nevertheless, values are NOT changed */
/*-----*/

/*-----*/
/* make a DOF vector from F, if not nil */
/*-----*/
if (F)
{
    dof_F.fe_space = data->fe_space;
    dof_F.size     = dim;
    dof_F.vec      = F;
}

/*-----*/
/* and now assemble DF(v) and/or F(v) */
/*-----*/
op_info->k      = data->k;
op_info->sigma   = data->sigma;

if (up_DF)
{
/*--- if v_vec[i] >= 0 for all i => matrix is pos.definite (p=1) ---*/
    data->solver = dof_min(&dof_v) >= 0 ? CG : BiCGStab;
    clear_dof_matrix(data->DF);
}

if (F)
{
    dof_set(0.0, &dof_F);
    L2scp_fct_bas(data->f, op_info->quad_fast->quad, &dof_F);
    dof_scal(-1.0, &dof_F);
}

fill_flag = CALL_LEAF_EL|FILL_COORDS|FILL_BOUND;
el_info = traverse_first(stack, data->fe_space->mesh, -1,
                        fill_flag);
while (el_info)

```

```

{
    const REAL    *v_loc = (*get_v_loc)(el_info->el, &dof_v, nil);
    const DOF     *dof   = (*get_dof)(el_info->el, admin, nil);
    const S_CHAR  *bound = (*get_bound)(el_info, nil);
    const REAL    **a_mat, **c_mat;

    /*-----*/
    /* initialization of values used by LALt and c */
    /*-----*/
    op_info->det   = el_grd_lambda(el_info, op_info->Lambda);
    op_info->v_qp   = uh_at_qp(op_info->quad_fast, v_loc, nil);

    a_mat = fill_a(el_info, a_info);
    c_mat = fill_c(el_info, c_info);

    if (up_DF) /*--- add element contribution to matrix DF(v) ---*/
    {
        /*-----*/
        /* add a(phi_i,phi_j) + 4*m(u^3*phi_i,phi_j) to matrix */
        /*-----*/
        add_element_matrix(data->DF, 1.0, n_phi, n_phi, dof, dof,
                           a_mat, bound);
        add_element_matrix(data->DF, 4.0, n_phi, n_phi, dof, dof,
                           c_mat, bound);
    }

    if (F) /*--- add element contribution to F(v) -----*/
    {
        int i, j;
        /*-----*/
        /* F(v) += a(v, phi_i) + m(v^4, phi_i) */
        /*-----*/
        for (i = 0; i < n_phi; i++)
        {
            if (bound[i] < DIRICHLET)
            {
                REAL val = 0.0;
                for (j = 0; j < n_phi; j++)
                    val += (a_mat[i][j] + c_mat[i][j])*v_loc[j];
                F[dof[i]] += val;
            }
            else
                F[dof[i]] = 0.0; /*- zero Dirichlet boundary values! */
        }
    }

    el_info = traverse_next(stack, el_info);
}

```

```

    free_traverse_stack(stack);

    return;
}

```

The linear sub-solver

For the solution of the linearized problem we use the `oem_solve_s()` function, which is also used in the solver for the linear Poisson equation (compare Section 2.1.8). Similar to the `update()` function, we have to convert the right hand side vector `F` and the solution vector `d` to DOF vectors. Information about the system matrix and parameters for the solver are passed by `ud`. The member `data->solver` is initialized in `update()`.

```

static int solve(void *ud, int dim, const REAL *F, REAL *d)
{
    NEWTON_DATA    *data = ud;
    int            iter;
    DOF_REAL_VEC   dof_F = {nil, nil, "F"};
    DOF_REAL_VEC   dof_d = {nil, nil, "d"};

    /*-----*/
    /* make DOF vectors from F and d                                */
    /*-----*/
    dof_F.fe_space = dof_d.fe_space = data->fe_space;
    dof_F.size     = dof_d.size     = dim;
    dof_F.vec       = (REAL *) F; /* cast needed ... */
    dof_d.vec       = d;

    iter = oem_solve_s(data->DF, &dof_F, &dof_d, data->solver,
                       data->tolerance, data->icon, data->restart,
                       data->max_iter, data->info);

    return(iter);
}

```

The computation of the H^1 semi norm

The H^1 semi norm can easily be calculated by converting the input vector `v` into a DOF vector and then calling the ALBERTA routine `H1_norm_uh()` (compare Section 3.10).

```

static REAL norm(void *ud, int dim, const REAL *v)
{
    NEWTON_DATA    *data = ud;
    DOF_REAL_VEC   dof_v = {nil, nil, "v"};

```

```

dof_v.fe_space = data->fe_space;
dof_v.size     = dim;
dof_v.vec      = (REAL *) v; /* cast needed ... */

return(H1_norm_uh(nil, &dof_v));
}

```

The nonlinear solver

The function `nlsolve()` initializes the structure `NEWTON_DATA` with problem dependent information. Here, we have to allocate a DOF matrix for storing the system matrix (only on the first call), and initialize parameters for the linear sub-solver and problem dependent data (like heat conductivity k , etc.)

The structure `NLS_DATA` is filled with information for the general Newton solver (the problem dependent routines `update()`, `solve()`, and `norm()` described above). All these routines use the same structure `NEWTON_DATA` for problem dependent information. For getting access to the definition of `NLS_DATA` and prototypes for the Newton solvers, we have to include the `nls.h` header file.

The dimension of the discrete equation is

```
dim = u0->fe_space->admin->size_used;
```

where `u0` is a pointer to a DOF vector storing the initial guess. Note, that after the call to `dof_compress()` in the `build()` routine, `dim` holds the true dimension of the discrete equation. Without a `dof_compress()` there may be holes in DOF vectors, and `u0->fe_space->admin->size_used` bigger than the last *used* index, and again `dim` is the dimension of the discrete equation for the Newton solver. The `ALBERTA` routines do not operate on unused indices, whereas the Newton solvers do operate on unused indices too, because they do not know about used and unused indices. In this situation, all unused DOFs would have to be cleared for the initial solution `u0` by

```
FOR_ALL_FREE_DOFs(u0->fe_space->admin, u0->vec[dof] = 0.0);
```

The same applies to the vector storing the right hand side in `update()`. The `dof_set()` function only initializes used indices.

Finally, we reallocate the workspace used by the Newton solvers (compare Section 3.15.6) and start the Newton method.

```

#include <nls.h>

int nlsolve(DOF_REAL_VEC *u0, REAL k, REAL sigma,
            REAL (*f)(const REAL_D))
{
    FUNCNAME("nlsolve");
    static NEWTON_DATA data = {nil,0,0,nil,nil,CG,0,500,2,10,1};
    static NLS_DATA    nls_data = {nil};

```

```

int                iter, dim = u0->fe_space->admin->size_used;

if (!data.fe_space)
{
/*-----*/
/*--  init parameters for newton  -----*/
/*-----*/
    nls_data.update      = update;
    nls_data.update_data = &data;
    nls_data.solve       = solve;
    nls_data.solve_data  = &data;
    nls_data.norm        = norm;
    nls_data.norm_data   = &data;

    nls_data.tolerance = 1.e-4;
    GET_PARAMETER(1, "newton tolerance", "%e", &nls_data.tolerance);
    nls_data.max_iter = 50;
    GET_PARAMETER(1, "newton max. iter", "%d", &nls_data.max_iter);
    nls_data.info = 8;
    GET_PARAMETER(1, "newton info", "%d", &nls_data.info);
    nls_data.restart = 0;
    GET_PARAMETER(1, "newton restart", "%d", &nls_data.restart);

/*-----*/
/*--  init data for update and solve  -----*/
/*-----*/

    data.fe_space = u0->fe_space;
    data.DF       = get_dof_matrix("DF(v)", u0->fe_space);

    data.tolerance = 1.e-2*nls_data.tolerance;
    GET_PARAMETER(1, "linear solver tolerance", "%f",
        &data.tolerance);
    GET_PARAMETER(1, "linear solver max iteration", "%d",
        &data.max_iter);
    GET_PARAMETER(1, "linear solver info", "%d", &data.info);
    GET_PARAMETER(1, "linear solver precon", "%d", &data.icon);
    GET_PARAMETER(1, "linear solver restart", "%d", &data.restart);
}
TEST_EXIT(data.fe_space == u0->fe_space)
    ("can't change f.e. spaces\n");

/*-----*/
/*--  init problem dependent parameters  -----*/
/*-----*/
    data.k      = k;
    data.sigma = sigma;
    data.f      = f;

```

```

/*-----*/
/*--  enlarge workspace used by newton(_fs), and solve by Newton --*/
/*-----*/
    if (nls_data.restart)
    {
        nls_data.ws = REALLOC_WORKSPACE(nls_data.ws, 4*dim*sizeof(REAL));
        iter = nls_newton_fs(&nls_data, dim, u0->vec);
    }
    else
    {
        nls_data.ws = REALLOC_WORKSPACE(nls_data.ws, 2*dim*sizeof(REAL));
        iter = nls_newton(&nls_data, dim, u0->vec);
    }

    return(iter);
}

```

2.3 Heat equation

In this section we describe a model implementation for the (linear) heat equation

$$\begin{aligned}
 \partial_t u - \Delta u &= f && \text{in } \Omega \subset \mathbb{R}^d \times (0, T), \\
 u &= g && \text{on } \partial\Omega \times (0, T), \\
 u &= u_0 && \text{on } \Omega \times \{0\}.
 \end{aligned}$$

We describe here only differences to the implementation of the linear Poisson problem. For common (or similar) routines we refer to Section 2.1.

2.3.1 Global variables

Additionally to the finite element space `fe_space`, the matrix `matrix` and the vectors `u_h` and `f_h`, we need a vector for storage of the solution U_n from the last time step. This one is implemented as a global variable, too. All these global variables are initialized in `main()`.

```
static DOF_REAL_VEC  *u_old = nil;
```

A global pointer to the `ADAPT_INSTAT` structure is used for access in the `build()` and `estimate()` routines, see below.

```
static ADAPT_INSTAT  *adapt_instat = nil;
```

Finally, a global variable `theta` is used for storing the parameter θ and `err_L2` for storing the actual L^2 error between true and discrete solution in the actual time step.

```
static REAL theta = 0.5;  /*- parameter of time discretization --*/
static REAL err_L2  = 0.0; /*- spatial error in single time step --*/

```

2.3.2 The main program for the heat equation

The main program initializes all program parameters from file and command line (compare Section 2.2.9), generates a mesh and finite element space, the DOF matrix and vectors, and allocates and fill the parameter structure `ADAPT_INSTAT` for the adaptive method for time dependent problems. This structure is accessed by `get_adapt_instat()` which already initializes besides the function pointers all members of this structure from the program parameters, compare Sections 3.13.4 and 2.1.2.

The (initial) time step size, read from the parameter file, is reduced when an initial global mesh refinement is performed. This reduction is automatically adapted to the order of time discretization (2nd order when $\theta = 0.5$, 1st order otherwise) and space discretization. For stability reasons, the time step size is scaled by a factor 10^{-3} if $\theta < 0.5$, see also Section 2.3.6.

Finally, the function pointers for the `adapt_instat()` structure are adjusted to the problem dependent routines for the heat equation and the complete numerical simulation is performed by a call to `adapt_method_instat()`.

```
int main(int argc, char **argv)
{
    FUNCNAME("main");
    MESH    *mesh;
    int      n_refine = 0, k, p = 1;
    char     filename[128];
    REAL     fac = 1.0;

    /*-----*/
    /*  first of all, init parameters of the init file          */
    /*-----*/

    init_parameters(0, "INIT/heat.dat");
    for (k = 1; k+1 < argc; k += 2)
        ADD_PARAMETER(0, argv[k], argv[k+1]);

    /*-----*/
    /*  get a mesh, and read the macro triangulation from file */
    /*-----*/

    GET_PARAMETER(1, "macro file name", "%s", filename);
    GET_PARAMETER(1, "global refinements", "%d", &n_refine);

    mesh = GET_MESH("ALBERTA mesh", init_dof_admin, init_leaf_data);
    read_macro(mesh, filename, nil);
    global_refine(mesh, n_refine*DIM);
    graphics(mesh, nil, nil);

    GET_PARAMETER(1, "parameter theta", "%e", &theta);
    if (theta < 0.5)
```



```

{
    WARNING("You are using the explicit Euler scheme\n");
    WARNING("Use a sufficiently small time step size!!!\n");
    fac = 1.0e-3;
}

matrix = get_dof_matrix("A", fe_space);
f_h     = get_dof_real_vec("f_h", fe_space);
u_h     = get_dof_real_vec("u_h", fe_space);
u_h->refine_interpol = fe_space->bas_fcts->real_refine_inter;
u_h->coarse_restrict = fe_space->bas_fcts->real_coarse_inter;
u_old   = get_dof_real_vec("u_old", fe_space);
u_old->refine_interpol = fe_space->bas_fcts->real_refine_inter;
u_old->coarse_restrict = fe_space->bas_fcts->real_coarse_inter;
dof_set(0.0, u_h);      /* initialize u_h ! */

/*-----*/
/*  init adapt structure and start adaptive method */
/*-----*/

adapt_instat = get_adapt_instat("heat", "adapt", 2, adapt_instat);

/*-----*/
/*  adapt time step size to refinement level and polynomial degree*/
/*-----*/

GET_PARAMETER(1, "polynomial degree", "%d", &p);

if (theta == 0.5)
    adapt_instat->timestep *= fac*pow(2, -(REAL)(p*(n_refine))/2.0);
else
    adapt_instat->timestep *= fac*pow(2, -(REAL)(p*(n_refine)));
MSG("using initial timestep size = %.4le\n",
    adapt_instat->timestep);

eval_time_u0 = adapt_instat->start_time;

adapt_instat->adapt_initial->get_el_est = get_el_est;
adapt_instat->adapt_initial->estimate = est_initial;
adapt_instat->adapt_initial->solve = interpol_u0;

adapt_instat->adapt_space->get_el_est  = get_el_est;
adapt_instat->adapt_space->get_el_estc = get_el_estc;
adapt_instat->adapt_space->estimate = estimate;
adapt_instat->adapt_space->build_after_coarsen = build;
adapt_instat->adapt_space->solve = solve;

adapt_instat->init_timestep = init_timestep;
adapt_instat->set_time      = set_time;

```

```

adapt_instat->get_time_est    = get_time_est;
adapt_instat->close_timestep = close_timestep;

adapt_method_instat(mesh, adapt_instat);

WAIT REALLY;
return(0);
}

```

2.3.3 The parameter file for the heat equation

The parameter file for the heat equation `INIT/heat.dat` (here for the 2d simulations) is similar to the parameter file for the Poisson problem. The main differences are additional parameters for the adaptive procedure, see Section 3.13.3. These additional parameters may also be optimized for 1d, 2d, and 3d.

Via the parameter `write finite element data` storage of meshes and finite element solution for post-processing purposes can be done. The parameter `write statistical data` selects storage of files containing number of DOFs, estimate, error, etc. versus time. Finally, `data path` can prescribe an existing path for storing such data.

```

macro file name:      Macro/macro.amc
global refinements:   0
polynomial degree:    1

% graphic windows: solution, estimate, and mesh if size > 0
graphic windows:      300 300 300
% for graphics you can specify the range for the values of
% discrete solution for displaying: min max
% automatical scaling by display routine if min >= max
graphic range:        -1.0 1.0

solver:                2 % 1: BICGSTAB 2: CG 3: GMRES 4: ODIR 5: ORES
solver max iteration:  1000
solver restart:        10 % only used for GMRES
solver tolerance:      1.e-12
solver info:           2
solver precon:         1 % 0: no precon 1: diag precon
                        % 2: HB precon 3: BPX precon

parameter theta:       1.0
adapt->start_time:      0.0
adapt->end_time:        5.0

adapt->tolerance:        1.0e-4
adapt->timestep:         1.0e-2

```

```

adapt->rel_initial_error:      0.5
adapt->rel_space_error:        0.5
adapt->rel_time_error:         0.5
adapt->strategy:                1 % 0=explicit, 1=implicit
adapt->max_iteration:          10
adapt->info:                    2

adapt->initial->strategy:       2 % 0=none, 1=GR, 2=MS, 3=ES, 4=GERS
adapt->initial->MS_gamma:       0.5
adapt->initial->max_iteration:  10
adapt->initial->info:           2

adapt->space->strategy:         3 % 0=none, 1=GR, 2=MS, 3=ES, 4=GERS
adapt->space->ES_theta:         0.9
adapt->space->ES_theta_c:       0.05
adapt->space->max_iteration:    10
adapt->space->coarsen_allowed:  1 % 0|1
adapt->space->info:             2

estimator C0:                  0.1
estimator C1:                  0.1
estimator C2:                  0.1
estimator C3:                  0.1

write finite element data:      0 % write data for post-processing ?
write statistical data:         0 % write statistical data or not
data path:                     ./data % path for data to be written

WAIT:                           0

```




Fig. 2.4. Time step size (left) and number of DOFs for different polynomial degrees (right) over time in 2d.

Figs. 2.4 and 2.5 show the variation of time step sizes and number of DOFs over time, automatically generated by the adaptive method in two and three space dimensions for a problem with time-periodic data. The number of DOFs is depicted for different spatial discretization order and shows the strong benefit from using a higher order method. The size of time steps was




Fig. 2.5. Time step size (left) and number of DOFs for different polynomial degrees (right) over time in 3d.

nearly the same for all spatial discretizations. Parameters for the adaptive procedure can be taken from the corresponding parameter files in 2d and 3d in the distribution.

2.3.4 Functions for leaf data

For time dependent problems, mesh adaption usually also includes coarsening of previously (for smaller t) refined parts of the mesh. For storage of local coarsening error estimates, the leaf data structure is enlarged by a second `REAL`. Functions `rw_el_estc()` and `get_el_estc()` are provided for access to that storage location.

```

struct heat_leaf_data
{
    REAL estimate;          /* one real for the element indicator */
    REAL est_c;            /* one real for the coarsening indicator */
};

static REAL *rw_el_estc(EL *el)
{
    if (IS_LEAF_EL(el))
        return(&((struct heat_leaf_data *)LEAF_DATA(el))->est_c);
    else
        return(nil);
}

static REAL get_el_estc(EL *el)
{
    if (IS_LEAF_EL(el))
        return(((struct heat_leaf_data *)LEAF_DATA(el))->est_c);
    else
        return(0.0);
}

```

2.3.5 Data of the differential equation

Data for the heat equation are the initial values u_0 , right hand side f , and boundary values g . When the true solution u is known, it can be used for computing the true error between discrete and exact solution.

The sample problem is defined such that the exact solution is

$$u(x, t) = \sin(\pi t) e^{-10|x|^2} \quad \text{on } (0, 1)^d \times [0, 1].$$

All library subroutines which evaluate a given data function (for integration, e.g.) are defined for space dependent functions only and do not know about a time variable. Thus, such a ‘simple’ space dependent function $f_{\text{space}}(x)$ has to be derived from a space–time dependent function $f(x, t)$. We do this by keeping the time in a global variable, and setting

$$f_{\text{space}}(x) := f(x, t).$$

```
static REAL eval_time_u = 0.0;
static REAL u(const REAL_D x)
{
    return(sin(M_PI*eval_time_u)*exp(-10.0*SCP_DOW(x,x)));
}

static REAL eval_time_u0 = 0.0;
static REAL u0(const REAL_D x)
{
    eval_time_u = eval_time_u0;
    return(u(x));
}

static REAL eval_time_g = 0.0;
static REAL g(const REAL_D x)      /* boundary values, not optional */
{
    eval_time_u = eval_time_g;
    return(u(x));
}

static REAL eval_time_f = 0.0;
static REAL f(const REAL_D x)      /* u_t - Delta u, not optional */
{
    REAL r2 = SCP_DOW(x,x), ux = sin(M_PI*eval_time_f)*exp(-10.0*r2);
    REAL ut = M_PI*cos(M_PI*eval_time_f)*exp(-10.0*r2);
    return(ut - (400.0*r2 - 20.0*DIM)*ux);
}
```

As indicated, the times for evaluation of boundary data and right hand side may be chosen independent of each other depending on the kind of time

discretization. The value of `eval_time_f` and `eval_time_g` are set by the function `set_time()`. Similarly, the evaluation time for the exact solution is set by `estimate()` where also the evaluation time of f is set for the evaluation of the element residual. In order to start the simulation not only at $t = 0$, we have introduced a variable `eval_time_u0`, which is set in `main()` at the beginning of the program to the value of `adapt_instat->start_time`.

2.3.6 Time discretization

The model implementation uses a variable time discretization scheme. Initial data is interpolated on the initial mesh,

$$U_0 = I_0 u_0.$$

For $\theta \in [0, 1]$, the finite element solution $U_{n+1} \approx u(\cdot, t_{n+1})$ is given by $U_{n+1} \in I_{n+1}g(\cdot, t_{n+1}) + \dot{X}_{n+1}$ such that

$$\begin{aligned} \frac{1}{\tau_{n+1}}(U_{n+1}, \Phi) + \theta(\nabla U_{n+1}, \nabla \Phi) &= \frac{1}{\tau_{n+1}}(I_{n+1}U_n, \Phi) \\ &\quad - (1 - \theta)(\nabla I_{n+1}U_n, \nabla \Phi) + (f(\cdot, t_n + \theta\tau_{n+1}), \Phi) \end{aligned} \quad (2.5)$$

for all $\Phi \in \dot{X}_{n+1}$. Using $\theta = 0$, this gives the forward (explicit) Euler scheme, for $\theta = 1$ the backward (implicit) Euler scheme. For $\theta = 0.5$, we obtain the Crank–Nicholson scheme, which is of second order in time. For $\theta \in [0.5, 1.0]$, the scheme is unconditionally stable, while for $\theta < 0.5$ stability is only guaranteed if the time step size is small enough. For that reason, the time step size is scaled by an additional factor of 10^{-3} in the main program if $\theta < 0.5$. But this might not be enough for guaranteeing stability of the scheme! We do recommend to use $\theta = 0.5, 1$ only.

2.3.7 Initial data interpolation

Initial data u_0 is just interpolated on the initial mesh, thus the `solve()` entry in `adapt_instat->adapt_initial` will point to a routine `interpol_u0()` which implements this by the library interpolation routine. No `build()` routine is needed by the initial mesh adaption procedure.

```
static void interpol_u0(MESH *mesh)
{
    dof_compress(mesh);
    interpol(u0, u_h);

    return;
}
```

2.3.8 The assemblage of the discrete system

Using a matrix notation, the discrete problem (2.5) can be written as

$$\left(\frac{1}{\tau_{n+1}}\mathbf{M} + \theta\mathbf{A}\right)\mathbf{U}_{n+1} = \left(\frac{1}{\tau_{n+1}}\mathbf{M} - (1 - \theta)\mathbf{A}\right)\mathbf{U}_n + \mathbf{F}_{n+1}.$$

Here, $\mathbf{M} = (\Phi_i, \Phi_j)$ denotes the mass matrix and $\mathbf{A} = (\nabla\Phi_i, \nabla\Phi_j)$ the stiffness matrix (up to Dirichlet boundary DOFs). The system matrix on the left hand side is not the same as the one applied to the old solution on the right hand side. But we want to compute the contribution of the solution from the old time step \mathbf{U}_n to the right hand side vector efficiently by a simple matrix–vector multiplication and thus avoiding additional element-wise integration. For doing this without storing both matrices \mathbf{M} and \mathbf{A} we are using the element-wise strategy explained and used in Section 2.2.6 when assembling the linearized equation in the Newton iteration for solving the nonlinear reaction–diffusion equation.

The subroutine `assemble()` generates both the system matrix and the right hand side at the same time. The mesh elements are visited via the non-recursive mesh traversal routines. On every leaf element, both the element mass matrix `c_mat` and the element stiffness matrix `a_mat` are calculated using the `el_matrix_fct()` provided by `fill_matrix_info()`. For this purpose, two different operators (the mass and stiffness operators) are defined and applied on each element. The stiffness operator uses the same `LALt()` function for the second order term as described in Section 2.1.7; the mass operator implements only the constant zero order coefficient $c = 1/\tau_{n+1}$, which is passed in `struct op_info` and evaluated in the function `c()`. The initialization and access of these operators is done in the same way as in Section 2.2.6 where this is described in detail. During the non-recursive mesh traversal, element stiffness matrix and the mass matrix are computed and added to the global system matrix. Then, the contribution to the right hand side vector of the solution from the old time step is computed by a matrix–vector product of these element matrices with the local coefficient vector on the element of \mathbf{U}_n and added to the global load vector.

After this step, the the right hand side f and Dirichlet boundary values g are treated by the standard routines.

```
struct op_info
{
    REAL_D  Lambda[DIM+1];
    REAL    det;

    REAL    tau_1;
};
```

```

static REAL c(const EL_INFO *el_info, const QUAD *quad, int iq,
              void *ud)
{
    struct op_info *info = ud;

    return(info->tau_1*info->det);
}

static void assemble(DOF_REAL_VEC *u_old, DOF_MATRIX *matrix,
                    DOF_REAL_VEC *fh, DOF_REAL_VEC *u_h, REAL theta,
                    REAL tau, REAL (*f)(const REAL_D),
                    REAL (*g)(const REAL_D))
{
    FUNCNAME("assemble");
    static struct op_info *op_info = nil;
    static const REAL      **(*fill_a)(const EL_INFO *, void *) = nil;
    static void            *a_info = nil;
    static const REAL      **(*fill_c)(const EL_INFO *, void *) = nil;
    static void            *c_info = nil;

    static const DOF_ADMIN *admin = nil;
    static int             n;
    static const REAL      *(*get_u_loc)(const EL *, const DOF_REAL_VEC *,
                                         REAL *);
    static const S_CHAR *(*get_bound)(const EL_INFO *, S_CHAR *);
    static const DOF *(*get_dof)(const EL *, const DOF_ADMIN *, DOF *);

    TRAVERSE_STACK *stack = get_traverse_stack();
    const EL_INFO *el_info;
    FLAGS          fill_flag;
    const REAL     **a_mat, **c_mat;
    REAL           *f_vec;
    const QUAD     *quad;
    int             i, j;

    quad = get_quadrature(DIM, 2*u_h->fe_space->bas_fcts->degree);

    /*-----*/
    /*  init functions for matrix assembling          */
    /*-----*/

    if (admin != u_h->fe_space->admin)
    {
        OPERATOR_INFO      o_info2 = {nil}, o_info0 = {nil};
        const EL_MATRIX_INFO *matrix_info;
        const BAS_FCTS      *bas_fcts = u_h->fe_space->bas_fcts;

        admin                = u_h->fe_space->admin;
    }

```



```

n = bas_fcts->n_bas_fcts;
get_dof      = bas_fcts->get_dof_indices;
get_bound    = bas_fcts->get_bound;
get_u_loc    = bas_fcts->get_real_vec;

if (!op_info)
    op_info = MEM_ALLOC(1, struct op_info);

o_info2.row_fe_space = o_info2.col_fe_space = u_h->fe_space;

o_info2.quad[2]      = quad;
o_info2.LALt        = LALt;
o_info2.LALt_pw_const = true;
o_info2.LALt_symmetric = true;
o_info2.user_data    = op_info;

matrix_info = fill_matrix_info(&o_info2, nil);
fill_a = matrix_info->el_matrix_fct;
a_info = matrix_info->fill_info;

o_info0.row_fe_space = o_info0.col_fe_space = u_h->fe_space;

o_info0.quad[0]      = quad;
o_info0.c            = c;
o_info0.c_pw_const   = true;
o_info0.user_data    = op_info;

matrix_info = fill_matrix_info(&o_info0, nil);
fill_c = matrix_info->el_matrix_fct;
c_info = matrix_info->fill_info;
}

op_info->tau_1 = 1.0/tau;

/*-----*/
/* and now assemble the matrix and right hand side */
/*-----*/

clear_dof_matrix(matrix);
dof_set(0.0, fh);
f_vec = fh->vec;

fill_flag = CALL_LEAF_EL|FILL_COORDS|FILL_BOUND;
el_info = traverse_first(stack, u_h->fe_space->mesh,-1,fill_flag);
while (el_info)
{
    const REAL *u_old_loc = (*get_u_loc)(el_info->el, u_old, nil);
    const DOF *dof         = (*get_dof)(el_info->el, admin, nil);
    const S_CHAR *bound    = (*get_bound)(el_info, nil);

```

```

/*-----*/
/* initialization of values used by LALt and c */
/*-----*/
    op_info->det    = el_grd_lambda(el_info, op_info->Lambda);

    a_mat = fill_a(el_info, a_info);
    c_mat = fill_c(el_info, c_info);

/*-----*/
/* add theta*a(psi_i,psi_j) + 1/tau*m(4*u^3*psi_i,psi_j) */
/*-----*/

    if (theta)
    {
        add_element_matrix(matrix, theta, n, n, dof, dof, a_mat, bound);
    }
    add_element_matrix(matrix, 1.0, n, n, dof, dof, c_mat, bound);

/*-----*/
/* f += -(1-theta)*a(u_old,psi_i) + 1/tau*m(u_old,psi_i) */
/*-----*/

    if (1.0 - theta)
    {
        REAL theta1 = 1.0 - theta;
        for (i = 0; i < n; i++)
        {
            if (bound[i] < DIRICHLET)
            {
                REAL val = 0.0;
                for (j = 0; j < n; j++)
                {
                    val += (-theta1*a_mat[i][j] + c_mat[i][j])*u_old_loc[j];
                    f_vec[dof[i]] += val;
                }
            }
        }
    }
    else
    {
        for (i = 0; i < n; i++)
        {
            if (bound[i] < DIRICHLET)
            {
                REAL val = 0.0;
                for (j = 0; j < n; j++)
                {
                    val += c_mat[i][j]*u_old_loc[j];
                    f_vec[dof[i]] += val;
                }
            }
        }
    }
}

```

```

    }
    el_info = traverse_next(stack, el_info);
}

free_traverse_stack(stack);

L2scp_fct_bas(f, quad, fh);
dirichlet_bound(g, fh, u_h, nil);

return;
}

```

The `build()` routine for one time step of the heat equation is nearly a dummy routine and just calls the `assemble()` routine described above. In order to avoid holes in vectors and matrices, as a first step, the mesh is compressed. This guarantees optimal performance of the BLAS1 routines used in the iterative solvers.

```

static void build(MESH *mesh, U_CHAR flag)
{
    FUNCNAME("build");

    dof_compress(mesh);

    INFO(adapt_instat->adapt_space->info, 2)
        ("%d DOFs for %s\n", fe_space->admin->size_used, fe_space->name);

    assemble(u_old, matrix, f_h, u_h, theta, adapt_instat->timestep,
            f, g);
    return;
}

```

The resulting linear system is solved by calling the `oem_solve_s()` library routine. This is done via the `solve()` subroutine described in Section 2.1.8.

2.3.9 Error estimation

The initial error $\|U_0 - u_0\|_{L^2(\Omega)}$ is calculated exactly (up to quadrature error) by a call to `L2_err()`. Local error contributions are written via `rw_el_est()` to the `estimate` value in `struct heat_leaf_data`. The `err_max` and `err_sum` of the `ADAPT_STAT` structure (which will be `adapt_instat->adapt_initial`, see below) are set accordingly.

```

static REAL est_initial(MESH *mesh, ADAPT_STAT *adapt)
{
    err_L2 = L2_err(u0, u_h, nil, 0, rw_el_est, &adapt->err_max);
    return(adapt->err_sum = err_L2);
}

```

In each time step, error estimation is done by the library routine `heat_est()`, which generates both time and space discretization indicators, compare Section 2.3.9. Similar to the estimator for elliptic problems, a function `r()` is needed for computing contributions of lower order terms and the right hand side. The flag for passing information about the discrete solution U_{n+1} or its gradient to `r()` is set to zero in `estimate()` since no lower order term is involved.

Local element indicators are stored to the `estimate` or `est_c` entries inside the data structure `struct heat_leaf_data` via the function pointers `rw_el_est()` and `rw_el_estc()`. The `err_max` and `err_sum` entries of `adapt->adapt_space` are set accordingly. The temporal error indicator is the return value by `heat_est()` and is stored in a global variable for later access by `get_time_est()`. In this example, the true solution is known and thus the true error $\|u(\cdot, t_{n+1}) - U_{n+1}\|_{L^2(\Omega)}$ is calculated additionally for comparison in the function `close_timestep()`.

```
static REAL r(const EL_INFO *el_info, const QUAD *quad, int iq,
              REAL t, REAL uh_iq, const REAL_D grd_uh_iq)
{
    REAL_D      x;
    coord_to_world(el_info, quad->lambda[iq], x);
    eval_time_f = t;
    return(-f(x));
}

static REAL estimate(MESH *mesh, ADAPT_STAT *adapt)
{
    FUNCNAME("estimate");
    static int      degree;
    static REAL     C[4] = {-1.0, 1.0, 1.0, 1.0};
    REAL_DD        A = {{0.0}};
    FLAGS          r_flag = 0;
    int             n;
    REAL            space_est;

    for (n = 0; n < DIM_OF_WORLD; n++)
        A[n][n] = 1.0;

    eval_time_u = adapt_instat->time;

    if (C[0] < 0)
    {
        C[0] = 1.0;
        GET_PARAMETER(1, "estimator C0", "%f", &C[0]);
        GET_PARAMETER(1, "estimator C1", "%f", &C[1]);
        GET_PARAMETER(1, "estimator C2", "%f", &C[2]);
        GET_PARAMETER(1, "estimator C3", "%f", &C[3]);
    }
}
```

```

degree = 2*u_h->fe_space->bas_fcts->degree;
time_est = heat_est(u_h, adapt_instat, rw_el_est, rw_el_estc,
                    degree, C, u_old, (const REAL_D *) A, r, r_flag);

space_est = adapt_instat->adapt_space->err_sum;
err_L2 = L2_err(u, u_h, nil, 0, nil, nil);

INFO(adapt_instat->info,2)
  ("---8<-----\n");
INFO(adapt_instat->info, 2)("time = %.4le with timestep = %.4le\n",
                           adapt_instat->time, adapt_instat->timestep);
INFO(adapt_instat->info, 2)("estimate   = %.4le, max = %.4le\n",
                           space_est, sqrt(adapt_instat->adapt_space->err_max));
INFO(adapt_instat->info, 2)("||u-uh||L2 = %.4le, ratio = %.2lf\n",
                           err_L2, err_L2/MAX(space_est,1.e-20));

return(adapt_instat->adapt_space->err_sum);
}

```

2.3.10 Time steps

Time dependent problems are calculated step by step in single time steps. In a fully implicit time-adaptive strategy, each time step includes an adaptation of the time step size as well as an adaptation of the corresponding spatial discretization. First, the time step size is adapted and then the mesh adaptation procedure is performed. This second part may again push the estimate for the time discretization error over the corresponding tolerance. In this case, the time step size is again reduced and the whole procedure is iterated until both, time and space discretization error estimates meet the prescribed tolerances (or until a maximal number of iterations is performed). For details and other time-adaptive strategies see Section 3.13.3.

Besides the `build()`, `solve()`, and `estimate()` routines for the adaptation of the initial grid and the grids in each time steps, additional routines for initializing time steps, setting time step size of the actual time step, and finalizing a time step are needed. For adaptive control of the time step size, the function `get_time_est()` gives information about the size of the time discretization error. The actual time discretization error is stored in the global variable `time_est` and its value is set in the function `estimate()`.

During the initialization of a new time step in `init_timestep()`, the discrete solution `u_h` from the old time step (or from interpolation of initial data) is copied to `u_old`. In the function `set_time()` evaluation times for the right hand side f and Dirichlet boundary data g are set accordingly to the chosen time discretization scheme. Since a time step can be rejected by the adaptive method by a too large time discretization error, this function can be called

several times during the computation of a single time step. On each call, information about the actual time and time step size is accessed via the entries `time` and `timestep` of the `adapt_instat` structure.

After accepting the current time step size and current grid by the adaptive method, the time step is completed by `close_time_step()`. The variables `space_est`, `time_est`, and `err_L2` now hold the final estimates resp. error, and `u_h` the accepted finite element solution for this time step. The final mesh and discrete solution can now be written to file for post-processing purposes, depending on the parameter value of `write finite element data`. The file name for the mesh/solution consists of the data path, the base name `mesh/u_h`, and the iteration counter of the actual time step. Such a composition can be easily constructed by the function `generate_filename()`, described in Section 3.1.6. Mesh and finite element solution are then exported to file by the `write_*_xdr()` routines in a portable binary format. Using this procedure, the sequence of discrete solutions can easily be visualized by the program `alberta_movi` which is an interface to GRAPE and comes with the distribution of ALBERTA, compare Section 3.16.3.

Depending on the parameter value of `write statistical data`, the evolution of estimates, error, number of DOFs, size of time step size, etc. are written to files by the function `write_statistical_data()`, which is included in `heat.c` but not described here. It produces for each quantity a two-column data file where the first column contains time and the second column estimate, error, etc. Such data can easily be evaluated by standard (graphic) tools.

Finally, a graphical output of the solution and the mesh is generated via the `graphics()` routine already used in the previous examples.

```
static REAL time_est = 0.0;

static REAL get_time_est(MESH *mesh, ADAPT_INSTAT *adapt)
{
    return(time_est);
}

static void init_timestep(MESH *mesh, ADAPT_INSTAT *adapt)
{
    FUNCNAME("init_timestep");

    INFO(adapt_instat->info,1)
        ("---8<-----\n");
    INFO(adapt_instat->info, 1)("starting new timestep\n");

    dof_copy(u_h, u_old);
    return;
}
```

```

static void set_time(MESH *mesh, ADAPT_INSTAT *adapt)
{
    FUNCNAME("set_time");

    INFO(adapt->info,1)
        ("---8<-----\n");
    if (adapt->time == adapt->start_time)
    {
        INFO(adapt->info, 1)("start time: %.4le\n", adapt->time);
    }
    else
    {
        INFO(adapt->info, 1)("timestep for (%.4le %.4le), tau = %.4le\n",
            adapt->time-adapt->timestep, adapt->time,
            adapt->timestep);
    }

    eval_time_f = adapt->time - (1 - theta)*adapt->timestep;
    eval_time_g = adapt->time;

    return;
}

static void close_timestep(MESH *mesh, ADAPT_INSTAT *adapt)
{
    FUNCNAME("close_timestep");
    static REAL err_max = 0.0;          /* max space-time error */
    static REAL est_max = 0.0;          /* max space-time estimate */
    static int  write_fe_data = 0, write_stat_data = 0;
    static int  step = 0;
    static char path[256] = "./";

    REAL        space_est = adapt->adapt_space->err_sum;
    REAL        tolerance = adapt->rel_time_error*adapt->timestep;

    err_max = MAX(err_max, err_L2);
    est_max = MAX(est_max, space_est + time_est);

    INFO(adapt->info,1)
        ("---8<-----\n");

    if (adapt->time == adapt->start_time)
    {
        tolerance = adapt->adapt_initial->timestep;
        INFO(adapt->info,1)("start time: %.4le\n", adapt->time);
    }
    else
    {
        tolerance += adapt->adapt_space->timestep;
    }
}

```

```

        INFO(adapt->info,1)("timestep for (%.4le %.4le), tau = %.4le\n",
                           adapt->time-adapt->timestep, adapt->time,
                           adapt->timestep);
    }
    INFO(adapt->info,2)("max. est. = %.4le, tolerance = %.4le\n",
                      est_max, tolerance);
    INFO(adapt->info,2)("max. error = %.4le, ratio = %.2lf\n",
                      err_max, err_max/MAX(est_max,1.0e-20));

    if (!step)
    {
        GET_PARAMETER(1, "write finite element data", "%d",
                      &write_fe_data);
        GET_PARAMETER(1, "write statistical data", "%d",
                      &write_stat_data);
        GET_PARAMETER(1, "data path", "%s", path);
    }

    /*---8<-----*/
    /*--- write mesh and discrete solution to file for post-processing ---*/
    /*----->8---*/

    if (write_fe_data)
    {
        const char *fn;

        fn= generate_filename(path, "mesh", step);
        write_mesh_xdr(mesh, fn, adapt->time);
        fn= generate_filename(path, "u_h", step);
        write_dof_real_vec(u_h, fn);
    }

    step++;

    /*---8<-----*/
    /*--- write data about estimate, error, time step size, etc. ---*/
    /*----->8---*/

    if (write_stat_data)
    {
        int n_dof = fe_space->admin->size_used;
        write_statistics(path, adapt, n_dof, space_est, time_est, err_L2);
    }

    graphics(mesh, u_h, get_el_est);

    return;
}

```


2.4 Installation of ALBERTA and file organization

2.4.1 Installation

Enclosed on CD you will find a gzipped archive file `alberta-1.2.tar.gz` for a UNIX/Linux operating system. It includes all sources of ALBERTA, the model implementations of the examples described in Chapter 2, and tools for installation. System requirements for ALBERTA are a Unix/Linux environment with C and FORTRAN Compilers, OpenGL graphics and GNU make.

The file `alberta-1.2.tar.gz` has to be unpacked which directly creates the sub-directory `albert-1.2` in the current directory with all data of ALBERTA. Changing to this sub-directory, the installation procedure is fully explained in the `README`. For a platform independent installation the GNU configure tools are used, documented in the file `INSTALL`. Installation options `configure` script can be added on the command line and are described in the `README` file or printed with command `configure --help`.

When ALBERTA should use the visualization packages `gltools` or `GRAPE`, these have to be installed first, see the corresponding web sites

```
http://www.wias-berlin.de/software/gltools/
http://www.iam.uni-bonn.de/sfb256/grape/
```

for obtaining the software. Paths to their installation directories have to be arguments to the `configure` script.

2.4.2 File organization

Using the ALBERTA library, the dimension enters in an application only by the parameters `DIM` and `DIM_OF_WORLD`. Thus, the code is usually the same for 1d, 2d, and 3d simulations. Nevertheless, the object files do depend on the dimension, since `DIM` and `DIM_OF_WORLD` are symbolic constants (which define the length of vectors, e.g.). Hence, all object files have to be rebuilt, when changing the dimension. To make sure that this is done automatically we use the following file organization, which is also reflected in the structure of `DEMO/src` sub-directory with the implementation of the model problems. We use the four sub-directories

```
./1d      ./2d      ./3d      ./Common
```

for organizing files. The directory `Common` contains all source files and header files that do not (or only slightly) depend on the dimension. The directories `1d`, `2d`, `3d` contain dimension-dependent data, like macro triangulations files and parameter files. Finally, a dimension-dependent `Makefile` is automatically created in `DEMO/src/*d` during installation of ALBERTA. These `Makefiles` contain all information about ALBERTA header files and all libraries used. In the 1d, 2d, or 3d sub-directories, the programs of the model problems for the corresponding space dimension are then generated by executing `make ellipt`, `make nonlin`, and `make heat`.

Data structures and implementation

ALBERTA provides two header files `alberta_util.h` and `alberta.h`, which contain the definitions of all data structures, macros, and subroutine prototypes. The file `alberta_util.h` is included in the header file `alberta.h`.

3.1 Basic types, utilities, and parameter handling

The file `alberta_util.h` contains some type definitions and macro definitions for memory (de-) allocation and messages, which we describe briefly in this section. The following system header files are included in `alberta_util.h`

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>
#include <string.h>
#include <float.h>
```

3.1.1 Basic types

ALBERTA uses the following basis symbolic constants and macro definition:

```
#define true      1
#define false    0
#define nil      NULL
#define MAX(a, b) ((a) > (b) ? (a) : (b))
#define MIN(a, b) ((a) < (b) ? (a) : (b))
#define ABS(a)    ((a) >= 0 ? (a) : -(a))
#define SQR(a)    ((a)*(a))
```

In order to store information in a compact way, we define two bit fields `U_CHAR` and `S_CHAR`:

```
typedef unsigned char    U_CHAR;
typedef signed char      S_CHAR;
```

The mesh traversal routines need flags which are stored in the data type **FLAGS**:

```
typedef unsigned long    FLAGS;
```

By the data type **REAL** the user can specify to store floating point values in the type **float** or **double**. All pointers to variables or vectors of floating point values have to be defined as **REAL**!

```
typedef double           REAL;
```

The use of **float** is also possible, but it is not guaranteed to work and may lead to problems when using other libraries (like libraries for linear solver or graphics, e.g.).

3.1.2 Message macros

There are several macros to write messages and error messages. Especially for error messages the exact location of the error is of interest. Thus, error messages are preceded by the name of the source file and the line number at that this error was detected. Such information is produced by the C-preprocessor. Additionally, the name of the function is printed. Since there is no symbolic constant defined by the C-preprocessor holding the function name, in each function a string **funcName** containing the name of the function has to be defined. This is usually done by the macro **FUNCNAME**

```
#define FUNCNAME(nn)    const char *funcName = nn
```

as the first declaration:

Example 3.1 (FUNCNAME).

```
static void refine_element(EL *el)
{
    FUNCNAME("refine_element");

    ...
}
```

All message macros use this local variable **funcName** and it has to be defined in each function.

Usual output to **stdout** is done by the macro **MSG()** which has the same arguments as **printf()**:

```
MSG(const char *format, ...);
```

The format string should be ended with the newline character **'\n'**. **MSG()** usually precedes the message by the function's name. If the previous message

was produced by the same function, the function's name is omitted and the space of the name is filled with blanks.

If the format string of `MSG()` does not end with the newline character, and one wants to print more information to the same line, this can be done by `print_msg()` which again has the same arguments as `printf()`:

```
print_msg(const char *format, ...);
```

Example 3.2 (`MSG()`, `print_msg()`).

```
static void refine_element(EL *el)
{
    FUNCNAME("refine_element");

    ...

    MSG("refining element %d\n", INDEX(el));
    MSG("neighbours of element: ");
    for (i = 0; i < N_VERTICES-1; i++)
        print_msg("%d, ", INDEX(NEIGH(el)[i]));
    print_msg("%d\n", INDEX(NEIGH(el)[N_VERTICES-1]));
}
```

produces for instance output

```
refine_element:    refining element 10
                  neighbours of element: 0, 14, 42
```

A simpler way to print vectors of integer or real numbers is provided by the macros `PRINT_INT_VEC` and `PRINT_REAL_VEC`.

```
PRINT_INT_VEC(const char *s, const int *vec, int no);
PRINT_REAL_VEC(const char *s, const REAL *vec, int no);
```

Based on the `MSG()` and `print_msg()` mechanisms, a comma-separated list of the `no` vector elements is produced.

Example 3.3 (`PRINT_REAL_VEC()`).

```
{
    FUNCNAME("test_routine");
    REAL_D point;
    ...

    PRINT_REAL_VEC("point coordinates", point, DIM_OF_WORLD);
}
```

generates for the second unit vector in 3D the output

```
test_routine:      point coordinates = (0.00000, 1.00000, 0.00000)
```

Often it is useful to suppress messages or to give only information up to a suitable level. There are two ways for defining such a level of information. The first one is a local level, which is determined by some variable in a function; the other one is a global restriction for information. For this global restriction a global variable

```
int      msg_info = 10;
```

is defined with an default value of 10. Using one of the macros

```
#define INFO(info,noinfo)\
    if (!(msg_info&&(MIN(msg_info,(info))>=(noinfo))));else MSG
#define PRINT_INFO(info,noinfo)\
    if (!(msg_info&&(MIN(msg_info,(info))>=(noinfo))));else print_msg
```

only messages are produced by `INFO()` or `PRINT_INFO()` if `msg_info` is non zero and the value `MIN(msg_info, info)` is greater or equal `noinfo`, where `noinfo` denotes some local level of information. Thus after setting `msg_info = 0`, no further messages are produced. Changing the value of this variable via a parameter file is described below in Section 3.1.5.

Example 3.4 (`INFO()`, `PRINT_INFO()`).

```
static void refine_element(EL *el)
{
    FUNCNAME("refine_element");

    ...

    INFO(info,4)("refining element %d\n", INDEX(el));
    INFO(info,6)("neighbours of element: ");
    for (i = 0; i < N_VERTICES-1; i++)
        PRINT_INFO(info,6)("%d, ", INDEX(NEIGH(el)[i]));
    PRINT_INFO(info,6)("%d\n", INDEX(NEIGH(el)[N_VERTICES-1]));
}
```

will print the element index, if the value of the global variable `info` ≥ 4 and additionally the indices of neighbours if `info` ≥ 6 .

For error messages macros `ERROR` and `ERROR_EXIT` are defined. `ERROR` has the same functionality as the `MSG` macro but the output is piped to `stderr`. `ERROR_EXIT` exits the program after using the `ERROR` macro with return value 1:

```
ERROR(const char *format, ...);
ERROR_EXIT(const char *format, ...);
```

Furthermore, two macros for testing boolean values are available:

```
#define TEST(test)      if ((test)); else ERROR
#define TEST_EXIT(test) if ((test)); else ERROR_EXIT
```

If `test` is not true both macros will print an error message. `TEST` will continue the program afterwards, meanwhile `TEST_EXIT` will exit the program with return value 1.

Error messages can not be suppressed and the information variable `msg_info` does not influence these error functions.

Example 3.5 (`TEST()`, `TEST_EXIT()`).

```
static void refine_element(EL *el)
{
    FUNCNAME("refine_element");

    TEST_EXIT(el)("no element for refinement\n");
    ...
}
```

Finally, there exists a macro `WARNING` for writing warnings to the same stream as for messages. Each warning is preceded by the word `WARNING`. Warnings can not be suppressed by the information variable `msg_info`.

```
WARNING(const char *format, ...);
```

Sometimes it may be very useful to write messages to file, or write parts of messages to file. By the functions

```
void change_msg_out(FILE *fp);
void open_msg_file(const char *filename, const char *type);
```

the user can select a new stream or file for message output. Using the first routine, the user directly specifies the new stream `fp`. If this stream is non nil, all messages are flushed to this stream, otherwise ALBERTA will use the old stream furthermore. The function `open_msg_file()` acts like `change_msg_out(fopen(filename, type))`.

Similar functions are available for error messages and they act in the same manner on the output stream for error messages:

```
void change_error_out(FILE *fp);
void open_error_file(const char *filename, const char *type);
```

For setting breakpoints in the program two macros

```
WAIT
WAIT REALLY
```

are defined.

`WAIT`: this macro uses a global variable `msg_wait` and if the value of this variable is not zero the statement `WAIT;` will produce the message

```
wait for <enter> ...
```

and will continue after pressing the `enter` or `return` key. If the value of `msg_wait` is zero, no message is produced and the program continues. The value of `msg_wait` can be modified by the parameter tools (see Section 3.1.5).

WAIT REALLY: the statement **WAIT REALLY** will always produce the above message and will wait for pressing the **enter** or **return** key.

3.1.3 Memory allocation and deallocation

Various functions and macros for memory allocation and deallocation are implemented. The basic ones are

```
void *alberta_alloc(size_t, const char *, const char *, int);
void *alberta_realloc(void *, size_t, size_t, const char *,
                    const char *, int);
void *alberta_calloc(size_t, size_t, const char *, const char *,int);
void alberta_free(void *, size_t);
```

In the following **name** is a pointer to the string holding the function name of the calling function (defined by the **FUNCNAME** macro, e.g.), **file** a pointer to the string holding the name of the source file (generated by the **__FILE__** CPP macro) and **line** is the line number of the call (generated by the **__LINE__** CPP macro). All functions will exit the running program with an error message, if the size to be allocated is 0 or the memory allocation by the system functions fails.

alberta_alloc(size, name, file, line): returns a pointer to a block of memory of at least the number of bytes specified by **size**.

alberta_realloc(ptr, old, new, name, file, line): adjusts the size of the block of memory pointed to by the pointer **ptr** to the number of bytes specified by **new**, and returns a pointer to the block. The contents of the block remain unchanged up to the lesser of the **old** and **new**; if necessary, a new block is allocated, and data is copied to it; if the **ptr** is a **nil** pointer, the **alberta_realloc()** function allocates a new block of the requested size.

alberta_calloc(n_el, el_size, name, file, line): returns a pointer to a vector with the **n_el** number of elements, where each element is of the size **el_size**; the space is initialized to zeros.

alberta_free(ptr, size): frees the block of memory pointed to by the argument **ptr** for further allocation; **ptr** must have been previously allocated by either **alberta_alloc()**, **alberta_realloc()**, or **alberta_calloc()**.

A more comfortable way to use these functions, is the use of the following macros:

```
TYPE* MEM_ALLOC(size_t, TYPE);
TYPE* MEM_REALLOC(TYPE *, size_t, size_t, TYPE);
TYPE* MEM_CALLOC(size_t, TYPE);
TYPE* MEM_FREE(TYPE *, size_t, TYPE);
```

They supply the above described functions with the function name, file name and line number automatically. For an allocation by these macros, the number of elements and the data type have to be specified; the actual size in

bytes is computed automatically. Additionally, casting to the correct type is performed.

MEM_ALLOC(*n*, *TYPE*): returns a pointer to a vector of type *TYPE* with the *n* number of elements.

MEM_REALLOC(*ptr*, *size_old*, *size_new*, *TYPE*): reallocates the vector of type *TYPE* at pointer *ptr* with *size_old* elements for *size_new* elements; values of the vector are not changed for all elements up to the minimum of *size_old* and *size_new*; returns a pointer to the new vector.

MEM_CALLOC(*n*, *TYPE*): returns a pointer to a vector of type *TYPE* with the *n* number of elements; the elements are initialized to zeros.

MEM_FREE(*ptr*, *n*, *TYPE*): frees a vector *ptr* of type *TYPE* with *n* number of elements which was previously allocated by either **MEM_ALLOC**(), **MEM_REALLOC**(), or **MEM_CALLOC**().

Example 3.6 (**MEM_ALLOC**(), **MEM_FREE**()).

```
REAL *u = MEM_ALLOC(10, REAL);
...
MEM_FREE(u, 10, REAL);
```

allocates a vector of 10 REALs and frees this vector again.

For some applications matrices are needed too. Matrices can be allocated and freed by the functions

```
void **alberta_matrix(size_t, size_t, size_t, const char *,
                     const char *, int);
void free_alberta_matrix(void **, size_t, size_t, size_t);
```

alberta_matrix(*nr*, *nc*, *el_size*, *name*, *file*, *line*): allocates a new matrix (in C-style) with *nr* number of rows and *nc* number of columns, where each element is of size *el_size* and returns a pointer to this matrix; *name* is a string holding the name of the calling function, *file* a string holding the name of the source file and *line* the line number of the call.

free_alberta_matrix(*ptr*, *nr*, *n_col*, *el_size*): frees the matrix at *ptr* which was previously allocated by **alberta_matrix**().

Again, the following macros simplify the use of the above functions:

```
TYPE** MAT_ALLOC(size_t, size_t, TYPE);
void MAT_FREE(TYPE **, size_t, size_t, TYPE);
```

They supply the above described functions with the function name, file name and line number automatically. These macros need the type of the matrix elements instead of the size. Casting to the correct type is performed.

MAT_ALLOC(*nr*, *nc*, *type*): returns a pointer ****ptr** to a matrix with elements **ptr[i][j]** of type *TYPE* and indices in the range $0 \leq i < nr$ and $0 \leq j < nc$.

MAT_FREE(ptr, nr, nc, type): frees a matrix allocated by **MAT_ALLOC()**.

Many subroutines need additional workspace for storing vectors, etc. (linear solvers like conjugate gradient methods, e.g.). Many applications need such kinds of workspaces for several functions. In order to make handling of such workspaces easy, a data structure **WORKSPACE** is available. In this data structure a pointer to the workspace and the actual size of the workspace is stored.

```
typedef struct workspace    WORKSPACE;

struct workspace
{
    size_t  size;
    void    *work;
};
```

The members yield following information:

size: actual size of the workspace in bytes.

work: pointer to the workspace.

The following functions access and enlarge workspaces:

```
WORKSPACE *get_workspace(size_t, const char *, const char *, int);
WORKSPACE *realloc_workspace(WORKSPACE *, size_t, const char *,
                             const char *, int);
```

Description:

get_workspace(size, name, file, line): return value is a pointer to a **WORKSPACE** structure holding a vector of length **size** bytes; **name** is a string holding the name of the calling function, **file** a string holding the name of the source file and **line** the line number of the call.

realloc_workspace(work_space, size, name, file, line): return value is a pointer to a **WORKSPACE** structure holding a vector of at least length **size** bytes; the member **size** holds the true length of the vector **work**; if **work_space** is a **nil** pointer, a new **WORKSPACE** structure is allocated; **name** is a string holding the name of the calling function, **file** a string holding the name of the source file and **line** the line number of the call.

The macros

```
WORKSPACE* GET_WORKSPACE(size_t)
WORKSPACE* REALLOC_WORKSPACE(WORKSPACE*, size_t)
```

simplify the use of **get_workspace()** and **realloc_workspace()** by supplying the function with **name**, **file**, and **line**.

GET_WORKSPACE(ws, size): returns a pointer to **WORKSPACE** structure holding a vector of length **size** bytes.

REALLOC_WORKSPACE(ws, size): returns a pointer to **WORKSPACE** structure holding a vector of at least length **size** bytes; the member **size** holds the

true length of the vector `work`; if `ws` is a `nil` pointer, a new `WORKSPACE` structure is allocated.

The functions

```
void clear_workspace(WORKSPACE *);
void free_workspace(WORKSPACE *);
```

are used for `WORKSPACE` deallocation. Description:

`clear_workspace(ws)`: frees the vector `ws->work` and sets `ws->work` to `nil` and `ws->size` to 0; the structure `ws` is not freed.

`free_workspace(ws)`: frees the vector `ws->work` and then the structure `ws`.

For convenience, the corresponding macros are defined as well.

```
void CLEAR_WORKSPACE(WORKSPACE *)
void FREE_WORKSPACE(WORKSPACE *)
```

For the handling of general linked lists, data structures and memory allocation functions are defined.

```
typedef struct void_list_element VOID_LIST_ELEMENT;

struct void_list_element
{
    void          *data;
    VOID_LIST_ELEMENT *next;
};
```

The members yield following information:

`data`: pointer to data section of list element.

`next`: pointer to next list element.

Such list elements can be accessed by the function

```
VOID_LIST_ELEMENT *get_void_list_element(void);
```

and subsequently be deallocated by a call to the function

```
void free_void_list_element(VOID_LIST_ELEMENT *);
```

Description:

`get_void_list_element()`: returns a pointer to a new list element with `list->data = list->next = nil`.

`free_void_list_element(list)`: frees the list element `list`. After return, the pointers `list->data` and `list->next` will be changed (the list element will be linked into a list of currently unused elements).

ALBERTA is recording the size of memory which is allocated and deallocated by the routines described above. Information about the actually used size of allocated memory can be obtained by a call of the function

```
void print_mem_use();
```

3.1.4 Parameters and parameter files

Many procedures need parameters, for example the maximal number of iterations for an iterative solver, the tolerance for the error in the adaptive procedure, etc. It is often very helpful to change the values of these parameters without recompiling the program by initializing them from a parameter file.

In order to avoid a fixed list of parameters, we use the following concept: Every parameter consists of two strings: a key string by which the parameter is identified, and a second string containing the parameter values. These values are stored as **ASCII**-characters and can be converted to **int**, **REAL**, etc. according to a format specified by the user (see below). Using this concept, parameters can be accessed at any point of the program.

Usually parameters are initialized from parameter files. Each line of the file describes either a single parameter: the key definition terminated by a ':' character followed by the parameter values, or specifies another parameter file to be included at that point (this can also be done recursively). The syntax of these files is described below and an example is given at the end of this section.

Parameter files

The definition of a parameter has the following syntax:

```
key: parameter values % optional comment
```

Lines are only read up to the first occurrence of the comment sign '%'. All characters behind this sign in the same line are ignored. The comment sign may be a character of the specified filename in an include statement (see below). In this case, '%' is treated as a usual character.

The definition of a new parameter consists out of a key string and a string containing the parameter values. The definition of the key for a new parameter has to be placed in one line before the first comment sign. For the parameter values a continuation line can be used (see below). The key string is a sequence of arbitrary characters except ':' and the comment character. It is terminated by ':', which does not belong to the key string. A key may contain blanks. Optional white space characters as blanks, tabs, etc. in front of a key and in between ':' and the first character of the parameter value are discarded.

Each parameter definition must have at least one parameter value, but it can have more than one. If there are no parameter values specified, i.e. the rest of the line (and all continuation lines) contain(s) only white-space characters (and the continuation character(s)). Such a parameter definition is ignored and the line(s) is (are) skipped.

One parameter value is a sequence of non white-space characters. We will call such a sequence of non white-space characters a word. Two parameter values are separated by at least one white-space character. A string as a

parameter value must not contain white-space characters. Strings enclosed in single or double quotes are not supported at the moment. These quotes are treated as usual characters.

Parameter values are stored as a sequence of words in one string. The words are separated by exactly one blank, although parameter values in the parameter file may be separated by more than one white-space character.

The key definition must be placed in one line. Parameter values can also be specified in so called continuation lines. A line is a continuation line if the last two characters in the preceding line are a `'\'` followed directly by the newline character. The `'\'` and the newline character are removed and the line is continued at the beginning of the next line. No additional blank is inserted.

Lines containing only white-space characters (if they are not continuation lines!) are skipped.

Besides a parameter definition we can include another parameter file with name `filename`:

```
#include "filename"
```

The effect of an include statement is the similar to an include statement in a C-program. Using the function `init_parameters()` (see below) for reading the parameter file, the named file is read by a recursive call of the function `init_parameters()`. Thus, the included parameter file may also contain an include statement. The rest of line behind the closing `"` is skipped. Initialization then is continued from the next line on. An include statement must not have a continuation line.

If a parameter file can not be opened for reading, an error message is produced and the reading of the file is skipped.

Errors occur and are reported if a key definition is not terminated in the same line by `':'`, no parameter values are specified, `filename` for include files are not specified correctly in between `" "`. The corresponding lines are ignored. No parameter is defined, or no file is included.

A parameter can be defined more than once but only the latest definition is valid. All previous definitions are ignored.

Reading of parameter files

Initializing parameters from such files is done by

```
void init_parameters(int, const char *);
```

Description:

`init_parameters(info, filename)`: initializes a set of parameters from the file `filename`; if values of the argument `info` and the global variable `msg_info` are not zero, a list of all defined parameters is printed to the message stream; if `init_parameters()` can not open the input file, or `filename` is a pointer to `nil`, no parameters are defined.

One call of this function should be the first executable statement in the main program. Several calls of `init_parameters()` are possible. If a key is defined more than once, parameter values from the latest definition are valid. Parameter values from previous definition(s) are ignored.

Adding of parameters or changing of parameter values

Several calls of `init_parameters()` are possible. This may add new parameters or change the value of an existing parameter since only the values from the latest definition are valid. Examples for giving parameter values from the command line and integrating them into the set of parameters are shown in Sections 2.2.3 and 2.3.2.

Parameters can also be defined or modified by the function or the macro

```
void add_parameter(int, const char *, const char *);
ADD_PARAMETER(int, const char *, const char *)
```

Description:

`add_parameter(info, key, value)`: initializes a parameter identified by `key` with values `value`; if the parameter already exists, the old values are replaced by the new one; if `info` is not zero information about the initialization is printed; This message can be suppressed by a global level of parameter information (see the parameter `parameter information` in Section 3.1.5).

`ADD_PARAMETER(info, key, value)`: is the same as the above described call `add_parameter(info, key, value)` but the function is additionally supplied with the name of the calling function, source file and line, which results in more detailed messages during parameter definition.

Display and saving of parameters and parameter values

All a list of all parameters together with the actual parameter values can be printed to `stdout` using the function

```
void print_parameters(void);
```

For long time simulations it is important to write all parameters and their values to file; using this file the simulation can be re-done with exactly the same parameter values although the original parameter file was changed. Thus, after the initialization of parameters in a long time simulation, they should be written to a file by the following function:

```
void save_parameters(const char *, int);
```

Description:

`save_parameters(file, info)`: writes all successfully initialized parameters to `file` according to the above described parameter file format; if the value of `info` is different from zero, the location of the initialization is supplied for

each parameter as a comment; no original comment is written, since these are not stored.

Getting parameter values

After initializing parameters by `init_parameters()` we can access the values of a parameter by a call of

```
int get_parameter(int, const char *, const char *, ...);
int GET_PARAMETER(int, const char *, const char *, ...)
```

Description:

`get_parameter(info, key, format, ...)`: looks for a parameter which matches the identifying key string `key` and converts the values of the corresponding string containing the parameter values according to the control string `format`. Pointers to variable(s) of suitable types are placed in the unnamed argument list (compare the syntax of `scanf()`). The first argument `info` defines the level of information during the initialization of parameters with a range of 0 to 4: no to full information. The return value is the number of successfully matched and assigned input items.

If there is no parameter key matching `key`, `get_parameter()` returns without an initialization. The return value is zero. It will also return without an initialization and return value zero if no parameter has been defined by `init_parameters()`.

In the case that a parameter matching the key is found, `get_parameter()` acts like a simplified version of `sscanf()`. The input string is the string containing the parameter values. The function reads characters from this string, interprets them according to a format, and stores the results in its arguments. It expects, as arguments, a control string, `format` (described below) and a set of pointer arguments indicating where the converted input should be stored. If there are insufficient arguments for the format, the behavior is undefined. If the format is exhausted while arguments remain, the excess arguments are simply ignored. The return value is the number of converted arguments.

The control string must only contain the following characters used as conversion specification: `%s`, `%c`, `%d`, `%e`, `%f`, `%g`, `%U`, `%S`, or `%*`. All other characters are ignored. In contrast to `scanf()`, a numerical value for a field width is not allowed. For each element of the control string the next word of the parameter string is converted as follows:

%s: a character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating `'\0'`, which will be added automatically; the string is one single word of the parameter string; as mentioned above strings enclosed in single or double quotes are not supported at the moment.

%c: matches a single character; the corresponding argument should be a pointer to a `char` variable; if the corresponding word of the parameter string consists of more than one character, the rest of the word is ignored; no space character is possible as argument.

%d: matches a decimal integer, whose format is the same as expected for the subject sequence of the `atoi()` function; the corresponding argument should be a pointer to an `int` variable.

%e,%f,%g: matches an optionally signed floating point number, whose format is the same as expected for the subject string of the `atof()` function; the corresponding argument should be a pointer to a `REAL` variable.

%U: matches an unsigned decimal integer in the range $[0,255]$, whose format is the same as expected for the subject sequence of the `atoi()` function; the corresponding argument should be a pointer to an `U_CHAR` variable.

%S: matches an optionally signed decimal integer in the range $[-128,127]$, whose format is the same as expected for the subject sequence of the `atoi()` function; the corresponding argument should be a pointer to an `S_CHAR` variable.

%*: next word of parameter string should be skipped; there must not be a corresponding argument.

`get_parameter()` will always finish its work, successfully or not. It may fail if a misspelled key is handed over or there are not so many parameter values as format specifiers (the remaining variables are not initialized!). If `info` is zero, `get_parameter()` works silently; no error message is produced. Otherwise the key and the initialized values and error messages are printed. The second way to influence messages produced by `get_parameter()` namely by is a parameter `parameter information` specified in a parameter file, see Section 3.1.5.

`GET_PARAMETER(info, key, format, ...)`: is a macro and acts in the same way as the function `get_parameter(info, key, format, ...)` but the function is additionally supplied with the name of the calling function, source file and line, which results in more detailed messages during parameter definition.

In order to prevent the program from working with uninitialized variables, all parameters should be initialized before! By the return value the number of converted arguments can be checked.

Example 3.7 (`init_parameters()`, `GET_PARAMETER()`). Consider the following parameter file `init.dat`:

```
adapt info: 3          % level of information of the adaptive method
adapt tolerance: 0.001 % tolerance for the error
```


Then

```
init_parameters(0, "init.dat");
...
tolerance = 0.1;
GET_PARAMETER(0, "adapt tolerance", "%e", &tolerance);
```

initializes the REAL variable `tolerance` with the value 0.001.

3.1.5 Parameters used by the utilities

The utility tools use the following parameters initialized with default values given in `()`:

level of information (10): the global level of information; can restrict the local level of information (compare Section 3.1.2).

parameter information (1): enforces more/less information than specified by the argument `info` of the routine `get_parameter(info, ...)`:

- 0: no message at all is produced, although the value `info` may be non zero;
- 1: gives only messages if the value of `info` is non zero;
- 2: all error messages are printed, although the value of `info` may be zero;
- 4: all messages are printed, although the value of `info` may be zero.

WAIT (1): sets the value of the global variable `msg_wait` and changes by that the behaviour of the macro `WAIT` (see Section 3.1.2).

3.1.6 Generating filenames for meshes and finite element data

During simulation of time-dependent problems one often wants to store meshes and finite element data for the sequence of time steps. A routine is provided to automatically generate file names composed from a given data path, a base-name for the file and a number which is iteration counter of the actual time step in time-dependent problems. Such a function simplifies the handling of a sequence of data for reading and writing. It also ensures that files are listed alphabetically in the given path (up to 1 million files with the same base-name).

```
const char *generate_filename(const char *, const char *, int);
```

Description:

`generate_filename(path, file, index)`: composes a filename from the given `path`, the base-name `file` of the file and the (iteration counter) `index`. When no path is given, the current directory `"/"` is used, if the first character of `path` is `'~'`, `path` is assumed to be located in the home directory and the name of the path is expanded accordingly, using the environment variable `HOME`. A pointer to a string containing the full filename is the return value; this string is overwritten on the next call to `generate_filename()`.

For instance, a call `generate_filename("./output", "mesh", 1)` constructs the filename `./output/mesh000001` and returns a pointer to it. An example how to use `generate_filename()` in a time dependent problem is given in Section 2.3.10.

3.2 Data structures for the hierarchical mesh

3.2.1 Constants describing the dimension of the mesh

The symbolic constant `DIM` defines the dimension of the triangulation:

```
#define DIM 1
```

for one dimensional grids,

```
#define DIM 2
```

for two dimensional grids, and

```
#define DIM 3
```

for three dimensional ones. Most parts of the finite element code do not depend on the particular choice of `DIM`. But especially the refinement/coarsening and hierarchical information processing routines strongly depend on `DIM`.

For finite element methods on one, two, or three dimensional curves and surfaces embedded in \mathbb{R}^n (like mean curvature flow [35]), the vertex coordinates of the simplices have n components. Thus, we need another symbolic constant `DIM_OF_WORLD` to define the number of these components. For applications in domains in \mathbb{R}^{DIM} , `DIM_OF_WORLD` equals `DIM` otherwise the value of `DIM_OF_WORLD` is the dimension of the \mathbb{R}^n where the `DIM`-dimensional surface is embedded.

As a convention all `typedef` definitions with the ending `_D` define a new data type for vectors of length `DIM_OF_WORLD`, definitions with the ending `_DD` a matrix of size `DIM_OF_WORLD` \times `DIM_OF_WORLD`, for example

```
typedef REAL          REAL_D[DIM_OF_WORLD];
typedef REAL          REAL_DD[DIM_OF_WORLD][DIM_OF_WORLD];
```

For `DIM_OF_WORLD` ≤ 3 , several macros are defined for working on `REAL_D` vectors:

```
void COPY_DOW(REAL_D, REAL_D)
REAL DIST_DOW(REAL_D, REAL_D)
REAL NORM_DOW(REAL_D)
REAL SCAL_DOW(REAL, REAL_D)
REAL SCP_DOW(REAL_D, REAL_D)
void SET_DOW(REAL, REAL_D)
```

Description:

`COPY_DOW(x, y)`: copies all elements of vector `x` to `y`.

`DIST_DOW(x, y)`: returns the Euclidean distance of the two vectors \mathbf{x} , \mathbf{y} .

`NORM_DOW(x)`: returns the Euclidean norm of the vector \mathbf{x} .

`SCAL_DOW(a, x)`: scales all elements of vector \mathbf{x} with \mathbf{a} .

`SCP_DOW(x, y)`: returns the Euclidean scalar product of the two vectors \mathbf{x} , \mathbf{y} .

`SET_DOW(a, x)`: set all elements of vector \mathbf{x} to \mathbf{a} .

For $\text{DIM_OF_WORLD} \leq 3$, the vector product of two vectors is computed by the routine

```
void vector_product(const REAL_D, const REAL_D, REAL_D);
```

Description:

`vector_product(x, y, z)`: calculates the vector product $z = x \wedge y$.

3.2.2 Constants describing the elements of the mesh

For a simplex we have to define the number of vertices, adjacent simplices, in two or three dimensions the number of edges, and in three dimensions the number of element faces:

```
#define N_VERTICES      2
#define N_NEIGH         2
```

for one dimension,

```
#define N_VERTICES      3
#define N_EDGES         3
#define N_NEIGH         3
```

for two dimensions, and

```
#define N_VERTICES      4
#define N_EDGES         6
#define N_FACES         4
#define N_NEIGH         4
```

for three dimensions.

3.2.3 Neighbour information

As mentioned in Section 1.2 neighbour information is usually not stored explicitly, but such an explicit storing may speed up special applications. **ALBERTA** supports both kinds of managing neighbour information. Via the symbolic constant `NEIGH_IN_EL` the user can enforce to store neighbour information explicitly in the elements. Setting

```
#define NEIGH_IN_EL    0
```

will result in calculating neighbour information (if requested) by the traversal routines, while

```
#define NEIGH_IN_EL    1
```

enforces that neighbour information is stored explicitly on each element and has to be set by the refinement/coarsening routines. This consumes additional computer memory, but will speed up the traversal routines that usually have to compute neighbour information. Macros

```
NEIGH(e1, el_info),  OPP_VERTEX(e1, el_info),  EL_TYPE(e1,el_info)
```

are defined to access information independently of the value of `NEIGH_IN_EL`, compare Section 3.2.10.

Currently, some routines may not work correctly for `NEIGH_IN_EL==1`. We recommend *not* to use this.

3.2.4 Element indices

It is often very helpful — especially during program development — for every element to have a unique global index. This requires an entry in the element data structure which adds to the needed computer memory.

On the other hand this additional amount of computer memory may be a disadvantage in real applications where a big number of elements is needed, and — after program development — element index information is no longer of interest.

After setting the value of the symbolic constant `EL_INDEX` to be 1 an element index is available and adds to the amount of computer memory, and setting it to be 0 no element index information is available which results in less computer memory usage. The macro

```
INDEX(e1)
```

is defined to access element indices independently of the value of `EL_INDEX`. If no indices are available, the macro returns -1, compare Section 3.2.11.

3.2.5 The BOUNDARY data structure

Degrees of freedom (DOFs) which describe the finite element space are located at a node which is either a vertex, an edge, a face, or the barycenter of an element (see Section 1.3). For each node we have to specify if (Dirichlet) boundary conditions are prescribed for DOFs at this node or if the DOFs are true degrees of freedom. At the moment we support three kinds of nodes:

- interior node where DOFs are handled as real degrees of freedom;
- nodes on the Dirichlet boundary where Dirichlet boundary conditions are prescribed;
- nodes on the Neumann boundary which are treated like interior nodes.

The boundary type is specified by a `S_CHAR`. We use the following symbolic constants and convention:

```

#define INTERIOR      0
#define DIRICHLET     1
#define NEUMANN       -1

```

An interior node has the boundary type `INTERIOR`, a node on a Dirichlet boundary has positive boundary type ($\geq \text{DIRICHLET}$), and a node on the Neumann boundary has negative type ($\leq \text{NEUMANN}$).

For the approximation of curved boundaries, the new nodes at the midpoint of the refinement edge are projected onto the curved boundary. Most likely, the boundary of the domain is split in several parts, where different functions for these projections are needed. For applying such a function to a node we have to know which part of the domain's boundary the refinement edge belongs.

In order to handle curved boundaries and to classify the type of nodes located at edges and faces we define the following structure:

```

typedef struct boundary  BOUNDARY;

struct boundary
{
    void          (*param_bound)(REAL_D  *);
    S_CHAR        bound;
};

```

The members yield following information:

param_bound: pointer to a function which projects a newly generated vertex onto a curved boundary;

param_bound(coord) projects the midpoint of the refinement edge with world coordinates **coord** onto the curved boundary; after the call, **coord** contains the coordinates of the projection, compare Section 3.4.1, Fig. 3.3; if **param_bound** is a pointer to **nil**, the edge or face does not belong to a curved boundary.

bound: boundary type of a node at the edge or face; the boundary type is uniquely defined by the sign of **bound**; by the value of **bound**, the part of the domain's boundary to which this object belongs can be stored (assuming, that these parts are enumerated).

There are some useful macros for the `BOUNDARY` structure defined in the header file `alberta.h`. They can be applied to any `BOUNDARY` pointer and give a default return value if the value of that pointer is **nil**.

```

#define GET_BOUND(boundary)\
    ((boundary) ? (boundary)->bound : INTERIOR)
#define IS_INTERIOR(boundary)\
    ((boundary) ? (boundary)->bound == INTERIOR : true)
#define IS_DIRICHLET(boundary)\
    ((boundary) ? (boundary)->bound >= DIRICHLET : false)
#define IS_NEUMANN(boundary)\
    ((boundary) ? (boundary)->bound <= NEUMANN : false)

```

3.2.6 The local indexing on elements

For the handling of higher order discretizations where besides vertices DOFs can be located at edges (in 2d and 3d), faces (in 3d), or center, we also need a local numbering for edges, and faces. Finally, a local numbering of neighbours for handling neighbour information is needed, used for instance in the refinement algorithm itself and for error estimator calculation.

The i -th neighbour is always the element opposite the i -th vertex. The i -th edge/face is the edge/face opposite the i -th vertex in 2d respectively 3d; edges in 3d are numbered in the following way (compare Fig. 3.1):

edge 0: between vertex 0 and 1, **edge 3:** between vertex 1 and 2,
edge 1: between vertex 0 and 2, **edge 4:** between vertex 1 and 3,
edge 2: between vertex 0 and 3, **edge 5:** between vertex 2 and 3.




Fig. 3.1. Local indices of edges/neighbours in 2d and local indices of edges in 3d.

The data structures described in the subsequent sections are based on this local numbering of vertices, edges, faces, and neighbours.

3.2.7 The MACRO_EL data structure

We now describe the macro triangulation and data type for an element of the macro triangulation. The macro triangulation is stored in a doubly linked list of macro elements:

```
typedef struct macro_el  MACRO_EL;

struct macro_el
{
    MACRO_EL    *next, *last;
    EL          *el;
    REAL        *coord[N_VERTICES];
    S_CHAR      bound[N_VERTICES];
#ifdef DIM == 2
    const BOUNDARY *boundary[N_EDGES];
#endif
}
```

```

#if DIM == 3
    const BOUNDARY *boundary[N_FACES+N_EDGES];
#endif
    MACRO_EL      *neigh[N_NEIGH];
    U_CHAR        opp_vertex[N_NEIGH];

    int           index;

#if DIM == 3
    U_CHAR        el_type;
    S_CHAR        orientation;
#endif
};

```

The members yield following information (alternatives for different dimensions are separated by ‘/’):

next: pointer to next macro element; if **next** is **nil** then the macro element is the last macro element.

last: pointer to previous macro element; if **last** is **nil** the macro element is the first macro element.

el: root of the binary tree located at this macro element.

coord: pointer to the world coordinates of the element’s vertices.

bound: **bound[i]** boundary type of vertex **i**.

boundary: **boundary[i]** is a pointer to a boundary structure of the **i**-th edge/face for $0 \leq i < N_NEIGH$ (in 2d and 3d); in 3d **boundary[N_FACES+i]** is a pointer to a boundary structure of the **i**-th edge for $0 \leq i < N_EDGES$; for an interior edge/face, it is a pointer to **nil**.

neigh: **neigh[i]** pointer to the macro element opposite the **i**-th local vertex; it is a pointer to **nil** if the vertex/edges/faces opposite the **i**-th local vertex belongs to the boundary.

opp-vertex: **opp-vertex[i]** is undefined if **neigh[i]** is a pointer to **nil**; otherwise it is the local index of the neighbour’s vertex opposite the common vertex/edge/face.

index: a unique index of all macro elements.

el_type: type of the element $\in [0, 1, 2]$ used for refinement and coarsening (for the definition of the element type see Section 1.1.1, only 3d).

orientation: orientation of a tetrahedron — depending on the vertex numbering, this is +1 or -1 (only 3d). The entry is calculated during the first mesh traversal with **FILL_ORIENTATION** flag set. As the calculation of orientation is based on the vertex coordinates stored in the macro elements, it may give wrong results in case of parametric meshes, when the parameterization is not orientation preserving.

3.2.8 The EL data structure

The elements of the binary trees and information that should be present for tree elements are stored in the data structure:

```
typedef struct el    EL;

struct el
{
    EL          *child[2];
    DOF         **dof;
#ifdef EL_INDEX
    int          index;
#endif
    S_CHAR       mark;
#ifdef DIM > 1
    REAL         *new_coord;
#endif

#ifdef NEIGH_IN_EL
    EL          *neigh[N_NEIGH];
    U_CHAR       opp_vertex[N_NEIGH];
#ifdef DIM == 3
    U_CHAR       el_type;
#endif
#endif
};
```

The members yield following information:

child: pointers to the two children for interior elements of the tree; **child[0]** is a pointer to **nil** for leaf elements; **child[1]** is a pointer to user data on leaf elements if **leaf_data_size** is bigger than zero, otherwise **child[1]** is also a pointer to **nil** for leaf elements (see Section 3.2.12).

dof: vector of pointers to DOFs; these pointers must be available for the element vertices (for the geometric description of the mesh); there may be pointers at the edges (in 2d and 3d), at the faces (in 3d), and at the barycenter; they are ordered in the following way: the first **N_VERTICES** entries correspond to the DOFs at the vertices; the next one are those at the edges, if present, then those at the faces, if present, and finally those at the barycenter, if present; the offsets are defined in the **MESH** structure (see Sections 3.2.14, 3.4.1, 3.4.2).

index: unique global index of the element; these indices are not strictly ordered and may be larger than the number of elements in the binary tree (the list of indices may have holes after coarsening); the index is available only if **EL_INDEX** is **true**.

mark: marker for refinement and coarsening: if **mark** is positive for a leaf element this element is refined **mark** times; if it is negative for a leaf element the element may be coarsened **-mark** times; (see Sections 3.4.1, 3.4.2).

new_coord: if the element has a boundary edge on a curved boundary this is a pointer to the coordinates of the new vertex that is created due to the refinement of the element, otherwise it is a `nil` pointer; thus, coordinate information can also be produced by the traversal routines in the case of a curved boundary.

If neighbour information should be present in the element structure (if `NEIGH_IN_EL` is `true`) then we have the additional entries:

neigh: `neigh[i]` pointer to the element opposite the *i*-th local vertex; it is a pointer to `nil` if the vertex/edges/faces opposite the *i*-th local vertex belongs to the boundary.

opp_vertex: `opp_vertex[i]` is undefined if `neigh[i]` is a pointer to `nil`; otherwise it is the local index of the neighbour's vertex opposite the common vertex/edge/face.

el_type: the element's type (see Section 3.4.1); has to be available on the element if neighbour information is located at the element, since such information then can not be produced by the traversal routines going from one element to another by the neighbour pointers.

3.2.9 The `EL_INFO` data structure

The `EL_INFO` data structure has entries for all information which is not stored on elements explicitly, but may be generated by the mesh traversal routines; most entries of the `EL_INFO` structure are only filled if requested (see Section 3.2.19).

```
typedef struct el_info  EL_INFO;

struct el_info
{
    MESH            *mesh;
    REAL_D          coord[N_VERTICES];
    const MACRO_EL  *macro_el;
    EL              *el, *parent;
    FLAGS           fill_flag;
    S_CHAR          bound[N_VERTICES];
#if DIM == 2
    const BOUNDARY  *boundary[N_EDGES];
#endif

#if DIM == 3
    const BOUNDARY  *boundary[N_FACES+N_EDGES];
#endif

    U_CHAR          level;
```

```

    #if ! NEIGH_IN_EL
        EL                *neigh[N_NEIGH];
        U_CHAR            opp_vertex[N_NEIGH];
    #if DIM == 3
        U_CHAR            el_type;
    #endif
    #endif
        REAL_D            opp_coord[N_NEIGH];

    #if DIM == 3
        S_CHAR            orientation;
    #endif
};

```

The members yield following information:

mesh: a pointer to the current mesh.

coord: `coord[i]` is a `DIM_OF_WORLD` vector storing world coordinates of the *i*-th vertex.

macro_el: the current element belongs to the binary tree located at the macro element `macro_el`.

el: pointer to the current element.

parent: `el` is a child of element `parent`.

fill_flag: a flag which indicates which elements are called and which information should be present (see Section 3.2.19).

bound: `bound[i]` boundary type of vertex *i*.

boundary: `boundary[i]` is a pointer to a boundary structure of the *i*-th edge/face for $i = 0, \dots, N_NEIGH - 1$ (in 2d and 3d); additionally in 3d, `boundary[N_FACES+i]` is a pointer to a boundary structure of the *i*-th edge for $i = 0, \dots, N_EDGES - 1$; it is a pointer to `nil` for an interior edge/face.

level: level of the current element; the level is zero for macro elements and the level of the children is (level of the parent + 1); the level is always filled by the traversal routines.

opp_coord: `opp_coord[i]` coordinates of the *i*-th neighbour's vertex opposite the common vertex/edge/face.

orientation: ± 1 : sign of the determinant of the transformation to the reference element with vertices (0, 0, 0), (1, 1, 1), (1, 1, 0), (1, 0, 0) (see Fig. 1.7).

If neighbour information is not present in the element structure (`NEIGH_IN_EL` is `false`), then we have the additional entries:

neigh: `neigh[i]` pointer to the element opposite the *i*-th local vertex; it is a pointer to `nil` if the vertex/edges/faces opposite the *i*-th local vertex belongs to the boundary.

opp_vertex: `opp_vertex[i]` is undefined if `neigh[i]` is a pointer to `nil`; otherwise it is the local index of the neighbour's vertex opposite the common vertex/edge/face.

`el_type`: the element's type (see Section 3.4.1); is filled automatically by the traversal routines (only 3d).

3.2.10 The NEIGH, OPP_VERTEX and EL_TYPE macros

If neighbour information is produced by the traversal routines (`NEIGH_IN_EL == 0`) we have to access neighbour information by the corresponding `el_info` structure:

```
neigh = el_info->neigh[i];
opp_v = el_info->opp_vertex[i];
```

If such information is stored explicitly at each element (`NEIGH_IN_EL == 1`) we get a pointer to the *i*-th neighbour of an element `el` and the corresponding `opp_vertex` by

```
neigh = el->neigh[i];
opp_v = el->opp_vertex[i];
```

To have same access for both situations there are two macros `NEIGH` and `OPP_VERTEX` defined in `alberta.h`:

```
#if NEIGH_IN_EL
#define NEIGH(el,el_info)      (el)->neigh
#define OPP_VERTEX(el,el_info) (el)->opp_vertex
#else
#define NEIGH(el,el_info)      (el_info)->neigh
#define OPP_VERTEX(el,el_info) (el_info)->opp_vertex
#endif
```

Similarly, the element type (only in 3d) is stored either in the `EL` structure or is generated during traversal in `EL_INFO`. A macro `EL_TYPE` is defined to access such information:

```
#if NEIGH_IN_EL
#define EL_TYPE(el,el_info) (el)->el_type
#else
#define EL_TYPE(el,el_info) (el_info)->el_type
#endif
```

Using these macros we always get information about the *i*-th neighbour or the element type by

```
neigh = NEIGH(el,el_info)[i];
opp_v = OPP_VERTEX(el,el_info)[i];
type  = EL_TYPE(el,el_info);
```

independently of the value of `NEIGH_IN_EL`.

3.2.11 The INDEX macro

In order to avoid to eliminate all lines where the element index is accessed after recompiling the source with `EL_INDEX == 0`, a macro `INDEX` is defined:

```

#if EL_INDEX
#define INDEX(e1)  ((e1) ? (e1)->index : -1)
#else
#define INDEX(e1)  -1
#endif

```

If element indices are stored at the elements `INDEX(e1)` gives the index of the element, if the `e1` is not a pointer to `nil`. If `e1` is a pointer to `nil`, the value of `INDEX(e1)` is `-1`. For instance, this allows a construction like `INDEX(NEIGH(e1,el_info)[i])` without testing whether this neighbour exists or not (in the case of a boundary edge/face). If element indices are not stored, the value of `INDEX(e1)` is always `-1`.

3.2.12 The LEAF_DATA_INFO data structure

As mentioned in Section 1.2, it is often necessary to provide access to special user data which is needed only on leaf elements. Error indicators give examples for such data.

Information for leaf elements depends strongly on the application and so it seems not to be appropriate to define a fixed data type for storing this information. Thus, we implemented the following general concept: The user can define his own type for data that should be present on leaf elements. ALBERTA only needs the size of memory that is required to store leaf data (one entry in a structure `leaf_data_info` described below in detail). During refinement and coarsening ALBERTA automatically allocates and deallocates memory for user data on leaf elements if the data size is bigger than zero. Thus, after grid modifications each leaf element possesses a memory area which is big enough to take leaf data.

To access leaf data we must have for each leaf element a pointer to the provided memory area. This would need an additional pointer on leaf elements. To make the element data structure as small as possible and in order to avoid different element types for leaf and interior elements we “hide” leaf data at the pointer of the second child on leaf elements:

By definition, a leaf element is an element without children. For a leaf element the pointers to the first *and* second child are pointers to `nil`, but since we use a binary tree the pointer to the second child must be `nil` if the pointer to the first child is a `nil` pointer and vice versa. Thus, only testing the first child will give correct information whether an element is a leaf element or not, and we do not have to use the pointer of the second child for this test. As consequence we can use the pointer of the second child as a pointer to the allocated area for leaf data and the user can write or read leaf data via this pointer (using casting to a data type defined by himself).

The consequence is that a pointer to the second child is only a pointer to an element if the pointer to the first child is not a `nil` pointer. Thus testing whether an element is a leaf element or not must only be done using the pointer to the first child.

If the leaf data size equals zero then the pointer to the second child is also a `nil` pointer for leaf elements.

Finally, the user may supply routines for transforming user data from parent to children during refinement and for transforming user data from children to parent during coarsening. If these routines are not supplied, information stored for the parent or the children respectively is lost.

In order to handle an arbitrary kind of user data on leaf elements we define a data structure `LEAF_DATA_INFO` which gives information about size of user data and may give access to functions transforming user data during refinement and coarsening. A pointer to such a structure is an entry in the `MESH` data structure.

```
typedef struct leaf_data_info  LEAF_DATA_INFO;

struct leaf_data_info
{
    char            *name;
    unsigned        leaf_data_size;
    void            (*refine_leaf_data)(EL *, EL *[2]);
    void            (*coarsen_leaf_data)(EL *, EL *[2]);
};
```

Following information is provided via this structure:

name: textual description of leaf data.

leaf_data_size: size of the memory area which is used for storing leaf data; if `leaf_data_size == 0` no memory for leaf data is allocated and `child[1]` is also pointer to `nil` for leaf elements; if `leaf_data_size > 0` size of leaf data will be \geq `leaf_data_size` (ALBERTA may increase the size of leaf data in order to guarantee an aligned memory access).

refine_leaf_data: pointer to a function for transformation of leaf data during refinement; first, `refine_leaf_data(parent, child)` transforms leaf data from the parent to the two children if `refine_leaf_data` is not `nil`; then leaf data of the parent is destroyed. Transformation only takes place if `leaf_data_size > 0`.

coarsen_leaf_data: pointer to a function for transformation of leaf data during coarsening; first, `coarsen_leaf_data(parent, child)` transforms leaf data from the two children to the parent if `refine_leaf_data` is not `nil`; then leaf data of the children is destroyed. Transformation only takes place if `leaf_data_size > 0`.

The following macros for testing leaf elements and accessing leaf data are provided:

```
#define IS_LEAF_EL(el) (!(el)->child[0])
#define LEAF_DATA(el) ((void *) (el)->child[1])
```

The first macro `IS_LEAF_EL(e1)` is true for leaf elements and false for elements inside the binary tree; for leaf elements, `LEAF_DATA(e1)` returns a pointer to leaf data hidden at the pointer to the second child.

3.2.13 The `RC_LIST_EL` data structure

For refining and coarsening we need information of the elements at the refinement and coarsening edge (compare Sections 1.1.1 and 1.1.2). Thus, we have to collect all elements at this edge. In 1d the patch is built from the current element only, in 2d we have at most the current element and its neighbour across this edge, if the edge is not part of the boundary. In 3d we have to loop around this edge to collect all the elements. Every element at the edge has at most two neighbours sharing the same edge. Defining an orientation for this edge, we can define the *right* and *left* neighbour in 3d.

For every element at the refinement/coarsening edge we have an entry in a vector. The elements of this vector build the refinement/coarsening patch. In 1d the vector has length 1, in 2d length 2, and in 3d length `mesh->max_no_edge_neigh` since this is the maximal number of elements sharing the same edge in the mesh `mesh`.

```
typedef struct rc_list_el  RC_LIST_EL;

struct rc_list_el
{
    EL          *el;
    int         no;
    int         flag;
#ifdef DIM == 3
    RC_LIST_EL  *neigh[2];
    int         opp_vertex[2];
    U_CHAR      el_type;
#endif
};
```

Information that is provided for every element in this `RC_LIST_EL` vector:

`el`: pointer to the element of the `RC_LIST_EL`.

`no`: this is the `no`-th entry in the vector.

`flag`: only used in the coarsening module: `flag` is `true` if the coarsening edge of the element is the coarsening edge of the patch, otherwise `flag` is `false`.

`neigh`: `neigh[0/1]` neighbour of element to the right/left in the orientation of the edge, or a `nil` pointer in the case of a boundary face (only 3d).

`opp-vertex`: `opp-vertex[0/1]` the opposite vertex of `neigh[0/1]` (only 3d).

`el_type`: the element type; this value is set during looping around the refinement/coarsening edge; if neighbour information is produced by the traversal routines, information about the type of an element can not be accessed via `el->el_type` and thus has to be stored in the `RC_LIST_EL` vector (only 3d).

This `RC_LIST_EL` vector is one argument to the interpolation and restriction routines for DOF vectors (see Section 3.3.3).

3.2.14 The MESH data structure

All information about a triangulation is accessible via the `MESH` data structure:

```
typedef struct mesh  MESH;

struct mesh
{
    const char      *name;
    int             n_vertices;
    int             n_elements;
    int             n_hier_elements;

#ifdef DIM == 2
    int             n_edges;
#endif

#ifdef DIM == 3
    int             n_edges;
    int             n_faces;
    int             max_edge_neigh;
#endif

    int             n_macro_el;
    MACRO_EL        *first_macro_el;

    REAL            diam[DIM_OF_WORLD];
    PARAMETRIC      *parametric;

    LEAF_DATA_INFO  leaf_data_info[1];
    U_CHAR          preserve_coarse_dofs;

    DOF_ADMIN       **dof_admin;
    int             n_dof_admin;

    int             n_dof_el;
    int             n_dof[DIM+1];
    int             n_node_el;
    int             node[DIM+1];

    /*-----*/
    /*---  pointer for administration; don't touch!  ---*/
    /*-----*/
    void            *mem_info;
};
```

The members yield following information:

name: string with a textual description for the mesh, or `nil`.
n_vertices: number of vertices of the mesh.
n_elements: number of leaf elements of the mesh.
n_hier_elements: number of all elements of the mesh.
n_edges: number of edges of the mesh (2d and 3d).
n_faces: number of faces of the mesh (3d).
max_edge_neigh: maximal number of elements that share one edge; used to allocate memory to store pointers to the neighbour at the refinement/coarsening edge (3d).
n_macro_el: number of macro elements.
first_macro_el: pointer to the first macro element.
diam: diameter of the mesh in the `DIM_OF_WORLD` directions.
parametric: is a pointer to `nil` if the mesh contains no parametric elements; otherwise it is a pointer to a `PARAMETRIC` structure containing coefficients of the parameterization and related information; the current version of `ALBERTA` includes only a preliminary definition of the `PARAMETRIC` data structure, which is not described here and is subject to change.
leaf_data_info: a structure with information about data on the leaf elements.
preserve_coarse_dofs: if the value is non zero then preserve all DOFs on all levels (can be used for multigrid, e.g.); otherwise all DOFs on the parent that are not handed over to a child are removed during refinement and added again on the parent during coarsening, compare Section 3.4.

The last entries are used for the administration of DOFs and are explained in Section 3.3 in detail.

dof_admin: vector of `dof_admins`.
n_dof_admin: number of `dof_admins`.
n_node_el: number of nodes on a single element where DOFs are located; needed for the (de-) allocation of the `dof`-vector on the element.
n_dof_el: number of all DOFs on a single element.
n_dof: number of DOFs at the different positions `VERTEX`, `EDGE`, `(FACE,)` `CENTER` on an element:
n_dof[VERTEX]: number of DOFs at a vertex (≥ 1);
n_dof[EDGE]: number of DOFs at an edge; if no DOFs are associated to edges, then this value is 0 (2d and 3d);
n_dof[FACE]: number of DOFs at a face; if no DOFs are associated to faces, then this value is 0 (3d);
n_dof[CENTER]: number of DOFs at the barycenter; if no DOFs are associated to the barycenter, then this value is 0.

node: gives the index of the first node at vertex, edge (2d and 3d), face (3d), and barycenter:

node[VERTEX]: has always value 0; **dof[0]**, ..., **dof[N_VERTICES-1]** are always DOFs at the vertices;

node[EDGE]: **dof[node[EDGE]]**, ..., **dof[node[EDGE]+N_EDGES-1]** are the DOFs at the **N_EDGES** edges, if DOFs are located at edges (2d and 3d);

node[FACE]: **dof[node[FACE]]**, ..., **dof[node[FACE]+N_FACES-1]** are the DOFs at the **N_FACES** faces, if DOFs are located at faces (3d);

node[CENTER]: **dof[node[CENTER]]** are the DOFs at the barycenter, if DOFs are located at the barycenter of elements.

Finally, the pointer **mem_info** is used for internal memory management and must not be changed.

3.2.15 Initialization of meshes

It is possible to handle more than one mesh at the same time. A mesh must be accessed by one of the following functions or macro

```
MESH *get_mesh(const char *, void (*)(MESH *),
               void (*)(LEAF_DATA_INFO *));
MESH *check_and_get_mesh(int, int, int, int, const char *,
                        const char *, void (*)(MESH *),
                        void (*)(LEAF_DATA_INFO *));
MESH *GET_MESH(const char *, void (*)(MESH *),
               void (*)(LEAF_DATA_INFO *))
```

Description:

get_mesh(name, init_dof_admins, init_leaf_data): returns a pointer to a filled mesh structure; **name** is a string holding a textual description of mesh and is duplicated at the member **name** of the mesh;

init_dof_admins is a pointer to a user defined function for the initialization of the required DOFs on the mesh (see Section 3.6.2 for an example of such a function); if this argument is a **nil** pointer a warning is given and a **DOF_ADMIN** structure for DOFs at the vertices of the mesh is generated. This is also done if none of the users **DOF_ADMINs** uses DOFs at vertices (as mentioned above, DOFs at vertices have to be present for the geometric description of the mesh);

init_leaf_data is a pointer to a function for the initialization of the **leaf_data_info** structure, i.e. the size of leaf data can be set and pointers to transformation functions for leaf data can be adjusted. If this argument is a **nil** pointer, the member **leaf_data_info** structure is initialized with zero.

The return value of the function is a pointer to the filled **mesh** data structure. The user must not change any entry in this structure aside from the

`max_edge_neigh`, `n_macro_el`, `first_macro_el`, `diam`, and `parametric` entries.

Currently no function for adding further DOFs after invoking `get_mesh()` is implemented. Such a function would have to perform something like a so called *p*-refinement.

There is no other possibility to define new meshes inside ALBERTA.

`check_and_get_mesh(d, dow, n, i, v, name, ida, ild)`: returns also a pointer to a filled mesh structure; the last three arguments are the same as of `get_mesh()`; in addition several checks about the used ALBERTA library are performed: `d` is DIM, `dow` is DIM_OF_WORLD, `n` is NEIGH_IN_EL, `i` is EL_INDEX, and `v` is VERSION in the user program; these values are checked against the constants in the used library; if these values are identical, the mesh is accessed by `get_mesh(name, ida, ild)`; otherwise an error message is produced and the program stops.

`GET_MESH(name, ida, ild)`: returns also pointer to a filled mesh structure; this macro calls `check_and_get_mesh()` and automatically supplies this function with the first five (missing) arguments; this macro should always be used for accessing a mesh.

After this initialization data for macro elements can be specified for example by reading it from file (see Section 3.2.16).

A mesh that is not needed any more can be freed by a call of the function

```
void free_mesh(MESH *);
```

Description:

`free_mesh(mesh)`: will de-allocate all memory used by `mesh` (elements, DOFs, etc.), and finally the data structure `mesh` too.

3.2.16 Reading macro triangulations

Data for macro triangulations can easily be stored in an ASCII-file (for binary macro files, see the end of this section). For the macro triangulation file we use a similar key-data format like the parameter initialization (see Section 3.1.4):

```
DIM:          dim
DIM_OF_WORLD: dow

number of vertices: nv
number of elements: ne

vertex coordinates:
DIM_OF_WORLD coordinates of vertex[0]
...
DIM_OF_WORLD coordinates of vertex[nv-1]
```

```

element vertices:
N_VERTICES indices of vertices of simplex[0]
...
N_VERTICES indices of vertices of simplex[ne-1]

element boundaries:
N_NEIGH boundary descriptions of simplex[0]
...
N_NEIGH boundary descriptions of simplex[ne-1]

element neighbours:
N_NEIGH neighbour indices of simplex[0]
...
N_NEIGH neighbour indices of of simplex[ne-1]

element type:
element type of simplex[0]
...
element type of simplex[ne-1]

```

All lines closed by the ‘:’ character are keys for following data (more or less self explaining). Data for elements and vertices are read and stored in vectors for the macro triangulation. Index information given in the file correspond to this vector oriented storage of data. Thus, index information must be in the range $0, \dots, ne-1$ for elements and $0, \dots, nv-1$ for vertices. Although a vertex may be a common vertex of several macro elements the coordinates are only stored once.

An important point is that the refinement edges are determined by the local numbering of the vertices on each element! This is always the edge in between the vertices with local index 0 and 1. In 1d the local vertex 0 has to be smaller than the local vertex 1; in 2d the local ordering of vertex indices must be in counterclockwise sense for each triangle; in 3d by this numbering and by the element type the distribution of the refinement edges for the children is determined.

Information about element boundaries and element neighbours is optional. Information about the element type must only be given in 3d and such information is optional too. Given values must be in the range 0,1,2. If such information is not given in the file, all elements are assumed to be of type 0.

There are only few restrictions on the ordering of data in the file: On the first two lines DIM and DIM_OF_WORLD have to be specified (in an arbitrary order); number of elements has to be specified before element vertices, element boundaries element neighbours (the last two optional), and element type (only 3d and optional); number of vertices has to be specified before vertex coordinates and element vertices. Besides these restrictions, ordering of data is left to the user.

For the 1d version of ALBERTA the parameter `DIM` must equal 1, for the 2d version `DIM` must equal 2, and for the 3d version `DIM` must equal 3. The parameter `DIM_OF_WORLD` must be $\geq \text{DIM}$. In the current version `DIM_OF_WORLD` ≤ 3 is only supported. By these values it is checked whether data matches to the actual version of ALBERTA.

If information about boundaries is not given in the file, all boundaries are assumed to be of Dirichlet type with boundary type `DIRICHLET`, compare Section 3.2.5.

If information about neighbours is supplied in the file the index is `-1` corresponds to a non existing neighbour at a boundary vertex (1d), edge (2d), or face (3d). If information about neighbours is not present or not correctly specified in the file, it is, as the local indices of opposite vertices, generated automatically.

Reading data of the macro grid from these files can be done by

```
void read_macro(MESH *, const char *,
               const BOUNDARY *(*)(MESH *, int ));
```

Description:

`read_macro(mesh, name, ibdry)`: reads data of the macro triangulation for `mesh` from the ASCII-file `name`; adjustment of `BOUNDARY` structures can be done via the function pointer `ibdry`; the `mesh` structure must have been initialized by a call of `get_mesh()` or `check_and_get_mesh()`; this is important for setting the correct DOF pointers on the macro triangulation.

Using index information from the file, all information concerning element vertices, neighbour relations can be calculated directly.

If the macro file gives information about boundary types, boundary types of vertices in 1d, of edges in 2d, and of faces in 3d are prescribed. Zero values correspond to an interior vertex/edge/face, positive values to an vertex/edge/face on the Dirichlet boundary and negative values to an vertex/edge/face on the Neumann boundary in 1d/2d/3d. Information for vertices in 2d, and for vertices and edges in 3d is derived from that information; we use the conventions that the Dirichlet boundary is a closed subset of the boundary. This leads to the following assignments in 2d and 3d:

1. a vertex belongs to the Dirichlet boundary if it is a vertex of one edge/face belonging to the Dirichlet boundary;
2. a vertex belongs to the Neumann boundary if it is a vertex of an edge/face belonging to the Neumann boundary and it is not a vertex of any edge/face on the Dirichlet boundary;
3. all other vertices belong to the interior.

The same holds accordingly for edges in 3d.

For boundary edges/faces a boundary structure has to be filled additionally; this is done by the third argument `ibdry` which is a pointer to a user defined function; it returns a pointer to a filled boundary structure, compare Section 3.2.5. The function is called for each boundary edge/face

```
(*ibdry)(mesh, bound)
```

where `mesh` is the structure to be filled and `bound` the boundary value read from the file for this edge/face; using this value, the user can choose inside `ibdry` the correct function for projecting nodes that will be created in this edge/face; if `ibdry` is a `nil` pointer, it is assumed that the domain is polygonal, no projection has to be done; the corresponding boundary pointers are adjusted to two default structures

```
const BOUNDARY dirichlet_boundary = {nil, DIRICHLET};
const BOUNDARY neumann_boundary = {nil, NEUMANN};
```

for positive respectively. negative values of `bound`.

In 3d we have to adjust a pointer to such a `BOUNDARY` structure also for edges. This is done by setting this pointer to the `BOUNDARY` structure of *one* of the meeting faces at that edge; if a `BOUNDARY` structure of a face supplies a function for the projection of a node onto a curved boundary during refinement, this function must be able to project any point in the closure of that face (because it may be used for any edge of that face); if the boundary type of the edge is `DIRICHLET`, then at least one the boundary types of the meeting faces is `DIRICHLET` also, and we will use the `BOUNDARY` structure of a `DIRICHLET` face.

During the initialization of the macro triangulation, other entries like `n_edges`, `n_faces`, and `max_edge_neigh` in the mesh data structure are calculated.

Example 3.8 (The standard triangulation of the unit interval in \mathbb{R}^1).


The easiest example is the macro triangulation for the interval $(0, 1)$ in 1d. We just have one element and two vertices.

```
DIM: 1
DIM_OF_WORLD: 1
```

```
number of elements: 1
number of vertices: 2
```

```
element vertices:
0 1
```

```
vertex coordinates:
0.0 0.0
1.0 0.0
```



Macro triangulation of the unit interval.

Example 3.9 (The standard triangulation of the unit square in \mathbb{R}^2).

Still rather simple is the macro triangulation for the unit square $(0, 1) \times (0, 1)$ in 2d. Here, we have two elements and four vertices. The refinement edge is the diagonal for both elements.

```


DIM: 2
DIM_OF_WORLD: 2

number of elements: 2
number of vertices: 4

element vertices:
2 0 1
0 2 3

vertex coordinates:
0.0 0.0
1.0 0.0
1.0 1.0
0.0 1.0

```



Macro triangulation of the unit square.

Example 3.10 (The standard triangulation of the unit cube in \mathbb{R}^3). More involved is already the macro triangulation for the unit cube $(0,1)^3$ in 3d. Here, we have eight vertices and six elements, all meeting at one diagonal; the shown specification of **element vertices** prescribes this diagonal as the refinement edge is for all elements.

```

DIM: 3
DIM_OF_WORLD: 3

number of vertices: 8
vertex coordinates:
-1.0 -1.0 -1.0
1.0 -1.0 -1.0
-1.0 -1.0 1.0
1.0 -1.0 1.0
1.0 1.0 -1.0
1.0 1.0 1.0
-1.0 1.0 -1.0
-1.0 1.0 1.0

number of elements: 6

element vertices:
0 5 4 1
0 5 3 1
0 5 3 2
0 5 4 6
0 5 7 6
0 5 7 2

```

Example 3.11 (A triangulation of three quarters of the unit disc).

Here, we describe a more complex example where we are dealing with a curved boundary and mixed type boundary condition. Due to the curved boundary, we have to provide a function for the projection during refinement in the boundary data structure. The actual projection is easy to implement, since we only have to normalize the coordinates for nodes belonging to the curved boundary. A pointer to this function is accessible inside `read_mesh()` by the `ibdry` function, which is the last argument to `read_mesh()`. Furthermore, we assume that the two straight edges belong to the Neumann boundary, and the curved boundary is the Dirichlet boundary. For handling mixed boundary types we have to specify `element boundaries` in the macro triangulation file. Information about `element boundaries` is also used inside the function `ibdry` for initializing routines for projecting the midpoints of the refinement edges onto the curved boundary.

```

DIM: 2
DIM_OF_WORLD: 2


number of vertices: 5
number of elements: 3

vertex coordinates:
0.0 0.0
1.0 0.0
0.0 1.0
-1.0 0.0
0.0 -1.0

element vertices:
1 2 0
2 3 0
3 4 0

element boundaries:
0 -1 2
0 0 2
-1 0 2

```



Macro triangulation of a 3/4 disc.

The functions `ball_proj()` for the projection and `ibdry()` for the initialization can be implemented as

```

static void ball_project(REAL_D p)
{
    REAL    norm = NORM_DOW(p);

    norm = 1.0/MAX(1.0E-15, norm);
    SCAL_DOW(norm, p);
}

```

```

    return;
}

const BOUNDARY *ibdry(MESH *mesh, int bound)
{
    FUNCNAME("ibdry");
    static const BOUNDARY curved_dirichlet = {ball_project, DIRICHLET};
    static const BOUNDARY straight_neumann = {nil, NEUMANN};

    switch(bound)
    {
        case 2: return(&curved_dirichlet);
        case -1: return(&straight_neumann);
        default: ERROR_EXIT("no boundary %d\n", bound);
    }
}

```

A binary data format allows faster import of a macro triangulation, especially when the macro triangulation consists already of many elements. Macro data written previously by binary `write_macro` routines (see below) can be read in native or machine independent binary format by the two routines

```

void read_macro_bin(MESH *, const char *,
                   const BOUNDARY (*)(MESH *, int ));
void read_macro_xdr(MESH *, const char *,
                   const BOUNDARY (*)(MESH *, int ));

```

Description:

`read_macro_bin(mesh, name, ibdry)`: reads data of the macro triangulation for `mesh` from the native binary file `name`; the file `name` was previously generated by the function `write_macro_bin()`, see below.

`read_macro_xdr(mesh, name, ibdry)`: reads data of the macro triangulation for `mesh` from the machine independent binary file `name`, the file `name` was previously generated by the function `write_macro_xdr()`, see below.

3.2.17 Writing macro triangulations

The counterpart of functions for reading macro triangulations are functions for writing macro triangulations to file. To be more general, it is possible to create a macro triangulation from the triangulation given by the leaf elements of a mesh. As mentioned above, it can be faster to use a binary format than the textual formal for writing and reading macro triangulations with many elements.

```

int write_macro(MESH *, const char *);
int write_macro_bin(MESH *, const char *);
int write_macro_xdr(MESH *, const char *);

```


Description:

`write_macro(mesh, name)`: writes the triangulation given by the leaf elements of `mesh` as a macro triangulation to the file specified by `name` in the above described format; if the file could be written, the return value is 1, otherwise an error message is produced and the return value is 0.

`write_macro_bin(mesh, name)`: writes the triangulation given by the leaf elements of `mesh` as a macro triangulation to the file specified by `name` in native binary format.

`write_macro_xdr(mesh, name)`: writes the triangulation given by the leaf elements of `mesh` as a macro triangulation to the file specified by `name` in machine independent binary format.

For exporting meshes including the whole hierarchy, see Section 3.3.8

3.2.18 Import and export of macro triangulations from/to other formats

When meshes are created using a simplicial grid generation tool, then data will usually not be in the ALBERTA macro triangulation format described above in Section 3.2.16. In order to simplify the import of such meshes, a vector-based data structure `MACRO_DATA` is provided. A vector-based data structure can easily be filled by an import routine; the filled data structure can then be converted into an ALBERTA mesh. The `MACRO_DATA` structure is defined as

```
typedef struct macro_data MACRO_DATA;

struct macro_data
{
    int n_total_vertices;
    int n_macro_elements;

    REAL_D *coords;
    int (*mel_vertices)[N_VERTICES];
    int (*neigh)[N_NEIGH];
    S_CHAR (*boundary)[N_NEIGH];
#ifdef DIM == 3
    U_CHAR *el_type;
#endif
};
```

The members yield following information:

`n_total_vertices`: number of vertices.

`n_macro_elements`: number of mesh elements.

`coords`: REAL array of size `[n_total_vertices][DIM_OF_WORLD]` holding the point coordinates of all vertices.

mel_vertices: integer array of size `[n_macro_elements][N_VERTICES]` storing element index information; `mel_vertices[i][j]` is the index of the *j*th vertex of element *i*.

neigh: integer array of size `[n_macro_elements][N_NEIGH]`, whereas element `neigh[i][j]` is the index of the *j*th neighbour element of element *i*, or -1 in case of a boundary.

boundary: `S_CHAR` array of size `[n_macro_elements][N_NEIGH]`, whereas element `boundary[i][j]` is the boundary type of the *j*th vertex/edge/face of element *i* (in 1d/2d/3d).

el_type: a `U_CHAR` vector of size `[n_macro_elements]` holding the element type of each mesh element (only 3d).

A `MACRO_DATA` structure can be allocated and freed by

```
MACRO_DATA *alloc_macro_data(int, int, FLAGS);
void free_macro_data(MACRO_DATA *);
```

Description:

`alloc_macro_data(n_vertices, n_elements, flags)`: allocates a structure `MACRO_DATA` together with all arrays needed to hold `n_vertices` vertices and `n_elements` mesh elements. The `coords` and `mel_vertices` arrays are allocated in any case, while `neigh`, `boundary` and `el_type` arrays are allocated only when requested as indicated by the corresponding flags `FILL_NEIGH`, `FILL_BOUNDARY`, and `FILL_EL_TYPE` set by a bitwise OR in `flags`.

`free_macro_data(data)`: completely frees all previously allocated storage for `MACRO_DATA` structure `data` including all vectors/arrays in it.

Once `MACRO_DATA` structure is filled, it can be saved to file in the ALBERTA macro triangulation format, or it can be directly be converted into a `MESH`.

```
void macro_data2mesh(MESH *, MACRO_DATA *,
                    const BOUNDARY *(*)(MESH *, int));
int write_macro_data(MACRO_DATA *, const char *);
int write_macro_data_bin(MACRO_DATA *, const char *);
int write_macro_data_xdr(MACRO_DATA *, const char *);
```

Description:

`macro_data2mesh(mesh, macro_data, bdry)`: converts data of a triangulation given in `macro_data` into a `MESH` structure. It sets most entries in `mesh`, allocates macro elements needed, assigns DOFs according to `mesh->n_dof`, and calculates `mesh->diam`. The coordinates in `macro_data->coords` are copied to a newly allocated array, thus the entire `MACRO_DATA` structure can thus be freed after calling this routine. When not `nil`, the `bdry` function is used to define the element boundaries.

`write_macro_data(macro_data, name)`: writes the macro triangulation with data stored in `macro_data` in the ALBERTA format described in Section

3.2.16 to file **name**. The return value is 0 when an error occurred and 1 in case the file was written successfully.

write_macro_data_bin(**macro_data**, **name**): writes data of the macro triangulation stored in **macro_data** in native binary format to file **name**; the return value is 0 when an error occurred and 1 in case the file was written successfully.

write_macro_data_xdr(**macro_data**, **name**): writes data of the macro triangulation stored in **macro_data** in machine independent binary format to file **name**; the return value is 0 when an error occurred and 1 in case the file was written successfully.

It is appropriate to check whether a macro triangulation given in a **MACRO_DATA** structure allows for recursive refinement, by testing for possible recursion cycles. An automatic correction by choosing other refinement edges may be done, currently implemented only in 2d.

```
void macro_test(MACRO_DATA *, const char *);
```

Description:

macro_test(**macro_data**, **name**): checks the triangulation given by the argument **macro_data** for potential cycles during recursive refinement. In the case that such a cycle is detected, the routine tries to correct this by renumbering element vertices (which is currently implemented only in 2d) and then writes the new, changed triangulation using **write_macro_data()** to a file **name**, when the second parameter is not **nil**.

3.2.19 Mesh traversal routines

As described before, the mesh is organized in a binary tree, and most local information is not stored at leaf element level, but is generated from hierarchical information and macro element data. The generation of such local information is done during tree traversal routines.

When some work has to be done at each tree element or leaf element, such a tree traversal is most easily done in a recursive way, calling some special subroutine at each (leaf) element which implements the operation that currently has to be done. For some other applications, it is necessary to operate on the (leaf) elements in another fashion, where a recursive traversal is not possible. To provide access for both situations, there exist recursive and non-recursive mesh traversal routines.

For both styles, selection criteria are available to indicate on which elements the operation should take place. The following constants are defined:

```
CALL EVERY_EL_PREORDER
CALL EVERY_EL_INORDER
CALL EVERY_EL_POSTORDER
CALL LEAF_EL
```

```

CALL_LEAF_EL_LEVEL
CALL_EL_LEVEL
CALL_MG_LEVEL

```

Choosing one of the flags `CALL_EVERY_EL_PREORDER`, `CALL_EVERY_EL_INORDER`, and `CALL_EVERY_EL_POSTORDER` operations are performed on *all* hierarchical elements of the mesh. These three differ in the sequence of operation on elements: `CALL_EVERY_EL_PREORDER` operates first on a parent element before operating on both children, `CALL_EVERY_EL_POSTORDER` operates first on both children before operating on their parent, and `CALL_EVERY_EL_INORDER` first operates on `child[0]`, then on the parent element, and last on `child[1]`.

`CALL_LEAF_EL` indicates to operate on *all* leaf elements of the tree, whereas `CALL_LEAF_EL_LEVEL` indicates to operate only on leaf elements which are exactly at a specified tree depth. `CALL_EL_LEVEL` operates on all tree elements at a specified tree depth. The option `CALL_MG_LEVEL` is special for multigrid operations. It provides the operation on all hierarchy elements on a specified multigrid level (which is usually `el->level/DIM`).

Additional flags are defined that specify which local information in `EL_INFO` has to be generated during the hierarchical mesh traversal. A bitwise OR of some of these constants is given as a parameter to the traversal routines. These flags are more or less self explaining:

`FILL_NOTHING`: no information needed at all.

`FILL_COORDS`: vertex coordinates `EL_INFO.coord` are filled.

`FILL_BOUND`: boundary information is filled in the entries `EL_INFO.bound` and `EL_INFO.boundary` (in 2d and 3d).

`FILL_NEIGH`: neighbour element information `EL_INFO.neigh` and corresponding `EL_INFO.opp_vertex` is generated (for `NEIGH_IN_EL == 0`).

`FILL_OPP_COORDS`: information about coordinates of the opposite vertex of neighbours is filled in `EL_INFO.opp_coords`; the flag `FILL_OPP_COORDS` can only be selected in combination with `FILL_COORDS|FILL_NEIGH`.

`FILL_ORIENTATION`: the element orientation info `EL_INFO.orientation` is generated (3d only). As mentioned in Section 3.2.7, the calculation of orientation may give wrong results in case of parametric meshes, when the parameterization is not preserving orientation.

`FILL_EL_TYPE`: the element type info is *always* generated, this flag is defined for compatible use as parameter to `get_macro_data()` or similar applications (3d only).

`FILL_ANY`: macro definition for a bitwise OR of any possible fill flags, used for separating the fill flags from the `CALL...` flags.

During mesh traversal, such information is generated hierarchically using the two subroutines

```

void fill_macro_info(MESH *, const MACRO_EL *, EL_INFO *);
void fill_elinfo(int, const EL_INFO *, EL_INFO *);

```

Description:

`fill_macro_info(mesh, mel, el_info)`: fills the `el_info` structure with macro element information of `mel` required by `el_info->flag` and sets `el_info->mesh` to `mesh`;

`fill_elinfo(ichild, parent_info, el_info)`: fills a given `el_info` structure for the child `ichild` using hierarchy information and parent data `parent_info` depending on `parent_info->flag`.

Sequence of visited elements

The sequence of elements which are visited during the traversal is given by the following rules:

- All elements in the binary mesh tree of one `MACRO_EL mel` are visited prior to any element in the tree of `mel->next`.
- For every `EL el`, all elements in the subtree `el->child[0]` are visited before any element in the subtree `el->child[1]`.
- The traversal order of an element and its two child trees is determined by the flags `CALL EVERY_EL_PREORDER`, `CALL EVERY_EL_INORDER`, and `CALL EVERY_EL_POSTORDER`, as defined above in Section 3.2.19.

Only during non-recursive traversal, this order may be changed by calling explicitly the `traverse_neighbour()` routine, see below.

Recursive mesh traversal routines

Recursive traversal of mesh elements is done by the routine

```
void mesh_traverse(MESH *, int, FLAGS, void (*)(const EL_INFO *));
```

Description:

`mesh_traverse(mesh, level, fill_flag, el_fct)`: traverses the triangulation `mesh`; the argument `level` specifies the element level if `CALL_EL_LEVEL` or `CALL_LEAF_EL_LEVEL`, or the multigrid level if `CALL_MG_LEVEL` is set in the `fill_flag`; otherwise this variable is ignored; by the argument `fill_flag` the elements to be traversed and data to be filled into `EL_INFO` is selected, using bitwise OR of one `CALL_...` flag and several `FILL_...` flags; the argument `el_fct` is a pointer to a function which is called on every element selected by the `CALL_...` part of `fill_flag`.

It is possible to use the recursive mesh traversal recursively, by calling `mesh_traverse()` from `el_fct`.

Example 3.12. An example of a mesh traversal is the computation of the measure of the computational domain. On each leaf element, the volume of the element is computed by the library function `el_volume()` and added to a global variable `measure_omega`, which finally holds the measure of the domain after the mesh traversal.

```

static REAL measure_omega;
static void measure_el(const EL_INFO *el_info)
{
    measure_omega += el_volume(el_info);
    return;
}

...

measure_omega = 0.0;
mesh_traverse(mesh, -1, CALL_LEAF_EL|FILL_COORDS, measure_el);
MSG("|Omega| = %e\n", measure_omega);

```

`el_volume()` computes the element volume and thus needs information about the elements vertex coordinates.

The only information passed to an element function like `measure_el()` is a pointer to the filled `el_info` structure. All other information needed by the element function has to be given by global variables, like `measure_omega` in this example.

Example 3.13. We give an implementation of the `CALL EVERY_EL...` routines to show the simple structure of all recursive traversal routines. A data structure `TRAVERSE_INFO`, only used by the traversal routines, holds the traversal flag and a pointer to the element function `el_fct()`:

```

void recursive_traverse(EL_INFO *el_info, TRAVERSE_INFO *trinfo)
{
    EL      *el = el_info->el;
    EL_INFO el_info_new;

    if (el->child[0])
    {
        if (trinfo->flag & CALL EVERY_EL_PREORDER)
            (trinfo->el_fct)(el_info);

        fill_elinfo(0, el_info, &el_info_new);
        recursive_traverse(&el_info_new, trinfo);

        if (trinfo->flag & CALL EVERY_EL_INORDER)
            (trinfo->el_fct)(el_info);

        fill_elinfo(1, el_info, &el_info_new);
        recursive_traverse(&el_info_new, trinfo);

        if (trinfo->flag & CALL EVERY_EL_POSTORDER)
            (trinfo->el_fct)(el_info);
    }
    else
    {

```

```

    (trinfo->el_fct)(el_info);
}
return;
}

void mesh_traverse_every_el(MESH *mesh, FLAGS fill_flag,
                           void (*el_fct)(const EL_INFO *))
{
    MACRO_EL      *mel;
    EL_INFO        el_info;
    TRAVERSE_INFO traverse_info;

    el_info.fill_flag = (flag & FILL_ANY);
    el_info.mesh = mesh;

    traverse_info.mesh = mesh;
    traverse_info.el_fct = el_fct;
    traverse_info.flag = flag;

    for (mel = mesh->first_macro_el; mel; mel = mel->next)
    {
        fill_macro_info(mel, &el_info);
        recursive_traverse(&el_info, &traverse_info);
    }
    return;
}

```

Non-recursive mesh traversal routines

As mentioned above in Example 3.12, all information needed by the element function `el_fct()`, besides data in the `el_info` structure, has to be given by global variables when using the recursive mesh traversal routines. Such a procedure can be done easier by using a non-recursive mesh traversal, where the element routine gets pointers to visited elements, one after another.

Additionally, mesh refinement and coarsening routines (for `NEIGH_IN_EL == 0`, see Sections 3.4.1 and 3.4.2), the `gltools` and `GRAPE` graphic interface (see Sections 3.16.2 and 3.16.3) are functions which need a non-recursive access to the mesh elements.

The implementation of the non-recursive mesh traversal routines uses a stack to save the tree path from a macro element to the current element. A data structure `TRAVERSE_STACK` holds such information. Before calling the non-recursive mesh traversal routines, such a stack must be allocated (and passed to the traversal routines).

```
typedef struct traverse_stack TRAVERSE_STACK;
```

By allocating a new stack, it is even possible to recursively call the non-recursive mesh traversal during another mesh traversal without destroying the stack which is already in use. For the non-recursive mesh traversal no pointer to an element function `el_fct()` has to be provided, because all operations are done by the routines which call the traversal functions. A mesh traversal is launched by each call to `traverse_first()` which also initializes the traverse stack. Advancing to the next element is done by the function `traverse_next()`. The following non-recursive routines are provided:

```
TRAVERSE_STACK *get_traverse_stack(void);
void free_traverse_stack(TRAVERSE_STACK *);
const EL_INFO *traverse_first(TRAVERSE_STACK *, MESH *, int, FLAGS);
const EL_INFO *traverse_next(TRAVERSE_STACK *, const EL_INFO *);
```

Description:

`get_traverse_stack()`: the return value is a pointer to a data structure `TRAVERSE_STACK`.

`free_traverse_stack(stack)`: frees the traverse stack `stack` previously accessed by `get_traverse_stack()`.

`traverse_first(stack, mesh, level, fill_flag)`: launches the non-recursive mesh traversal; the return value is a pointer to an `el_info` structure of the first element to be visited;

`stack` is a traverse stack previously accessed by `get_traverse_stack()`;

`mesh` is a pointer to a mesh to be traversed, `level` specifies the element level if `CALL_EL_LEVEL` or `CALL_LEAF_EL_LEVEL`, or the multigrid level if `CALL_MG_LEVEL` is set; otherwise this variable is ignored;

`fill_flag` specifies the elements to be traversed and data to be filled into `EL_INFO` is selected, using bitwise OR of one `CALL_...` flag and several `FILL_...` flags;

`traverse_next(stack, el_info)`: returns an `EL_INFO` structure with data about the next element of the mesh traversal or a pointer to `nil`, if `el_info->el` is the last element to be visited;

information which elements are visited and which data has to be filled is accessible via the traverse stack `stack`, initialized by `traverse_first()`. After calling `traverse_next()`, all `EL_INFO` information about previous elements is invalid, the structure may be overwritten with new data.

Usually, the interface to a graphical environment uses the non-recursive mesh traversal, compare the gltools (Section 3.16.2) and GRAPE interfaces (Section 3.16.3).

Example 3.14. The computation of the measure of the computational domain with the non-recursive mesh traversal routines is shown in the following code segment. No global variable is needed for storing information which is needed on elements.


```

REAL measure_omega(MESH *mesh)
{
    TRAVERSE_STACK *stack = get_traverse_stack();
    const EL_INFO *el_info;
    FLAGS          fill_flag;
    REAL           measure_omega = 0.0;

    el_info = traverse_first(stack, mesh, -1, CALL_LEAF_EL|FILL_COORDS);
    while (el_info)
    {
        measure_omega += el_volume(el_info);
        el_info = traverse_next(stack, el_info);
    }
    free_traverse_stack(stack);

    return(measure_omega);
}

```

Neighbour traversal

Some applications, like the recursive refinement algorithm, need the possibility to jump from one element to another element using neighbour relations. Such a traversal can not be performed by the recursive traversal routines and thus needs the non-recursive mesh traversal. The traversal routine for going from one element to a neighbour is

```
EL_INFO *traverse_neighbour(TRAVERSE_STACK *, EL_INFO *, int);
```

Description:

traverse_neighbour(stack, el_info, i): returns a pointer to a structure with **EL_INFO** information about the *i*-th neighbour opposite the *i*-th vertex of **el_info->el**;

The function can be called at any time during the non-recursive mesh traversal after initializing the first element by **traverse_first()**.

Calling **traverse_neighbour()**, all **EL_INFO** information about a previous element is completely lost. It can only be regenerated by calling **traverse_neighbour()** again with the *old* **OPP_VERTEX** value. If called at the boundary, when no adjacent element is available, then the routine returns **nil**; nevertheless, information from the old **EL_INFO** may be overwritten and lost. To avoid such behavior, one should check for boundary vertices/edges/faces (1d/2d/3d) before calling **traverse_neighbour()**.

Access to an element at world coordinates x

Some applications need the access to elements at a special location in world coordinates. Examples are characteristic methods for convection problems,

or the implementation of a special right hand side like point evaluations or curve integrals. In a characteristic method, the point x is usually given by $x = x_0 - \mathbf{V}\tau$, where x_0 is the starting point, \mathbf{V} the advection and τ the time step size. For points x_0 close to the boundary it may happen that x does not belong to the computational domain. In this situation it is convenient to know the point on the domain's boundary which lies on the line segment between the old point x_0 and the new point x . This point is uniquely determined by the scalar value s such that $x_0 + s(x - x_0) \in \partial\text{Domain}$.

The following function accesses an element at world coordinates x :

```
int find_el_at_pt(MESH *, const REAL_D, EL_INFO **, FLAGS,
                 REAL [DIM+1], const MACRO_EL *, const REAL_D,
                 REAL *);
```

Description:

`find_el_at_pt(mesh,x,el_info_p,fill_flag,bary,start_mel,x0,sp):`

fills element information in an `EL_INFO` structure and corresponding barycentric coordinates of the element where the point \mathbf{x} is located; the return value is `true` if \mathbf{x} is inside the domain, or `false` otherwise. Arguments of the function are:

`mesh` is the mesh to be traversed;

`x` are the world coordinates of the point (should be in the domain occupied by `mesh`);

`el_info_p` is the return address for a pointer to the `EL_INFO` for the element at \mathbf{x} (or when \mathbf{x} is outside the domain but `x0` was given, of the element containing the point $x_0 + s(x - x_0) \in \partial\text{Domain}$);

`fill_flag` are the flags which specify which information should be filled in the `EL_INFO` structure, coordinates are included in any case as they are needed by the routine itself;

`bary` is a pointer where to return the barycentric coordinates of \mathbf{x} on `*el_info_p->el` (or, when \mathbf{x} is outside the domain but `x0` was given, of the point $x_0 + s(x - x_0) \in \partial\text{Domain}$);

`start_mel` an initial guess for the macro element containing \mathbf{x} , or `nil`;

`x0` starting point of a characteristic method, see above, or `nil`;

`sp` return address for the relative distance to domain boundary in a characteristic method if `x0 != nil`, see above, or `nil`.

The implementation of `find_el_at_pt()` is based on the transformation from world to local coordinates, available via the routine `world_to_coord()`, compare Section 3.7. At the moment, `find_el_at_pt()` works correctly only for domains with non-curved boundary. This is due to the fact that the implementation first looks for the macro-element containing \mathbf{x} and then finds its path through the corresponding element tree based on the macro barycentric coordinates. For domains with curved boundary, it is possible that in some cases a point inside the domain is considered as external.

3.3 Administration of degrees of freedom

Degrees of freedom (DOFs) give connection between local and global finite element functions, compare Sections 1.4.2 and 1.3. We want to be able to have several finite element spaces and corresponding sets of DOFs at the same time. One set of DOFs may be shared between different finite element spaces, when appropriate.

During adaptive refinement and coarsening of a triangulation, not only elements of the mesh are created and deleted, but also degrees of freedom. The geometry is handled dynamically in a hierarchical binary tree structure, using pointers from parent elements to their children. For data corresponding to DOFs, which are usually involved with matrix-vector operations, simpler storage and access methods are more efficient. For that reason every DOF is realized just as an integer index, which can easily be used to access data from a vector or to build matrices that operate on vectors of DOF data.

During coarsening of the mesh, DOFs are deleted. In general, the deleted DOF is not the one which corresponds to the largest integer index. “Holes” with unused indices appear in the total range of used indices. One of the main aspects of the DOF administration is to keep track of all used and unused indices. One possibility to remove holes from vectors is the compression of DOFs, i.e. the renumbering of all DOFs such that all unused indices are shifted to the end of the index range, thus removing holes of unused indices. While the global index corresponding to a DOF may change, the *relative* order of DOF indices remains unchanged during compression.

During refinement of the mesh, new DOFs are added, and additional indices are needed. If a deletion of DOFs created some unused indices before, some of these can be reused for the new DOFs. Otherwise, the total range of used indices has to be enlarged, and the new indices are taken from this new range. At the same time, all vectors and matrices which are supposed to use these DOF indices have to be adjusted in size, too. This is the next major aspect of the DOF administration. To be able to do this, lists of vectors and matrices are included in the `DOF_ADMIN` data structure. Entries are added to or removed from these lists via special subroutines, see Section 3.3.2.

In ALBERTA, every abstract DOF is realized as an integer index into vectors:

```
typedef signed int    DOF;
```

These indices are administrated via the `DOF_ADMIN` data structure (see 3.3.1) and some subroutines. For each set of DOFs, one `DOF_ADMIN` structure is created. Degrees of freedom are directly connected with the mesh. The `MESH` data structure contains a reference to all sets of DOFs which are used on a mesh, compare Section 3.2.14. The `FE_SPACE` structure describing a finite element space references the corresponding set of DOFs, compare Sections 1.4.2, 3.5.1. Several `FE_SPACES` may share the same set of DOFs, thus reference the same `DOF_ADMIN` structure. Usually, a `DOF_ADMIN` structure is created during def-

inition of a finite element space by `get_fe_space()`, see Section 3.6.2. For special applications, additional DOF sets, that are not connected to any finite element space may also be defined (compare Section 3.6.2).

In Sections 3.3.5 and 3.3.6, we describe storage and access methods for global DOFs and local DOFs on single mesh elements.

As already mentioned above, special data types for data vectors and matrices are defined, see Sections 3.3.2 and 3.3.4. Several BLAS routines are available for such data, see Section 3.3.7.

3.3.1 The DOF_ADMIN data structure

The following data structure holds all data about one set of DOFs. It includes information about used and unused DOF indices, as well as linked lists of matrices and vectors of different data types, that are automatically resized and resorted during mesh changes. Currently, only an automatic *enlargement* of vectors is implemented, but no automatic shrinking. The actual implementation of used and unused DOFs is not described here in detail — it uses only one bit of storage for every integer index.

```
typedef struct dof_admin DOF_ADMIN;
typedef unsigned int DOF_FREE_UNIT;

struct dof_admin
{
    MESH          *mesh;
    const char    *name;

    DOF_FREE_UNIT *dof_free; /* flag bit vector */
    unsigned int  dof_free_size; /* flag bit vector size */
    unsigned int  first_hole; /* first non-zero dof_free entry */

    DOF size; /* allocated size of dof_list vector */
    DOF used_count; /* number of used dof indices */
    DOF hole_count; /* number of FREED dof indices */
    DOF size_used; /* > max. index of a used entry */

    int n_dof[DIM+1]; /* dofs from THIS dof_admin */
    int n0_dof[DIM+1]; /* dofs from previous dof_admins */

    DOF_INT_VEC *dof_int_vec; /* linked list of int vectors */
    DOF_DOF_VEC *dof_dof_vec; /* linked list of dof vectors */
    DOF_DOF_VEC *int_dof_vec; /* linked list of dof vectors */
    DOF_UCHAR_VEC *dof_uchar_vec; /* linked list of u_char vectors*/
    DOF_SCHAR_VEC *dof_schar_vec; /* linked list of s_char vectors*/
    DOF_REAL_VEC *dof_real_vec; /* linked list of real vectors */
    DOF_REAL_D_VEC *dof_real_d_vec; /* linked list of real_d vectors*/
    DOF_MATRIX *dof_matrix; /* linked list of matrices */
}
```

```

    void          *mem_info; /*--- don't touch! -----*/
};

```

The entries yield following information:

mesh: this is a `dof_admin` on `mesh`;

name: a string holding a textual description of this `dof_admin`;

dof_free, dof_free_size, first_hole: internally used variables for administration of used and free DOF indices;

size: current size of vectors in `dof*_vec` and `dof_matrix` lists;

used_count: number of used dof indices;

hole_count: number of *freed* dof indices (*not* `size-used_count`);

size_used: \geq largest used DOF index;

n_dof: numbers of degrees of freedom defined by this `dof_admin` structure; `n_dof[VERTEX]`, `n_dof[EDGE]`, `n_dof[FACE]`, and `n_dof[CENTER]` are the DOF counts at vertices, edges, faces (only in 3d) and element interiors, compare Section 3.3.6. These values are usually set by `get_fe_space()` as a copy from `bas_fcts->n_dof` (compare Section 3.5.1);

n0_dof: start indices `n0_dof[VERTEX/EDGE/FACE/CENTER]` of the first dofs defined by this `dof_admin` at vertices, edges (2d and 3d), faces (only 3d), respectively center in the element's `dof[VERTEX/EDGE/FACE/CENTER]` vectors. The values are the sums of all degrees of freedom defined by previous `dof_admin` structures on the same mesh, see Section 3.3.6 for details and usage; the values of `n0_dof[VERTEX/EDGE/FACE/CENTER]` are set automatically by the function `get_fe_space()`, compare Section 3.6.2;

dof*_vec, dof_matrix: pointers to linked lists of all used `DOF*_VEC` and `DOF_MATRIX` structures which are associated with the DOFs administrated by this `DOF_ADMIN` and whose size is automatically adjusted during mesh refinements, compare Section 3.3.2;

mem_info: used internally for memory management.

Deletion of DOFs occurs not only when the mesh is (locally) coarsened, but also during refinement of a mesh with higher order elements. This is due to the fact, that during local interpolation operations, both coarse-grid and fine-grid DOFs must be present, so deletion of coarse-grid DOFs that are no longer used is done after allocation of new fine-grid DOFs. Usually, all operations concerning DOFs are done automatically by routines doing mesh adaption or handling finite element spaces. The removal of “holes” in the range of used DOF indices is not done automatically. It is actually not *needed* to be done, but may speed up the access in loops over global DOFs; When there are no holes, then a simple `for`-loop can be used without checking for each index, whether it is currently in use or not. The `FOR_ALL_DOFS()`-macro described in Section 3.3.5 checks this case. Hole removal is done for all `DOF_ADMIN`s of a mesh by the function

```

void dof_compress(MESH *);

```

Description:

`dof_compress(mesh)`: remove all holes of unused DOF indices by compressing the used range of indices (it does *not* resize the vectors). While the global index corresponding to a DOF may change, the *relative* order of DOF indices remains unchanged during compression.

This routine is usually called after a mesh adaption involving higher order elements or coarsening.

Remark 3.15. Currently, the use of DOF matrices which combine two different sets of DOFs may produce errors during `dof_compress()`. Such matrices should be cleared by calling `clear_dof_matrix()` before a call to `dof_compress()`.

Usually, the range of DOF indices is enlarged in fixed increments given by the symbolic constant `SIZE_INCREMENT`, defined in `dof_admin.c`. If an estimate of the finally needed number of DOFs is available, then a direct enlargement of the DOF range to that number can be forced by calling:

```
void enlarge_dof_lists(DOF_ADMIN *, int);
```

Description:

`enlarge_dof_lists(admin, minsize)`: enlarges the range of the indices of `admin` to `minsize`.

3.3.2 Vectors indexed by DOFs: The `DOF*_VEC` data structures

The DOFs described above are just integers that can be used as indices into vectors and matrices. During refinement and coarsening of the mesh, the number of used DOFs, the meaning of one integer index, and even the total range of DOFs change. To be able to handle these changes automatically for all vectors, which are indexed by the DOFs, special data structures are used which contain such vector data. Lists of these structures are kept in the `DOF_ADMIN` structure, so that all vectors in the lists can be resized together with the range of DOFs. During refinement and coarsening of elements, values can be interpolated automatically to new DOFs, and restricted from old DOFs, see Section 3.3.3.

ALBERTA includes data types for vectors of type `REAL`, `REAL_D`, `S_CHAR`, `U_CHAR`, and `int`. Below, the `DOF_REAL_VEC` structure is described in full detail. Structures `DOF_REAL_D_VEC`, `DOF_SCHAR_VEC`, `DOF_UCHAR_VEC`, and `DOF_INT_VEC` are declared similarly, the only difference between them is the type of the structure entry `vec`.

Although the administration of such vectors is done completely by the DOF administration which needs `DOF_ADMIN` data, the following data structures include a reference to a `FE_SPACE`, which includes additionally the `MESH` and `BAS_FCTS`. In this way, complete information about a finite element function given by a `REAL`- and `REAL_D`-valued vector is directly accessible.

```

typedef struct dof_real_vec DOF_REAL_VEC;

struct dof_real_vec
{
    DOF_REAL_VEC  *next;
    const FE_SPACE *fe_space;

    const char    *name;
    DOF           size;
    REAL          *vec;    /* different type in DOF_INT_VEC, ... */

    void (*refine_interpol)(DOF_REAL_VEC *, RC_LIST_EL *, int n);
    void (*coarse_restrict)(DOF_REAL_VEC *, RC_LIST_EL *, int n);
};

```

The members yield following information:

next: linked list of DOF_REAL_VEC structures in `fe_space->admin`;

fe_space: FE_SPACE structure with information about DOFs and basis functions;

name: string with a textual description of vector values, or `nil`;

size: current size of `vec`;

vec: pointer to REAL vector of size `size`;

refine_interpol, coarse_restrict: interpolation and restriction routines, see Section 3.3.3. For REAL and REAL_D vectors, these usually point to the corresponding routines from `fe_space->bas_fcts`, compare Section 3.5.1. While we distinguish there between *restriction* and *interpolation* during coarsening, only one such operation is appropriate for a given vector, as it either represents a finite element function or values of a functional applied to basis functions.

All DOF vectors linked in the corresponding `dof_admin->dof*_vec` list are automatically adjusted in size and reordered during mesh changes. Values are transformed during local mesh changes, if the `refine_interpol` and/or `coarse_restrict` entries are not `nil`, compare Section 3.3.3.

Integer DOF vectors can be used in several ways: They may either hold an `int` value for each DOF, or reference a DOF value for each DOF. In both cases, the vectors should be automatically resized and rearranged during mesh changes. Additionally, values should be automatically changed in the second case. Such vectors are referenced in the `dof_admin->dof_int_vec` and `dof_admin->dof_dof_vec` lists.

On the other hand, DOF_INT_VECs provide a way to implement for special applications a vector of DOF values, which is *not indexed* by DOFs. For such vectors, only the values are automatically changed during mesh changes, but not the size or order. The user program is responsible for allocating memory for the `vec` vector. Such DOF vectors are referenced in the `dof_admin->int_dof_vec` list.

A macro `GET_DOF_VEC` is defined to simplify the secure access to a `DOF*_VEC`'s data. It assigns `dof_vec->vec` to `ptr`, if both `dof_vec` and `dof_vec->vec` are not `nil`, and generates an error in other cases:

```
#define GET_DOF_VEC(ptr, dof_vec)\
    TEST_EXIT((dof_vec)&&(ptr = (dof_vec)->vec))\
    ("%s == nil", (dof_vec) ? (dof_vec)->name : #dof_vec)
```

The following subroutines are provided to handle DOF vectors. Allocation of a new `DOF*_VEC` and freeing of a `DOF*_VEC` (together with its `vec`) are done with:

```
DOF_REAL_VEC    *get_dof_real_vec(const char *, const FE_SPACE *);
DOF_REAL_D_VEC  *get_dof_real_d_vec(const char *, const FE_SPACE *);
DOF_INT_VEC     *get_dof_int_vec(const char *, const FE_SPACE *);
DOF_INT_VEC     *get_dof_dof_vec(const char *, const FE_SPACE *);
DOF_INT_VEC     *get_int_dof_vec(const char *, const FE_SPACE *);
DOF_SCHAR_VEC   *get_dof_schar_vec(const char *, const FE_SPACE *);
DOF_UCHAR_VEC   *get_dof_uchar_vec(const char *, const FE_SPACE *);
void            free_dof_real_vec(DOF_REAL_VEC *);
void            free_dof_real_d_vec(DOF_REAL_D_VEC *);
void            free_dof_int_vec(DOF_INT_VEC *);
void            free_dof_dof_vec(DOF_INT_VEC *);
void            free_int_dof_vec(DOF_INT_VEC *);
void            free_dof_schar_vec(DOF_SCHAR_VEC *);
void            free_dof_uchar_vec(DOF_UCHAR_VEC *);
```

By specifying a finite element space for a `DOF*_VEC`, the corresponding set of DOFs is implicitly specified by `fe_space->admin`. The `DOF*_VEC` is linked into `DOF_ADMIN`'s appropriate `dof*_vec` list for automatic handling during mesh changes. The `DOF*_VEC` structure entries `next` and `admin` are set during creation and must not be changed otherwise! The size of the `dof_vec->vec` vector is automatically adjusted to the range of DOF indices controlled by `fe_space->admin`.

There is a special list for each type of DOF vectors in the `DOF_ADMIN` structure. All used `DOF_REAL_VECs`, `DOF_REAL_D_VECs`, `DOF_UCHAR_VECs`, and `DOF_SCHAR_VECs` are added to the respective lists, whereas a `DOF_INT_VEC` may be added to one of three lists in `DOF_ADMIN`: `dof_int_vec`, `dof_dof_vec`, and `int_dof_vec`. The difference between these three lists is their handling during a resize or compress of the DOF range. In contrast to all other cases, for a vector in `admin`'s `int_dof_vec` list, the `size` is NOT changed with `admin->size`. But the values `vec[i]`, $i = 1, \dots, \text{size}$ are adjusted when `admin` is compressed, for example. For vectors in the `dof_dof_vec` list, both adjustments in `size` and adjustment of values is done.

The `get*_vec()` routines automatically allocate enough memory for the data vector `vec` as indicated by `fe_space->admin->size`. Pointers to the routines `refine_interpol` and `coarse_restrict` are set to `nil`. They must be set explicitly after the call to `get*_vec()` for an interpolation during refinement and/or interpolation/restriction during coarsening. The `free*_vec()`

routines remove the vector from a `vec->fe_space->admin->dof_*_vec` list and free the memory used by `vec->vec` and `*vec`.

A printed output of DOF vector is produced by the routines:

```
void print_dof_int_vec(const DOF_INT_VEC *);
void print_dof_real_vec(const DOF_REAL_VEC *);
void print_dof_real_d_vec(const DOF_REAL_D_VEC *);
void print_dof_schar_vec(const DOF_SCHAR_VEC *);
void print_dof_uchar_vec(const DOF_UCHAR_VEC *);
```

Description:

`print_dof_*_vec(dof_vec)`: prints the elements of the DOF vector `dof_vec` together with its name to the message stream.

3.3.3 Interpolation and restriction of DOF vectors during mesh refinement and coarsening

During mesh refinement and coarsening, new DOFs are produced, or old ones are deleted. In many cases, information stored in `DOF_*_VECs` has to be adjusted to the new distribution of DOFs. To do this automatically during the refinement and coarsening process, each `DOF_*_VEC` can provide pointers to subroutines `refine_interpol` and `coarse_restrict`, that implements these operations on data. During refinement and coarsening of a `mesh`, these routines are called for all `DOF_*_VECs` with non-`nil` pointers in all `DOF_ADMINs` in `mesh->dof_admin`.

Before doing the mesh operations, it is checked whether any automatic interpolations or restrictions during refinement or coarsening are requested. If yes, then the corresponding operations will be performed during local mesh changes.

As described in Sections 3.4.1 and 3.4.2, interpolation resp. restriction of values is done during the mesh refinement and coarsening locally on every refined resp. coarsened patch of elements. Which of the local DOFs are created new, and which ones are kept from parent/children elements, is described in these other sections, too. All necessary interpolations or restrictions are done by looping through all `DOF_ADMINs` in `mesh` and calling the `DOF_*_VEC's` routines

```
struct dof_real_vec
{
    ...

    void (*refine_interpol)(DOF_REAL_VEC *, RC_LIST_EL *, int);
    void (*coarse_restrict)(DOF_REAL_VEC *, RC_LIST_EL *, int);
}
```

Those implement interpolation and restriction on one patch of mesh elements for this `DOF_*_VEC`. Only these have to know about the actual meaning of the DOFs. Here, `RC_LIST_EL` is a vector holding pointers to all `n` parent elements

which build the patch (and thus have a common refinement edge). Usually, the interpolation and restriction routines for `REAL` or `REAL_D` vectors are defined in the corresponding `dof_vec->fe_space->bas_fcts` structures. Interpolation or restriction of non-real values (`int` or `CHAR`) is usually application dependent and is not provided in the `BAS_FCTS` structure.

Examples of these routines are shown in Sections 3.5.4-3.5.7.

3.3.4 The `DOF_MATRIX` data structure

Not only vectors indexed by DOFs are available in `ALBERTA`, but also matrices which operate on these `DOF_*_VECS`. For finite element calculations, these matrices are usually sparse, and should be stored in a way that reflects this sparseness. We use a storage method which is similar to the one used in [53]. Every row of a matrix is realized as a linked list of `MATRIX_ROW` structures, each holding a maximum of `ROW_LENGTH` matrix entries from that row. Each entry consists of a column DOF index and the corresponding `REAL` matrix entry. Unused entries in a `MATRIX_ROW` are marked with a negative column index. The `ROW_LENGTH` is a symbolic preprocessor constant defined in `alberta.h`. For `DIM=2` meshes built from triangles, the refinement by bisection generates usually at most eight elements meeting at a common vertex, more elements may meet only at macro vertices. Thus, for piecewise linear (Lagrange) elements on triangles, up to nine entries are non-zero in most rows of a mass or stiffness matrix. This motivates the choice `ROW_LENGTH = 9`. For higher order elements or tetrahedra, there are much more non-zero entries in each row. Thus, a split of rows into short `MATRIX_ROW` parts should not produce too much overhead.

```
typedef struct matrix_row  MATRIX_ROW;
#define ROW_LENGTH 9

struct matrix_row
{
    MATRIX_ROW *next;
    DOF         col[ROW_LENGTH];
    REAL        entry[ROW_LENGTH];
};

#define ENTRY_USED(col)      ((col) >= 0)
#define ENTRY_NOT_USED(col) ((col) < 0)
#define UNUSED_ENTRY       -1
#define NO_MORE_ENTRIES    -2
```

Such a `DOF_MATRIX` structure is usually filled by local operations on single elements, using the `update_matrix()` routine, compare Section 3.12.1, which automatically generates space for new matrix entries by adding new `MATRIX_ROWS`, if needed.

An automatic adjustment of matrix entry values during mesh refinement and coarsening is not possible in the current implementation.

A `DOF_MATRIX` may also be used to combine two different sets of DOFs, compare Section 3.12.1. In the moment, the use of such matrices may produce errors during `dof_compress()`. Such a matrix should be cleared by calling `clear_dof_matrix()` before a call to `dof_compress()`.

```
typedef struct dof_matrix DOF_MATRIX;

struct dof_matrix
{
    DOF_MATRIX    *next;
    const FE_SPACE *fe_space;
    const char     *name;

    MATRIX_ROW    **matrix_row;
    DOF            size;
};
```

The entries yield following information:

`next`: linked list of `DOF_MATRIX` structures in `fe_space->admin`;
`fe_space`: `FE_SPACE` structure with information about corresponding row DOFs and basis functions;
`name`: a textual description for the matrix, or `nil`;
`matrix_row`: vector of pointers to `MATRIX_ROWS`, one for each row;
`size`: current size of the `matrix_row` vector.

The following routines are available for DOF-matrices:

```
DOF_MATRIX *get_dof_matrix(const char *, const FE_SPACE *);
void        free_dof_matrix(DOF_MATRIX *);
void        clear_dof_matrix(DOF_MATRIX *);
void        print_dof_matrix(const DOF_MATRIX *);
```

Description:

`get_dof_matrix(name, fe_space)`: allocates a new `DOF_MATRIX` structure for the finite element space `fe_space`; `name` is a textual description for the name of the new matrix; the new matrix is automatically linked into the `fe_space->admin->dof_matrix` list; a `matrix_row` vector of length `fe_space->admin->size` is allocated and all entries are set to `nil`;
`free_dof_matrix(matrix)`: frees the DOF matrix `matrix` previously accessed by the function `get_dof_matrix()`; in a first step, all `MATRIX_ROWS` in `matrix->matrix_row` are freed, then `matrix->matrix_row`, and finally the structure `*matrix`;
`clear_dof_matrix(matrix)`: clears all entries of the DOF matrix `matrix`; this is done by *removing* all entries from the DOF matrix, which means that all `MATRIX_ROWS` in `matrix->matrix_row` are freed and all entries in `matrix->matrix_row` are set to `nil`;
`print_dof_matrix(dof_matrix)`: prints the elements of `dof_matrix` together with its name to the message stream.

3.3.5 Access to global DOFs: Macros for iterations using DOF indices

For loops over all used (or free) DOFs, the following macros are defined:

```
FOR_ALL_DOFS(const DOF_ADMIN *, todo);
FOR_ALL_FREE_DOFS(const DOF_ADMIN *, todo);
```

Description:

FOR_ALL_DOFS(admin, todo): loops over all used DOFs of `admin`; `todo` is a list of C-statements which are to be executed for every used DOF index. During `todo`, the local variable `int dof` holds the current index of the used entry; it must not be altered by `todo`;

FOR_ALL_FREE_DOFS(admin, todo): loops over all unused DOFs of `admin`; `todo` is a list of C-statements which are to be executed for every unused DOF index. During `todo`, the local variable `int dof` holds the current index of the unused entry; it must not be altered by `todo`.

Two examples illustrate the usage of these macros.

Example 3.16 (Initialization of vectors). This BLAS-1 routine `dset()` initializes all elements of a vector with a given value; for `DOF_REAL_VECs` we have to set this value for all *used* DOFs. All used entries of the `DOF_REAL_VEC *drv` are set to a value `alpha` by:

```
FOR_ALL_DOFS(drv->fe_space->admin, drv->vec[dof] = alpha);
```

The BLAS-1 routine `dof_set()` is written this way, compare Section 3.3.7.

Example 3.17 (Matrix-vector multiplication). As a more complex example we give the main loop from an implementation of the matrix-vector product in `dof_mv()`, compare Sections 3.3.4 and 3.3.7:

```
FOR_ALL_DOFS(admin,
    sum = 0.0;
    for (row = a->matrix_row[dof]; row; row = row->next)
    {
        for (j=0; j<ROW_LENGTH; j++)
        {
            jcol = row->col[j];
            if (ENTRY_USED(jcol))
            {
                sum += row->entry[j] * xvec[jcol];
            }
            else
            {
                if (jcol == NO_MORE_ENTRIES)
                    break;
            }
        }
    }
}
```

```

        yvec[dof] = sum;
    );

```

3.3.6 Access to local DOFs on elements

As shown by the examples in Fig. 1.17, the DOF administration is able to handle different sets of DOFs, defined by different `DOF_ADMIN` structures, at the same time. All operations with finite element functions, like evaluation or integration, are done locally on the level of single elements. Thus, access on element level to DOFs from a single `DOF_ADMIN` has to be provided in a way that is independent from all other finite element spaces which might be defined on the mesh.

As described in Section 3.2.8, the `EL` data structure holds a vector of pointers to DOF vectors, that contain data for *all* DOFs on the element from all `DOF_ADMIN`s:

```

struct el
{
    ...
    DOF                **dof;
    ...
};

```

During initialization of a mesh, the lengths of these vectors are computed by collecting data from all `DOF_ADMIN`s associated with the mesh; details are given below. Information about all DOFs associated with a mesh is collected and accessible in the `MESH` data structure (compare Section 3.2.14):

```

struct mesh
{
    ...
    DOF_ADMIN **dof_admin;
    int        n_dof_admin;

    int        n_dof_el;
    int        n_dof [DIM+1];
    int        n_node_el;
    int        node [DIM+1];
    ...
};

```

The meaning of these entries is:

`dof_admin`: a vector of pointers to all `DOF_ADMIN` structures for the mesh;

`n_dof_admin`: number of all `DOF_ADMIN` structures for the mesh;

`n_dof_el`: total number of DOFs on one element from all `DOF_ADMIN` structures;

`n_dof`: total number of VERTEX, EDGE, FACE, and CENTER DOFs from all `DOF_ADMIN` structures;

n_node_el: number of used nodes on each element (vertices, edges, faces, and center), this gives the dimension of **el->dof**;

node: The entry **node[i]**, $i \in \{\text{VERTEX}, \text{EDGE}, \text{FACE}, \text{CENTER}\}$ gives the index of the first *i*-node in **el->dof**.

All these variables must not be changed by a user routine — they are set during the **init_dof_admins()** routine given as a parameter to **GET_MESH()**, compare Section 3.2.15. Actually, only the subroutine **get_fe_space()** is allowed to change such information (compare Section 3.6.2).

We denote the different locations of DOFs on an element by *nodes*. As there are DOFs connected with different-dimensional (sub-) simplices, there are *vertex*, *edge*, *face*, and *center* nodes. Using the symbolic constants from Section 3.2.2, there may be **N_VERTICES** vertex nodes, **N_EDGES** edge nodes (2d or 3d), **N_FACES** face nodes (in 3d), and one center node. Depending on the finite element spaces in use, not all possible nodes must be associated with DOFs, but some nodes may be associated with DOFs from several different finite element spaces (and several **DOF_ADMINs**). In order to minimize the memory usage for pointers and DOF vectors, the elements store data only for such nodes where DOFs are used. Thus, the number of nodes on an element is determined during mesh initialization, when all finite element spaces and DOFs are defined. The total number of nodes is stored in **mesh->n_node_el**, which will be the length of the **el->dof** vector for all elements.

In order to access the DOFs for one node, **mesh->node[1]** contains the index of the first 1-node in **el->dof**, where 1 is either **VERTEX**, **EDGE**, **FACE**, or **CENTER** (compare Fig. 3.2). So, a pointer to DOFs from the *i*-th edge node is stored at **el->dof[mesh->node[EDGE]+i]** ($0 \leq i < \text{N_EDGES}$), and these DOFs (and the vector holding them) are shared by all elements meeting at this edge.




Fig. 3.2. DOF vector indices in **el->dof** for DOFs at vertices, vertices and edges, vertices, edges and center, and vertices and center (in 2d). Corresponding **mesh->node** values are $\{0,0,0\}$, $\{0,3,0\}$, $\{0,3,6\}$, and $\{0,0,3\}$.

The total number of DOFs at an 1-node is available in **mesh->n_dof[1]**. This number is larger than zero, iff the node is in use. All DOFs from different **DOF_ADMINs** are stored together in one vector. In order to access DOFs from a given finite element space (and its associated **DOF_ADMIN**), the start index

for DOFs from this `DOF_ADMIN` must be known. This start index is generated during mesh initialization and stored in `admin->n0_dof[1]`. The number of DOFs from this `DOF_ADMIN` is given in `admin->n_dof[1]`. Thus, a loop over all DOFs associated with the i -th edge node can be done by:

```
DOF *dof_ptr = el->dof[mesh->node[EDGE]+i] + admin->n0_dof[EDGE];
for (j = 0 ; j < admin->n_dof[EDGE]; j++)
{
    dof = dof_ptr[j];
    ...
}
```

In order to simplify the access of DOFs for a finite element space on an element, the `BAS_FCTS` structure provides a routine

```
const DOF *(*get_dof_indices)(const EL *, const DOF_ADMIN *, DOF *);
```

which returns a vector containing all global DOFs associated with basis functions, in the correct order: the k -th DOF is associated with the k -th local basis function (compare Section 3.5.1).

3.3.7 BLAS routines for DOF vectors and matrices

Several basic linear algebra subroutines (BLAS [45, 28]) are implemented for DOF vectors and DOF matrices, see Table 3.1. Some non-standard routines are added: `dof_xpay()` is a variant of `dof_axpy()`, `dof_min()` and `dof_max()` calculate minimum and maximum values, and `dof_mv()` is a simplified version of the general `dof_gemv()` matrix-vector multiplication routine. The BLAS-2 routines `dof_gemv()` and `dof_mv()` use a `MatrixTranspose` flag: `transpose = NoTranspose = 0` indicates the original matrix, while `transpose = Transpose = 1` indicates the transposed matrix. We use the C-BLAS definition,

```
typedef enum {NoTranspose, Transpose, ConjugateTranspose}
            MatrixTranspose;
```

Similar routines are provided for `DOF_REAL_D` vectors, see Table 3.2.

3.3.8 Reading and writing of meshes and vectors

Section 3.2.16 described the input and output of ASCII files for macro triangulations. Locally refined triangulations including the mesh hierarchy and corresponding DOFs are saved in binary formats. Finite element data is saved (and restored) in binary format, too, in order to keep the full data precision. As the binary data and file format does usually depend on hardware and operating system, the interchange of data between different platforms needs a machine independent format. The XDR (External Data Representation) library provides a widely used interface for such a format. The `_xdr` routines

Table 3.1. Implemented BLAS routines for DOF vectors and matrices

REAL dof_nrm2(const DOF_REAL_VEC *x)	$nrm2 = (\sum X_i^2)^{1/2}$
REAL dof_asum(const DOF_REAL_VEC *x)	$asum = \sum X_i $
REAL dof_min(const DOF_REAL_VEC *x)	$min = \min X_i$
REAL dof_max(const DOF_REAL_VEC *x)	$max = \max X_i$
void dof_set(REAL alpha, DOF_REAL_VEC *x)	$X = (\alpha, \dots, \alpha)$
void dof_scal(REAL alpha, DOF_REAL_VEC *x)	$X = \alpha * X$
REAL dof_dot(const DOF_REAL_VEC *x, const DOF_REAL_VEC *y)	$dot = \sum X_i Y_i$
void dof_copy(const DOF_REAL_VEC *x, DOF_REAL_VEC *y)	$Y = X$
void dof_axpy(REAL alpha, const DOF_REAL_VEC *x, DOF_REAL_VEC *y)	$Y = \alpha * X + Y$
void dof_xpay(REAL alpha, const DOF_REAL_VEC *x, DOF_REAL_VEC *y)	$Y = X + \alpha * Y$
void dof_gemv(MatrixTranspose transpose, REAL alpha, const DOF_MATRIX *a, const DOF_REAL_VEC *x, REAL beta, DOF_REAL_VEC *y)	$Y = \alpha * A * X + \beta * Y$ or $Y = \alpha * A^t * X + \beta * Y$
void dof_mv(MatrixTranspose transpose, const DOF_MATRIX *a, const DOF_REAL_VEC *x, DOF_REAL_VEC *y)	$Y = A * X$ or $Y = A^t * X$

should be used whenever data must be transfered between different computer platforms.

```
int write_mesh(MESH *mesh, const char *name, REAL time);

MESH *read_mesh(const char *name, REAL *timeptr,
               void (*init_leaf_data)(LEAF_DATA_INFO *),
               const BOUNDARY *(*init_boundary)(MESH *, int));
```

The routine `write_mesh` stores information about the mesh in a file named `name`. Written data includes the corresponding `time` (only important for time dependent problems), macro elements, mesh elements including the parent/child hierarchy information, DOF administration and element DOFs. The return value is 1 if an error occurs, otherwise 0.

Routine `read_mesh` reads a complete mesh from file `name`, which was created by `write_mesh`. The corresponding time, if any, is stored at `*timeptr`. The arguments `init_leaf_data` and `init_boundary` are used in the same way as in `read_macro`, compare Section 3.2.16.

For input and output of finite element data, the following routines are provided which read or write files containing binary DOF vectors.

Table 3.2. Implemented BLAS routines for DOF_REAL_D vectors

REAL dof_nrm2_d(const DOF_REAL_D_VEC *x)	$nrm2 = (\sum X_i ^2)^{1/2}$
REAL dof_min_d(const DOF_REAL_D_VEC *x)	$min = \min X_i $
REAL dof_max_d(const DOF_REAL_D_VEC *x)	$max = \max X_i $
void dof_set_d(REAL alpha, DOF_REAL_D_VEC *x)	$X = ((\alpha, \dots, \alpha), \dots)$
void dof_scal_d(REAL alpha, DOF_REAL_D_VEC *x)	$X = \alpha * X$
REAL dof_dot_d(const DOF_REAL_D_VEC *x, const DOF_REAL_D_VEC *y)	$dot = \sum (X_i, Y_i)$
void dof_copy_d(const DOF_REAL_D_VEC *x, DOF_REAL_D_VEC *y)	$Y = X$
void dof_axpy_d(REAL alpha, const DOF_REAL_D_VEC *x, DOF_REAL_D_VEC *y)	$Y = \alpha * X + Y$
void dof_xpay_d(REAL alpha, const DOF_REAL_D_VEC *x, DOF_REAL_D_VEC *y)	$Y = X + \alpha * Y$
void dof_gemv(MatrixTranspose transpose, REAL alpha, const DOF_MATRIX *a, const DOF_REAL_D_VEC *x, REAL beta, DOF_REAL_D_VEC *y)	$Y = \alpha * A * X + \beta * Y$ or $Y = \alpha * A^t * X + \beta * Y$
void dof_mv_d(MatrixTranspose transpose, const DOF_MATRIX *a, const DOF_REAL_D_VEC *x, DOF_REAL_D_VEC *y)	$Y = A * X$ or $Y = A^t * X$

```

int write_dof_int_vec(const DOF_INT_VEC *div, const char *name);
int write_dof_real_vec(const DOF_REAL_VEC *drv, const char *name);
int write_dof_real_d_vec(const DOF_REAL_D_VEC *drdv, const char *);
int write_dof_schar_vec(const DOF_SCHAR_VEC *dsv, const char *);
int write_dof_uchar_vec(const DOF_UCHAR_VEC *duv, const char *);

DOF_INT_VEC    *read_dof_int_vec(const char *name, MESH *, FE_SPACE*);
DOF_REAL_VEC   *read_dof_real_vec(const char *, MESH *, FE_SPACE *);
DOF_REAL_D_VEC *read_dof_real_d_vec(const char *, MESH *, FE_SPACE *);
DOF_SCHAR_VEC  *read_dof_schar_vec(const char *, MESH *, FE_SPACE *);
DOF_UCHAR_VEC  *read_dof_uchar_vec(const char *, MESH *, FE_SPACE *);

```

For the output and input of machine independent data files, similar routines are provided. The XDR library is used, and all routine names end with `_xdr`:

```

int write_mesh_xdr(MESH *mesh, const char *name, REAL time);

MESH *read_mesh_xdr(const char *name, REAL *timeptr,
void (*init_leaf_data)(LEAF_DATA_INFO *),
const BOUNDARY *(*init_boundary)(MESH *, int));

```

```

int write_dof_int_vec_xdr(const DOF_INT_VEC *, const char *name);
int write_dof_real_vec_xdr(const DOF_REAL_VEC *, const char *name);
int write_dof_real_d_vec_xdr(const DOF_REAL_D_VEC *, const char *);
int write_dof_schar_vec_xdr(const DOF_SCHAR_VEC *, const char *name);
int write_dof_uchar_vec_xdr(const DOF_UCHAR_VEC *, const char *name);

DOF_INT_VEC      *read_dof_int_vec_xdr(const char *, MESH *, FE_SPACE*);
DOF_REAL_VEC     *read_dof_real_vec_xdr(const char *, MESH *,
                                         FE_SPACE *);
DOF_REAL_D_VEC   *read_dof_real_d_vec_xdr(const char *, MESH *,
                                         FE_SPACE *);
DOF_SCHAR_VEC    *read_dof_schar_vec_xdr(const char *, MESH *,
                                         FE_SPACE *);
DOF_UCHAR_VEC    *read_dof_uchar_vec_xdr(const char *, MESH *,
                                         FE_SPACE *);

```

Remark 3.18. Currently, the functions for reading and writing meshes in a binary fashion, `read_mesh(_xdr)()` and `write_mesh(_xdr)()`, are only available for `NEIGH_IN_EL==0`.

3.4 The refinement and coarsening implementation

3.4.1 The refinement routines

For the refinement of a mesh the following symbolic constant is defined and the refinement is done by the functions

```

#define MESH_REFINED    1

U_CHAR refine(MESH *);
U_CHAR global_refine(MESH *, int);

```

Description:

refine(mesh): refines all leaf elements with a *positive* element marker **mark** times (this mark is usually set by some adaptive procedure); the routine loops over all leaf elements and refines the elements with a positive marker until there is no element left with a positive marker; the return value is `MESH_REFINED`, if at least one element was refined, and 0 otherwise. Every refinement has to be done via this routine. The basic steps of this routine are described below.

global_refine(mesh, mark): sets all element markers for leaf elements of **mesh** to **mark**; the mesh is then refined by **refine()** which results in a **mark** global refinement of the mesh; the return value is `MESH_REFINED`, if **mark** is positive, and 0 otherwise.

Basic steps of the refinement algorithm

The refinement of a mesh is principally done in two steps. In the first step no coordinate information is available on the elements. In the case that neighbour information is stored at the elements such coordinate information can not be produced when going from one neighbour to another by the neighbour pointers. Thus, only a topological refinement is performed. If new nodes are created on the boundary these can be projected onto a curved boundary in the second step when coordinate information is available.

Again using the notion of “refinement edge” for the element itself in 1d, the algorithm performs the following steps:

1. The whole mesh is refined only topologically. This part consists of
 - the collection of a compatible refinement patch; this includes the recursive refinement of adjacent elements with an incompatible refinement edge;
 - the topological bisection of the patch elements;
 - the transformation of leaf data from parent to child, if such a function is available in the `leaf_data_info` structure;
 - allocation of new DOFs;
 - handing on of DOFs from parent to the children;
 - interpolation of DOF vectors from the coarse grid to the fine one on the whole refinement patch, if the function `refine_interpol()` is available for these DOF vectors (compare Section 3.3.3); these routines must not use coordinate information;
 - a deallocation of DOFs on the parent when `preserve_coarse_dofs == 0`.

This process is described in detail below.

2. New nodes which belong to the curved part of the boundary are now projected onto the curved boundary via the `param_bound()` function in the `BOUNDARY` structure of the refinement edge. The coordinates of the projected node are stored in a `REAL_D`-vector and the pointers `e1->new_coord` of all parents `e1` which belong to the refinement patch are set to this vector. This step is only done in 2d and 3d. Fig. 3.3 shows some refinements of a triangle with one edge on the curved boundary. The projections of refinement edge midpoints (small circles) to the curved boundary are shown by the black dots.

The topological refinement is done by the recursive refinement Algorithm 1.5. In 1d, no recursion is needed. In 2d and 3d, all elements at the refinement edge of a marked element are collected. If a neighbour with an incompatible refinement edge is found, this neighbour is refined first by a recursive call of the refinement function. Thus, after looping around the refinement edge, the patch of simplices at this edge is always a compatible refinement patch. The elements of this patch are stored in a vector `ref_list` with elements of type `RC_LIST_EL`, compare Section 3.2.13. This vector is an argument for the




Fig. 3.3. Refinement at curved boundary: refinement edge midpoints \circ are projected by `param_bound()` to the curved boundary \bullet

functions for interpolation of DOF vectors during refinement, compare Section 3.3.3.

In 1d the vector has length 1. In 2d the length is 2 if the refinement edge is an interior edge; for a boundary edge the length is 1 since only the element itself has to be refined. For 1d and 2d, only the `el` entry of the components is set and used.

In 3d this vector is allocated with length `mesh->max_edge_neigh`. As mentioned in Section 3.2.13 we can define an orientation of the edge and by this orientation we can define the right and left neighbours (inside the patch) of an element at this edge.

The patch is bisected by first inserting a new vertex at the midpoint of the refinement edge. Then all elements of the refinement patch are bisected. This includes the allocation of new DOFs, the adjustment of DOF pointers, and the memory allocation for leaf data (if the leaf data size is positive) and transformation of leaf data from parent to child (if a pointer to a function `refine_leaf_data()` is provided by the user in the `mesh->leaf_data_info` structure). Then memory for parents' leaf data is freed and information stored there is definitely lost.

In the case of higher order elements we also have to add new DOFs on the patch and if we do not need information about the higher order DOFs on coarser levels they are removed from the parents. There are some basic rules for adding and removing DOFs which are important for the prolongation and restriction of data (see Section 3.3.3):

1. Only DOFs of the same kind (i.e. `VERTEX`, `EDGE`, or `FACE`) and whose nodes have the same geometrical position on parent and child are handed on to this child from the parent;

2. DOFs at a vertex, an edge or a face belong to all elements sharing this vertex, edge, face, respectively;
3. DOFs on the parent are only removed if the entry `preserve_coarse_dofs` in the mesh data structure is `false`; in that case only DOFs which are not handed on to a child are removed on the parent.

A direct consequence of 1. is that only DOFs inside the patch are added or removed; DOFs on the patch boundary stay untouched. **CENTER** DOFs can not be handed from parent to child since the centers of the parent and the children are always at different positions.

Using standard Lagrange finite elements, only DOFs that are not handed from parent to child have to be set while interpolating a finite element function to the finer grid; all values of the other DOFs stay the same (the same holds during coarsening and interpolating to the coarser grid).

Due to 2. it is clear that DOFs shared by more than one element have to be allocated only once and pointers to these DOFs are set correctly for all elements sharing it.

Now, we take a closer look at DOFs that are handed on by the parents and those that have to be allocated: In 1d we have

```
child[0]->dof[0] = el->dof[0];
child[1]->dof[1] = el->dof[1];
```

in 2d

```
child[0]->dof[0] = el->dof[2];
child[0]->dof[1] = el->dof[0];
child[1]->dof[0] = el->dof[1];
child[1]->dof[1] = el->dof[2];
```

In 3d for `child[1]` this passing of DOFs additionally depends on the element type `el_type` of the parent. For `child[0]` we always have

```
child[0]->dof[0] = el->dof[0];
child[0]->dof[1] = el->dof[2];
child[0]->dof[2] = el->dof[3];
```

For `child[1]` and a parent of type 0 we have

```
child[1]->dof[0] = el->dof[1];
child[1]->dof[1] = el->dof[3];
child[1]->dof[2] = el->dof[2];
```

and for a parent of type 1 or 2

```
child[1]->dof[0] = el->dof[1];
child[1]->dof[1] = el->dof[2];
child[1]->dof[2] = el->dof[3];
```

In 1d

```
child[0]->dof[1] = child[1]->dof[0]
```

and in 2d and 3d

```
child[0]->dof[DIM] = child[1]->dof[DIM]
```

is the newly allocated DOF at the midpoint of the refinement edge (compare Fig. 1.4 on page 13 for the 1d and 2d situation and Fig. 1.5 on page 13 for the 3d situation).

In the case that we have DOFs at the midpoint of edges (only 2d and 3d) the following DOFs are passed on (let `enode = mesh->node[EDGE]` be the offset for DOFs at edges): for 2d

```
child[0]->dof[enode+2] = el->dof[enode+1];
child[1]->dof[enode+2] = el->dof[enode+0];
```

and for 3d

```
child[0]->dof[enode+0] = el->dof[enode+1];
child[0]->dof[enode+1] = el->dof[enode+2];
child[0]->dof[enode+3] = el->dof[enode+5];
```

for `child[0]` a for `child[1]` of a parent of type 0

```
child[1]->dof[enode+0] = el->dof[enode+4];
child[1]->dof[enode+1] = el->dof[enode+3];
child[1]->dof[enode+3] = el->dof[enode+5];
```

and finally for `child[1]` of a parent of type 1 or 2

```
child[1]->dof[enode+0] = el->dof[enode+3];
child[1]->dof[enode+1] = el->dof[enode+4];
child[1]->dof[enode+3] = el->dof[enode+5];
```




Fig. 3.4. Edge DOFs that are freed \bullet , passed on \circ , and newly allocated \square

We also have to create new DOFs (compare Fig. 3.4). Two additional DOFs are created in the refinement edge which are shared by all patch elements. Pointers to these DOFs are adjusted for

```
child[0]->dof[enode+0],
child[1]->dof[enode+1]
```

in 2d and

```
child[0]->dof[enode+2],
child[1]->dof[enode+2]
```

in 3d for all patch elements. In 3d, for each interior face of the refinement patch there is a new edge where we have to add a new DOF vector. These DOFs are shared by two children in the case of a boundary face; otherwise it is shared by four children and pointers of

```
child[0]->dof[enode+4] = child[1]->dof[enode+{5,4,4}],
child[0]->dof[enode+5] = child[1]->dof[enode+{4,4,5}]
```

are adjusted for those elements.

In 3d, there may be also DOFs at faces; the face DOFs in the boundary of the patch are passed on (let `fnode = mesh->node[FACE]` be the offset for DOFs at faces):

```
child[0]->dof[fnode+3] = el->dof[fnode+1];
child[1]->dof[fnode+3] = el->dof[fnode+0];
```

For the common face of `child[0]` and `child[1]` we have to allocate a new face DOF vector which is located at

```
child[0]->dof[fnode+0] = child[1]->dof[fnode+0]
```

and finally for each interior face of the patch two new face DOF vectors are created and pointers for adjacent children are adjusted:

```
child[0]->dof[fnode+1],
child[0]->dof[fnode+2],
child[1]->dof[fnode+1],
child[1]->dof[fnode+2]
```

Each of these DOF vectors may be shared with another child of a patch element.

If DOFs are located at the barycenter they have to be allocated for both children in 2d and 3d (let `cnode = mesh->node[CENTER]` be the offset for DOFs at the center)

```
child[0]->dof[cnode],
child[1]->dof[cnode].
```

After adding and passing on of DOFs on the patch we can interpolate data from the coarse to the fine grid on the whole patch. This is an operation on the whole patch since new DOFs can be shared by more than one patch element and usually the value(s) of such a DOF should only be calculated once.

All DOF vectors where a pointer to a function `refine_interpol()` in the corresponding `DOF_*_VEC` data structure is provided are interpolated to the fine grid. Such a function does essentially depend on the described passing on and new allocation of DOFs. An abstract description of such functions can be found in Section 1.4.4 and a more detailed one for Lagrange elements in Section 3.5.2.

After such an interpolation, DOFs of higher degree on parent elements may be no longer of interest (when not using a higher order multigrid method). In such a case the entry `preserve_coarse_dofs` in the mesh data structure

has to be **false** and all DOFs on the parent that are not handed over to the children will be removed. The following DOFs are removed on the parent for all patch elements (some DOFs are shared by several elements): The DOFs at the center

```
e1->dof[mesh->node[CENTER]]
```

are removed in all dimensions. In 2d, additionally DOFs in the refinement edge

```
e1->dof[mesh->node[EDGE]+2]
```

are removed and in 3d the DOFs in the refinement edge and the DOFs in the two faces adjacent to the refinement edge

```
e1->dof[mesh->node[EDGE]+0],
e1->dof[mesh->node[FACE]+2],
e1->dof[mesh->node[FACE]+3],
e1->dof[mesh->node[CENTER]]
```

are deleted on the parent. All pointers to DOFs at edges, faces and centers are set to **nil** on the parent. No information about these DOFs is available on interior tree elements in this case. Again we want to point out that for the geometrical description of the mesh we do not free vertex DOFs and all pointers to vertex DOFs stay untouched on the parent elements.

This setting of DOF pointers and pointers to children (and adjustment of adjacency relations when **NEIGH_IN_EL == 1**) is the main part of the refinement module.

If neighbour information is produced by the traversal routines, then it is valid for all tree elements. If it is stored explicitly, then neighbour information is valid for all leaf elements with all neighbours; but for interior elements of the tree this is not the case, it is only valid for those neighbours that belong to the common refinement patch!

3.4.2 The coarsening routines

For the coarsening of a mesh the following symbolic constant is defined and the coarsening is done by the functions

```
#define MESH_COARSENEED 2

U_CHAR coarsen(MESH *);
U_CHAR global_coarsen(MESH *, int);
```

Description:

coarsen(mesh): tries to coarsen all leaf element with a *negative* element marker **|mark|** times (again, this mark is usually set by an adaptive procedure); the return value is **MESH_COARSENEED** if any element was coarsened, and 0 otherwise.

`global_coarsen(mesh, mark)`: sets all element markers for leaf elements of `mesh` to `mark`; the mesh is then coarsened by `coarsen()`; depending on the actual distribution of coarsening edges on the mesh, this may not result in a `|mark|` global coarsening; the return value is `coarsen(mesh)` if `mark` is negative, and 0 otherwise.

The function `coarsen()` implements Algorithm 1.10. For a marked element, the coarsening patch is collected first. This is done in the same manner as in the refinement procedure. If such a patch can *definitely* not be coarsened (if one element of the patch may not be coarsened, e.g.) all coarsening markers for all patch elements are reset. If we can not coarsen the patch immediately, because one of the elements has not a common coarsening edge but is allowed to be coarsened more than once, then nothing is done in the moment and we try to coarsen this patch later on (compare Remark 1.11).

The coarsening of a patch is the “inverse” of the refinement of a compatible patch. If DOFs of the parents were removed during refinement, these are now added on the parents. Pointers to those that have been passed on to the children are now adjusted back on the parents (see Section 3.4.1 which DOFs of children are now assigned to the parent, just swap the left and right hand sides of the assignments). DOFs that were freed have to be newly allocated.

If leaf data is stored at the pointer of `child[1]`, then memory for the parent’s leaf data is allocated. If a function `coarsen_leaf_data` is provided in the `mesh->leaf_data_info` structure then leaf data is transformed from children to parent. Finally, leaf data on both children is freed.

Similar to the interpolation of data during refinement, we now can restrict or interpolate data from children to parent. This is done by the `coarse_restrict()` functions for all those DOF vectors where such a function is available in the corresponding `DOF_*_VEC` data structure. Since it does not make sense to both interpolate and restrict data, `coarse_restrict()` may be a pointer to a function either for interpolation or restriction. An abstract description of those functions can be found in Section 1.4.4 and a more detailed one for Lagrange elements in Section 3.5.2.

After these preliminaries the main part of the coarsening can be performed. DOFs that have been created in the refinement step are now freed again, and the children of all patch elements are freed and the pointer to the first child is set to `nil` and the pointer to the second child is adjusted to the `leaf_data` of the parent, or also set to `nil`. Thus, all fine grid information is lost at that moment, which makes clear that a restriction of data has to be done in advance.

3.5 Implementation of basis functions

In order to construct a finite element space, we have to specify a set of local basis functions. We follow the concept of finite elements which are given on a single element S in local coordinates: Finite element functions on an element

S are defined by a finite dimensional function space $\bar{\mathbb{P}}$ on a reference element \bar{S} and the (one to one) mapping $\lambda^S : \bar{S} \rightarrow S$ from the reference element \bar{S} to the element S . In this situation the non vanishing basis functions on an arbitrary element are given by the set of basis functions of $\bar{\mathbb{P}}$ in local coordinates λ^S . Also, derivatives are given by the derivatives of basis functions on $\bar{\mathbb{P}}$ and derivatives of λ^S .

Each local basis function on S is uniquely connected to a global degree of freedom, which can be accessed from S via the DOF administration. Together with this DOF administration and the underlying mesh, the finite element space is given. In the following section we describe the basic data structures for storing basis function information.

In ALBERTA Lagrange finite elements up to order four are implemented; they are presented in the subsequent sections.

3.5.1 Data structures for basis functions

For the handling of local basis functions, i.e. a basis of the function space $\bar{\mathbb{P}}$ (compare Section 1.4.2) we use the following data structures:

```
typedef REAL      BAS_FCT(const REAL [DIM+1]);
typedef const REAL *GRD_BAS_FCT(const REAL [DIM+1]);
typedef const REAL (*D2_BAS_FCT(const REAL [DIM+1])) [DIM+1];
```

Description:

BAS_FCT: the data type for a local finite element function, i.e. a function $\bar{\varphi} \in \bar{\mathbb{P}}$, evaluated at barycentric coordinates $\lambda \in \mathbb{R}^{\text{DIM}+1}$ and its return value $\bar{\varphi}(\lambda)$ is of type **REAL**.

GRD_BAS_FCT: the data type for the gradient (with respect to λ) of a local finite element function, i.e. a function returning a pointer to $\nabla_\lambda \bar{\varphi}$ for some function $\bar{\varphi} \in \bar{\mathbb{P}}$:

$$\nabla_\lambda \bar{\varphi}(\lambda) = \left(\frac{\partial \bar{\varphi}(\lambda)}{\partial \lambda_0}, \dots, \frac{\partial \bar{\varphi}(\lambda)}{\partial \lambda_{\text{DIM}}} \right);$$

the arguments of such a function are barycentric coordinates and the return value is a pointer to a **const REAL** vector of length **[DIM+1]** storing $\nabla_\lambda \bar{\varphi}(\lambda)$; this vector will be overwritten during the next call of the function.

D2_BAS_FCT: the data type for the second derivatives (with respect to λ) of a local finite element function, i.e. a function returning a pointer to the matrix $D_\lambda^2 \bar{\varphi}$ for some function $\bar{\varphi} \in \bar{\mathbb{P}}$:

$$D_\lambda^2 \bar{\varphi} = \begin{pmatrix} \frac{\partial^2 \bar{\varphi}(\lambda)}{\partial \lambda_0 \partial \lambda_0} & \cdots & \frac{\partial^2 \bar{\varphi}(\lambda)}{\partial \lambda_0 \partial \lambda_{\text{DIM}}} \\ \vdots & & \vdots \\ \frac{\partial^2 \bar{\varphi}(\lambda)}{\partial \lambda_{\text{DIM}} \partial \lambda_0} & \cdots & \frac{\partial^2 \bar{\varphi}(\lambda)}{\partial \lambda_{\text{DIM}} \partial \lambda_{\text{DIM}}} \end{pmatrix};$$

the arguments of such a function are barycentric coordinates and the return value is a pointer to a $(\text{DIM} + 1) \times (\text{DIM} + 1)$ matrix storing $D_\lambda^2 \bar{\varphi}$; this matrix will be overwritten during the next call of the function.

For the implementation of a finite element space, we need a basis of the function space $\bar{\mathbb{P}}$. For such a basis we need the connection of *local* and *global* DOFs on each element (compare Section 1.4.3), information about the interpolation of a given function on an element, and information about interpolation/restriction of finite element functions during refinement/coarsening (compare Section 1.4.4). Such information is stored in the BAS_FCTS data structure:

```
typedef struct bas_fcts BAS_FCTS;

struct bas_fcts
{
    char          *name;
    int           n_bas_fcts;
    int           degree;
    int           n_dof[DIM+1];

    void          (*init_element)(const EL_INFO *, const FE_SPACE *,
                                  U_CHAR);

    BAS_FCT       **phi;
    GRD_BAS_FCT   **grd_phi;
    D2_BAS_FCT    **D2_phi;

    const DOF      *(*get_dof_indices)(const EL *, const DOF_ADMIN *,
                                       DOF *);
    const S_CHAR   *(*get_bound)(const EL_INFO *, S_CHAR *);

    /*----- entries must be set for interpolation -----*/

    const REAL     *(*interpol)(const EL_INFO *, int, const int *,
                                REAL (*)(const REAL_D), REAL *);
    const REAL_D   *(*interpol_d)(const EL_INFO *, int, const int *b_no,
                                  const REAL (*)(const REAL_D, REAL_D),
                                  REAL_D *);

    /*----- optional entries -----*/

    const int      *(*get_int_vec)(const EL *, const DOF_INT_VEC *,
                                   int *);
    const REAL     *(*get_real_vec)(const EL *, const DOF_REAL_VEC *,
                                    REAL *);
    const REAL_D   *(*get_real_d_vec)(const EL *, const DOF_REAL_D_VEC *,
                                      REAL_D *);
    const U_CHAR   *(*get_uchar_vec)(const EL *, const DOF_UCHAR_VEC *,
                                     U_CHAR *);
}
```

```

const S_CHAR *(*get_schar_vec)(const EL *, const DOF_SCHAR_VEC *,
                               S_CHAR *);

void (*real_refine_inter)(DOF_REAL_VEC *, RC_LIST_EL *, int);
void (*real_coarse_inter)(DOF_REAL_VEC *, RC_LIST_EL *, int);
void (*real_coarse_restr)(DOF_REAL_VEC *, RC_LIST_EL *, int);

void (*real_d_refine_inter)(DOF_REAL_D_VEC *, RC_LIST_EL *, int);
void (*real_d_coarse_inter)(DOF_REAL_D_VEC *, RC_LIST_EL *, int);
void (*real_d_coarse_restr)(DOF_REAL_D_VEC *, RC_LIST_EL *, int);

void *bas_fcts_data;
};

```

The entries yield following information:

name: string containing a textual description or `nil`.

n_bas_fcts: number of local basis functions.

degree: maximal polynomial degree of the basis functions; this entry is used by routines using numerical quadrature where no `QUAD` structure is provided; in such a case via **degree** some default numerical quadrature is chosen (see Section 3.8.1); additionally, **degree** is used by some graphics routines (see Section 251).

n_dof: vector with the count of DOFs for this set of basis functions;

n_dof[VERTEX(,EDGE(,FACE)),CENTER] count of DOFs at the vertices, edges (only 2d and 3d), faces (only 3d), and the center of an element; the corresponding DOF administration of the finite element space uses such information.

init_element: is used for the initialization of element dependent local finite element spaces; is not used or supported by `ALBERTA` in this version;

phi: vector of function pointers for the evaluation of local basis functions in barycentric coordinates;

(*phi[i])(lambda) returns the value $\bar{\varphi}^i(\lambda)$ of the i -th basis function at **lambda** for $0 \leq i < \text{n_bas_fcts}$.

grd_phi: vector of function pointers for the evaluation of gradients of the basis functions in barycentric coordinates;

(*grd_phi[i])(lambda) returns a pointer to a vector of length `DIM+1` containing all first derivatives (with respect to the barycentric coordinates) of the i -th basis function at **lambda**, i.e. **(*grd_phi[i])(lambda)[k]** = $\bar{\varphi}_{,\lambda_k}^i(\lambda)$ for $0 \leq k \leq \text{DIM}$, $0 \leq i < \text{n_bas_fcts}$; this vector is overwritten on the next call of **(*grd_phi[i])()**.

D2_phi: vector of function pointers for the evaluation of second derivatives of the basis functions in barycentric coordinates;

(*D2_phi[i])(lambda) returns a pointer to a matrix of size $(\text{DIM}+1) \times (\text{DIM}+1)$ containing all second derivatives (with respect to the barycentric coordinates) of the i -th local basis function at point **lambda**, i.e.

`(*D2_phi[i])(lambda)[k][l] = $\bar{\varphi}_{\lambda_k \lambda_l}^i(\lambda)$` $0 \leq k, l \leq \text{DIM}$, $0 \leq i < \text{n_bas_fcts}$; this matrix is overwritten on the next call of `(*D2_phi[i])()`.

`get_dof_indices`: pointer to a function which connects the set of local basis functions with its global DOFs (an implementation of the function j_S in Section 1.4.3);

`get_dof_indices(el, admin, dof)` returns a pointer to a `const` DOF vector of length `n_bas_fcts` where the i -th entry is the index of the DOF associated to the i -th basis function; arguments are the actual element `el` and the DOF admin `admin` of the corresponding finite element space `fe_space` (these indices depend on all defined DOF admins and thus on the corresponding `admin`); if the last argument `dof` is `nil`, `get_dof_indices()` has to provide memory for storing this vector, which is overwritten on the next call of `get_dof_indices()`; if `dof` is not `nil`, `dof` is a pointer to a vector which has to be filled.

`get_bound`: pointer to a function which fills a vector with the boundary types of the basis functions;

`get_bound(el_info, bound)` returns a pointer to this vector of length `n_bas_fcts` where the i -th entry is the boundary type of the i -th basis function; `bound` may be a pointer to a vector which has to be filled (compare the `dof` argument of `get_dof_indices()`);

such a function needs boundary information; thus, all routines using this function on the elements need the `FILL_BOUND` flag during mesh traversal.

When using ALBERTA routines for the interpolation of `REAL(D)` valued functions the `interpol(D)` function pointer must be set (for example the calculation of Dirichlet boundary values by `dirichlet_bound()` described in Section 3.12.5):

`interpol(D)`: pointer to a function which performs the local interpolation of a `REAL(D)` valued function on an element;

`interpol(D)(el_info, n, indices, f, coeff)` returns a pointer to a `const REAL(D)` vector with interpolation coefficients of the `REAL(D)` valued function `f`; if `indices` is a pointer to `nil`, the coefficients for all basis functions are calculated and the i -th entry in the vector is the coefficient of the i -th basis function; if `indices` is non `nil`, only the coefficients for a subset of the local basis functions have to be calculated; `n` is the number of those basis functions, `indices[0], ..., indices[n-1]` are the local indices of the basis functions where the coefficients have to be calculated, and the i -th entry in the return vector is then the coefficient of the `indices[i]`-th basis function; `coeff` may be a pointer to a vector which has to be filled (compare the `dof` argument of `get_dof_indices()`);

such a function usually needs vertex coordinate information; thus, all routines using this function on the elements need the `FILL_COORDS` flag during mesh traversal.

Optional entries:

get*_vec: pointer to a function which fills a local vector with values of a **DOF*_VEC** at the DOFs of the basis functions;

get*_vec(*el*, *dof*_vec*, * *values*) returns a pointer to a **const *** vector where the *i*-th entry is the value of **dof*_vec** at the DOF of the *i*-th basis function on element *el*; *values* may be a pointer to a vector which has to be filled (compare the **dof** argument of **get_dof_indices()**), when *values*==**nil** then the vector will be overwritten during the next call of **get*_vec()**.

Since the interpolation of finite element functions during refinement and coarsening, as well as the restriction of functionals during coarsening, strongly depend on the basis functions and its DOFs (compare Section 1.4.4), pointers for functions which perform those operations can be stored at the following entries:

real(_d)_refine_inter: pointer to a function for interpolating a **REAL(_D)** valued function during refinement;

real(_d)_refine_inter(*vec*, *refine_list*, *n*) interpolates a given vector *vec* of type **DOF_REAL(_D)_VEC** on the refinement patch onto the finer grid; information about all parents of the refinement patch is accessible in the vector *refine_list* of length *n*.

real(_d)_coarse_inter: pointer to a function for interpolating a **REAL(_D)** valued function during coarsening;

real(_d)_coarse_inter(*vec*, *refine_list*, *n*) interpolates a given vector *vec* of type **DOF_REAL(_D)_VEC** on the coarsening patch onto the coarser grid; information about all parents of the refinement patch is accessible in the vector *refine_list* of length *n*.

real(_d)_coarse_restr: pointer to a function for restriction of **REAL(_D)** valued linear functionals during coarsening;

real(_d)_coarse_restr(*vec*, *refine_list*, *n*) restricts a given vector *vec* of type **DOF_REAL(_D)_VEC** on the coarsening patch onto the coarser grid; information about all parents of the refinement patch is accessible in the vector *refine_list* of length *n*.

Finally, there is an optional pointer for storing internal data:

bas_fcts_data: pointer to internal data used by the basis functions, like Lagrange nodes in barycentric coordinates for Lagrange elements, e. g.

In Section 3.5.4 and 3.5.5 examples for the implementation of those functions are given.

Remark 3.19. The access of local element vectors via the **get*_vec()** routines can also be done in a standard way by using the **get_dof_indices()** function which must be supplied; if some of the **get*_vec()** are **nil** pointer in a new basis-functions data structure, ALBERTA fills in pointers to some

standard functions using `get_dof_indices()`. But a specialized function may be faster. An example of such a standard routine is:

```
const int *get_int_vec(const EL *el, const DOF_INT_VEC *vec,
                      int *ivec)
{
    FUNCNAME("get_int_vec");
    int i, n_bas_fcts;
    const DOF *dof;
    int *v = nil, *rvec;
    const FE_SPACE *fe_space = vec->fe_space;
    static int *local_vec = nil;
    static int max_size = 0;

    GET_DOF_VEC(v, vec);
    n_bas_fcts = fe_space->bas_fcts->n_bas_fcts;
    if (ivec)
    {
        rvec = ivec;
    }
    else
    {
        if (max_size < n_bas_fcts)
        {
            local_vec = MEM_REALLOC(local_vec, max_size, n_bas_fcts, int);
            max_size = n_bas_fcts;
        }
        rvec = local_vec;
    }
    dof = fe_space->bas_fcts->get_dof_indices(el, fe_space->admin, nil);

    for (i = 0; i < n_bas_fcts; i++)
        rvec[i] = v[dof[i]];

    return((const int *) rvec);
}
```

A specialized implementation for linear finite elements e.g. is more efficient:

```
const int *get_int_vec(const EL *el, const DOF_INT_VEC *vec,
                      int *ivec)
{
    FUNCNAME("get_int_vec");
    int i, n0;
    static int local_vec[N_VERTICES];
    int *v = vec->vec, *rvec = ivec ? ivec : local_vec;
    DOF **dof = el->dof;

    n0 = vec->fe_space->admin->n0_dof[VERTEX];
```

```

    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = v[dof[i][n0]];

    return((const int *) rvec);
}

```

Any kind of basis functions can be implemented by filling the above described structure for basis functions. All non optional entries have to be defined. Since in the functions for reading and writing of meshes, the basis functions are identified by their names, all used basis functions have to be registered before using these functions (compare Section 3.3.8). All Lagrange finite elements described below are already registered, with names "**lagrange0**" to "**lagrange4**"; newly defined basis functions must use different names.

```
int new_bas_fcts(const BAS_FCTS * bas_fcts);
```

Description:

new_bas_fcts(bas_fcts): puts the new set of basis functions **bas_fcts** to an internal list of all used basis functions; different sets of basis functions are identified by their **name**; thus, the member **name** of **bas_fcts** must be a string with positive length holding a description; if an existing set of basis functions with the same name is found, the program stops with an error; if the entries **phi**, **grd_phi**, **get_dof_indices**, and **get_bound** are not set, this also result in an error and the program stops.

Basis functions can be accessed from that list by

```
const BAS_FCTS *get_bas_fcts(const char *name)
```

Description:

get_bas_fcts(name): looks for a set of basis functions with name **name** in the internal list of all registered basis functions; if such a set is found, the return value is a pointer to the corresponding **BAS_FCTS** structure, otherwise the return value is **nil**.

Lagrange elements can be accessed by **get_lagrange()**, see Section 3.5.8.

3.5.2 Lagrange finite elements

ALBERTA provides Lagrange finite elements up to order four which are described in the following sections. Lagrange finite elements are given by $\bar{\mathbb{P}} = \mathbb{P}_p(\bar{S})$ (polynomials of degree $p \in \mathbb{N}$ on \bar{S}) and they are globally continuous. They are uniquely determined by the values at the associated Lagrange nodes $\{x_i\}$. The Lagrange basis functions $\{\phi_i\}$ satisfy

$$\phi_i(x_j) = \delta_{ij} \quad \text{for } i, j = 1, \dots, N = \dim X_h.$$

Now, consider the basis functions $\{\bar{\varphi}^i\}_{i=1}^m$ of $\bar{\mathbb{P}}$ with the associated Lagrange nodes $\{\lambda_i\}_{i=1}^m$ given in barycentric coordinates:

$$\bar{\varphi}^i(\lambda_j) = \delta_{ij} \quad \text{for } i, j = 1, \dots, m.$$

Basis functions are located at the vertices, (edges, faces,) or at the center of an element. The corresponding DOF is a vertex, (edge, face,) or center DOF, respectively. The boundary type of a basis function is the boundary type of the associated vertex (or edge or face). Basis functions at the center are always INTERIOR. Such boundary information is filled by the `get_bound()` function in the `BAS_FCTS` structure and is straight forward.

The interpolation coefficient for a function f for basis function $\bar{\varphi}^i$ on element S is the value of f at the Lagrange node: $f(x(\lambda_i))$. These coefficients are calculated by the `interpol(_d)()` function in the `BAS_FCTS` structure. Examples for both functions are given below for linear finite elements.

3.5.3 Piecewise constant finite elements

Piecewise constant, discontinuous finite elements are uniquely defined by their values on the elements of the triangulation. For all dimensions we have exactly one (constant) basis function on each element, where the corresponding Lagrange node is the barycenter.

3.5.4 Piecewise linear finite elements

Table 3.3. Local basis functions for linear finite elements in 1d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \lambda_0$	vertex 0	$\lambda^0 = (1, 0)$
$\bar{\varphi}^1(\lambda) = \lambda_1$	vertex 1	$\lambda^1 = (0, 1)$

Table 3.4. Local basis functions for linear finite elements in 2d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0)$
$\bar{\varphi}^1(\lambda) = \lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0)$
$\bar{\varphi}^2(\lambda) = \lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1)$

Piecewise linear, continuous finite elements are uniquely defined by their values at the vertices of the triangulation. On each element we have `N_VERTICES` basis functions which are the barycentric coordinates of the element. Thus, in 1d we have two, in 2d we have three, and in 3d four basis functions for Lagrange elements of first order; the basis functions and the corresponding Lagrange nodes in barycentric coordinates are shown in Tables 3.3, 3.4 and 3.5. The calculation of derivatives is straight forward.


Table 3.5. Local basis functions for linear finite elements in 3d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0, 0)$
$\bar{\varphi}^1(\lambda) = \lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0, 0)$
$\bar{\varphi}^2(\lambda) = \lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1, 0)$
$\bar{\varphi}^3(\lambda) = \lambda_3$	vertex 3	$\lambda^3 = (0, 0, 0, 1)$

The global DOF index of the i -th basis functions on element `e1` is stored for linear finite elements at

```
e1->dof[i][admin->n0_dof[VERTEX]]
```

Setting `nv = admin->n0_dof[VERTEX]` the associated DOFs are shown in Fig. 3.5.

**Fig. 3.5.** DOFs and local numbering of the basis functions for linear elements in 2d and 3d.

For linear finite elements we want to give examples for the implementation of some routines in the corresponding `BAS_FCTS` structure.

Example 3.20 (Accessing DOFs for piecewise linear finite elements).

The implementation of `get_dof_indices()` can be done in the following way, compare Fig. 3.5 and Remark 3.19 with the implementation of the function `get_int_vec()` for accessing a local element vector from a global `DOF_INT_VEC` for piecewise linear finite elements.

```
const DOF *get_dof_indices(const EL *el, const DOF_ADMIN *admin,
                           DOF *idof)
{
    FUNCNAME("get_dof_indices");
    static DOF index_vec[N_VERTICES];
    DOF      *rvec = idof ? idof : index_vec;
    int      i, n0 = admin->n0_dof[VERTEX];
    DOF      **dof = el->dof;
```

```

    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = dof[i][n0];

    return((const DOF *) rvec);
}

```

Example 3.21 (Accessing the boundary types of DOFs for piecewise linear finite elements). The `get_bound()` function fills the bound vector with the boundary type of the vertices:

```

const S_CHAR *get_bound(const EL_INFO *el_info, S_CHAR *bound)
{
    FUNCNAME("get_bound");
    static S_CHAR bound_vec[N_VERTICES];
    S_CHAR      *rvec = bound ? bound : bound_vec;
    int          i;

    TEST_FLAG(FILL_BOUND, el_info);

    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = el_info->bound[i];

    return((const S_CHAR *) rvec);
}

```

Example 3.22 (Interpolation for piecewise linear finite elements). For the interpolation `interpol()` routine we have to evaluate the given function at the vertices. Thus, interpolation can be implemented as follows:

```

const REAL *interpol(const EL_INFO *el_info, int n, const int *dofs,
                    REAL (*f)(const REAL_D), REAL *vec)
{
    FUNCNAME("interpol");
    static REAL inter[N_VERTICES];
    REAL      *rvec = vec ? vec : inter;
    int        i;

    TEST_FLAG(FILL_COORDS, el_info);

    if (dofs)
    {
        if (n <= 0 || n > N_VERTICES)
        {
            MSG("something is wrong, doing nothing\n");
            rvec[0] = 0.0;
            return((const REAL *) rvec);
        }
    }
}

```

```

    for (i = 0; i < n; i++)
        rvec[i] = f(el_info->coord[dofs[i]]);
}
else
{
    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = f(el_info->coord[i]);
}
return((const REAL *) rvec);
}

```

Example 3.23 (Interpolation and restriction routines for piecewise linear finite elements). The implementation of functions for interpolation during refinement and restriction of linear functionals during coarsening is very simple for linear elements; we do not have to loop over the refinement patch since only the vertices at the refinement/coarsening edge and the new DOF at the midpoint are involved in this process. No interpolation during coarsening has to be done since all values at the remaining vertices stay the same; no function has to be defined.

```

void real_refine_inter(DOF_REAL_VEC *drv, RC_LIST_EL *list, int n)
{
    FUNCNAME("real_refine_inter");
    EL      *el;
    REAL     *vec = nil;
    DOF      dof_new, dof0, dof1;
    int      n0;

    if (n < 1) return;

    GET_DOF_VEC(vec, drv);
    n0 = drv->fe_space->admin->n0_dof[VERTEX];
    el = list->el;
    dof0 = el->dof[0][n0];      /* 1st endpoint of refinement edge */
    dof1 = el->dof[1][n0];      /* 2nd endpoint of refinement edge */
    dof_new = el->child[0]->dof[DIM][n0]; /* newest vertex is DIM */
    vec[dof_new] = 0.5*(vec[dof0] + vec[dof1]);

    return;
}

void real_coarse_restr(DOF_REAL_VEC *drv, RC_LIST_EL *list, int n)
{
    FUNCNAME("real_coarse_restr");
    EL      *el;
    REAL     *vec = nil;
    DOF      dof_new, dof0, dof1;
    int      n0;

```

```

if (n < 1) return;

GET_DOF_VEC(vec, drv);
n0 = drv->fe_space->admin->n0_dof[VERTEX];
el = list->el;
dof0 = el->dof[0][n0];      /* 1st endpoint of refinement edge */
dof1 = el->dof[1][n0];      /* 2nd endpoint of refinement edge */
dof_new = el->child[0]->dof[DIM][n0]; /* newest vertex is DIM */
vec[dof0] += 0.5*vec[dof_new];
vec[dof1] += 0.5*vec[dof_new];

return;
}

```

3.5.5 Piecewise quadratic finite elements

Table 3.6. Local basis functions for quadratic finite elements in 1d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \lambda_0(2\lambda_0 - 1)$	vertex 0	$\lambda^0 = (1, 0)$
$\bar{\varphi}^1(\lambda) = \lambda_1(2\lambda_1 - 1)$	vertex 1	$\lambda^1 = (0, 1)$
$\bar{\varphi}^2(\lambda) = 4\lambda_0 \lambda_1$	center	$\lambda^2 = (\frac{1}{2}, \frac{1}{2})$


Table 3.7. Local basis functions for quadratic finite elements in 2d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \lambda_0(2\lambda_0 - 1)$	vertex 0	$\lambda^0 = (1, 0, 0)$
$\bar{\varphi}^1(\lambda) = \lambda_1(2\lambda_1 - 1)$	vertex 1	$\lambda^1 = (0, 1, 0)$
$\bar{\varphi}^2(\lambda) = \lambda_2(2\lambda_2 - 1)$	vertex 2	$\lambda^2 = (0, 0, 1)$
$\bar{\varphi}^3(\lambda) = 4\lambda_1 \lambda_2$	edge 0	$\lambda^3 = (0, \frac{1}{2}, \frac{1}{2})$
$\bar{\varphi}^4(\lambda) = 4\lambda_2 \lambda_0$	edge 1	$\lambda^4 = (\frac{1}{2}, 0, \frac{1}{2})$
$\bar{\varphi}^5(\lambda) = 4\lambda_0 \lambda_1$	edge 2	$\lambda^5 = (\frac{1}{2}, \frac{1}{2}, 0)$

Piecewise quadratic, continuous finite elements are uniquely defined by their values at the vertices and the edges' midpoints (center in 1d) of the triangulation. In 1d we have three, in 2d we have six, and in 3d we have ten basis functions for Lagrange elements of second order; the basis functions and

Table 3.8. Local basis functions for quadratic finite elements in 3d.

function	position	Lagrange node
$\varphi^0(\lambda) = \lambda_0(2\lambda_0 - 1)$	vertex 0	$\lambda^0 = (1, 0, 0, 0)$
$\varphi^1(\lambda) = \lambda_1(2\lambda_1 - 1)$	vertex 1	$\lambda^1 = (0, 1, 0, 0)$
$\varphi^2(\lambda) = \lambda_2(2\lambda_2 - 1)$	vertex 2	$\lambda^2 = (0, 0, 1, 0)$
$\varphi^3(\lambda) = \lambda_3(2\lambda_3 - 1)$	vertex 3	$\lambda^3 = (0, 0, 0, 1)$
$\varphi^4(\lambda) = 4\lambda_0 \lambda_1$	edge 0	$\lambda^4 = (\frac{1}{2}, \frac{1}{2}, 0, 0)$
$\varphi^5(\lambda) = 4\lambda_0 \lambda_2$	edge 1	$\lambda^5 = (\frac{1}{2}, 0, \frac{1}{2}, 0)$
$\varphi^6(\lambda) = 4\lambda_0 \lambda_3$	edge 2	$\lambda^6 = (\frac{1}{2}, 0, 0, \frac{1}{2})$
$\varphi^7(\lambda) = 4\lambda_1 \lambda_2$	edge 3	$\lambda^7 = (0, \frac{1}{2}, \frac{1}{2}, 0)$
$\varphi^8(\lambda) = 4\lambda_1 \lambda_3$	edge 4	$\lambda^8 = (0, \frac{1}{2}, 0, \frac{1}{2})$
$\varphi^9(\lambda) = 4\lambda_2 \lambda_3$	edge 5	$\lambda^9 = (0, 0, \frac{1}{2}, \frac{1}{2})$

**Fig. 3.6.** DOFs and local numbering of the basis functions for quadratic elements in 2d and 3d.

the corresponding Lagrange nodes in barycentric coordinates are shown in Tables 3.6, 3.7, and 3.8.

The associated DOFs for basis functions at vertices/edges are located at the vertices/edges of the element; the entry in the vector of DOF indices at the vertices/edges is determined by the vertex/edge offset in the corresponding `admin` of the finite element space: the DOF index of the i -th basis functions on element `e1` is

```
e1->dof[i][admin->n0_dof[VERTEX]]
```

for $i = 0, \dots, N_VERTICES-1$ and

```
e1->dof[i][admin->n0_dof[EDGE]]
```

for $i = N_VERTICES, \dots, N_VERTICES+N_EDGES-1$. Here we used the fact, that for quadratic elements DOFs are located at the vertices and the edges on the mesh. Thus, regardless of any other set of DOFs, the offset `mesh->node[VERTEX]` is zero and `mesh->node[EDGE]` is `N_VERTICES`.

Setting `nv = admin->n0_dof[VERTEX]` and `ne = admin->n0_dof[EDGE]`, the associated DOFs are shown in Fig. 3.6.

Example 3.24 (Accessing DOFs for piecewise quadratic finite elements). The function `get_dof_indices()` for quadratic finite elements can be implemented in 2d and 3d by (compare Fig. 3.6):

```
const DOF *get_dof_indices(const EL *el, const DOF_ADMIN *admin,
                          DOF *idof)
{
    static DOF index_vec[N_VERTICES+N_EDGES];
    DOF        *rvec = idof ? idof : index_vec, **dof = el->dof;
    int         i, n0 = admin->n0_dof[VERTEX];

    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = dof[i][n0];
    n0 = admin->n0_dof[EDGE];
    for (i = N_VERTICES; i < N_VERTICES+N_EDGES; i++)
        rvec[i] = dof[i][n0];

    return((const DOF *) rvec);
}
```

The boundary type of a basis functions at a vertex is the the boundary type of the vertex, and the boundary type of a basis function at an edge is the boundary type of the edge. The i -th interpolation coefficient of a function f on element S is just $f(x(\lambda_i))$. The implementation is similar to that for linear finite elements and is not shown here.

The implementation of functions for interpolation during refinement and coarsening and the restriction during coarsening becomes more complicated and differs between the dimensions. Here we have to set values for all elements of the refinement patch. The interpolation during coarsening is not trivial anymore. As an example of such implementations we present the interpolation during refinement for 2d and 3d.

Example 3.25 (Interpolation during refinement for piecewise quadratic finite elements in 2d). We have to set values for the new vertex in the refinement edge, and for the two midpoints of the bisected edge. Then we have to set the value for the midpoint of the common edge of the two children of the bisected triangle and we have to set the corresponding value on the neighbor in the case that the refinement edge is not a boundary edge:

```
void real_refine_inter(DOF_REAL_VEC *drv, RC_LIST_EL *list, int n)
{
    FUNC_NAME("real_refine_inter");
    EL        *el = list->el;
    int        node, n0;
    DOF        cdof;
```

```

const DOF      *pdof;
const DOF      *(*get_dof_indices)(const EL *, const DOF_ADMIN *,
                                   DOF *);

const DOF_ADMIN *admin = drv->fe_space->admin;
REAL            *v      = drv->vec;

if (n < 1) return;

get_dof_indices = drv->fe_space->bas_fcts->get_dof_indices;
pdof = (*get_dof_indices)(el, admin, nil);

node = drv->fe_space->mesh->node[VERTEX];
n0 = admin->n0_dof[VERTEX];
/*-----*/
/* newest vertex of child[0] and child[1] */
/*-----*/
cdof = el->child[0]->dof[node+DIM][n0];
v[cdof] = v[pdof[5]];
/*-----*/
/* midpoint of edge on child[0] at the refinement edge */
/*-----*/
node = drv->fe_space->mesh->node[EDGE];
n0 = admin->n0_dof[EDGE];

cdof = el->child[0]->dof[node][n0];
v[cdof] = 0.375*v[pdof[0]] - 0.125*v[pdof[1]] + 0.75*v[pdof[5]];
/*-----*/
/* node in the common edge of child[0] and child[1] */
/*-----*/
cdof = el->child[0]->dof[node+1][n0];
v[cdof] = -0.125*(v[pdof[0]] + v[pdof[1]]) + 0.25*v[pdof[5]]
          + 0.5*(v[pdof[3]] + v[pdof[4]]);
/*-----*/
/* midpoint of edge on child[1] at the refinement edge */
/*-----*/
cdof = el->child[1]->dof[node+1][n0];
v[cdof] = -0.125*v[pdof[0]] + 0.375*v[pdof[1]] + 0.75*v[pdof[5]];

if (n > 1)
{
/*-----*/
/* adjust value at midpoint of the common edge of neigh's children */
/*-----*/
el = list[1].el;
pdof = (*get_dof_indices)(el, admin, nil);

cdof = el->child[0]->dof[node+1][n0];
v[cdof] = -0.125*(v[pdof[0]] + v[pdof[1]]) + 0.25*v[pdof[5]]
          + 0.5*(v[pdof[3]] + v[pdof[4]]);

```



```

    }
    return;
}

```

Example 3.26 (Interpolation during refinement for piecewise quadratic finite elements in 3d). Here, we first have to set values for all DOFs that belong to the first element of the refinement patch. Then we have to loop over the refinement patch and set all DOFs that have not previously been set on another patch element. In order to set values only once, by the variable `lr_set` we check, if a common DOFs with a left or right neighbor is set by the neighbor. Such values are already set if the neighbor is a prior element in the list. Since all values are set on the first element for all subsequent elements there must be DOFs which have been set by another element.

```

void real_refine_inter(DOF_REAL_VEC *drv, RC_LIST_EL *list, int n)
{
    FUNCNAME("real_refine_inter");
    EL      *el = list->el;
    const DOF *cdof;
    DOF      pdof[10], cdofi;
    int      i, lr_set;
    int      node0, n0;
    const DOF *(*get_dof_indices)(const EL *, const DOF_ADMIN *,
                                  DOF *);

    REAL      *v      = drv->vec;
    const DOF_ADMIN *admin = drv->fe_space->admin;

    if (n < 1) return;

    get_dof_indices = drv->fe_space->bas_fcts->get_dof_indices;
    (*get_dof_indices)(el, admin, pdof);
    /*-----*/
    /* values on child[0] */
    /*-----*/
    cdof = (*get_dof_indices)(el->child[0], admin, nil);

    v[cdof[3]] = (v[pdof[4]]);
    v[cdof[6]] = 0.375*v[pdof[0]] - 0.125*v[pdof[1]] + 0.75*v[pdof[4]];
    v[cdof[8]] = (0.125*(-v[pdof[0]] - v[pdof[1]]) + 0.25*v[pdof[4]]
                  + 0.5*(v[pdof[5]] + v[pdof[7]]));
    v[cdof[9]] = (0.125*(-v[pdof[0]] - v[pdof[1]]) + 0.25*v[pdof[4]]
                  + 0.5*(v[pdof[6]] + v[pdof[8]]));
    /*-----*/
    /* values on child[1] */
    /*-----*/
    node0 = drv->fe_space->mesh->node[EDGE];
    n0 = admin->n0_dof[EDGE];

    cdofi = el->child[1]->dof[node0+2][n0];
}

```

```

    v[cdofi] = -0.125*v[pdof[0]] + 0.375*v[pdof[1]] + 0.75*v[pdof[4]];
/*-----*/
/*  adjust neighbour values                                */
/*-----*/
for (i = 1; i < n; i++)
{
    el = list[i].el;
    (*get_dof_indices)(el, admin, pdof);

    lr_set = 0;
    if (list[i].neigh[0] && list[i].neigh[0]->no < i)
        lr_set = 1;

    if (list[i].neigh[1] && list[i].neigh[1]->no < i)
        lr_set += 2;

    TEST_EXIT(lr_set)("no values set on both neighbours\n");

    switch (lr_set)
    {
    case 1:
        cdofi = el->child[0]->dof[node0+4][n0];
        v[cdofi] = (0.125*(-v[pdof[0]] - v[pdof[1]]) + 0.25*v[pdof[4]]
                    + 0.5*(v[pdof[5]] + v[pdof[7]]));
        break;
    case 2:
        cdofi = el->child[0]->dof[node0+5][n0];
        v[cdofi] = (0.125*(-v[pdof[0]] - v[pdof[1]]) + 0.25*v[pdof[4]]
                    + 0.5*(v[pdof[6]] + v[pdof[8]]));
    }
}
return;
}

```

3.5.6 Piecewise cubic finite elements

Table 3.9. Local basis functions for cubic finite elements in 1d.


function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{2}(3\lambda_0 - 1)(3\lambda_0 - 2)\lambda_0$	vertex 0	$\lambda^0 = (1, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{2}(3\lambda_1 - 1)(3\lambda_1 - 2)\lambda_1$	vertex 1	$\lambda^1 = (0, 1)$
$\bar{\varphi}^2(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_1$	center	$\lambda^2 = (\frac{2}{3}, \frac{1}{3})$
$\bar{\varphi}^3(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_0$	center	$\lambda^3 = (\frac{1}{3}, \frac{2}{3})$

Table 3.10. Local basis functions for cubic finite elements in 2d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{2}(3\lambda_0 - 1)(3\lambda_0 - 2)\lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{2}(3\lambda_1 - 1)(3\lambda_1 - 2)\lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0)$
$\bar{\varphi}^2(\lambda) = \frac{1}{2}(3\lambda_2 - 1)(3\lambda_2 - 2)\lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1)$
$\bar{\varphi}^3(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_2$	edge 0	$\lambda^3 = (0, \frac{2}{3}, \frac{1}{3})$
$\bar{\varphi}^4(\lambda) = \frac{9}{2}(3\lambda_2 - 1)\lambda_2\lambda_1$	edge 0	$\lambda^4 = (0, \frac{1}{3}, \frac{2}{3})$
$\bar{\varphi}^5(\lambda) = \frac{9}{2}(3\lambda_2 - 1)\lambda_2\lambda_0$	edge 1	$\lambda^5 = (\frac{1}{3}, 0, \frac{2}{3})$
$\bar{\varphi}^6(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^6 = (\frac{2}{3}, 0, \frac{1}{3})$
$\bar{\varphi}^7(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_1$	edge 2	$\lambda^7 = (\frac{2}{3}, \frac{1}{3}, 0)$
$\bar{\varphi}^8(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_0$	edge 2	$\lambda^8 = (\frac{1}{3}, \frac{2}{3}, 0)$
$\bar{\varphi}^9(\lambda) = 27\lambda_0\lambda_1\lambda_2$	center	$\lambda^9 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

For Lagrange elements of third order we have four basis functions in 1d, ten basis functions in 2d, and 20 in 3d; the basis functions and the corresponding Lagrange nodes in barycentric coordinates are shown in Tables 3.9, 3.10, and 3.11.

For cubic elements we have to face a further difficulty. At each edge two basis functions are located. The two DOFs of the i -th edge are subsequent entries in the vector `e1->dof[i]`. For two neighboring triangles the common edge has a different orientation with respect to the local numbering of vertices on the two triangles. In Fig. 3.7 the 3rd local basis function on the left and the 4th on the right triangle built up the global basis function, e.g.; thus, both local basis function must have access to the same global DOF.

**Fig. 3.7.** Cubic DOFs on a patch of two triangles.

In order to combine the global DOF with the local basis function in the implementation of the `get_dof_indices()` function, we have to give every edge a *global orientation*, i.e., every edge has a unique beginning and end

Table 3.11. Local basis functions for cubic finite elements in 3d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{2}(3\lambda_0 - 1)(3\lambda_0 - 2)\lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{2}(3\lambda_1 - 1)(3\lambda_1 - 2)\lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0, 0)$
$\bar{\varphi}^2(\lambda) = \frac{1}{2}(3\lambda_2 - 1)(3\lambda_2 - 2)\lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1, 0)$
$\bar{\varphi}^3(\lambda) = \frac{1}{2}(3\lambda_3 - 1)(3\lambda_3 - 2)\lambda_3$	vertex 3	$\lambda^3 = (0, 0, 0, 1)$
$\bar{\varphi}^4(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_1$	edge 0	$\lambda^4 = (\frac{2}{3}, \frac{1}{3}, 0, 0)$
$\bar{\varphi}^5(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_0$	edge 0	$\lambda^5 = (\frac{1}{3}, \frac{2}{3}, 0, 0)$
$\bar{\varphi}^6(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^6 = (\frac{2}{3}, 0, \frac{1}{3}, 0)$
$\bar{\varphi}^7(\lambda) = \frac{9}{2}(3\lambda_2 - 1)\lambda_2\lambda_0$	edge 1	$\lambda^7 = (\frac{1}{3}, 0, \frac{2}{3}, 0)$
$\bar{\varphi}^8(\lambda) = \frac{9}{2}(3\lambda_0 - 1)\lambda_0\lambda_3$	edge 2	$\lambda^8 = (\frac{2}{3}, 0, 0, \frac{1}{3})$
$\bar{\varphi}^9(\lambda) = \frac{9}{2}(3\lambda_3 - 1)\lambda_3\lambda_0$	edge 2	$\lambda^9 = (\frac{1}{3}, 0, 0, \frac{2}{3})$
$\bar{\varphi}^{10}(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_2$	edge 3	$\lambda^{10} = (0, \frac{2}{3}, \frac{1}{3}, 0)$
$\bar{\varphi}^{11}(\lambda) = \frac{9}{2}(3\lambda_2 - 1)\lambda_2\lambda_1$	edge 3	$\lambda^{11} = (0, \frac{1}{3}, \frac{2}{3}, 0)$
$\bar{\varphi}^{12}(\lambda) = \frac{9}{2}(3\lambda_1 - 1)\lambda_1\lambda_3$	edge 4	$\lambda^{12} = (0, \frac{2}{3}, 0, \frac{1}{3})$
$\bar{\varphi}^{13}(\lambda) = \frac{9}{2}(3\lambda_3 - 1)\lambda_3\lambda_1$	edge 4	$\lambda^{13} = (0, \frac{1}{3}, 0, \frac{2}{3})$
$\bar{\varphi}^{14}(\lambda) = \frac{9}{2}(3\lambda_2 - 1)\lambda_2\lambda_3$	edge 5	$\lambda^{14} = (0, 0, \frac{2}{3}, \frac{1}{3})$
$\bar{\varphi}^{15}(\lambda) = \frac{9}{2}(3\lambda_3 - 1)\lambda_3\lambda_2$	edge 5	$\lambda^{15} = (0, 0, \frac{1}{3}, \frac{2}{3})$
$\bar{\varphi}^{16}(\lambda) = 27\lambda_1\lambda_2\lambda_3$	face 0	$\lambda^{16} = (0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
$\bar{\varphi}^{17}(\lambda) = 27\lambda_2\lambda_3\lambda_0$	face 1	$\lambda^{17} = (\frac{1}{3}, 0, \frac{1}{3}, \frac{1}{3})$
$\bar{\varphi}^{18}(\lambda) = 27\lambda_3\lambda_0\lambda_1$	face 2	$\lambda^{18} = (\frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3})$
$\bar{\varphi}^{19}(\lambda) = 27\lambda_0\lambda_1\lambda_2$	face 3	$\lambda^{19} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$

point. Using the orientation of an edge we are able to order the DOFs stored at this edge. Let for example the common edge in Fig. 3.7 be oriented from bottom to top. The global DOF corresponding to 3rd local DOF on the left and the 4th local DOF on the right is then

```
e1->dof[N_VERTICES+0][admin->n0_dof[EDGE]]
```

and for the 4th local DOF on the left and 3rd local DOF on the right

```
e1->dof[N_VERTICES+0][admin->n0_dof[EDGE]+1]
```

The global orientation gives a unique access of local DOFs from global ones.

Example 3.27 (Accessing DOFs for piecewise cubic finite elements).

For the implementation, we use in 2d as well as in 3d an orientation defined by the DOF indices at the edges' vertices. The vertex with the smaller (global) DOF index is the beginning point, the vertex with the higher index the end point. For cubics the implementation differs between 2d and 3d. In 2d we have one degree of freedom at the center and in 3d on degree of freedom at each

face and no one at the center. The DOFs at an edge are accessed according to the orientation of the edge.

```

#if DIM == 2
#define N_BAS N_VERTICES+2*N_EDGES+1
#endif
#if DIM == 3
#define N_BAS N_VERTICES+2*N_EDGES+N_FACES
#endif

const DOF *get_dof_indices(const EL *el, const DOF_ADMIN *admin,
                          DOF *idof)
{
    static DOF index_vec[N_BAS];
    DOF *rvec = idof ? idof : index_vec, **dof = el->dof;
    int i, j;

    /*--- vertex DOFs -----*/
    for (i = 0; i < N_VERTICES; i++)
        rvec[i] = dof[i][admin->n0_dof[VERTEX]];

    /*--- edge DOFs -----*/
    for (i = 0, j = N_VERTICES; i < N_EDGES; i++)
    {
        if (dof[vertex_of_edge[i][0]][0] < dof[vertex_of_edge[i][1]][0])
        {
            rvec[j++] = dof[N_VERTICES+i][admin->n0_dof[EDGE]];
            rvec[j++] = dof[N_VERTICES+i][admin->n0_dof[EDGE]+1];
        }
        else
        {
            rvec[j++] = dof[N_VERTICES+i][admin->n0_dof[EDGE]+1];
            rvec[j++] = dof[N_VERTICES+i][admin->n0_dof[EDGE]];
        }
    }

    #if DIM == 3
    /*--- face DOFs, only 3d -----*/
    for (i = 0; i < N_FACES; i++)
        rvec[j++] = dof[N_VERTICES+N_EDGES+i][admin->n0_dof[FACE]];
    #endif

    #if DIM == 2
    /*--- center DOF, only 2d -----*/
    rvec[j] = dof[N_VERTICES+N_EDGES][admin->n0_dof[CENTER]];
    #endif

    return((const DOF *) rvec);
}

```

3.5.7 Piecewise quartic finite elements

For Lagrange elements of fourth order we have 5 basis functions in 1d, 15 in 2d, and 35 in 3d; the basis functions and the corresponding Lagrange nodes in barycentric coordinates are shown in Tables 3.12, 3.13, and 3.14.

For the implementation of `get_dof_indices()` for quartics, we again need a global orientation of the edges on the mesh. At every edge three DOFs are located, which then can be ordered with respect to the orientation of the corresponding edge. In 3d, we also need a global orientation of faces for a one to one mapping of global DOFs located at a face to local DOFs on an element at that face. Such an orientation can again be defined by DOF indices at the face's vertices.

Table 3.12. Local basis functions for quartic finite elements in 1d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)(4\lambda_0 - 3)\lambda_0$	vertex 0	$\lambda^0 = (1, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)(4\lambda_1 - 3)\lambda_1$	vertex 1	$\lambda^1 = (0, 1)$
$\bar{\varphi}^2(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_1$	center	$\lambda^2 = (\frac{3}{4}, \frac{1}{4})$
$\bar{\varphi}^3(\lambda) = 4(4\lambda_0 - 1)(4\lambda_1 - 1)\lambda_0\lambda_1$	center	$\lambda^3 = (\frac{1}{2}, \frac{1}{2})$
$\bar{\varphi}^4(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_0\lambda_1$	center	$\lambda^4 = (\frac{1}{4}, \frac{3}{4})$

Table 3.13. Local basis functions for quartic finite elements in 2d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)(4\lambda_0 - 3)\lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)(4\lambda_1 - 3)\lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0)$
$\bar{\varphi}^2(\lambda) = \frac{1}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)(4\lambda_2 - 3)\lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1)$
$\bar{\varphi}^3(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_1\lambda_2$	edge 0	$\lambda^3 = (0, \frac{3}{4}, \frac{1}{4})$
$\bar{\varphi}^4(\lambda) = 4(4\lambda_1 - 1)(4\lambda_2 - 1)\lambda_1\lambda_2$	edge 0	$\lambda^4 = (0, \frac{1}{2}, \frac{1}{2})$
$\bar{\varphi}^5(\lambda) = \frac{16}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)\lambda_1\lambda_2$	edge 0	$\lambda^5 = (0, \frac{1}{4}, \frac{3}{4})$
$\bar{\varphi}^6(\lambda) = \frac{16}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^6 = (\frac{1}{4}, 0, \frac{3}{4})$
$\bar{\varphi}^7(\lambda) = 4(4\lambda_2 - 1)(4\lambda_0 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^7 = (\frac{1}{2}, 0, \frac{1}{2})$
$\bar{\varphi}^8(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^8 = (\frac{3}{4}, 0, \frac{1}{4})$
$\bar{\varphi}^9(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_1$	edge 2	$\lambda^9 = (\frac{3}{4}, \frac{1}{4}, 0)$
$\bar{\varphi}^{10}(\lambda) = 4(4\lambda_0 - 1)(4\lambda_1 - 1)\lambda_0\lambda_1$	edge 2	$\lambda^{10} = (\frac{1}{2}, \frac{1}{2}, 0)$
$\bar{\varphi}^{11}(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_0\lambda_1$	edge 2	$\lambda^{11} = (\frac{1}{4}, \frac{3}{4}, 0)$
$\bar{\varphi}^{12}(\lambda) = 32(4\lambda_0 - 1)\lambda_0\lambda_1\lambda_2$	center	$\lambda^{12} = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$
$\bar{\varphi}^{13}(\lambda) = 32(4\lambda_1 - 1)\lambda_0\lambda_1\lambda_2$	center	$\lambda^{13} = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$
$\bar{\varphi}^{14}(\lambda) = 32(4\lambda_2 - 1)\lambda_0\lambda_1\lambda_2$	center	$\lambda^{14} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$

Table 3.14. Local basis functions for quartic finite elements in 3d.

function	position	Lagrange node
$\bar{\varphi}^0(\lambda) = \frac{1}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)(4\lambda_0 - 3)\lambda_0$	vertex 0	$\lambda^0 = (1, 0, 0, 0)$
$\bar{\varphi}^1(\lambda) = \frac{1}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)(4\lambda_1 - 3)\lambda_1$	vertex 1	$\lambda^1 = (0, 1, 0, 0)$
$\bar{\varphi}^2(\lambda) = \frac{1}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)(4\lambda_2 - 3)\lambda_2$	vertex 2	$\lambda^2 = (0, 0, 1, 0)$
$\bar{\varphi}^3(\lambda) = \frac{1}{3}(4\lambda_3 - 1)(2\lambda_3 - 1)(4\lambda_3 - 3)\lambda_3$	vertex 3	$\lambda^3 = (0, 0, 0, 1)$
$\bar{\varphi}^4(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_1$	edge 0	$\lambda^4 = (\frac{3}{4}, \frac{1}{4}, 0, 0)$
$\bar{\varphi}^5(\lambda) = 4(4\lambda_0 - 1)(4\lambda_1 - 1)\lambda_0\lambda_1$	edge 0	$\lambda^5 = (\frac{1}{2}, \frac{1}{2}, 0, 0)$
$\bar{\varphi}^6(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_0\lambda_1$	edge 0	$\lambda^6 = (\frac{1}{4}, \frac{3}{4}, 0, 0)$
$\bar{\varphi}^7(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^7 = (\frac{3}{4}, 0, \frac{1}{4}, 0)$
$\bar{\varphi}^8(\lambda) = 4(4\lambda_0 - 1)(4\lambda_2 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^8 = (\frac{1}{2}, 0, \frac{1}{2}, 0)$
$\bar{\varphi}^9(\lambda) = \frac{16}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)\lambda_0\lambda_2$	edge 1	$\lambda^9 = (\frac{1}{4}, 0, \frac{3}{4}, 0)$
$\bar{\varphi}^{10}(\lambda) = \frac{16}{3}(4\lambda_0 - 1)(2\lambda_0 - 1)\lambda_0\lambda_3$	edge 2	$\lambda^{10} = (\frac{3}{4}, 0, 0, \frac{1}{4})$
$\bar{\varphi}^{11}(\lambda) = 4(4\lambda_0 - 1)(4\lambda_3 - 1)\lambda_0\lambda_3$	edge 2	$\lambda^{11} = (\frac{1}{2}, 0, 0, \frac{1}{2})$
$\bar{\varphi}^{12}(\lambda) = \frac{16}{3}(4\lambda_3 - 1)(2\lambda_3 - 1)\lambda_0\lambda_3$	edge 2	$\lambda^{12} = (\frac{1}{4}, 0, 0, \frac{3}{4})$
$\bar{\varphi}^{13}(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_1\lambda_2$	edge 3	$\lambda^{13} = (0, \frac{3}{4}, \frac{1}{4}, 0)$
$\bar{\varphi}^{14}(\lambda) = 4(4\lambda_1 - 1)(4\lambda_2 - 1)\lambda_1\lambda_2$	edge 3	$\lambda^{14} = (0, \frac{1}{2}, \frac{1}{2}, 0)$
$\bar{\varphi}^{15}(\lambda) = \frac{16}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)\lambda_1\lambda_2$	edge 3	$\lambda^{15} = (0, \frac{1}{4}, \frac{3}{4}, 0)$
$\bar{\varphi}^{16}(\lambda) = \frac{16}{3}(4\lambda_1 - 1)(2\lambda_1 - 1)\lambda_1\lambda_3$	edge 4	$\lambda^{16} = (0, \frac{3}{4}, 0, \frac{1}{4})$
$\bar{\varphi}^{17}(\lambda) = 4(4\lambda_1 - 1)(4\lambda_3 - 1)\lambda_1\lambda_3$	edge 4	$\lambda^{17} = (0, \frac{1}{2}, 0, \frac{1}{2})$
$\bar{\varphi}^{18}(\lambda) = \frac{16}{3}(4\lambda_3 - 1)(2\lambda_3 - 1)\lambda_1\lambda_3$	edge 4	$\lambda^{18} = (0, \frac{1}{4}, 0, \frac{3}{4})$
$\bar{\varphi}^{19}(\lambda) = \frac{16}{3}(4\lambda_2 - 1)(2\lambda_2 - 1)\lambda_2\lambda_3$	edge 5	$\lambda^{19} = (0, 0, \frac{3}{4}, \frac{1}{4})$
$\bar{\varphi}^{20}(\lambda) = 4(4\lambda_2 - 1)(4\lambda_3 - 1)\lambda_2\lambda_3$	edge 5	$\lambda^{20} = (0, 0, \frac{1}{2}, \frac{1}{2})$
$\bar{\varphi}^{21}(\lambda) = \frac{16}{3}(4\lambda_3 - 1)(2\lambda_3 - 1)\lambda_2\lambda_3$	edge 5	$\lambda^{21} = (0, 0, \frac{1}{4}, \frac{3}{4})$
$\bar{\varphi}^{22}(\lambda) = 32(4\lambda_1 - 1)\lambda_1\lambda_2\lambda_3$	face 0	$\lambda^{22} = (0, \frac{1}{2}, \frac{1}{4}, \frac{1}{4})$
$\bar{\varphi}^{23}(\lambda) = 32(4\lambda_2 - 1)\lambda_1\lambda_2\lambda_3$	face 0	$\lambda^{23} = (0, \frac{1}{4}, \frac{1}{2}, \frac{1}{4})$
$\bar{\varphi}^{24}(\lambda) = 32(4\lambda_3 - 1)\lambda_1\lambda_2\lambda_3$	face 0	$\lambda^{24} = (0, \frac{1}{4}, \frac{1}{4}, \frac{1}{2})$
$\bar{\varphi}^{25}(\lambda) = 32(4\lambda_0 - 1)\lambda_0\lambda_2\lambda_3$	face 1	$\lambda^{25} = (\frac{1}{2}, 0, \frac{1}{4}, \frac{1}{4})$
$\bar{\varphi}^{26}(\lambda) = 32(4\lambda_2 - 1)\lambda_0\lambda_2\lambda_3$	face 1	$\lambda^{26} = (\frac{1}{4}, 0, \frac{1}{2}, \frac{1}{4})$
$\bar{\varphi}^{27}(\lambda) = 32(4\lambda_3 - 1)\lambda_0\lambda_2\lambda_3$	face 1	$\lambda^{27} = (\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{2})$
$\bar{\varphi}^{28}(\lambda) = 32(4\lambda_0 - 1)\lambda_0\lambda_1\lambda_3$	face 2	$\lambda^{28} = (\frac{1}{2}, \frac{1}{4}, 0, \frac{1}{4})$
$\bar{\varphi}^{29}(\lambda) = 32(4\lambda_1 - 1)\lambda_0\lambda_1\lambda_3$	face 2	$\lambda^{29} = (\frac{1}{4}, \frac{1}{2}, 0, \frac{1}{4})$
$\bar{\varphi}^{30}(\lambda) = 32(4\lambda_3 - 1)\lambda_0\lambda_1\lambda_3$	face 2	$\lambda^{30} = (\frac{1}{4}, \frac{1}{4}, 0, \frac{1}{2})$
$\bar{\varphi}^{31}(\lambda) = 32(4\lambda_0 - 1)\lambda_0\lambda_1\lambda_2$	face 3	$\lambda^{31} = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4}, 0)$
$\bar{\varphi}^{32}(\lambda) = 32(4\lambda_1 - 1)\lambda_0\lambda_1\lambda_2$	face 3	$\lambda^{32} = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4}, 0)$
$\bar{\varphi}^{33}(\lambda) = 32(4\lambda_2 - 1)\lambda_0\lambda_1\lambda_2$	face 3	$\lambda^{33} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{2}, 0)$
$\bar{\varphi}^{34}(\lambda) = 256\lambda_0\lambda_1\lambda_2\lambda_3$	center	$\lambda^{34} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$

3.5.8 Access to Lagrange elements

The Lagrange elements described above are already implemented in ALBERTA; access to Lagrange elements is given by the function

```
const BAS_FCTS *get_lagrange(int);
```

Description:

get_lagrange(degree): returns a pointer to a filled **BAS_FCTS** structure for Lagrange elements of order **degree**, where $0 \leq \text{degree} \leq 4$ for all dimensions; no additional call of **new_bas_fcts()** is needed. The entry **bas_fcts_data** of the **BAS_FCTS** structure is a pointer to a **(*)[DIM+1]** vector of length **n_bas_fcts** storing the Lagrange nodes in barycentric coordinates.

3.6 Implementation of finite element spaces

3.6.1 The finite element space data structure

All information about the underlying mesh, the local basis functions, and the DOFs are collected in the following data structure which defines one single finite element space:

```
typedef struct fe_space      FE_SPACE;

struct fe_space
{
    const char      *name;
    const DOF_ADMIN *admin;
    const BAS_FCTS  *bas_fcts;
    MESH            *mesh;
};
```

Description:

name: holds a textual description of the finite element space.

admin: pointer to the DOF administration for the DOFs of this finite element space, see Section 3.3.1.

bas_fcts: pointer to the local basis functions, see Section 3.5.1.

mesh: pointer to the underlying mesh, see Section 3.2.14.

Several finite element spaces can be handled on the same mesh. Different finite element spaces can use the same DOF administration, if they share exactly the same DOFs.

3.6.2 Access to finite element spaces

A finite element space can only be accessed by the function

```
const FE_SPACE *get_fe_space(MESH *, const char *, const int [DIM+1],
                             const BAS_FCTS *);
```

Description:

`get_fe_space(mesh, name, n_dof, bas_fcts)`: defines a new finite element space on `mesh`; it looks for an existing `dof_admin` defined on `mesh`, which manages DOFs uniquely defined by `bas_fcts->dof_admin->n_dof` or by `n_dof` in case that `bas_fcts` is `nil` (compare Section 3.3.1); if such a `dof_admin` is not found, a new `dof_admin` is created;

the return value is a newly created `FE_SPACE` structure, where `name` is duplicated, and the members `mesh`, `bas_fcts` and `admin` are adjusted correctly.

The selection of finite element spaces defines the DOFs that must be present on the mesh elements. For each finite element space there must be a corresponding DOF administration, having information about the used DOFs. For the access of a mesh, via the `GET_MESH()` routine, all used DOFs, i.e. all DOF administrations, have to be specified in advance. After the access of a mesh, no further `DOF_ADMIN` can be defined (this would correspond to a so called p -refinement of the mesh, which is not implemented yet). The specification of the used DOFs is done via a user defined routine

```
void init_dof_admins(MESH *);
```

which is the second argument of the `GET_MESH()` routine and the function is called during the initialization of the mesh. Inside the function `init_dof_admins()`, the finite element spaces are accessed, which thereby determine the used DOFs. Since a mesh gives only access to the `DOF_ADMIN`s defined on it, the user has to store pointers to the `FE_SPACE` structures in some global variable; no access to `FE_SPACES` is possible via the underlying mesh.

Example 3.28 (Initializing DOFs for Stokes and Navier–Stokes).

Now, as an example we look at a user defined function `init_dof_admins()`. In the example we want to define two finite element spaces on the mesh, for a mixed finite element formulation of the Stokes or Navier–Stokes equations with the Taylor–Hood element, e.g. We want to use Lagrange finite elements of order `degree` (for the velocity) and `degree – 1` (for the pressure). Pointers to the corresponding `FE_SPACE` structures are stored in the global variables `u_fe` and `p_fe`.

```
static FE_SPACE *u_fe, *p_fe;
static int      degree;

void init_dof_admins(MESH *mesh)
{
    FUNCNAME("init_dof_admins");
    const BAS_FCTS *lagrange;
```

```

TEST_EXIT(mesh)("no MESH\n");
TEST_EXIT(degree > 1)("degree must be greater than 1\n");

lagrange = get_lagrange(degree);
u_fe = get_fe_space(mesh, lagrange->name, nil, lagrange);

lagrange = get_lagrange(degree-1);
p_fe = get_fe_space(mesh, lagrange->name, nil, lagrange);

return;
}

```

An initialization of a mesh with this `init_dof_admins()` function, will provide all DOFs for the two finite element spaces on the elements and the corresponding `DOF_ADMINs` will have information about the access of local DOFs for both finite element spaces.

It is also possible to define only one or even more finite element spaces; the use of special user defined basis functions is possible too. These should be added to the list of all used basis functions by a call of `new_bas_fcts()` inside `init_dof_admins()`.

Remark 3.29. For some applications, additional DOF sets, that are not directly connected to any finite element space may also be defined by `get_fe_space(mesh, name, n_dof, nil)`. Here, `n_dof` is a vector storing the number of DOFs at a `VERTEX`, `EDGE` (2d and 3d), `FACE` (3d), and `CENTER` that should be present. An additional DOF set with 1 DOF at each edge in 2d and face in 3d can be defined in the following way inside the `init_dof_admins()` function:

```

#if DIM == 2
    int n_dof[DIM+1] = {0,1,0};
#endif
#if DIM == 3
    int n_dof[DIM+1] = {0,0,1,0};
#endif
    ...

#if DIM > 1
    face_dof = get_fe_space(mesh, "face dofs", n_dof, nil);
#endif
    ...

```

3.7 Routines for barycentric coordinates

Operations on single elements are performed using barycentric coordinates. In many applications, the world coordinates x of the local barycentric coordinates

λ have to be calculated (see Section 3.12, e.g.). Some other applications will need the calculation of barycentric coordinates for given world coordinates (see Section 3.2.19, e.g.). Finally, derivatives of finite element functions on elements involve the Jacobian of the barycentric coordinates (see Section 3.9, e.g.).

In case of a grid with parametric elements, these operations strongly depend on the element parameterization and no general routines can be supplied. For non-parametric simplices, ALBERTA supplies functions to perform these basic tasks:

```
const REAL *coord_to_world(const EL_INFO *, const REAL *, REAL_D);
int world_to_coord(const EL_INFO *, const REAL *, REAL [DIM+1]);
REAL el_grd_lambda(const EL_INFO *, REAL [DIM+1][DIM_OF_WORLD]);
REAL el_det(const EL_INFO *);
REAL el_volume(const EL_INFO *);
```

Description:

`coord_to_world(el_info, lambda, world)`: returns a pointer to a vector, which contains the world coordinates of a point in barycentric coordinates `lambda` with respect to the element `el_info->el`;

if `world` is not `nil` the world coordinates are stored in this vector; otherwise the function itself provides memory for this vector; in this case the vector is overwritten during the next call of `coord_to_world()`;

`world_to_coord(el_info, world, lambda)`: calculates the barycentric coordinates with respect to the element `el_info->el` of a point with world coordinates `world` and stores them in the vector given by `lambda`. The return value is `-1` when the point is inside the simplex (or on its boundary), otherwise the index of the barycentric coordinate with largest negative value (between 0 and DIM);

`el_grd_lambda(el_info, Lambda)`: calculates the Jacobian of the barycentric coordinates on `el_info->el` and stores the matrix in `Lambda`; the return value of the function is the absolute value of the determinant of the affine linear parameterization's Jacobian;

`el_det(el_info)`: returns the the absolute value of the determinant of the affine linear parameterization's Jacobian;

the function `el_det()` needs vertex coordinates information; thus, the flag `FILL_COORDS` has to be set during mesh traversal when calling this routine on elements.

`el_volume(el_info)`: returns the the volume of the simplex.

All functions need vertex coordinates information; thus, the flag `FILL_COORDS` has to be set during mesh traversal when calling one of the above described routine on elements.

3.8 Data structures for numerical quadrature

For the numerical calculation of general integrals

$$\int_S f(x) dx$$

we use quadrature formulas described in 1.4.6. ALBERTA supports numerical integration in one, two, and three dimensions on the standard simplex \hat{S} in barycentric coordinates.

3.8.1 The QUAD data structure

A quadrature formula is described by the following structure, which is defined both as type QUAD and QUADRATURE:

```
typedef struct quadrature    QUAD;
typedef struct quadrature    QUADRATURE;

struct quadrature
{
    char          *name;
    int           degree;

    int           dim;
    int           n_points;
    const double  **lambda;
    const double  *w;
};
```

Description:

name: textual description of the quadrature;

degree: quadrature is exact of degree **degree**;

dim: quadrature for dimension **dim**;

n_points: number of quadrature points;

lambda: vector `lambda[0], ..., lambda[n_points-1]` of quadrature points given in barycentric coordinates (thus having DIM+1 components);

w: vector `w[0], ..., w[n_points-1]` of quadrature weights.

Currently, numerical quadrature formulas exact up to order 19 in one (Gauss formulas), up to order 17 in two, up to order 7 in three dimensions are implemented. We only use stable formulas; this results in more quadrature points for some formulas (for example in 3d the formula which is exact of degree 3). A compilation of quadrature formulas on triangles and tetrahedra is given in [26]. The implemented quadrature formulas are taken from [34, 39, 42, 68].

Functions for numerical quadrature are

```
const QUAD *get_quadrature(unsigned dim, unsigned degree);
const QUAD *get_lumping_quadrature(unsigned dim);
REAL integrate_std_simp(const QUAD *quad, REAL (*f)(const REAL *));
```

Description:

get_quadrature(dim, degree): returns a pointer to a filled QUAD data structure for numerical integration in `dim` dimensions which is exact of degree `min(19, degree)` for `dim==1`, `min(17, degree)` for `dim==2`, and `min(7, degree)` for `dim==3`.

get_lumping_quadrature(dim): returns a pointer to a QUAD structure for numerical integration in `dim` dimensions which implements the mass lumping (or vertex) quadrature rule.

integrate_std_simp(quad, f): approximates an integral by the numerical quadrature described by `quad`;

`f` is a pointer to a function to be integrated, evaluated in barycentric coordinates; the return value is

$$\sum_{k=0}^{\text{quad} \rightarrow \text{n_points} - 1} \text{quad} \rightarrow \text{w}[k] * (*f)(\text{quad} \rightarrow \text{lambda}[k]);$$

for the approximation of $\int_S f$ we have to multiply this value with $d!|S|$ for a simplex S ; for a parametric simplex, `f` should be a pointer to a function which calculates $f(\lambda)|\det DF_S(\hat{x}(\lambda))|$.

The following functions initialize values and gradients of functions at the quadrature nodes:

```
const REAL *f_at_qp(const QUAD *, REAL (*)(const REAL [DIM+1]),
                    REAL *);
const REAL_D *grd_f_at_qp(const QUAD *,
                           const REAL (*)(const REAL [DIM+1]),
                           REAL_D *vec);
const REAL_D *f_d_at_qp(const QUAD *,
                        const REAL (*)(const REAL [DIM+1]),
                        REAL_D *);
const REAL_DD *grd_f_d_at_qp(const QUAD *,
                              const REAL_D (*)(const REAL [DIM+1]),
                              REAL_DD *);
```

Description:

f_at_qp(quad, f, vec): returns a pointer `ptr` to a vector storing the values of a REAL valued function at all quadrature points of `quad`; the length of this vector is `quad->n_points`; `f` is a pointer to that function, evaluated in barycentric coordinates; if `vec` is not `nil`, the values are stored in this vector, otherwise the values are stored in some static local vector, which is overwritten on the next call;

`ptr[i]=(*f)(quad->lambda[i])` for $0 \leq i < \text{quad} \rightarrow \text{n_points}$.

grd_f_at_qp(quad, **grd_f**, **vec**): returns a pointer **ptr** to a vector of length **quad->n_points** storing the gradient (with respect to world coordinates) of a **REAL** valued function at all quadrature points of **quad**;

grd_f is a pointer to a function, evaluated in barycentric coordinates and returning a pointer to a vector of length **DIM_OF_WORLD** storing the gradient; if **vec** is not **nil**, the values are stored in this vector, otherwise the values are stored in some local static vector, which is overwritten on the next call; **ptr[i][j]** = **(*grd_f)(quad->lambda[i])[j]**, for $0 \leq j < \text{DIM_OF_WORLD}$ and $0 \leq i < \text{quad->n_points}$,

f_d_at_qp(quad, **fd**, **vec**): returns a pointer **ptr** to a vector of length **quad->n_points** storing the values of a **REAL_D** valued function at all quadrature points of **quad**;

fd is a pointer to that function, evaluated in barycentric coordinates and returning a pointer to a vector of length **DIM_OF_WORLD** storing all components; if the second argument **val** of **(*fd)(lambda, val)** is not **nil**, the values have to be stored at **val**, otherwise **fd** has to provide memory for the vector which may be overwritten on the next call;

if **vec** is not **nil**, the values are stored in this vector, otherwise the values are stored in some static local vector, which is overwritten on the next call; **ptr[i][j]** = **(*fd)(quad->lambda[i], val)[j]**, for values $0 \leq j < \text{DIM_OF_WORLD}$ and $0 \leq i < \text{quad->n_points}$.

grd_f_d_at_qp(quad, **grd_fd**, **vec**): returns a pointer **ptr** to a vector of length **quad->n_points** storing the Jacobian (with respect to world coordinates) of a **REAL_D** valued function at all quadrature points of **quad**;

grd_fd is a pointer to a function, evaluated in barycentric coordinates and returning a pointer to a matrix of size **DIM_OF_WORLD** × **DIM_OF_WORLD** storing the Jacobian; if the second argument **val** of **(*grd_fd)(x, val)** is not **nil**, the Jacobian has to be stored at **val**, otherwise **grd_fd** has to provide memory for the matrix which may be overwritten on the next call;

if **vec** is not **nil**, the values are stored in this vector, otherwise the values are stored in some static local vector, which is overwritten on the next call; **ptr[i][j][k]** = **(*grd_fd)(quad->lambda[i], val)[j][k]**, for $0 \leq j, k < \text{DIM_OF_WORLD}$ and $0 \leq i < \text{quad->n_points}$,

3.8.2 The QUAD_FAST data structure

Often numerical integration involves basis functions, such as the assembling of the system matrix and right hand side, or the integration of finite element functions. Since numerical quadrature involves only the values at the quadrature points and the values of basis functions and its derivatives (with respect to barycentric coordinates) are the same at these points for all elements of the grid, such routines can be much more efficient, if they can use pre-computed values of the basis functions at the quadrature points. In this case the basis

functions do not have to be evaluated for each quadrature point newly on every element.

Information that should be pre-computed can be specified by the following symbolic constants:

```
INIT_PHI
INIT_GRD_PHI
INIT_D2_PHI
```

Description:

INIT_PHI: pre-compute the values of all basis functions at all quadrature nodes;

INIT_GRD_PHI: pre-compute the gradients (with respect to the barycentric coordinates) of all basis functions at all quadrature nodes;

INIT_D2_PHI: pre-compute all 2nd derivatives (with respect to the barycentric coordinates) of all basis functions at all quadrature nodes.

In order to store such information for one set of basis functions we define the data structure

```
typedef struct quad_fast      QUAD_FAST;

struct quad_fast
{
    const QUAD      *quad;
    const BAS_FCTS  *bas_fcts;

    int             n_points;
    int             n_bas_fcts;
    const double    *w;

    U_CHAR          init_flag;

    REAL            **phi;
    REAL            (**grd_phi)[DIM+1];
    REAL            (**D2_phi)[DIM+1][DIM+1];
};
```

The entries yield following information:

quad: values stored for numerical quadrature **quad**;

bas_fcts: values stored for basis functions **bas_fcts**;

n_points: number of quadrature points; equals **quad->n_points**;

n_bas_fcts: number of basis functions; equals **bas_fcts->n_bas_fcts**;

w: vector of quadrature weights; **w** = **quad->w**;

init_flag: indicates which information is initialized; may be one of, or a bitwise OR of several of **INIT_PHI**, **INIT_GRD_PHI**, **INIT_D2_PHI**;

phi: matrix storing function values if the flag `INIT_PHI` is set;
`phi[i][j]` stores the value `bas_fcts->phi[j](quad->lambda[i])`, $0 \leq j < n_bas_fcts$ and $0 \leq i < n_points$;

grd_phi: matrix storing all gradients (with respect to the barycentric coordinates) if the flag `INIT_GRD_PHI` is set;
`grd_phi[i][j][k] = bas_fcts->grd_phi[j](quad->lambda[i])[k]` for $0 \leq j < n_bas_fcts$, $0 \leq i < n_points$, and $0 \leq k \leq DIM$;

D2_phi: matrix storing all second derivatives (with respect to the barycentric coordinates) if the flag `INIT_D2_PHI` is set;
`D2_phi[i][j][k][l] = bas_fcts->D2_phi[j](quad->lambda[i])[k][l]` for $0 \leq j < n_bas_fcts$, $0 \leq i < n_points$, and $0 \leq k, l \leq DIM$.

A filled structure can be accessed by a call of

```
const QUAD_FAST *get_quad_fast(const BAS_FCTS *, const QUAD *,
                               U_CHAR);
```

Description:

`get_quad_fast(bas_fcts, quad, init_flag)`: `bas_fcts` is a pointer to a filled `BAS_FCTS` structure, `quad` a pointer to some quadrature (accessed by `get_quadrature()`, e.g.) and `init_flag` indicates which information should be filled into the `QUAD_FAST` structure; it may be one of, or a bitwise OR of several of `INIT_PHI`, `INIT_GRD_PHI`, `INIT_D2_PHI`; the function returns a pointer to a filled `QUAD_FAST` structure where all demanded information is computed and stored.

All used `QUAD_FAST` structures are stored in a linked list and are identified uniquely by the members `quad` and `bas_fcts`; first, `get_quad_fast()` looks for a matching structure in the linked list; if no structure is found, a new structure is generated and linked to the list; thus for one combination `bas_fcts` and `quad` only one `QUAD_FAST` structure is created.

Then `get_quad_fast()` allocates memory for all information demanded by `init_flag` and which is not yet initialized for this structure; only such information is then computed and stored; on the first call for `bas_fcts` and `quad`, all information demanded `init_flag` is generated, on a subsequent call only missing information is generated.

`get_quad_fast()` will return a `nil` pointer, if `INIT_PHI` flag is set and `bas_fcts->phi` is `nil`, `INIT_GRD_PHI` flag is set and `bas_fcts->grd_phi` is `nil`, and `INIT_D2_PHI` flag is set and `bas_fcts->D2_phi` is `nil`.

There may be several `QUAD_FAST` structures in the list for the same set of basis functions for different quadratures, and there may be several `QUAD_FAST` structures for one quadrature for different sets of basis functions.

The function `get_quad_fast()` should not be called on each element during mesh traversal, because it has to look in a list for an existing entry for a set of basis functions and a quadrature; a pointer to the `QUAD_FAST` structure should be accessed before mesh traversal and stored in some global variable for instance.

Many functions using the `QUAD_FAST` structure need vectors for storing values at all quadrature points; for these functions it can be of interest to get the count of the maximal number of quadrature nodes used by the all initialized `quad_fast` structures in order to avoid several memory reallocations. This count can be accessed by the function

```
int max_quad_points(void);
```

Description:

`max_quad_points()`: returns the maximal number of quadrature points for all yet initialized `quad_fast` structures; this value may change after a new initialization of a `quad_fast` structures;
this count is *not* the maximal number of quadrature points of all used `QUAD` structures, since new quadratures can be used at any time without an initialization.

3.8.3 Integration over sub-simplices (edges/faces)

The weak formulation of non-homogeneous Neumann or Robin boundary values needs integration over $\text{DIM}-1$ dimensional boundary simplices of DIM dimensional mesh elements, and the evaluation of jump residuals for error estimators (compare Sections 1.5, 3.14) needs integration over all interior $\text{DIM}-1$ dimensional sub-simplices. The quadrature formulas and data structures described above are available for any d dimensional simplex, $d = 0, 1, 2, 3$. So, the above task can be accomplished by using a $\text{DIM}-1$ dimensional quadrature formula and augmenting the corresponding DIM dimensional barycentric coordinates of quadrature points on edges/faces to $\text{DIM}+1$ dimensional coordinates on adjacent mesh elements.

When an integral over an edge/face involves values from both adjacent elements (in the computation of jump residuals e.g.) it is necessary to have a common orientation of the edge/face from both elements. Only a common orientation of the edges/faces ensures that augmenting DIM dimensional barycentric coordinates of quadrature points on the edge/face to $\text{DIM}+1$ dimensional barycentric coordinates on the adjacent mesh elements results in the same points from both sides. Additionally, the calculation of Gram's determinant for the $\text{DIM}-1$ dimensional transformation as well as edge/face normals is needed. The following routines give such information.

```
const int *sort_face_indices(const EL *, int, int *);
REAL get_face_normal(const EL_INFO *, int, REAL *);
```

Description:

`sort_face_indices(el, i, vec)`: calculates a unique ordering of local vertex indices for the side (vertex/edge/face) opposite the i -th vertex of mesh element `el`. These indices give the same ordering of vertices from both adjacent mesh element, thus the barycentric coordinates of a quadrature point

(for $\text{DIM} - 1$ dimensional quadrature) may be transformed into the same vertex/edge/face points from both sides.

`get_face_normal(el_info, i, normal)`: calculates the *outer* unit normal vector for the side (vertex/edge/face) opposite the i -th vertex of element `el_info->el` and stores it in `normal`. `get_face_normal()` needs coordinate information filled in the `el_info` structure. The return value is Gram's determinant of the transformation from the $\text{DIM} - 1$ dimensional reference element.

3.9 Functions for the evaluation of finite elements

Finite element functions are evaluated locally on single elements using barycentric coordinates (compare Section 1.4.3). ALBERTA supplies several functions for calculating values and first and second derivatives of finite element functions on single elements. Functions for the calculation of derivatives are currently only implemented for (non-parametric) simplices.

Recalling (1.4.3) on page 30 we obtain for the value of a finite element function u_h on an element S

$$u_h(x(\lambda)) = \sum_{i=1}^m u_S^i \bar{\varphi}^i(\lambda) \quad \text{for all } \lambda \in \bar{S},$$

where $(\bar{\varphi}^1, \dots, \bar{\varphi}^m)$ is a basis of $\bar{\mathbb{P}}$ and (u_S^1, \dots, u_S^m) the local coefficient vector of u_h on S . Derivatives are evaluated on S by

$$\nabla u_h(x(\lambda)) = \Lambda^t \sum_{i=1}^m u_S^i \nabla_\lambda \bar{\varphi}^i(\lambda), \quad \lambda \in \bar{S}$$

and

$$D^2 u_h(x(\lambda)) = \Lambda^t \sum_{i=1}^m u_S^i D_\lambda^2 \bar{\varphi}^i(\lambda) \Lambda, \quad \lambda \in \bar{S},$$

where Λ is the Jacobian of the barycentric coordinates, compare Section 1.4.3.

These formulas are used for all evaluation routines. Information about values of basis functions and their derivatives can be calculated via function pointers in the `BAS_FCTS` structure. Additionally, the local coefficient vector and the Jacobian of the barycentric coordinates are needed (for the calculation of derivatives).

The following routines calculate values of a finite element function at a single point, given in barycentric coordinates:

```
REAL eval_uh(const REAL [DIM+1], const REAL *, const BAS_FCTS *);
const REAL *eval_grd_uh(const REAL [DIM+1], const REAL_D [DIM+1],
                        const REAL *, const BAS_FCTS *, REAL_D);
const REAL_D *eval_D2_uh(const REAL [DIM+1], const REAL_D [DIM+1],
                        const REAL *, const BAS_FCTS *, REAL_DD);
```

```

const REAL *eval_uh_d(const REAL [DIM+1], const REAL_D *,
                     const BAS_FCTS *, REAL_D);
const REAL_D *eval_grd_uh_d(const REAL [DIM+1], const REAL_D [DIM+1],
                           const REAL_D *, const BAS_FCTS *,
                           REAL_DD);
REAL eval_div_uh_d(const REAL [DIM+1], const REAL_D [DIM+1],
                  const REAL_D *, const BAS_FCTS *);
const REAL_DD *eval_D2_uh_d(const REAL [DIM+1], const REAL_D [DIM+1],
                           const REAL_D *, const BAS_FCTS *,
                           REAL_DD *);

```

Description:

In the following $\text{lambda} = \lambda$ are the barycentric coordinates at which the function is evaluated, $\text{Lambda} = \Lambda$ is the Jacobian of the barycentric coordinates, uh the local coefficient vector $(u_S^0, \dots, u_S^{m-1})$ (where u_S^i is a `REAL` or a `REAL_D`), and `bas_fcts` is a pointer to a `BAS_FCTS` structure, storing information about the set of local basis functions $(\bar{\varphi}^0, \dots, \bar{\varphi}^{m-1})$.

All functions returning a pointer to a vector or matrix provide memory for the vector or matrix in a static local variable. This area is overwritten during the next call. If the last argument of such a function is not `nil`, then memory is provided by the calling function, where values must be stored (optional memory pointer). These values are not overwritten. The memory area must be of correct size, no check is performed.

`eval_uh(lambda, uh, bas_fcts)`: the function returns $u_h(\lambda)$.

`eval_grd_uh(lambda, Lambda, uh, bas_fcts, grd)`: returns a pointer `ptr` to a vector of length `DIM_OF_WORLD` storing $\nabla u_h(\lambda)$, i.e.

$$\text{ptr}[i] = u_{h,x_i}(\lambda), \quad i = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`grd` is an optional memory pointer.

`eval_D2_uh(lambda, Lambda, uh, bas_fcts, D2)`: the function returns a pointer `ptr` to a matrix of size $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ storing $D^2 u_h(\lambda)$, i.e.

$$\text{ptr}[i][j] = u_{h,x_i x_j}(\lambda), \quad i, j = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`D2` is an optional memory pointer.

`eval_uh_d(lambda, uh, bas_fcts, val)`: the function returns a pointer `ptr` to a vector of length `DIM_OF_WORLD` storing $u_h(\lambda)$, i.e.

$$\text{ptr}[k] = u_{h_k}(\lambda), \quad k = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`val` is an optional memory pointer.

`eval_grd_uh_d(lambda, Lambda, uh, bas_fcts, grd)`: returns a pointer `ptr` to a matrix of size $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ storing $\nabla u_h(\lambda)$, i.e.

$$\text{ptr}[k][i] = u_{h_k, x_i}(\lambda), \quad k, i = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`grd` is an optional memory pointer.

`eval_div_uh_d(lambda, Lambda, uh, bas_fcts)`: returns $\text{div } u_h(\lambda)$.

`eval_D2_uh(lambda, Lambda, uh, bas_fcts, D2)`: the function returns a pointer `ptr` to a vector of $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ matrices of length DIM_OF_WORLD storing $D^2 u_h(\lambda)$, i.e.

$$\text{ptr}[k][i][j] = u_{h, x_i x_j}(\lambda), \quad k, i, j = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`D2` is an optional memory pointer;

Using pre-computed values of basis functions at the evaluation point, these routines can be implemented more efficiently.

```
REAL eval_uh_fast(const REAL *, const REAL *, int);
const REAL *eval_grd_uh_fast(const REAL_D [DIM+1], const REAL *,
                             const REAL (*)[DIM+1], int, REAL_D);
const REAL_D *eval_D2_uh_fast(const REAL_D [DIM+1], const REAL *,
                              const REAL (*)[DIM+1][DIM+1], int,
                              REAL_DD);
const REAL *eval_uh_d_fast(const REAL_D *, const REAL *, int,
                           REAL_D);
const REAL_D *eval_grd_uh_d_fast(const REAL_D [DIM+1],
                                 const REAL_D *,
                                 const REAL (*)[DIM+1], int,
                                 REAL_DD);
REAL eval_div_uh_d_fast(const REAL_D [DIM+1], const REAL_D *,
                       const REAL (*)[DIM+1], int);
const REAL_DD *eval_D2_uh_d_fast(const REAL_D [DIM+1],
                                 const REAL_D *,
                                 const REAL (*)[DIM+1][DIM+1],
                                 int, REAL_DD *);
```

Description:

In the following `Lambda = A` denotes the Jacobian of the barycentric coordinates, `uh` the local coefficient vector $(u_S^0, \dots, u_S^{m-1})$ (where u_S^i is a `REAL` or a `REAL_D`), and `m` the number of local basis functions on an element.

`eval_uh_fast(uh, phi, m)`: the function returns $u_h(\lambda)$;

`phi` is a vector storing the values $\bar{\varphi}^0(\lambda), \dots, \bar{\varphi}^{m-1}(\lambda)$.

`eval_grd_uh_fast(Lambda, uh, grd_phi, m, grd)`: the function returns a pointer `ptr` to a vector of length DIM_OF_WORLD storing $\nabla u_h(\lambda)$, i.e.

$$\text{ptr}[i] = u_{h, x_i}(\lambda), \quad i = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

`grd_phi` is a vector of $(\text{DIM} + 1)$ vectors storing $\nabla_\lambda \bar{\varphi}^0(\lambda), \dots, \nabla_\lambda \bar{\varphi}^{m-1}(\lambda)$; `grd` is an optional memory pointer.

`eval_D2_uh_fast(Lambda, uh, D2_phi, m, D2)`: returns a pointer `ptr` to a matrix of size $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ storing $D^2 u_h(\lambda)$, i.e.

$$\text{ptr}[i][j] = u_{h, x_i x_j}(\lambda), \quad i, j = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

D2_phi is a vector of $(\text{DIM}+1 \times \text{DIM}+1)$ matrices storing the second derivatives $D_\lambda^2 \bar{\varphi}^0(\lambda), \dots, D_\lambda^2 \bar{\varphi}^{m-1}(\lambda)$; **D2** is an optional memory pointer.

eval_uh_d_fast(uh, phi, m, val): the function returns a pointer **ptr** to a vector of **DIM_OF_WORLD** vectors of length **DIM_OF_WORLD** storing $\nabla u_h(\lambda)$, i.e.

$$\text{ptr}[\mathbf{k}][\mathbf{i}] = u_{h\mathbf{k},x_{\mathbf{i}}}(\lambda), \quad \mathbf{k}, \mathbf{i} = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

phi is a vector storing the values $\bar{\varphi}^0(\lambda), \dots, \bar{\varphi}^{m-1}(\lambda)$; **val** is an optional memory pointer.

eval_grd_uh_d_fast(Lambda, uh, grd_phi, m, grd): the function returns a pointer **ptr** to a vector of **DIM_OF_WORLD** vectors of length **DIM_OF_WORLD** storing $\nabla u_h(\lambda)$, i.e.

$$\text{ptr}[\mathbf{k}][\mathbf{i}] = u_{h\mathbf{k},x_{\mathbf{i}}}(\lambda), \quad \mathbf{k}, \mathbf{i} = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

grd_phi is a vector of $(\text{DIM}+1)$ vectors storing $\nabla_\lambda \bar{\varphi}^0(\lambda), \dots, \nabla_\lambda \bar{\varphi}^{m-1}(\lambda)$; **grd** is an optional memory pointer.

eval_div_uh_d_fast(Lambda, uh, grd_phi, m): returns $\text{div } u_h(\lambda)$;

grd_phi is a vector of $(\text{DIM}+1)$ vectors storing $\nabla_\lambda \bar{\varphi}^0(\lambda), \dots, \nabla_\lambda \bar{\varphi}^{m-1}(\lambda)$.

eval_D2_uh_d_fast(Lambda, uh, D2_phi, m, D2): the function returns a pointer **ptr** to a vector of $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ matrices of length **DIM_OF_WORLD** storing $D^2 u_h(\lambda)$, i.e.

$$\text{ptr}[\mathbf{k}][\mathbf{i}][\mathbf{j}] = u_{h\mathbf{k},x_{\mathbf{i}}x_{\mathbf{j}}}(\lambda), \quad \mathbf{k}, \mathbf{i}, \mathbf{j} = 0, \dots, \text{DIM_OF_WORLD} - 1;$$

D2_phi is a vector of $(\text{DIM}+1 \times \text{DIM}+1)$ matrices storing the second derivatives $D_\lambda^2 \bar{\varphi}^0(\lambda), \dots, D_\lambda^2 \bar{\varphi}^{m-1}(\lambda)$; **D2** is an optional memory pointer.

One important task is the evaluation of finite element functions at all quadrature nodes for a given quadrature formula. Using the **QUAD_FAST** data structures, the values of the basis functions are known at the quadrature nodes which results in an efficient calculation of values and derivatives of finite element functions at these quadrature points.

```
const REAL *uh_at_qp(const QUAD_FAST *, const REAL *, REAL *);
const REAL_D *grd_uh_at_qp(const QUAD_FAST *, const REAL_D [DIM+1],
                           const REAL *, REAL_D *);
const REAL_DD *D2_uh_at_qp(const QUAD_FAST *, const REAL_D [DIM+1],
                           const REAL *, REAL_DD *);

const REAL_D *uh_d_at_qp(const QUAD_FAST *, const REAL_D *,
                        REAL_D *);
const REAL_DD *grd_uh_d_at_qp(const QUAD_FAST *,
                              const REAL_D [DIM+1],
                              const REAL_D *, REAL_DD *);
const REAL *div_uh_d_at_qp(const QUAD_FAST *, const REAL_D [DIM+1],
                          const REAL_D *, REAL *);
```

```

const REAL_DD (*D2_uh_d_at_qp(const QUAD_FAST *,
                             const REAL_D [DIM+1], const REAL_D *,
                             REAL_DD (*)[DIM_OF_WORLD])
               ) [DIM_OF_WORLD];

```

Description: In the following, `uh` denotes a given local coefficient vector $(u_S^0, \dots, u_S^{m-1})$ (where u_S^i is a `REAL` or a `REAL_D`) on an element. We use `lambda` for the quadrature nodes of `quad_fast` and `n_points` for their number.

`uh_at_qp(quad_fast, uh, val)`: the function returns a pointer `ptr` to a vector of length `n_points` storing the values of u_h at all quadrature points of `quad_fast->quad`, i.e.

$$\text{ptr}[l] = u_h(\text{lambda}[l]), \quad l = 0, \dots, \text{n_points} - 1;$$

the `INIT_PHI` flag must be set in `quad_fast->init_flag`; `val` is an optional memory pointer.

`grd_uh_at_qp(quad_fast, Lambda, uh, grd)`: returns a pointer `ptr` to a vector of length `n_points` of `DIM_OF_WORLD` vectors storing ∇u_h at all quadrature points of `quad_fast->quad`, i.e.

$$\text{ptr}[l][i] = u_{h,x_i}(\text{lambda}[l])$$

for $l = 0, \dots, \text{n_points} - 1$, and $i = 0, \dots, \text{DIM_OF_WORLD} - 1$; the `INIT_GRD_PHI` flag must be set in `quad_fast->init_flag`; `grd` is an optional memory pointer.

`D2_uh_at_qp(quad_fast, Lambda, uh, D2)`: returns a pointer `ptr` to a vector of length `n_points` of $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ matrices storing $D^2 u_h$ at all quadrature points of `quad_fast->quad`, i.e.

$$\text{ptr}[l][i][j] = u_{h,x_i x_j}(\text{lambda}[l])$$

for indices $l = 0, \dots, \text{n_points} - 1$, and $i, j = 0, \dots, \text{DIM_OF_WORLD} - 1$; the `INIT_D2_PHI` flag must be set in `quad_fast->init_flag`; `D2` is an optional memory pointer.

`uh_d_at_qp(quad_fast, uh, val)`: the function returns a pointer `ptr` to a vector of length `n_points` of `DIM_OF_WORLD` vectors storing the values of u_h at all quadrature points of `quad_fast->quad`, i.e.

$$\text{ptr}[l][k] = u_{h,k}(\text{lambda}[l])$$

where $l = 0, \dots, \text{n_points} - 1$, and $k = 0, \dots, \text{DIM_OF_WORLD} - 1$; the `INIT_PHI` flag must be set in `quad_fast->init_flag`; `val` is an optional memory pointer.

`grd_uh_d_at_qp(quad_fast, Lambda, uh, grd)`: returns a pointer `ptr` to a vector of length `n_points` of $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$ matrices storing ∇u_h at all quadrature points of `quad_fast->quad`, i.e.

$$\text{ptr}[l][k][i] = u_{h_{k,x_i}}(\text{lambda}[l])$$

where $l = 0, \dots, \text{n_points} - 1$, and $k, i = 0, \dots, \text{DIM_OF_WORLD} - 1$; the `INIT_GRD_PHI` flag must be set in `quad_fast->init_flag`; `grd` is an optional memory pointer.

`D2_uh_d_at_qp(quad_fast, Lambda, uh, D2)`: returns a pointer `ptr` to a vector of tensors of length `n_points` storing $D^2 u_h$ at all quadrature points of `quad_fast->quad`; the tensors are of size $(\text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD} \times \text{DIM_OF_WORLD})$:

$$\text{ptr}[l][k][i][j] = u_{h_{k,x_i x_j}}(\text{lambda}[l])$$

where $l = 0, \dots, \text{n_points} - 1$, and $k, i, j = 0, \dots, \text{DIM_OF_WORLD} - 1$; the `INIT_D2_PHI` flag must be set in `quad_fast->init_flag`; `D2` is an optional memory pointer;

3.10 Calculation of norms for finite element functions

ALBERTA supplies functions for the calculation of the L^2 norm and H^1 semi-norm of a given scalar or vector valued finite element function, currently only implemented for non-parametric meshes.

```
REAL H1_norm_uh(const QUAD *, const DOF_REAL_VEC *);
REAL L2_norm_uh(const QUAD *, const DOF_REAL_VEC *);
REAL H1_norm_uh_d(const QUAD *, const DOF_REAL_D_VEC *);
REAL L2_norm_uh_d(const QUAD *, const DOF_REAL_D_VEC *);
```

Description:

`H1_norm_uh(quad, uh)`: returns an approximation to the H^1 semi norm of a finite element function, i.e. $(\int_{\Omega} |\nabla u_h|^2)^{1/2}$; the coefficient vector of the vector is stored in `uh`; the domain is given by `uh->fe_space->mesh`; the element integrals are approximated by the numerical quadrature `quad`, if `quad` is not `nil`; otherwise a quadrature which is exact of degree $2 * \text{uh->fe_space->bas_fcts->degree} - 2$ is used.

`L2_norm_uh(quad, uh)`: returns an approximation to the L^2 norm of a finite element function, i.e. $(\int_{\Omega} |u_h|^2)^{1/2}$; the coefficient vector of the vector is stored in `uh`; the domain is given by `uh->fe_space->mesh`; the element integrals are approximated by the numerical quadrature `quad`, if `quad` is not `nil`; otherwise a quadrature which is exact of degree $2 * \text{uh->fe_space->bas_fcts->degree}$ is used.

`H1_norm_uh_d(quad, uh_d)`: returns an approximation to the H^1 semi norm of a vector valued finite element function; the coefficient vector of the vector is stored in `uh_d`; the domain is given by `uh_d->fe_space->mesh`; the element integrals are approximated by the numerical quadrature `quad`, if `quad` is not `nil`; otherwise a quadrature which is exact of degree $2 * \text{uh_d->fe_space->bas_fcts->degree} - 2$ is used.

L2_norm_uh_d(quad, uh_d): returns an approximation to the L^2 norm of a vector valued finite element function; the coefficient vector of the vector is stored in **uh_d**; the domain is given by **uh_d->fe_space->mesh**; the element integrals are approximated by the numerical quadrature **quad**, if **quad** is not **nil**; otherwise a quadrature which is exact of degree $2*uh_d->fe_space->bas_fcts->degree$ is used;

3.11 Calculation of errors of finite element approximations

For test purposes it is convenient to calculate the “exact error” between a finite element approximation and the exact solution. ALBERTA supplies functions to calculate the error in several norms. For test purposes, the integral error routines may be used as “error estimators” in an adaptive method. The local element error $\int_S |\nabla(u - u_h)|^2$ or $\int_S |u - u_h|^2$ can be used as an error indicator and can be stored on the element leaf data, e.g.

```
REAL max_err_at_qp(REAL (*)(const REAL_D), const DOF_REAL_VEC *,
                  const QUAD *);
REAL H1_err(const REAL (*)(const REAL_D), const DOF_REAL_VEC *,
            const QUAD *, int, REAL (*)(EL *), REAL *);
REAL L2_err(REAL (*)(const REAL_D), const DOF_REAL_VEC *,
            const QUAD *, int, REAL (*)(EL *), REAL *);
REAL max_err_d_at_qp(const REAL (*)(const REAL_D, REAL_D),
                    const DOF_REAL_D_VEC *, const QUAD *);
REAL H1_err_d(const REAL_D (*)(const REAL_D, REAL_D),
              const DOF_REAL_D_VEC *, const QUAD *, int,
              REAL (*)(EL *), REAL *);
REAL L2_err_d(const REAL (*)(const REAL_D, REAL_D),
              const DOF_REAL_D_VEC *, const QUAD *, int,
              REAL (*)(EL *), REAL *);
```

Description:

max_err_at_qp(u, uh, quad): returns the maximal error, $\max |u - u_h|$, between the true solution and the approximation at all quadrature nodes on all elements of a mesh; **u** is a pointer to a function for the evaluation of the true solution, **uh** stores the coefficients of the approximation, **uh->fe_space->mesh** is the underlying mesh, and **quad** is the quadrature which gives the quadrature nodes; if **quad** is **nil**, a quadrature which is exact of degree $2*uh->fe_space->bas_fcts->degree-2$ is used.

H1_err(grd_u, uh, quad, rel_err, rw_el_err, max): returns an approximation to the absolute error $(\int_\Omega |\nabla(u - u_h)|^2)^{1/2}$ (if **rel_err** is **false**) or relative error $(\int_\Omega |\nabla(u - u_h)|^2 / \int_\Omega |\nabla u|^2)^{1/2}$ (if **rel_err** is **true**) between the true solution and the approximation in the H^1 semi norm; **grd_u** is a pointer to a function for the evaluation of the gradient of the true solution returning a **DIM_OF_WORLD** vector storing this gradient, **uh**

stores the coefficients of the approximation, `uh->fe_space->mesh` is the underlying mesh, and `quad` is the quadrature for the approximation of the element integrals; if `quad` is `nil`, a quadrature which is exact of degree $2*uh->fe_space->bas_fcts->degree-2$ is used;

if `rw_el_err` is not `nil`, the return value of `(*rw_el_err)(el)` provides for each mesh element `el` an address where the local error is stored; if `max` is not `nil`, `*max` is the maximal local error on an element on output.

`L2_err(u, uh, quad, rel_err, rw_el_err, max)`: the function returns an approximation to the absolute error $(\int_{\Omega} |u - u_h|^2)^{1/2}$ (if `rel_err` is `false`) or the relative error $(\int_{\Omega} |u - u_h|^2 / \int_{\Omega} |u|^2)^{1/2}$ (if `rel_err` is `true`) between the true solution and the approximation in the L^2 norm,

`u` is a pointer to a function for the evaluation of the true solution, `uh` stores the coefficients of the approximation, `uh->fe_space->mesh` is the underlying mesh, and `quad` is the quadrature for the approximation of the element integrals; if `quad` is `nil`, a quadrature which is exact of degree $2*uh->fe_space->bas_fcts->degree-2$ is used;

if `rw_el_err` is not `nil`, the return value of `(*rw_el_err)(el)` provides for each mesh element `el` an address where the local error is stored; if `max` is not `nil`, `*max` is the maximal local error on an element on output.

`max_err_at_qp_d(u_d, uh_d, quad)`: the function returns the maximal error between the true solution and the approximation at all quadrature nodes on all elements of a mesh; `u_d` is a pointer to a function for the evaluation of the true solution returning a `DIM_OF_WORLD` vector storing the value of the function, `uh_d` stores the coefficients of the approximation, `uh_d->fe_space->mesh` is the underlying mesh, and `quad` is the quadrature which gives the quadrature nodes; if `quad` is `nil`, a quadrature which is exact of degree $2*uh_d->fe_space->bas_fcts->degree-2$ is used.

`H1_err2_d(grd_u_d, uh_d, quad, rel_err, rw_el_err, max)`: returns an approximation to the absolute error (if `rel_err` is `false`) or relative error (if `rel_err` is `true`) between the true solution and the approximation in the H^1 semi norm;

`grd_u_d` is a pointer to a function for the evaluation of the Jacobian of the true solution returning a `DIM_OF_WORLD` \times `DIM_OF_WORLD` matrix storing this Jacobian, `uh_d` stores the coefficients of the approximation, `uh_d->fe_space->mesh` is the underlying mesh, and `quad` is the quadrature for the approximation of the element integrals; if `quad` is `nil`, a quadrature which is exact of degree $2*uh_d->fe_space->bas_fcts->degree-2$ is used; if `rw_el_err` is not `nil`, the return value of `(*rw_el_err)(el)` provides for each mesh element `el` an address where the local error is stored; if `max` is not `nil`, `*max` is the maximal local error on an element on output.

`L2_err2_d(u_d, uh_d, quad, rel_err, rw_el_err, max)`: the function returns an approximation to the absolute error (if `rel_err` is `false`), or the

relative error (if `rel_err` is `true`) between the true solution and the approximation in the L^2 norm;

`u_d` is a pointer to a function for the evaluation of the true solution returning a `DIM_OF_WORLD` vector storing the value of the function, `uh_d` stores the coefficients of the approximation, `uh_d->fe_space->mesh` is the underlying mesh, and `quad` is the quadrature for the approximation of the element integrals; if `quad` is `nil`, a quadrature which is exact of degree $2 * \text{uh_d->fe_space->bas_fcts->degree} - 2$ is used;

if `rw_el_err` is not `nil`, the return value of `(*rw_el_err)(e1)` provides for each mesh element `e1` an address where the local error is stored; if `max` is not `nil`, `*max` is the maximal local error on an element on output.

3.12 Tools for the assemblage of linear systems

This section describes data structures and subroutines for matrix and vector assembly. Section 3.12.1 presents basic routines for the update of global matrices and vectors by adding contributions from one single element. Data structures and routines for global matrix assembly are described in Section 3.12.2. This includes library routines for the efficient implementation of a general second order linear elliptic operator. Section 3.12.3 presents data structures and routines for the handling of pre-computed integrals, which are used to speed up calculations in the case of problems with constant coefficients. The assembly of (right hand side) vectors is described in Section 3.12.4. The incorporation of Dirichlet boundary values into the right hand side is presented in Section 3.12.5. Finally, routines for generation of interpolation coefficients are described in Section 3.12.6.

3.12.1 Assembling matrices and right hand sides

The usual way to assemble the system matrix and the load vector is to loop over all (leaf) elements, calculate the local element contributions and add these to the global system matrix and the global load vector. The updating of the load vector is rather easy. The contribution of a local degree of freedom is added to the value of the corresponding global degree of freedom. Here we have to use the function j_S defined on each element S in (1.4) on page 30. It combines uniquely the local DOFs with the global ones. The basis functions provide in the `BAS_FCTS` structure the entry `get_dof_indices()` which is an implementation of j_S , see Section 3.5.1.

The updating of the system matrix is not that easy. As mentioned in Section 1.4.7, the system matrix is usually sparse and we use special data structures for storing these matrices, compare Section 3.3.4. For sparse matrices we do not have for each DOF a matrix row storing values for all other DOFs; only the values for pairs of DOFs are stored, where the corresponding

global basis functions have a common support. Usually, the exact number of entries in one row of a sparse matrix is not known a priori and can change during grid modifications.

Thus, we use the following concept: A call of `clear_matrix()` will not set all matrix entries to zero, but will remove all matrix rows from the matrix, compare the description of this function on page 169. During the updating of a matrix for the value corresponding to a pair of local DOFs (i, j) , we look in the $j_S(i)$ th row of the matrix for a column $j_S(j)$ (the `col` member of `matrix_row`); if such an entry exists, we add the current contribution; if this entry does not yet exist we will create a new entry, set the current value and column number. This creation may include an enlargement of the row, by linking a new matrix row to the list of matrix rows, if no space for a new entry is left. After the assemblage we then have a sparse matrix, storing all values for pairs of global basis functions with common support.

The function which we describe now allows also to handle matrices where the DOFs indexing the rows can differ from the DOFs indexing the columns; this makes the combination of DOFs from different finite element spaces possible. Currently, only one reference to a `FE_SPACE` structure is included in the `DOF_MATRIX` structure. Thus, the handling of two different DOF sets is not yet fully implemented, especially the handling of matrices during `dof_compress()` may produce wrong results. Such matrices should be cleared by calling `clear_dof_matrix()` before a call to `dof_compress()`.

The following functions can be used on elements for updating matrices and vectors.

```
void add_element_matrix(DOF_MATRIX *, REAL, int row, int col,
                      const DOF *, const DOF *, const REAL **,
                      const S_CHAR *);
void add_element_vec(DOF_REAL_VEC *, REAL, int, const DOF *,
                   const REAL *, const S_CHAR *);
void add_element_d_vec(DOF_REAL_D_VEC *, REAL, int, const DOF *,
                     const REAL_D *, const S_CHAR *);
```

Description:

`add_element_matrix(mat, fac, nr, nc, rdof, cdof, el_mat, bound):`
 updates the `DOF_MATRIX` `mat` by adding element contributions;
`fac` is a multiplier for the element contributions; usually `fac` is 1 or -1;
`nr` is the number of rows of the element matrix;
`nc` is the number of columns of the element matrix; `nc` may be less or equal to zero if the DOFs indexing the columns are the same as the DOFs indexing the rows; in this case `nc = nr` is used;
`rdof` is a vector of length `nr` storing the global row indices;
`cdof` is a vector of length `nc` storing the global column indices, `cdof` may be a `nil` pointer if the DOFs indexing the columns are the same as the DOFs indexing the rows; in this case `cdof = rdof` is used;
`el_mat` is a matrix of size `nr × nc` storing the element contributions;

bound is an optional **S_CHAR** vector of length **n_row**; **bound** *must* be a **nil** pointer if row indices and column indices are not the same; if row indices and column indices are the same and **bound** is not **nil** it holds the boundary type of the global DOFs; contributions are not added in rows corresponding to a Dirichlet DOF; the diagonal matrix entry of a Dirichlet DOF is set to 1.0.

If row indices and column indices differ or if **bound** is a **nil** pointer, then for all **i** the values **fac*el_mat[i][j]** are added to the entries (**rdof[i]**, **cdof[j]**) in the global matrix **mat** ($0 \leq i < \mathbf{nr}$, $0 \leq j < \mathbf{nc}$); if such an entry exists in the **rdof[i]**-th row of the matrix the value is added; otherwise a new entry is created in the row, the value is set and the column number is set to **cdof[j]**; this may include an enlargement of the row by adding a new **MATRIX_ROW** structure to the list of matrix rows.

If row DOFs and column DOFs are the same and **bound** is not **nil** these values are only added for row indices **i** with **bound[i] < DIRICHLET**; for row indices **i** with **bound[i] >= DIRICHLET** only the diagonal element is set to 1.0; the values in the **i**-th row of **el_mat** are ignored.

If row indices and column indices are the same, the diagonal element is *always* the first entry in a matrix row; this makes the access to the diagonal element easy for a diagonal preconditioner, e.g.

add_element_vec(drv, fac, n_dof, dof, el_vec, bound): updates a given **DOF_REAL_VEC** **drv**;

fac is a multiplier for the element contributions; usually **fac** is 1 or -1;

n_dof is the number of local degrees of freedom;

dof is a vector of length **n_dof** storing the global DOFs, i.e. **dof[i]** is the global DOF of the **i**-th contribution;

el_vec is a **REAL** vector of length **n_dof** storing the element contributions;

bound is an optional vector of length **n_dof**; if **bound** is not **nil** it holds the boundary type of the global DOFs; contributions are only added for non-Dirichlet DOFs.

For all **i** or, if **bound** is not **nil**, for all **i** with **bound[i] < DIRICHLET** ($0 \leq i < \mathbf{n_dof}$) the value **fac*el_vec[i]** is added to **drv->vec[dof[i]]**; values in **drv->vec** are not changed for Dirichlet DOFs.

add_element_d_vec(drdv, fac, n_dof, dof, el_vec, bound): updates the **DOF_REAL_D_VEC** **drdv**;

fac is a multiplier for the element contributions; usually **fac** is 1 or -1;

n_dof is the number of local degrees of freedom;

dof is a vector of length **n_dof** storing the global DOFs, i.e. **dof[i]** is the global DOF of the **i**-th contribution;

el_vec is a **REAL_D** vector of length **n_dof** storing the element contributions;

bound is an optional vector of length **n_dof**; if **bound** is not **nil** it holds the boundary type of the global DOFs; contributions are only added for non-Dirichlet DOFs.

For all i or, if `bound` is not `nil`, for all i with `bound[i] < DIRICHLET` ($0 \leq i < n_dof$) the value `fac*el_vec[i]` is added to `drdv->vec[dof[i]]`; values in the `drdv->vec` are not changed for Dirichlet DOFs.

3.12.2 Data structures and function for matrix assemblage

The following structure holds full information for the assembling of element matrices. This structure is used by the function `update_matrix()` described below.

```
typedef struct el_matrix_info  EL_MATRIX_INFO;

struct el_matrix_info
{
    int                n_row;
    const DOF_ADMIN    *row_admin;
    const DOF          *(*get_row_dof)(const EL *, const DOF_ADMIN *,
                                      DOF *);

    int                n_col;
    const DOF_ADMIN    *col_admin;
    const DOF          *(*get_col_dof)(const EL *, const DOF_ADMIN *,
                                      DOF *);

    const S_CHAR       *(*get_bound)(const EL_INFO *, S_CHAR *);

    REAL               factor;

    const REAL         *(*el_matrix_fct)(const EL_INFO *, void *);
    void               *fill_info;

    FLAGS              fill_flag;
};
```

Description:

`n_row`: number of rows of the element matrix.

`row_admin`: pointer to a `DOF_ADMIN` structure for the administration of DOFs indexing the rows of the matrix.

`get_row_dof`: pointer to a function for the access of the global row DOFs on a single element; `get_row_dof(el, row_admin, row_dof)` returns a pointer to a vector of length `n_row` storing the global DOFs indexing the rows of the matrix; if `row_dof` is a `nil` pointer, `get_row_dof()` has to provide memory for storing this vector, which may be overwritten on the next call; otherwise the DOFs have to be stored at `row_dof`; usually, `get_row_dof()` is the `get_dof_indices()` function from a `BAS_FCTS` structure (compare Section 3.5.1).

n_col: number of columns of the element matrix; **n_col** may be less or equal to zero if the DOFs indexing the columns are the same as the DOFs indexing the rows; in this case **n_col = n_row** is used.

col_admin: pointer to a **DOF_ADMIN** structure for the administration of DOFs indexing the columns of the matrix; **col_admin** may be a **nil** pointer if the column DOFs are the same as the row DOFs.

get_col_dof: pointer to a function for the access of the global row DOFs on a single element; **get_col_dof** may be a **nil** pointer if the column DOFs are the same as the row DOFs; if row and column DOFs differ then the column DOFs are accessed by **get_col_dof()** otherwise the vector accessed by **get_row_dof()** is used for the columns also; **get_col_dof(el, col_admin, col_dof)** returns a pointer to a vector of length **n_col** storing the global DOFs indexing the rows of the matrix; if **col_dof** is a **nil** pointer, **get_col_dof()** has to provide memory for storing this vector, which may be overwritten on the next call; otherwise the DOFs have to be stored at **col_dof**; usually, **get_col_dof()** is the **get_dof_indices()** function from a **BAS_FCTS** structure (compare Section 3.5.1).

get_bound: is an optional pointer to a function which provides information about the boundary type of DOFs; **get_bound** *must* be a **nil** pointer if row indices and column indices are not the same; otherwise **get_bound(el_info, bound)** returns a pointer to a vector of length **n_row** storing the boundary type of the local DOFs; this pointer is the optional pointer to a vector holding boundary information of the function **add_element_matrix()** described above; if **bound** is a **nil** pointer, **get_bound()** has to provide memory for storing this vector, which may be overwritten on the next call; otherwise the DOFs have to be stored at **bound**; usually, **get_bound()** is the **get_bound()** function from a **BAS_FCTS** structure (compare Section 3.5.1).

factor: is a multiplier for the element contributions; usually **factor** is 1 or -1.

el_matrix_fct: is a pointer to a function for the computation of the element matrix; **el_matrix_fct(el_info, fill_info)** returns a pointer to a matrix of size $n_row \times n_col$ storing the element matrix on element **el_info->el**; **fill_info** is a pointer to data needed by **el_matrix_fct()**; the function has to provide memory for storing the element matrix, which can be overwritten on the next call.

fill_info: pointer to data needed by **el_matrix_fct()**; will be given as second argument to this function.

fill_flag: the flag for the mesh traversal for assembling the matrix.

The following function updates a matrix by assembling element contributions during mesh traversal; information for computing the element matrices is provided in an **EL_MATRIX_INFO** structure:

```
void update_matrix(DOF_MATRIX *matrix, const EL_MATRIX_INFO *);
```

Description:

update_matrix(matrix, info): updates the matrix **matrix** by traversing the underlying mesh and assembling the element contributions into the matrix; information about the computation of element matrices and connection of local and global DOFs is stored in **info**;

the flags for the mesh traversal of the mesh **matrix->fe_space->mesh** are stored at **info->fill_flag** which specifies the elements to be visited and information that should be present on the elements for the calculation of the element matrices and boundary information (if **info->get_bound** is not **nil**).

Information about row DOFs is accessed elementwise by the function **info->get_row_dof** using **info->row_admin**; this vector is also used for the column DOFs if **info->n_col** is less or equal to zero, or if one of **info->get_col_admin** or **info->get_col_dof** is a **nil** pointer; when row and column DOFs are the same, the boundary type of the DOFs is accessed by **info->get_bound** if **info->get_bound** is not a **nil** pointer; then the element matrix is computed by **info->el_matrix_fct(el_info, info->fill_info)**; these contributions, multiplied by **info->factor**, are eventually added to **matrix** by a call of **add_element_matrix()** with all information about row and column DOFs, the element matrix, and boundary types, if available;

update_matrix() only adds element contributions; this makes several calls for the assemblage of one matrix possible; before the first call, the matrix should be cleared by calling **clear_dof_matrix()**.

Now we want to describe some tools which enable an easy assemblage of the system matrix. For this we have to provide a function for the calculation of the element matrix. For a general elliptic problem the element matrix $\mathbf{L}_S = (L_S^{ij})_{i,j=1,\dots,m}$ is given by (recall (1.16) on page 40)

$$\begin{aligned} L_S^{ij} = & \int_{\hat{S}} \nabla_{\lambda} \bar{\varphi}^i(\lambda(\hat{x})) \cdot \bar{A}(\lambda(\hat{x})) \nabla_{\lambda} \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x} \\ & + \int_{\hat{S}} \bar{\varphi}^i(\lambda(\hat{x})) \bar{b}(\lambda(\hat{x})) \cdot \nabla_{\lambda} \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x} \\ & + \int_{\hat{S}} \bar{c}(\lambda(\hat{x})) \bar{\varphi}^i(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x}, \end{aligned}$$

where \bar{A} , \bar{b} , and \bar{c} are functions depending on given data and on the actual element, namely

$$\begin{aligned} \bar{A}(\lambda) &:= (\bar{a}_{kl}(\lambda))_{k,l=0,\dots,d} := |\det DF_S(\hat{x}(\lambda))| A(x(\lambda)) A(x(\lambda)) A^t(x(\lambda)), \\ \bar{b}(\lambda) &:= (\bar{b}_l(\lambda))_{l=0,\dots,d} := |\det DF_S(\hat{x}(\lambda))| A(x(\lambda)) b(x(\lambda)), \quad \text{and} \\ \bar{c}(\lambda) &:= |\det DF_S(\hat{x}(\lambda))| c(x(\lambda)). \end{aligned}$$

Having access to functions for the evaluation of \bar{A} , \bar{b} , and \bar{c} at given quadrature nodes, the above integrals can be computed by some general routine for

any set of local basis functions using quadrature. Additionally, if a coefficient is piecewise constant on the mesh, only an integration of basis functions has to be done (compare (1.16) on page 42) for this term. Here we can use pre-computed integrals of the basis functions on the standard element and transform them to the actual element. Such a computation is usually much faster than using quadrature on each single element. Data structures for storing such pre-computed values are described in Section 3.12.3.

For the assemblage routines which we will describe now, we use the following slight generalization: In the discretization of the first order term, sometimes integration by parts is used too. For a divergence free vector field b and purely Dirichlet boundary values this leads for instance to

$$\int_{\Omega} \varphi b \cdot \nabla u \, dx = \frac{1}{2} \left(\int_{\Omega} \varphi b \cdot \nabla u \, dx - \int_{\Omega} \nabla \varphi \cdot b u \, dx \right)$$

yielding a modified first order term for the element matrix

$$\int_{\hat{S}} \bar{\varphi}^i(\lambda(\hat{x})) \frac{1}{2} \bar{b}(\lambda(\hat{x})) \cdot \nabla_{\lambda} \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x} - \int_{\hat{S}} \nabla_{\lambda} \bar{\varphi}^i(\lambda(\hat{x})) \cdot \frac{1}{2} \bar{b}(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x}.$$

Secondly, we allow that we have two finite element spaces with local basis functions $\{\bar{\psi}_i\}_{i=1,\dots,n}$ and $\{\bar{\varphi}_i\}_{i=1,\dots,m}$.

In general, the following contributions of the element matrix $\mathbf{L}_S = (L_S^{ij})_{i=1,\dots,n, j=1,\dots,m}$ have to be computed:

$$\begin{aligned} & \int_{\hat{S}} \nabla_{\lambda} \bar{\psi}^i(\lambda(\hat{x})) \cdot \bar{A}(\lambda(\hat{x})) \nabla_{\lambda} \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x} && \text{second order term,} \\ & \int_{\hat{S}} \bar{\psi}^i(\lambda(\hat{x})) \bar{b}^0(\lambda(\hat{x})) \cdot \nabla_{\lambda} \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x} && \text{first order terms,} \\ & \int_{\hat{S}} \nabla_{\lambda} \bar{\psi}^i(\lambda(\hat{x})) \cdot \bar{b}^1(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x} && \\ & \int_{\hat{S}} \bar{c}(\lambda(\hat{x})) \bar{\psi}^i(\lambda(\hat{x})) \bar{\varphi}^j(\lambda(\hat{x})) \, d\hat{x} && \text{zero order term,} \end{aligned}$$

where for instance $\bar{b}^0 = \bar{b}$ and $\bar{b}^1 = 0$, or using integration by parts $\bar{b}^0 = \frac{1}{2} \bar{b}$ and $\bar{b}^1 = -\frac{1}{2} \bar{b}$.

In order to store information about the finite element spaces, the problem dependent functions \bar{A} , \bar{b}^0 , \bar{b}^1 , \bar{c} and the quadrature that should be used for the numerical integration of the element matrix, we define the following data structure:

```
typedef struct operator_info  OPERATOR_INFO;

struct operator_info
{
    const FE_SPACE *row_fe_space;
    const FE_SPACE *col_fe_space;
```



```

const QUAD *quad[3];

void      (*init_element)(const EL_INFO *, const QUAD *[3],
                          void *);
const REAL (*(*LALt)(const EL_INFO *, const QUAD *, int, void *)
              ) [DIM+1];
int        LALt_pw_const;
int        LALt_symmetric;
const REAL (*(*Lb0)(const EL_INFO *, const QUAD *, int, void *)
            ) [DIM+1];
int        Lb0_pw_const;
const REAL (*(*Lb1)(const EL_INFO *, const QUAD *, int, void *)
            ) [DIM+1];
int        Lb1_pw_const;
int        Lb0_Lb1_anti_symmetric;
REAL       (*c)(const EL_INFO *, const QUAD *, int, void *);
int        c_pw_const;

int        use_get_bound;
void       *user_data;
FLAGS      fill_flag;
};

```

Description:

row_fe_space: pointer to a finite element space connected to the row DOFs of the resulting matrix.

col_fe_space: pointer to a finite element space connected to the column DOFs of the resulting matrix.

quad: vector with pointers to quadratures; **quad[0]** is used for the integration of the zero order term, **quad[1]** for the first order term(s), and **quad[2]** for the second order term.

init_element: pointer to a function for doing an initialization step on each element; **init_element** may be a **nil** pointer;

if **init_element** is not **nil**, **init_element(el_info, quad, user_data)** is the first statement executed on each element **el_info->el** and may initialize data which is used by the functions **LALt()**, **Lb0()**, **Lb1()**, and/or **c()** (calculate the Jacobian of the barycentric coordinates in the 1st and 2nd order terms or the element volume for all order terms, e.g.); **quad** is a pointer to a vector of quadratures which is actually used for the integration of the various order terms and **user_data** may hold a pointer to user data, filled by **init_element()**, e.g.;

LALt: is a pointer to a function for the evaluation of \bar{A} at quadrature nodes on the element; **LALt** may be a **nil** pointer, if no second order term has to be integrated;

if **LALt** is not a **nil** pointer, **LALt(el_info, quad, iq, user_data)** returns a pointer to a matrix of size $\text{DIM}+1 \times \text{DIM}+1$ storing the value of \bar{A} at

`quad->lambda[iq]`; `quad` is the quadrature for the second order term and `user_data` is a pointer to user data;

LAlt_pw_const: should be `true` if \bar{A} is piecewise constant on the mesh (constant matrix A on a non-parametric mesh, e.g.); thus integration of the second order term can use pre-computed integrals of the basis functions on the standard element; otherwise integration is done by using quadrature on each element;

LAlt_symmetric: should be `true` if \bar{A} is a symmetric matrix; if the finite element spaces for rows and columns are the same, only the diagonal and the upper part of the element matrix for the second order term have to be computed; elements of the lower part can then be set using the symmetry; otherwise the complete element matrix has to be calculated;

Lb0: is a pointer to a function for the evaluation of \bar{b}^0 , at quadrature nodes on the element; **Lb0** may be a `nil` pointer, if this first order term has not to be integrated;

if **Lb0** is not `nil`, `Lb0(el_info, quad, iq, user_data)` returns a pointer to a vector of length `DIM+1` storing the value of \bar{b}^0 at `quad->lambda[iq]`; `quad` is the quadrature for the first order term and `user_data` is a pointer to user data;

Lb0_pw_const: should be `true` if \bar{b}^0 is piecewise constant on the mesh (constant vector b on a non-parametric mesh, e.g.); thus integration of the first order term can use pre-computed integrals of the basis functions on the standard element; otherwise integration is done by using quadrature on each element;

Lb1: is a pointer to a function for the evaluation of \bar{b}^1 , at quadrature nodes on the element; **Lb1** may be a `nil` pointer, if this first order term has not to be integrated;

if **Lb1** is not `nil`, `Lb1(el_info, quad, iq, user_data)` returns a pointer to a vector of length `DIM+1` storing the value of \bar{b}^1 at `quad->lambda[iq]`; `quad` is the quadrature for the first order term and `user_data` is a pointer to user data;

Lb1_pw_const: should be `true` if \bar{b}^1 is piecewise constant on the mesh (constant vector b on a non-parametric mesh, e.g.); thus integration of the first order term can use pre-computed integrals of the basis functions on the standard element; otherwise integration is done by using quadrature on each element;

Lb0_Lb1_anti_symmetric: should be `true` if the contributions of the complete first order term to the *local* element matrix are anti symmetric (only possible if both **Lb0** and **Lb1** are not `nil`, $\bar{b}^0 = -\bar{b}^1$, e.g.); if the finite element spaces for rows and columns are the same then only the upper part of the element matrix for the first order term has to be computed; elements of the lower part can then be set using the anti symmetry; otherwise the complete element matrix has to be calculated;

c: is a pointer to a function for the evaluation of \bar{c} at quadrature nodes on the element; **c** may be a `nil` pointer, if no zero order term has to be integrated; if **c** is not `nil`, `c(el_info, quad, iq, user_data)` returns the value of the function \bar{c} at `quad->lambda[iq]`; `quad` is the quadrature for the zero order term and `user_data` is a pointer to user data;

c_pw_const: should be `true` if the zero order term \bar{c} is piecewise constant on the mesh (constant function c on a non-parametric mesh, e.g.); thus integration of the zero order term can use pre-computed integrals of the basis functions on the standard element; otherwise integration is done by using quadrature on each element;

use_get_bound: if non-zero, then the `get_bound` entry in `EL_MATRIX_INFO` is set to `row_fe_space->bas_fcts->get_bound()`, and Dirichlet boundary DOFs are handled accordingly; **use_get_bound** must be zero if two different finite element spaces are used.

user_data: optional pointer to memory segment for user data used by `init_element()`, `LALt()`, `Lb0()`, `Lb1()`, and/or `c()` and is the last argument to these functions.

fill_flag: the flag for the mesh traversal routine indicating which elements should be visited and which information should be present in the `EL_INFO` structure for `init_element()`, `LALt()`, `Lb0()`, `Lb1()`, and/or `c()` on the visited elements.

Information stored in such a structure is used by the following function which returns a pointer to a filled `EL_MATRIX_INFO` structure; this structure can be used as an argument to the `update_matrix()` function which will then assemble the discrete matrix corresponding to the operator defined in the `OPERATOR_INFO`:

```
const EL_MATRIX_INFO *fill_matrix_info(const OPERATOR_INFO *,
                                       EL_MATRIX_INFO *);
```

Description:

fill_matrix_info(op_info, mat_info): the function returns a pointer to a filled `EL_MATRIX_INFO` structure for the assemblage of the system matrix for the operator defined in `op_info`.

If the second argument `mat_info` is a `nil` pointer, a new structure `mat_info` is allocated and filled; otherwise the structure `mat_info` is filled; all members are newly assigned.

`op_info->row_fe_space` and `op_info->col_fe_space` are pointers to the finite element spaces (and by this to the basis functions and DOFs) connected to the row DOFs and the column DOFs of the matrix to be assembled. If both pointers are `nil` pointers, an error message is given, and the program stops; if only one of these pointers is `nil`, row and column degrees of freedom are connected with the same finite element space (i.e. either `op_info->row_fe_space = op_info->col_fe_space`, or `op_info->col_fe_space = op_info->row_fe_space` is used).

The numbers of basis functions of the `row_fe_space` and `col_fe_space` determine the members `mat_info->n_row` and `mat_info->n_col`, the corresponding `dof_admin` structures define the entries `mat_info->row_admin` and `mat_info->col_admin`; the `get_dof_indices()` of the basis functions are used for `mat_info->get_row_dof` and `mat_info->get_col_dof`; the entries for the columns are only set if the finite element spaces are not the same; if the spaces are the same and the `use_get_bound` entry is not zero, then the `get_bound()` function of the basis functions is used for the member `mat_info->get_bound()`.

The most important member in the `EL_MATRIX_INFO` structure, namely `mat_info->el_matrix_fct`, is adjusted to some general routine for the integration of the element matrix for any set of local basis functions; `fill_matrix_info()` tries to use the fastest available function for the element integration for the operator defined in `op_info`, depending on `op_info->LAlt_pw_const` and similar hints;

Denote by `row_degree` and `col_degree` the degree of the basis functions connected to the rows and columns; the following vector `quad` of quadratures is used for the element integration, if not specified by `op_info->quad` using the following rule: pre-computed integrals of basis functions should be evaluated exactly, and all terms calculated by quadrature on the elements should use the same quadrature formula (this is more efficient than to use different quadratures). To be more specific:

If the 2nd order term has to be integrated and `op_info->quad[2]` is not `nil`, `quad[2] = op_info->quad[2]` is used, otherwise `quad[2]` is a quadrature which is exact of degree `row_degree+col_degree-2`. If the 2nd order term is not integrated then `quad[2]` is set to `nil`.

If the 1st order term has to be integrated and `op_info->quad[1]` is not a `nil` pointer, `quad[1] = op_info->quad[1]` is used; otherwise: if `op_info->Lb_pw_const` is zero and `quad[2]` is not `nil`, `quad[1] = quad[2]` is used, otherwise `quad[1]` is a quadrature which is exact of degree `row_degree+col_degree-1`. If the 1st order term is not integrated then `quad[1]` is set to `nil`.

If the zero order term has to be integrated and `op_info->quad[0]` is not a `nil` pointer, `quad[0] = op_info->quad[0]` is used; otherwise: if `op_info->c_pw_const` is zero and `quad[2]` is not `nil`, `quad[0] = quad[2]` is used, if `quad[2]` is `nil` and `quad[1]` is not `nil`, `quad[0] = quad[1]` is used, or if both quadratures are `nil`, `quad[0]` is a quadrature which is exact of degree `row_degree+col_degree`. If the zero order term is not integrated then `quad[0]` is set to `nil`.

If the function pointer `op_info->init_element` is not `nil` then a call of `op_info->init_element(el_info, quad, op_info->user_data)` is always the first statement of `mat_info->el_matrix_fct()` on each element; `el_info` is a pointer to the `EL_INFO` structure of the actual element, `quad` is the

quadrature vector described above (now giving information about the actually used quadratures), and the last argument is a pointer to user data.

If `op_info->LALt` is not `nil`, the 2nd order term is integrated using the quadrature `quad[2]`; if `op_info->LALt_pw_const` is not zero, the integrals of the product of gradients of the basis functions on the standard simplex are initialized (using the quadrature `quad[2]` for the integration) and used for the computation on the elements; `op_info->LALt()` is only called once with arguments `op_info->LALt(el_info, quad[2], 0, op_info->user_data)`, i.e. the matrix of the 2nd order term is evaluated only at the first quadrature node; otherwise the integrals are approximated by quadrature and `op_info->LALt()` is called for each quadrature node of `quad[2]`; if `op_info->LALt_symmetric` is not zero, the symmetry of the element matrix is used, if the finite element spaces are the same and this term is not integrated by the same quadrature as the first order term.

If `op_info->Lb0` is not `nil`, this 1st order term is integrated using the quadrature `quad[1]`; if `op_info->Lb0_pw_const` is not zero, the integrals of the product of basis functions with gradients of basis functions on the standard simplex are initialized (using the quadrature `quad[1]` for the integration) and used for the computation on the elements; `op_info->Lb0()` is only called once with arguments `op_info->Lb0(el_info, quad[1], 0, op_info->user_data)`, i.e. the vector of this 1st order term is evaluated only at the first quadrature node; otherwise the integrals are approximated by quadrature and `op_info->Lb0()` is called for each quadrature node of `quad[1]`;

If `op_info->Lb1` is not `nil`, this 1st order term is integrated also using the quadrature `quad[1]`; if `op_info->Lb1_pw_const` is not zero, the integrals of the product of gradients of basis functions with basis functions on the standard simplex are initialized (using the quadrature `quad[1]` for the integration) and used for the computation on the elements; `op_info->Lb1()` is only called once with arguments `op_info->Lb1(el_info, quad[1], 0, op_info->user_data)`, i.e. the vector of this 1st order term is evaluated only at the first quadrature node; otherwise the integrals are approximated by quadrature and `op_info->Lb1()` is called for each quadrature node of `quad[1]`.

If both `op_info->Lb0` and `op_info->Lb1` are not `nil`, the finite element spaces for rows and columns are the same and `Lb0_Lb1_anti_symmetric` is non-zero, then the contributions of the 1st order term are computed using this anti symmetry property.

If `op_info->c` is not `nil`, the zero order term is integrated using the quadrature `quad[0]`; if `op_info->c_pw_const` is not zero, the integrals of the product of basis functions on the standard simplex are initialized (using the quadrature `quad[0]` for the integration) and used for the computation on the elements; `op_info->c()` is only called once with arguments `op_info->c(el_info, quad[0], 0, op_info->user_data)`, i.e. the zero or-

der term is evaluated only at the first quadrature node; otherwise the integrals are approximated by quadrature and `op_info->c()` is called for each quadrature node of `quad[0]`.

The functions `LALt()`, `Lb0()`, `Lb1()`, and `c()`, can be called in an arbitrary order on the elements, if not `nil` (this depends on the type of integration, using pre-computed values, using same/different quadrature for the second, first, and/or zero order term, e.g.) but commonly used data for these functions is always initialized first by `op_info->init_element()`, if this function pointer is not `nil`.

Using all information about the operator and quadrature, an “optimal” routine for the assemblage is chosen. Information for this routine is stored at `mat_info` which includes the pointer to user data `op_info->user_data` (the last argument to `init_element()`, `LALt()`, `Lb0()`, `Lb1()`, and/or `c()`).

Finally, the flag for the mesh traversal which is used by the function `update_matrix()` is set in `mat_info->fill_flag` to `op_info->fill_flag`; it indicates which elements should be visited and which information should be present in the `EL_INFO` structure for `init_element()`, `LALt()`, `Lb0/1()`, and/or `c()` on the visited elements. If the finite element spaces are the same, and the entry `op_info->use_get_bound` is set, then the `FILL_BOUND` flag is added to `mat_info->fill_flag`.

The the access of a `MATRIX_INFO` structure for the automatic assemblage of the system matrix corresponding to the Laplace operator is given in Section 2.1.7 on 63.

3.12.3 Data structures for storing pre-computed integrals of basis functions

Assume a non-parametric triangulation and constant coefficient functions A , b , and c . Since the Jacobian of the barycentric coordinates is constant on S , the functions \bar{A}_S , \bar{b}_S^0 , \bar{b}_S^1 , and \bar{c}_S are constant on S also. Now, looking at the element matrix approximated by some quadrature \hat{Q} , we observe

$$\begin{aligned}\hat{Q}\left(\sum_{k,l=0}^d(\bar{a}_{S,kl}\bar{\psi}_{,\lambda_k}^i\bar{\varphi}_{,\lambda_l}^j)\right) &= \sum_{k,l=0}^d\bar{a}_{S,kl}\hat{Q}\left(\bar{\psi}_{,\lambda_k}^i\bar{\varphi}_{,\lambda_l}^j\right), \\ \hat{Q}\left(\sum_{l=0}^d(\bar{b}_{S,l}^0\bar{\psi}^i\bar{\varphi}_{,\lambda_l}^j)\right) &= \sum_{l=0}^d\bar{b}_{S,l}^0\hat{Q}\left(\bar{\psi}^i\bar{\varphi}_{,\lambda_l}^j\right), \\ \hat{Q}\left(\sum_{k=0}^d(\bar{b}_{S,k}^1\bar{\psi}_{,\lambda_k}^i\bar{\varphi}^j)\right) &= \sum_{k=0}^d\bar{b}_{S,k}^1\hat{Q}\left(\bar{\psi}_{,\lambda_k}^i\bar{\varphi}^j\right), \quad \text{and} \\ \hat{Q}\left((\bar{c}_S\bar{\psi}^i\bar{\varphi}^j)\right) &= \bar{c}_S\hat{Q}\left(\bar{\psi}^i\bar{\varphi}^j\right).\end{aligned}$$

The values of the quadrature applied to the basis functions do only depend on the basis functions and the standard element but not on the actual simplex S . All information about S is given by \bar{A}_S , \bar{b}_S^0 , \bar{b}_S^1 , and \bar{c}_S . Thus, these quadratures have only to be calculated once, and can then be used on each element during the assembling.

For this we have to store for the basis functions $\{\bar{\psi}_i\}_{i=1,\dots,n}$ and $\{\bar{\varphi}_j\}_{j=1,\dots,m}$ the values

$$\hat{Q}_{ij,kl}^{11} := \hat{Q}\left(\bar{\psi}_{,\lambda_k}^i \bar{\varphi}_{,\lambda_l}^j\right) \quad \text{for } 1 \leq i \leq n, 1 \leq j \leq m, 0 \leq k, l \leq \text{DIM},$$

if $A \neq 0$,

$$\hat{Q}_{ij,l}^{01} := \hat{Q}\left(\bar{\psi}^i \bar{\varphi}_{,\lambda_l}^j\right) \quad \text{for } 1 \leq i \leq n, 1 \leq j \leq m, 0 \leq l \leq \text{DIM},$$

if $b^0 \neq 0$,

$$\hat{Q}_{ij,k}^{10} := \hat{Q}\left(\bar{\psi}_{,\lambda_k}^i \bar{\varphi}^j\right) \quad \text{for } 1 \leq i \leq n, 1 \leq j \leq m, 0 \leq k \leq \text{DIM}$$

if $b^1 \neq 0$, and

$$\hat{Q}_{ij}^{00} := \hat{Q}\left(\bar{\psi}^i \bar{\varphi}^j\right) \quad \text{for } 1 \leq i \leq n, 1 \leq j \leq m,$$

if $c \neq 0$. Many of these values are zero, especially for the first and second order terms (if $\bar{\psi}^i$ and $\bar{\varphi}^j$ are the linear nodal basis functions $\hat{Q}_{ij,kl}^{11} = \delta_{ij}\delta_{kl}$). Thus, we use special data structures for a sparse storage of the non zero values for these terms. These are described now.

In order to “define” zero entries we use

```
static const REAL TOO_SMALL = 1.e-15;
```

and all computed values `val` with $|\text{val}| \leq \text{TOO_SMALL}$ are treated as zeros. As we are considering here integrals over the standard simplex, non-zero integrals are usually of order one, such that the above constant is of the order of roundoff errors for double precision.

The following data structure is used for storing values \hat{Q}^{11} for two sets of basis functions integrated with a given quadrature

```
typedef struct q11_psi_phi    Q11_PSI_PHI;
```

```
struct q11_psi_phi
{
    const BAS_FCTS    *psi;
    const BAS_FCTS    *phi;
    const QUAD        *quad;

    const int         **n_entries;
```

```

const REAL ***values;
const int ***k;
const int ***l;
};

```

Description:

psi: pointer to the first set of basis functions.

phi: pointer to the second set of basis functions.

quad: pointer to the quadrature which is used for the integration.

n_entries: matrix of size **psi**->**n_bas_fcts** × **phi**->**n_bas_fcts** storing the count of non zero integrals;

n_entries[i][j] is the count of non zero values of $\hat{Q}_{ij,kl}^{11}$ ($0 \leq k, l \leq \text{DIM}$) for the pair (**psi**[i], **phi**[j]), $0 \leq i < \text{psi}->\text{n_bas_fcts}$, $0 \leq j < \text{phi}->\text{n_bas_fcts}$.

values: tensor storing the non zero integrals;

values[i][j] is a vector of length **n_entries**[i][j] storing the non zero values for the pair (**psi**[i], **phi**[j]).

k, **l**: tensor storing the indices k, l of the non zero integrals;

k[i][j] and **l**[i][j] are vectors of length **n_entries**[i][j] storing at **k**[i][j][r] and **l**[i][j][r] the indices k and l of the value stored at **values**[i][j][r], i.e.

$$\text{values}[i][j][r] = \hat{Q}_{ij, k[i][j][r], l[i][j][r]}^{11} = \hat{Q}\left(\bar{\psi}_{\lambda_{k[i][j][r]}}^i, \bar{\varphi}_{\lambda_{l[i][j][r]}}^j\right),$$

for $0 \leq r < \text{n_entries}[i][j]$. Using these pre-computed values we have for $n_{ij} = \text{n_entries}[i][j]$ on all elements S

$$\sum_{k,l=0}^d \bar{a}_{S,kl} \hat{Q}\left(\bar{\psi}_{\lambda_k}^i, \bar{\varphi}_{\lambda_l}^j\right) = \sum_{r=0}^{n_{ij}-1} \bar{a}_{S, k[i][j][r], l[i][j][r]} * \text{values}[i][j][r].$$

The following function initializes a Q11_PSI_PHI structure:

```

const Q11_PSI_PHI *get_q11_psi_phi(const BAS_FCTS *,
                                   const BAS_FCTS *, const QUAD *);

```

Description:

get_q11_psi_phi(**psi**, **phi**, **quad**): the function returns a pointer to a filled Q11_PSI_PHI structure.

psi is a pointer to the first set of basis functions, **phi** is a pointer to the second set of basis functions; if both are **nil** pointers, nothing is done and the return value is **nil**; if one of the pointers is a **nil** pointer, the structure is initialized using the same set of basis functions for the first and second set, i.e. **phi** = **psi** or **psi** = **phi** is used.

`quad` is a pointer to a quadrature for the approximation of the integrals; if `quad` is `nil`, then a default quadrature which is exact of degree `psi->degree+phi->degree-2` is used.

All used `Q11_PSI_PHI` structures are stored in a linked list and are identified uniquely by the members `psi`, `phi`, and `quad`, and for such a combination only one `Q11_PSI_PHI` structure is created during runtime;

First, `get_q11_psi_phi()` looks for a matching structure in the linked list; if such a structure is found a pointer to this structure is returned; the values are not computed a second time. Otherwise a new structure is generated, linked to the list, and the values are computed using the quadrature `quad`; all values `val` with $|val| \leq \text{T00_SMALL}$ are treated as zeros.

Example 3.30. The following example shows how to use these pre-computed values for the integration of the 2nd order term

$$\int_S \nabla_\lambda \bar{\psi}^i(\lambda(\hat{x})) \cdot \bar{A}(\lambda(\hat{x})) \nabla_\lambda \bar{\varphi}^j(\lambda(\hat{x})) d\hat{x}$$

for all i, j . We only show the body of a function for the integration and assume that `LALt_fct` returns a matrix storing \bar{A} (compare the member `LALt` in the `OPERATOR_INFO` structure):

```
if (!q11_psi_phi) q11_psi_phi = get_q11_psi_phi(psi, phi, quad[2]);

LALt = LALt_fct(el_info, quad, 0, user_data);
n_entries = q11_psi_phi->n_entries;

for (i = 0; i < psi->n_bas_fcts; i++)
{
    for (j = 0; j < phi->n_bas_fcts; j++)
    {
        k      = q11_psi_phi->k[i][j];
        l      = q11_psi_phi->l[i][j];
        values = q11_psi_phi->values[i][j];
        for (val = m = 0; m < n_entries[i][j]; m++)
            val += values[m]*LALt[k[m]][l[m]];
        mat[i][j] += val;
    }
}
```

The values \hat{Q}^{01} for the set of basis functions `psi` and `phi` are stored in

```
typedef struct q01_psi_phi    Q01_PSI_PHI;
```

```
struct q01_psi_phi
{
    const BAS_FCTS    *psi;
    const BAS_FCTS    *phi;
    const QUAD         *quad;
```

```

const int    **n_entries;
const REAL   ***values;
const int    ***l;
};

```

Description:

psi: pointer to the first set of basis functions.

phi: pointer to the second set of basis functions.

quad: pointer to the quadrature which is used for the integration.

n_entries: matrix of size **psi**->**n_bas_fcts** \times **phi**->**n_bas_fcts** storing the count of non zero integrals;

n_entries[i][j] is the count of non zero values of $\hat{Q}_{ij,1}^{01}$ ($0 \leq i \leq \text{DIM}$) for the pair (**psi**[i],**phi**[j]), $0 \leq i < \text{psi}->\text{n_bas_fcts}$, $0 \leq j < \text{phi}->\text{n_bas_fcts}$.

values: tensor storing the non zero integrals;

values[i][j] is a vector of length **n_entries**[i][j] storing the non zero values for the pair (**psi**[i],**phi**[j]).

l: tensor storing the indices *l* of the non zero integrals;

l[i][j] is a vector of length **n_entries**[i][j] storing at **l**[i][j][r] the index *l* of the value stored at **values**[i][j][r], i.e.

$$\text{values}[i][j][r] = \hat{Q}_{ij,1[i][j][r]}^{01} = \hat{Q}(\bar{\psi}^i \bar{\varphi}_{\lambda_{1[i][j][r]}}^j),$$

for $0 \leq r < \text{n_entries}[i][j]$. Using these pre-computed values we have for all elements *S*

$$\sum_{l=0}^d \bar{b}_{S,l}^0 \hat{Q}(\bar{\psi}^i \bar{\varphi}_{\lambda_l}^j) = \sum_{r=0}^{\text{n_entries}[i][j]-1} \bar{b}_{S,1[i][j][r]}^0 * \text{values}[i][j][r].$$

The following function initializes a Q01_PSI_PHI structure:

```

const Q01_PSI_PHI *get_q01_psi_phi(const BAS_FCTS *,
                                   const BAS_FCTS *, const QUAD *);

```

Description:

get_q01_psi_phi(**psi**, **phi**, **quad**): the function returns a pointer to a filled Q01_PSI_PHI structure.

psi is a pointer to the first set of basis functions **phi** is a pointer to the second set of basis functions; if both are **nil** pointers, nothing is done and the return value is **nil**; if one of the pointers is a **nil** pointer, the structure is initialized using the same set of basis functions for the first and second set, i.e. **phi** = **psi** or **psi** = **phi** is used.

quad is a pointer to a quadrature for the approximation of the integrals; if **quad** is **nil**, a default quadrature which is exact of degree **psi**->**degree**+**phi**->**degree**-1 is used.

All used Q01_PSI_PHI structures are stored in a linked list and are identified uniquely by the members **psi**, **phi**, and **quad**, and for such a combination only one Q01_PSI_PHI structure is created during runtime;

First, `get_q01_psi_phi()` looks for a matching structure in the linked list; if such a structure is found a pointer to this structure is returned; the values are not computed a second time. Otherwise a new structure is generated, linked to the list, and the values are computed using the quadrature **quad**; all values **val** with $|\text{val}| \leq \text{T00_SMALL}$ are treated as zeros.

The values \hat{Q}^{10} for the set of basis functions **psi** and **phi** are stored in

```
typedef struct q10_psi_phi    Q10_PSI_PHI;

struct q10_psi_phi
{
    const BAS_FCTS    *psi;
    const BAS_FCTS    *phi;
    const QUAD        *quad;

    const int         **n_entries;
    const REAL        ***values;
    const int         ***k;
};
```

Description:

psi: pointer to the first set of basis functions.

phi: pointer to the second set of basis functions.

quad: pointer to the quadrature which is used for the integration.

n_entries: matrix of size **psi**->**n_bas_fcts** \times **phi**->**n_bas_fcts** storing the count of non zero integrals;

n_entries[i][j] is the count of non zero values of $\hat{Q}_{ij,k}^{10}$ ($0 \leq k \leq \text{DIM}$) for the pair (**psi**[i],**phi**[j]), $0 \leq i < \text{psi}-\text{n_bas_fcts}$, $0 \leq j < \text{phi}-\text{n_bas_fcts}$.

values: tensor storing the non zero integrals;

values[i][j] is a vector of length **n_entries**[i][j] storing the non zero values for the pair (**psi**[i],**phi**[j]).

k: tensor storing the indices k of the non zero integrals;

k[i][j] is a vector of length **n_entries**[i][j] storing at **k**[i][j][r] the index k of the value stored at **values**[i][j][r], i.e.

$$\text{values}[i][j][r] = \hat{Q}_{ij,k[i][j][r]}^{10} = \hat{Q}\left(\bar{\psi}_{\lambda_{k[i][j][r]}}^i \bar{\varphi}^j\right),$$

for $0 \leq r < \text{n_entries}[i][j]$. Using these pre-computed values we have for all elements S

$$\sum_{k=0}^d \bar{b}_{S,k}^1 \hat{Q}\left(\bar{\psi}_{\lambda_k}^i \bar{\varphi}^j\right) = \sum_{r=0}^{\text{n_entries}[i][j]-1} \bar{b}_{S,k[i][j][r]}^1 * \text{values}[i][j][r].$$

The following function initializes a Q10_PSI_PHI structure:

```
const Q10_PSI_PHI *get_q10_psi_phi(const BAS_FCTS *,
                                   const BAS_FCTS *, const QUAD *);
```

Description:

`get_q10_psi_phi(psi, phi, quad)`: the function returns a pointer to a filled Q10_PSI_PHI structure.

`psi` is a pointer to the first set of basis functions `phi` is a pointer to the second set of basis functions; if both are `nil` pointers, nothing is done and the return value is `nil`; if one of the pointers is a `nil` pointer, the structure is initialized using the same set of basis functions for the first and second set, i.e. `phi = psi` or `psi = phi` is used.

`quad` is a pointer to a quadrature for the approximation of the integrals; if `quad` is `nil`, a default quadrature which is exact of degree `psi->degree+phi->degree-1` is used.

All used Q10_PSI_PHI structures are stored in a linked list and are identified uniquely by the members `psi`, `phi`, and `quad`, and for such a combination only one Q10_PSI_PHI structure is created during runtime;

First, `get_q10_psi_phi()` looks for a matching structure in the linked list; if such a structure is found a pointer to this structure is returned; the values are not computed a second time. Otherwise a new structure is generated, linked to the list, and the values are computed using the quadrature `quad`; all values `val` with $|val| \leq \text{T00_SMALL}$ are treated as zeros.

Finally, the values \hat{Q}^{00} for the set of basis functions `psi` and `phi` are stored in

```
typedef struct q00_psi_phi    Q00_PSI_PHI;

struct q00_psi_phi
{
    const BAS_FCTS    *psi;
    const BAS_FCTS    *phi;
    const QUAD        *quad;

    const REAL    **values;
};
```

Description:

`psi`: pointer to the first set of basis functions.

`phi`: pointer to the second set of basis functions.

`quad`: pointer to the quadrature which is used for the integration.

`values`: matrix storing the integrals;

$$\text{values}[i][j] = \hat{Q}_{ij}^{00} = \hat{Q}(\bar{\psi}^i \bar{\varphi}^j),$$

for the pair $(\psi[i], \phi[j])$, $0 \leq i < \psi \rightarrow n_{\text{bas_fcts}}$, $0 \leq j < \phi \rightarrow n_{\text{bas_fcts}}$.

The following function initializes a Q00_PSI_PHI structure:

```
const Q00_PSI_PHI *get_q00_psi_phi(const BAS_FCTS *,
                                   const BAS_FCTS *, const QUAD *);
```

Description:

`get_q00_psi_phi(psi, phi, quad)`: the function returns a pointer to a filled Q00_PSI_PHI structure.

`psi` is a pointer to the first set of basis functions `phi` is a pointer to the second set of basis functions; if both are `nil` pointers, nothing is done and the return value is `nil`; if one of the pointers is a `nil` pointer, the structure is initialized using the same set of basis functions for the first and second set, i.e. `phi = psi` or `psi = phi` is used.

`quad` is a pointer to a quadrature for the approximation of the integrals; if `quad` is `nil`, a default quadrature which is exact of degree `psi->degree+phi->degree` is used.

All used Q00_PSI_PHI structures are stored in a linked list and are identified uniquely by the members `psi`, `phi`, and `quad`, and for such a combination only one Q00_PSI_PHI structure is created during runtime;

First, `get_q00_psi_phi()` looks for a matching structure in the linked list; if such a structure is found a pointer to this structure is returned; the values are not computed a second time. Otherwise a new structure is generated, linked to the list, and the values are computed using the quadrature `quad`.

3.12.4 Data structures and functions for vector update

Besides the general routines `update_real_vec()` and `update_real_d_vec()`, this section presents also easy to use routines for calculation of L^2 scalar products between a given function and all basis functions of a finite element space.

The following structures hold full information for the assembling of element vectors. They are used by the functions `update_real_vec()` and `update_real_d_vec()` described below.

```
typedef struct el_vec_info    EL_VEC_INFO;
typedef struct el_vec_d_info  EL_VEC_D_INFO;

struct el_vec_info
{
    int                n_dof;
    const DOF_ADMIN    *admin;
    const DOF           (*get_dof)(const EL *, const DOF_ADMIN *, DOF *);
    const S_CHAR        (*get_bound)(const EL_INFO *, S_CHAR *);

    REAL               factor;
```

```

    const REAL      *(*el_vec_fct)(const EL_INFO *, void *);
    void            *fill_info;

    FLAGS           fill_flag;
};

struct el_vec_d_info
{
    int              n_dof;
    const DOF_ADMIN *admin;
    const DOF        *(*get_dof)(const EL *, const DOF_ADMIN *, DOF *);
    const S_CHAR     *(*get_bound)(const EL_INFO *, S_CHAR *);

    REAL             factor;

    const REAL_D     *(*el_vec_fct)(const EL_INFO *, void *);
    void             *fill_info;

    FLAGS           fill_flag;
};

```

Description:

n_dof: size of the element vector.

admin: pointer to a **DOF_ADMIN** structure for the administration of DOFs of the vector to be filled.

get_dof: pointer to a function for the access of the global DOFs on a single element; **get_row_dof**(**el**, **admin**, **dof**) returns a pointer to a vector of length **n_dof** storing the global DOFs; if **dof** is a **nil** pointer, **get_dof**() has to provide memory for storing this vector, which may be overwritten on the next call; otherwise the DOFs have to be stored at **dof**; (usually, **get_dof**() is the **get_dof_indices**() function inside a **BAS_FCTS** structure (compare Section 3.5.1).

get_bound: optional pointer to a function providing information about the boundary type of DOFs; if **get_bound**() is not **nil**, **get_bound**(**el_info**, **bound**) returns a pointer to a vector of length **n_dof** storing the boundary type of the local DOFs; this pointer is the optional pointer to a vector holding boundary information of the function **add_element[_d]_vec**() described above; if **bound** is a **nil** pointer, **get_bound**() has to provide memory for storing this vector, which may be overwritten on the next call; otherwise the DOFs have to be stored at **bound**; (usually, **get_bound**() is the **get_bound**() function inside a **BAS_FCTS** structure (compare Section 3.5.1).

factor: is a multiplier for the element contributions; usually **factor** is 1 or -1.

el_vec_fct: is a pointer to a function for the computation of the element vector; **el_vec_fct**(**el_info**, **fill_info**) returns a pointer to a **REAL**

resp. `REAL_D` vector of length `n.dof` storing the element vector on element `el_info->el`; `fill_info` is a pointer to data needed by `el_vec_fct()`; the function has to provide memory for storing the element vector, which can be overwritten on the next call.

`fill_info`: pointer to data needed by `el_vec_fct()`; is the second argument of this function.

`fill_flag`: the flag for the mesh traversal for assembling the vector.

The following function does the update of vectors by assembling element contributions during mesh traversal; information for computing the element vectors is held in a `EL_VEC[_D]_INFO` structure:

```
void update_real_vec(DOF_REAL_VEC *, const EL_VEC_INFO *);
void update_real_d_vec(DOF_REAL_D_VEC *, const EL_VEC_D_INFO *);
```

`update_real[_d]_vec(dr[d]v, info)`: updates the vector `drv` resp. `drdv` by traversing the underlying mesh and assembling the element contributions into the vector; information about the computation of element vectors and connection of local and global DOFs is stored in `info`.

The flags for the mesh traversal of the mesh `dr[d]v->fe_space->mesh` are stored at `info->fill_flags` which specifies the elements to be visited and information that should be present on the elements for the calculation of the element vectors and boundary information (if `info->get_bound` is not `nil`); Information about global DOFs is accessed element-wise by `info->get_dof` using `info->admin`; in addition, the boundary type of the DOFs is accessed by `info->get_bound` if it is not a `nil` pointer; then the element vector is computed by `info->el_vec_fct(el_info, info->fill_info)`; these contributions are finally added to `dr[d]v` multiplied by `info->factor` by a call of `add_element[_d]_vec()` with all information about global DOFs, the element vector, and boundary types, if available.

`update_real[_d]_vec()` only adds element contributions; this makes several calls for the assemblage of one vector possible; before the first call, the vector should be set to zero by a call of `dof_set[_d](0.0, dr[d]v)`.

L^2 scalar products. In many applications, the load vector is just the L^2 scalar product of a given function with all basis functions of the finite element space or this scalar product is a part of the right hand side; such a scalar product can be directly assembled by the functions

```
void L2scp_fct_bas(REAL (*)(const REAL_D), const QUAD *,
                  DOF_REAL_VEC *);
void L2scp_fct_bas_d(const REAL (*)(const REAL_D, REAL_D),
                    const QUAD *, DOF_REAL_D_VEC *);
```

Description:

L2scp_fct_bas(f, quad, fh): approximates the L^2 scalar products of a given function with all basis functions by numerical quadrature and adds the corresponding values to a DOF vector;

f is a pointer for the evaluation of the given function in world coordinates x and returns the value of that function at x ; if **f** is a **nil** pointer, nothing is done;

fh is the DOF vector where at the i -th entry the approximation of the L^2 scalar product of the given function with the i -th global basis function of **fh->fe_space** is added;

quad is the quadrature for the approximation of the integral on each leaf element of **fh->fe_space->mesh**; if **quad** is a **nil** pointer, a default quadrature which is exact of degree $2*\text{fh->fe_space->bas_fcts->degree}-2$ is used.

The integrals are approximated by looping over all leaf elements, computing the approximations to the element contributions and adding these values to the vector **fh** by **add_element_vec()**.

The vector **fh** is *not* initialized with 0.0; only the new contributions are added.

L2scp_fct_bas_d(fd, quad, fhd): approximates the L^2 scalar products of a given vector valued function with all scalar valued basis functions by numerical quadrature and adds the corresponding values to a vector valued DOF vector;

fd is a pointer for the evaluation of the given function in world coordinates x ; **fd(x, fx)** returns a pointer to a vector storing the value at x ; if **fx** is not **nil**, the value is stored at **fx** otherwise the function has to provide memory for storing this vector, which can be overwritten on the next call; if **fd** is a **nil** pointer, nothing is done;

fhd is the DOF vector where at the i -th entry (a **REAL_D** vector) the approximation of the L^2 scalar product of the given vector valued function with the i -th global (scalar valued) basis function of **fhd->fe_space** is added;

quad is the quadrature for the approximation of the integral on each leaf element of **fhd->fe_space->mesh**; if **quad** is a **nil** pointer, a default quadrature which is exact of degree $2*\text{fhd->fe_space->bas_fcts->degree}-2$ is used.

The integrals are approximated by looping over all leaf elements, computing the approximations to the element contributions and adding these values to the vector **fhd** by **add_element_d_vec()**.

The vector **fhd** is *not* initialized with $(0.0, \dots, 0.0)$; only the new contributions are added.

3.12.5 Dirichlet boundary conditions

For the solution of the discrete system (1.13) on page 37 derived in Section 1.4.5, we have to set the Dirichlet boundary values for all Dirichlet DOFs. Usually, we take for the approximation g_h of g the interpolant of g , i.e. $g_h = I_h g$ and we have to copy the coefficients of g_h at the Dirichlet DOFs to the load vector (compare (1.12) on page 37). Additionally we know that the discrete solution has also the same coefficients at Dirichlet DOFs. Using an iterative solver we should use such information for the initial guess. Copying the coefficients of g_h at the Dirichlet DOFs to the initial guess will result in an initial residual (and then for all subsequent residuals) which is (are) zero at all Dirichlet DOFs.

Furthermore, the matrix we have derived in (1.11) on page 37 (and which is assembled in this way by the assemblage tools) is not symmetric even for the discretization of symmetric and elliptic operators. Applying directly the conjugate gradient method for solving (1.13) will not work, because the matrix is not symmetric. But setting the Dirichlet boundary values also in the initial guess at Dirichlet DOFs, all residuals are zero at Dirichlet DOFs and thus the conjugate gradient method will only apply to the non Dirichlet DOFs, which means that the conjugate gradient method will only “see” the symmetric and positive definite part of the matrix.

The following functions will set Dirichlet boundary values for all DOFs on the Dirichlet boundary, using an interpolation of the boundary values g :

```
void dirichlet_bound(REAL (*)(const REAL_D), DOF_REAL_VEC *,
                    DOF_REAL_VEC *, DOF_SCHAR_VEC *);
void dirichlet_bound_d(const REAL (*)(const REAL_D, REAL_D),
                      DOF_REAL_D_VEC *, DOF_REAL_D_VEC *,
                      DOF_SCHAR_VEC *);
```

Description:

`dirichlet_bound(g, fh, uh, bound)`: the function sets Dirichlet boundary values at all Dirichlet DOFs on leaf elements of `fh->fe_space->mesh` or `uh->fe_space->mesh`; values at DOFs not belonging to the Dirichlet boundary are not changed by this function.

`g` is a pointer to a function for the evaluation of the boundary data; if `g` is a `nil` pointer, all coefficients at Dirichlet DOFs are set to 0.0.

`fh` and `uh` are vectors where Dirichlet boundary values should be set (usually, `fh` stores the load vector and `uh` an initial guess for an iterative solver); one of `fh` and `uh` may be a `nil` pointer; if both arguments are `nil` pointers, nothing is done; if both arguments are not `nil`, `fh->fe_space` must equal `uh->fe_space`.

Boundary values are set element-wise on all leaf elements at Dirichlet DOFs, which are identified by `fe_space->bas_fcts->get_bound()`; the finite element space `fe_space` of either `fh` or `uh` is used. Interpolation is then done by `fe_space->bas_fcts->interpol()` solely for the element's Dirichlet DOFs;

the function `g` must meet the requirements of this function. For Lagrange elements, `(*g)()` is evaluated for all Lagrange nodes on the Dirichlet boundary and has to return the values at these nodes (compare Section 3.5.1); the flag of the mesh traversal is `CALL_LEAF_EL|FILL_BOUND|FILL_COORDS`.

`bound` is a vector for storing the boundary type for each used DOF; `bound` may be `nil`; if it is not `nil`, the i -th entry of the vector is filled with the boundary type of the i -th DOF. `bound->fe_space` must be the same as `fh`'s or `uh`'s `fe_space`.

`dirichlet_bound_d(gd, fhd, uhd, bound)`: sets Dirichlet boundary values for vector valued functions at all Dirichlet DOFs on leaf elements of `fhd->fe_space->mesh` or `uhd->fe_space->mesh`; values at DOFs not belonging to the Dirichlet boundary are not changed by this function.

`gd` is a pointer to a function for the evaluation of boundary data; the return value is a pointer to a vector storing the values; if the second argument `val` of `(*gd)(x, val)` is not `nil`, the values have to be stored at `val`, otherwise `gd` has to provide memory for the vector which may be overwritten on the next call; if `gd` is a `nil` pointer, all coefficients at Dirichlet DOFs are set to $(0.0, \dots, 0.0)$.

`fhd` and `uhd` are DOF vectors where Dirichlet boundary values should be set (usually, `fhd` stores the load vector and `uhd` an initial guess for an iterative solver); one of `fhd` and `uhd` may be a `nil` pointer; if both arguments are `nil` pointers, nothing has been done; if both arguments are not `nil` pointers, `fhd->fe_space` must equal `uhd->fe_space`.

Boundary values are set element-wise on all leaf elements at Dirichlet DOFs, which are identified by `fe_space->bas_fcts->get_bound()`; the finite element space `fe_space` of either `fhd` or `uhd` is used. Interpolation is then done by `fe_space->bas_fcts->interp_d()` solely for the element's Dirichlet DOFs; the function `gd` must meet the requirements of this function. For Lagrange elements, `(*gd)()` is evaluated for all Lagrange nodes on the Dirichlet boundary and has to return the values at these nodes (compare Section 3.5.1); the flag of the mesh traversal is `CALL_LEAF_EL|FILL_BOUND|FILL_COORDS`.

`bound` is a vector for storing the boundary type for each used DOF; `bound` may be `nil`; if it is not `nil`, the i -th entry of the vector is filled with the boundary type of the i -th DOF. `bound->fe_space` must be the same as `fhd`'s or `uhd`'s `fe_space`.

3.12.6 Interpolation into finite element spaces

In time dependent problems, usually the “solve” step in the adaptive method for the adaptation of the initial grid is an interpolation of initial data to the finite element space, i.e. a DOF vector is filled with the coefficient of the interpolant. The following functions are implemented for this task:

```

void interpol(REAL (*)(const REAL_D), DOF_REAL_VEC *);
void interpol_d(const REAL (*)(const REAL_D, REAL_D),
               DOF_REAL_D_VEC *);

```

Description:

interpol(f, fh): computes the coefficients of the interpolant of a function and stores these in a DOF vector;

f is a pointer to a function for the evaluation of the function to be interpolated; if **f** is a **nil** pointer, all coefficients are set to 0.0.

fh is a DOF vector for storing the coefficients; if **fh** is a **nil** pointer, nothing is done.

The actual interpolation is done element-wise on the leaf elements of **fh->fe_space->mesh** utilizing **fh->fe_space->bas_fcts->interpol()**; **f** must meet the requirements of this function, for instance the point-wise evaluation of **(*f)()** at all Lagrange nodes when using Lagrange finite elements elements (compare Section 3.5.1); the **fill_flag** of the mesh traversal is **CALL_LEAF_EL|FILL_COORDS**.

interpol_d(fd, fhd): computes the coefficients of the interpolant of a vector valued function and stores these in a DOF vector;

fd is a pointer to a function for the evaluation of the function to be interpolated; the return value is a pointer to a vector storing the values; if the second argument **fx** of **(*fd)(x, fx)** is not **nil**, the values have to be stored at **fx**, otherwise **fd** has to provide memory for the vector which may be overwritten on the next call; if **fd** is a **nil** pointer, all coefficients are set to (0.0, ..., 0.0).

fhd is a DOF vector for storing the coefficients; if **fhd** is a **nil** pointer, nothing is done.

The actual interpolation is done element-wise on the leaf elements of **fhd->fe_space->mesh** by **fhd->fe_space->bas_fcts->interpol_d()**; **fd** must meet the requirements of this function, for instance the point-wise evaluation of **(*fd)()** at all Lagrange nodes when using Lagrange finite elements elements (compare Section 3.5.1); the **fill_flag** of the mesh traversal is **CALL_LEAF_EL|FILL_COORDS**.

3.13 Data structures and procedures for adaptive methods

3.13.1 ALBERTA adaptive method for stationary problems

The basic data structure **ADAPT_STAT** for stationary adaptive methods contains pointers to problem dependent routines to build the linear or nonlinear system(s) of equations on an adapted mesh, and to a routine which solves the discrete problem and computes the new discrete solution(s). For flexibility

and efficiency reasons, building and solution of the system(s) may be split into several parts, which are called at various stages of the mesh adaption process.

ADAPT_STAT also holds parameters used for the adaptive procedure. Some of the parameters are optional or used only when a special marking strategy is selected.

```
typedef struct adapt_stat      ADAPT_STAT;

struct adapt_stat
{
    const char  *name;
    REAL        tolerance;
    REAL        p;                /* power of the estimator norm */
    int         max_iteration;
    int         info;

    REAL  (*estimate)(MESH *mesh, ADAPT_STAT *adapt);
    REAL  (*get_el_est)(EL *el); /* access local error indicator */
    REAL  (*get_el_estc)(EL *el); /* access local coarsening error */
    U_CHAR (*marking)(MESH *mesh, ADAPT_STAT *adapt);

    void  *est_info; /* data used for the estimator ---*/
    REAL  err_sum, err_max; /*--- sum and max of el_est ---*/

    void  (*build_before_refine)(MESH *mesh, U_CHAR flag);
    void  (*build_before_coarsen)(MESH *mesh, U_CHAR flag);
    void  (*build_after_coarsen)(MESH *mesh, U_CHAR flag);
    void  (*solve)(MESH *mesh);

    int  refine_bisections;
    int  coarsen_allowed; /*--- 0: false, 1:true ---*/
    int  coarse_bisections;

    int  strategy; /*--- 1: GR, 2: MS, 3: ES, 4:GERS ---*/
    /*----- parameters for the different strategies -----*/
    REAL  MS_gamma, MS_gamma_c;
    REAL  ES_theta, ES_theta_c;
    REAL  GERS_theta_star, GERS_nu, GERS_theta_c;
};
```

The entries yield following information:

name: textual description of the adaptive method, or `nil`.

tolerance: given tolerance for the (absolute or relative) error.

p: power p used in estimate (1.17), $1 \leq p < \infty$.

max_iteration: maximal allowed number of iterations of the adaptive procedure; if `max_iteration` ≤ 0 , no iteration bound is used.

info: level of information printed during the adaptive procedure; if **info** ≥ 2 , the iteration count and final error estimate are printed; if **info** ≥ 4 , then information is printed after each iteration of the adaptive procedure; if **info** ≥ 6 , additional information about the CPU time used for mesh adaption and building the linear systems is printed.

estimate: pointer to a problem dependent function for computing the global error estimate and the local error indicators; must not be **nil**;

estimate(mesh, adapt) computes the error estimate and fills the entries **adapt->err_sum** and **adapt->err_max** with

$$\begin{aligned}\mathbf{adapt}\mathbf{->err_sum} &= \left(\sum_{S \in \mathcal{S}_h} \eta_S(u_h)^p \right)^{1/p}, \\ \mathbf{adapt}\mathbf{->err_max} &= \max_{S \in \mathcal{S}_h} \eta_S(u_h)^p.\end{aligned}$$

The return value is the total error estimate **adapt->err_sum**. User data, like additional parameters for **estimate()**, can be passed via the **est_info** entry of the **ADAPT_STAT** structure to a (problem dependent) parameter structure. Usually, **estimate()** stores the local error indicator(s) $\eta_S(u_h)^p$ (and coarsening error indicator(s) $\eta_{c,S}(u_h)^p$) in **LEAF_DATA(e1)**.

For sample implementations of error estimators for quasi-linear elliptic and parabolic problems, see Section 3.14.

get_el_est: pointer to a problem dependent subroutine returning the value of the local error indicator; must not be **nil** if via the entry **strategy** adaptive refinement is selected and the specialized marking routine **marking** is **nil**;

get_el_est(e1) returns the value $\eta_S(u_h)^p$, of the local error indicator on leaf element **e1**; usually, local error indicators are computed by **estimate()** and stored in **LEAF_DATA(e1)**, which is problem dependent and thus not directly accessible by general-purpose routines. **get_el_est()** is needed by the ALBERTA marking strategies.

get_el_estc: pointer to a function which returns the local coarsening error indicator;

get_el_estc(e1) returns the value $\eta_{c,S}(u_h)^p$ of the local coarsening error indicator on leaf element **e1**, usually computed by **estimate()** and stored in **LEAF_DATA(e1)**; if not **nil**, **get_el_estc()** is called by the ALBERTA marking routines; this pointer may be **nil**, which means $\eta_{c,S}(u_h) = 0$.

marking: specialized marking strategy; if **nil**, a standard ALBERTA marking routine is selected via the entry **strategy**;

marking(mesh, adapt) selects and marks elements for refinement or coarsening; the return value is

0: no element is marked;

MESH_REFINED: elements are marked but only for refinement;

MESH_COARSENE: elements are marked but only for coarsening;

MESH_REFINED|MESH_COARSENED: elements are marked for refinement and coarsening.

est_info: pointer to (problem dependent) parameters for the **estimate()** routine; via this pointer the user can pass information to the estimate routine; this pointer may be **nil**.

err_sum: holds the sum of local error indicators $(\sum_{S \in \mathcal{S}} \eta_S(u_h)^p)^{1/p}$; the value for this entry must be set by the function **estimate()**.

err_max: holds the maximal local error indicator $\max_{S \in \mathcal{S}} \eta_S(u_h)^p$; the value for this entry must be set by the function **estimate()**.

build_before_refine: pointer to a subroutine that builds parts of the (non-)linear system(s) before any mesh adaptation; if it is **nil**, this assemblage stage omitted;

build_before_refine(mesh, flag) launches the assembling of the assembling of the discrete system on **mesh**; **flag** gives information which part of the system has to be built; the mesh will be refined if the **MESH_REFINED** bit is set in **flag** and it will be coarsened if the bit **MESH_COARSENE**D is set in **flag**.

build_before_coarsen: pointer to a subroutine that builds parts of the (non-)linear system(s) between the refinement and coarsening; if it is **nil**, this assemblage stage omitted;

build_before_coarsen(mesh, flag) performs an intermediate assembling step on **mesh** (compare Section 1.4.4 for an example when such a step is needed); **flag** gives information which part of the system has to be built; the mesh was refined if the **MESH_REFINED** bit is set in **flag** and it will be coarsened if the bit **MESH_COARSENE**D is set in **flag**.

build_after_coarsen: pointer to a subroutine that builds parts of the (non-)linear system(s) after all mesh adaptation; if it is **nil**, this assemblage stage omitted;

build_after_coarsen(mesh, flag) performs the final assembling step on **mesh**; **flag** gives information which part of the system has to be built; the mesh was refined if the **MESH_REFINED** bit is set in **flag** and it was coarsened if the bit **MESH_COARSENE**D is set in **flag**.

solve: pointer to a subroutine for solving the discrete (non-)linear system(s); if it is **nil**, the solution step is omitted;

solve(mesh) computes the new discrete solution(s) on **mesh**.

refine_bisections: number of bisection steps for the refinement of an element marked for refinement; used by the **ALBERTA** marking strategies; default value is **DIM**.

coarsen_allowed: flag used by the **ALBERTA** marking strategies to allow (**true**) or forbid (**false**) mesh coarsening;

coarse_bisections: number of bisection steps for the coarsening of an element marked for coarsening; used by the **ALBERTA** marking strategies; default value is **DIM**.

strategy: parameter to select an ALBERTA marking routine; possible values are:

- 0: no mesh adaption,
- 1: global refinement (GR),
- 2: maximum strategy (MS),
- 3: equidistribution strategy (ES),
- 4: guaranteed error reduction strategy (GERS),

see Section 3.13.2 for a description of these strategies.

MS_gamma, **MS_gamma_c**: parameters for the marking by the *maximum strategy*, see Sections 1.5.2 and 1.5.3.

ES_theta, **ES_theta_c**: parameters for the marking by the *equidistribution strategy*, see Sections 1.5.2 and 1.5.3.

GERS_theta_star, **GERS_nu**, **GERS_theta_c**: parameters for the marking by the *guaranteed error reduction strategy*, see Sections 1.5.2 and 1.5.3.

The routine `adapt_method_stat()` implements the whole adaptive procedure for a stationary problem, using the parameters given in **ADAPT_STAT**:

```
void adapt_method_stat(MESH *, ADAPT_STAT *);
```

Description:

`adapt_method_stat(mesh, adapt_stat)`: the adaptive procedure for a stationary problem; solves the problem adaptively on `mesh` by the method described in Section 1.5.1; `adapt_stat` is a pointer to a filled **ADAPT_STAT** data structure, holding all information about the problem to be solved and parameters for the adaptive method.

The main loop of the adaptive method is given in the following source fragment (assuming that `adapt->max_iteration` is non-negative):

```
void adapt_method_stat(MESH *mesh, ADAPT_STAT *adapt)
{
    int      iter;
    REAL     est;

    ...

    /*--- get solution on initial mesh -----*/
    if (adapt->build_before_refine)
        adapt->build_before_refine(mesh, 0);
    if (adapt->build_before_coarsen)
        adapt->build_before_coarsen(mesh, 0);
    if (adapt->build_after_coarsen)
        adapt->build_after_coarsen(mesh, 0);
    if (adapt->solve)
        adapt->solve(mesh);
```

```

    est = adapt->estimate(mesh, adapt);

    for (iter = 0;
        (est > adapt->tolerance) && iter < adapt->max_iteration;
        iter++)
    {
        if (adapt_mesh(mesh, adapt))
        {
            if (adapt->solve) adapt->solve(mesh);
            est = adapt->estimate(mesh, adapt);
        }
        ...
    }
}

```

The actual mesh adaption is done in a subroutine `adapt_mesh()`, which combines marking, refinement, coarsening and the linear system building routines:

```

static U_CHAR adapt_mesh(MESH *mesh, ADAPT_STAT *adapt)
{
    U_CHAR    flag = 0;
    U_CHAR    mark_flag;

    ...

    if (adapt->marking)
        mark_flag = adapt->marking(mesh, adapt);
    else
        mark_flag = marking(mesh, adapt); /*-- use standard marking --*/

    if (!adapt->coarsen_allowed)
        mark_flag &= MESH_REFINED; /*-- use only refinement --*/

    if (adapt->build_before_refine)
        adapt->build_before_refine(mesh, mark_flag);

    if (mark_flag & MESH_REFINED)
        flag = refine(mesh);

    if (adapt->build_before_coarsen)
        adapt->build_before_coarsen(mesh, mark_flag);

    if (mark_flag & MESH_COARSENEED)
        flag |= coarsen(mesh);

    if (adapt->build_after_coarsen)
        adapt->build_after_coarsen(mesh, flag);
}

```



```

...
return(flag);
}

```

Remark 3.31. As the same procedure is used for time dependent problems in single time steps, different pointers to routines for building parts of the (non-)linear systems make it possible, for example, to assemble the right hand side including a functional involving the solution from the old time step *before* coarsening the mesh, and then using the `DOF_VEC` restriction during coarsening to compute exactly the projection to the coarsened finite element space, without losing any information, compare Section 1.4.4.

Remark 3.32. For time dependent problems, usually the system matrices depend on the current time step size. Thus, matrices may have to be rebuilt even if meshes are not changed, but when the time step size was changed. Such changes can be detected in the `set_time()` routine, for example.

3.13.2 Standard ALBERTA marking routine

When the `marking` procedure pointer in the `ADAPT_STAT` structure is `nil`, then the standard ALBERTA marking routine is called. The `strategy` entry, allows the selection of one of five different marking strategies (compare Sections 1.5.2 and 1.5.3). Elements are only marked for coarsening and coarsening parameters are only used if the entry `coarsen_allowed` is `true`. The number of bisection steps for refinement and coarsening is selected by the entries `refine_bisections` and `coarse_bisections`.

`strategy=0`: no refinement or coarsening is performed;

`strategy=1`: Global Refinement (GR):

the mesh is refined globally, no coarsening is performed;

`strategy=2`: Maximum Strategy (MS):

the entries `MS_gamma`, `MS_gamma_c` are used as refinement and coarsening parameters;

`strategy=3`: Equidistribution strategy (ES):

the entries `ES_theta`, `ES_theta_c` are used as refinement and coarsening parameters;

`strategy=4`: Guaranteed error reduction strategy (GERS):

the entries `GERS_theta_star`, `GERS_nu`, and `GERS_theta_c` are used as refinement and coarsening parameters.

Remark 3.33. As `get_el_est()` and `get_el_estc()` return the p -th power of the local estimates, all algorithms are implemented to use the values η_S^p instead of η_S . This results in a small change to the coarsening tolerances

for the equidistribution strategy described in Section 1.5.3. The implemented equidistribution strategy uses the inequality

$$\eta_S^p + \eta_{c,S}^p \leq c^p \text{tol}^p / N_k$$

instead of

$$\eta_S + \eta_{c,S} \leq c \text{tol} / N_k^{1/p}.$$

3.13.3 ALBERTA adaptive method for time dependent problems

Similar to the data structure `ADAPT_STAT` for collecting information about the adaptive solution for a stationary problem, the data structure `ADAPT_INSTAT` is used for gather all information needed for the time and space adaptive solution of instationary problems. Using a time stepping scheme, in each time step a stationary problem is solved; the adaptive method for these stationary is based on the `adapt_method_stat()` routine described in Section 3.13.1, the `ADAPT_INSTAT` structure includes two `ADAPT_STAT` parameter structures. Additional entries give information about the time adaptive procedure.

```
typedef struct adapt_instat      ADAPT_INSTAT;
struct adapt_instat
{
    const char  *name;

    ADAPT_STAT adapt_initial[1];
    ADAPT_STAT adapt_space[1];

    REAL  time;
    REAL  start_time, end_time;
    REAL  timestep;

    void  (*init_timestep)(MESH *, ADAPT_INSTAT *);
    void  (*set_time)(MESH *, ADAPT_INSTAT *);
    void  (*one_timestep)(MESH *, ADAPT_INSTAT *);
    REAL  (*get_time_est)(MESH *, ADAPT_INSTAT *);
    void  (*close_timestep)(MESH *, ADAPT_INSTAT *);

    int    strategy;
    int    max_iteration;

    REAL  tolerance;
    REAL  rel_initial_error;
    REAL  rel_space_error;
    REAL  rel_time_error;
    REAL  time_theta_1;
    REAL  time_theta_2;
    REAL  time_delta_1;
```

```

    REAL    time_delta_2;
    int     info;
};

```

The entries yield following information:

name: textual description of the adaptive method, or **nil**.

adapt_initial: mesh adaption parameters for the initial mesh, compare Section 3.13.1.

adapt_space: mesh adaption parameters during time steps, compare Section 3.13.1.

time: actual time, end of time interval for current time step.

start_time: initial time for the adaptive simulation.

end_time: final time for the adaptive simulation.

timestep: current time step size, will be changed by the time adaptive method.

init_timestep: pointer to a routine called at beginning of each time step; if **nil**, initialization of a new time step is omitted;

init_timestep(mesh, adapt) initializes a new time step;

set_time: pointer to a routine called after changes of the time step size if not **nil**;

set_time(mesh, adapt) is called by the adaptive method each time when the actual time **adapt->time** has changed, i. e. at a new time step and after a change of the time step size **adapt->timestep**; information about actual time and time step size is available via **adapt**.

one_timestep: pointer to a routine which implements one (adaptive) time step, if **nil**, a default routine is called;

one_timestep(mesh, adapt) implements the (adaptive) solution of the problem in one single time step; information about the stationary problem of the time step is available in the **adapt->adapt_space** data structure.

get_time_est: pointer to a routine returning an estimate for the time error; if **nil**, no time step adaptation is done;

get_time_est(mesh, adapt) returns an estimate η_τ for the current time error at time **adapt->time** on **mesh**.

close_timestep: pointer to a routine called after finishing a time step, may be **nil**.

close_timestep(mesh, adapt) is called after accepting the solution(s) of the discrete problem on **mesh** at time **adapt->time** by the time-space adaptive method; can be used for visualization and export to file for post-processing of the **mesh** and discrete solution(s).

strategy: parameter for the default ALBERTA **one_timestep** routine; possible values are: 0: explicit strategy, 1: implicit strategy.

max_iteration: parameter for the default **one_timestep** routine; maximal number of time step size adaptation steps, only used by the implicit strategy.

tolerance: given total error tolerance tol .
rel_initial_error: portion Γ_0 of tolerance allowed for initial error, compare Section 1.5.4;
rel_space_error: portion Γ_h of tolerance allowed for error from spatial discretization in each time step, compare Section 1.5.4.
rel_time_error: portion Γ_τ of tolerance allowed for error from time discretization in each time step, compare Section 1.5.4.
time_theta_1: safety parameter θ_1 for the time adaptive method in the default ALBERTA `one_timestep()` routine; the tolerance for the time estimate η_τ is $\theta_1 \Gamma_\tau tol$, compare Algorithm 1.24.
time_theta_2: safety parameter θ_2 for the time adaptive method in the default ALBERTA `one_timestep()` routine; enlargement of the time step size is only allowed for $\eta_\tau \leq \theta_2 \Gamma_\tau tol$, compare Algorithm 1.24.
time_delta_1: factor δ_1 used for the reduction of the time step size in the default ALBERTA `one_timestep()` routine, compare Algorithm 1.24.
time_delta_2: factor δ_2 used for the enlargement of the time step size in the default ALBERTA `one_timestep()` routine, compare Algorithm 1.24.
info: level of information produced by the time-space adaptive procedure.
Using information given in the ADAPT_INSTAT data structure, the space and time adaptive procedure is performed by:

```
void adapt_method_instat(MESH *, ADAPT_INSTAT *);
```

Description:

adapt_method_instat(mesh, adapt_instat): the function solves an stationary problem on `mesh` by the space-time adaptive procedure described in Section 1.5.4; `adapt_instat` is a pointer to a filled ADAPT_INSTAT data structure, holding all information about the problem to be solved and parameters for the adaptive method.

Implementation of the routine is very simple. All essential work is done by calling `adapt_method_stat()` for the generation of the initial mesh, based on parameters given in `adapt->adapt_initial` with tolerance `adapt->tolerance*adapt->rel_space_error`, and in `one_timestep()` which solves the discrete problem and does mesh adaption and time step adjustment for one single time step.

```
void adapt_method_instat(MESH *mesh, ADAPT_INSTAT *adapt)
{
/*--8<-----*/
/* adaptation of the initial grid: done by adapt_method_stat() */
/*----->8---*/

adapt->time = adapt->start_time;
if (adapt->set_time) adapt->set_time(mesh, adapt);
```

```

adapt->adapt_initial->tolerance
    = adapt->tolerance * adapt->rel_initial_error;
adapt_method_stat(mesh, adapt->adapt_initial);

if (adapt->close_timestep)
    adapt->close_timestep(mesh, adapt);

/*---8<-----*/
/* adaptation of timestepsize and mesh: done by one_timestep() */
/*----->8---*/

while (adapt->time < adapt->end_time)
{
    if (adapt->init_timestep)
        adapt->init_timestep(mesh, adapt);

    if (adapt->one_timestep)
        adapt->one_timestep(mesh, adapt);
    else
        one_timestep(mesh, adapt);

    if (adapt->close_timestep)
        adapt->close_timestep(mesh, adapt);
}
}

```

The default **ALBERTA** `one_timestep()` routine

The default `one_timestep()` routine provided by **ALBERTA** implements both the explicit strategy and the implicit time strategy A. The semi-implicit strategy described in Section 1.5.4 is only a special case of the implicit strategy with a limited number of iterations (exactly one).

The routine uses the parameter `adapt->strategy` to select the strategy:

strategy 0: Explicit strategy, **strategy 1:** Implicit strategy.

Explicit strategy. The explicit strategy does one adaption of the mesh based on the error estimate computed from the last time step's discrete solution by using parameters given in `adapt->adapt_space` and with `tolerance` set to `adapt->tolerance*adapt->rel_space_error`. Then the current time step's discrete problem is solved, and the error estimators are computed. No time step size adjustment is done.

Implicit strategy. The implicit strategy starts with the old mesh given from the last time step. Using parameters given in `adapt->adapt_space`, the discrete problem is solved on the current mesh. Error estimates are computed,

and time step size and mesh are adjusted, as shown in Algorithm 1.25. The tolerances used for time and space estimate are set to

```
adapt->tolerance * adapt->rel_time_error
```

and

```
adapt->tolerance * adapt->rel_space_error,
```

respectively. This is iterated until the given error bounds are reached, or until `adapt->max_iteration` is reached.

With parameter `adapt->max_iteration==0`, this is equivalent to the semi-implicit strategy described in Section 1.5.4.

3.13.4 Initialization of data structures for adaptive methods

ALBERTA provides functions for the initialization of the data structures `ADAPT_STAT` and `ADAPT_INSTAT`. Both functions do *not* fill any function pointer entry in the structures! These function pointers have to be adjusted in the application.

Table 3.15. Initialized members of an `ADAPT_STAT` structure accessed by the function `get_adapt_stat()` with default values and key for the initialization by `GET_PARAMETER()`.

member	default	parameter key
<code>tolerance</code>	1.0	<code>pre->tolerance</code>
<code>p</code>	2	<code>pre->p</code>
<code>max_iteration</code>	30	<code>pre->max_iteration</code>
<code>info</code>	2	<code>pre->info</code>
<code>refine_bisections</code>	DIM	<code>pre->refine_bisections</code>
<code>coarsen_allowed</code>	0	<code>pre->coarsen_allowed</code>
<code>coarse_bisections</code>	DIM	<code>pre->coarse_bisections</code>
<code>strategy</code>	1	<code>pre->strategy</code>
<code>MS_gamma</code>	0.5	<code>pre->MS_gamma</code>
<code>MS_gamma_c</code>	0.1	<code>pre->MS_gamma_c</code>
<code>ES_theta</code>	0.9	<code>pre->ES_theta</code>
<code>ES_theta_c</code>	0.2	<code>pre->ES_theta_c</code>
<code>GERS_theta_star</code>	0.6	<code>pre->GERS_theta_star</code>
<code>GERS_nu</code>	0.1	<code>pre->GERS_nu</code>
<code>GERS_theta_c</code>	0.1	<code>pre->GERS_theta_c</code>

```
ADAPT_STAT *get_adapt_stat(const char *, const char *, int,
                           ADAPT_STAT *);
ADAPT_INSTAT *get_adapt_instat(const char *, const char *, int,
                               ADAPT_INSTAT *);
```

Description:

`get_adapt_stat(name, pre, info, adapt)`: returns a pointer to a partly initialized `ADAPT_STAT` structure; if the argument `adapt` is `nil`, a new structure is created, the name `name` is duplicated at the name entry of the structure, if `name` is not `nil`; if `name` is `nil`, and `pre` is not `nil`, this string is duplicated at the name entry; for a newly created structure, all function pointers of the structure are initialized with `nil`; all other members are initialized with some default value; if the argument `adapt` is not `nil`, this initialization part is skipped, the name and function pointers are not changed; if `pre` is not a `nil` pointer, `get_adapt_stat()` tries then to initialize members by a call of `GET_PARAMETER()`, where the key for each member is `value(pre)->member name`; the argument `info` is the first argument of `GET_PARAMETER()` giving the level of information for the initialization; only the parameters for the actually chosen strategy are initialized using the function `GET_PARAMETER()`: for `strategy == 2` only `MS_gamma` and `MS_gamma_c`, for `strategy == 3` only `ES_theta` and `ES_theta_c`, and for `strategy == 4` only `GERS_theta_star`, `GERS_nu`, and `GERS_theta_c`; since the parameter tools are used for the initialization, `get_adapt_stat()` should be called *after* the initialization of all parameters; there may be no initializer in the parameter file(s) for some member, if the default value should be used; if `info` is not zero and there is no initializer for some member this will result in an error message by `GET_PARAMETER()` which can be ignored; Table 3.15 shows the initialized members, the default values and the key used for the initialization by `GET_PARAMETER()`;

Table 3.16. Initialization of the main parameters in an `ADAPT_INSTAT` structure for the time-adaptive strategy by `get_adapt_instat()`; initialized members, the default values and keys used for the initialization by `GET_PARAMETER()`.

member	default	parameter key
<code>start_time</code>	0.0	<code>pre->start_time</code>
<code>end_time</code>	1.0	<code>pre->end_time</code>
<code>timestep</code>	0.01	<code>pre->timestep</code>
<code>strategy</code>	0	<code>pre->strategy</code>
<code>max_iteration</code>	0	<code>pre->max_iteration</code>
<code>tolerance</code>	1.0	<code>pre->tolerance</code>
<code>rel_initial_error</code>	0.1	<code>pre->rel_initial_error</code>
<code>rel_space_error</code>	0.4	<code>pre->rel_space_error</code>
<code>rel_time_error</code>	0.4	<code>pre->rel_time_error</code>
<code>time_theta_1</code>	1.0	<code>pre->time_theta_1</code>
<code>time_theta_2</code>	0.3	<code>pre->time_theta_2</code>
<code>time_delta_1</code>	0.7071	<code>pre->time_delta_1</code>
<code>time_delta_2</code>	1.4142	<code>pre->time_delta_2</code>
<code>info</code>	8	<code>pre->info</code>

`get_adapt_instat(name, pre, info, adapt)`: returns a pointer to a partially initialized `ADAPT_INSTAT` structure; if the argument `adapt` is `nil`, a new structure is created, the name `name` is duplicated at the name entry of the structure, if `name` is not `nil`; if `name` is `nil`, and `pre` is not `nil`, this string is duplicated at the name entry; for a newly created structure, all function pointers of the structure are initialized with `nil`; all other members are initialized with some default value; if the argument `adapt` is not `nil`, this default initialization part is skipped;

if `pre` is not `nil`, `get_adapt_instat()` tries then to initialize members by a call of `GET_PARAMETER()`, where the key for each member is `value(pre)->member name`; the argument `info` is the first argument of `GET_PARAMETER()` giving the level of information for the initialization;

Tables 3.16–3.18 show all initialized members, their default values and the key used for the initialization by `GET_PARAMETER()`. The tolerances in the two sub-structures `adapt_initial` and `adapt_space` are set automatically to the values `adapt->tolerance*adapt->rel_initial_error`, respectively `adapt->tolerance*adapt->rel_space_error`. A special initialization is done for the `info` parameters: when `adapt_initial->info` or `adapt_space->info` are negative, then they are set to `adapt->info-2`.

Table 3.17. Initialization of `adapt_initial` inside an `ADAPT_INSTAT` structure for the adaptation of the initial grid by `get_adapt_instat()`; initialized members, the default values and keys used for the initialization by `GET_PARAMETER()`.

member	default	parameter key
<code>adapt_initial->tolerance</code>	–	–
<code>adapt_initial->p</code>	2	<code>pre->initial->p</code>
<code>adapt_initial->max_iteration</code>	30	<code>pre->initial->max_iteration</code>
<code>adapt_initial->info</code>	2	<code>pre->initial->info</code>
<code>adapt_initial->refine_bisections</code>	DIM	<code>pre->initial->refine_bisections</code>
<code>adapt_initial->coarsen_allowed</code>	0	<code>pre->initial->coarsen_allowed</code>
<code>adapt_initial->coarse_bisections</code>	DIM	<code>pre->initial->coarse_bisections</code>
<code>adapt_initial->strategy</code>	1	<code>pre->initial->strategy</code>
<code>adapt_initial->MS_gamma</code>	0.5	<code>pre->initial->MS_gamma</code>
<code>adapt_initial->MS_gamma_c</code>	0.1	<code>pre->initial->MS_gamma_c</code>
<code>adapt_initial->ES_theta</code>	0.9	<code>pre->initial->ES_theta</code>
<code>adapt_initial->ES_theta_c</code>	0.2	<code>pre->initial->ES_theta_c</code>
<code>adapt_initial->GERS_theta_star</code>	0.6	<code>pre->initial->GERS_theta_star</code>
<code>adapt_initial->GERS_nu</code>	0.1	<code>pre->initial->GERS_nu</code>
<code>adapt_initial->GERS_theta_c</code>	0.1	<code>pre->initial->GERS_theta_c</code>

Table 3.18. Initialization of `adapt_space` inside an `ADAPT_INSTAT` structure for the adaptation of the grids during time-stepping by `get_adapt_instat()`; initialized members, the default values and keys used for the initialization by `GET_PARAMETER()`.

member	default	parameter key
<code>adapt_space->tolerance</code>	—	—
<code>adapt_space->p</code>	2	<code>pre->space->p</code>
<code>adapt_space->max_iteration</code>	30	<code>pre->space->max_iteration</code>
<code>adapt_space->info</code>	2	<code>pre->space->info</code>
<code>adapt_space->refine_bisections</code>	DIM	<code>pre->space->refine_bisections</code>
<code>adapt_space->coarsen_allowed</code>	1	<code>pre->space->coarsen_allowed</code>
<code>adapt_space->coarse_bisections</code>	DIM	<code>pre->space->coarse_bisections</code>
<code>adapt_space->strategy</code>	1	<code>pre->space->strategy</code>
<code>adapt_space->MS_gamma</code>	0.5	<code>pre->space->MS_gamma</code>
<code>adapt_space->MS_gamma_c</code>	0.1	<code>pre->space->MS_gamma_c</code>
<code>adapt_space->ES_theta</code>	0.9	<code>pre->space->ES_theta</code>
<code>adapt_space->ES_theta_c</code>	0.2	<code>pre->space->ES_theta_c</code>
<code>adapt_space->GERS_theta_star</code>	0.6	<code>pre->space->GERS_theta_star</code>
<code>adapt_space->GERS_nu</code>	0.1	<code>pre->space->GERS_nu</code>
<code>adapt_space->GERS_theta_c</code>	0.1	<code>pre->space->GERS_theta_c</code>

3.14 Implementation of error estimators

3.14.1 Error estimator for elliptic problems

ALBERTA provides a residual type error estimator for non-linear elliptic problems of the type

$$\begin{aligned}
 -\nabla \cdot A \nabla u(x) + f(x, u(x), \nabla u(x)) &= 0 & x \in \Omega, \\
 u(x) &= 0 & x \in \Gamma_D, \\
 \nu \cdot A \nabla u(x) &= 0 & x \in \Gamma_N,
 \end{aligned}$$

where $A \in \mathbb{R}^{d \times d}$ is a positive definite matrix and $\partial\Omega = \Gamma_D \cup \Gamma_N$.

Verfürth [71] proved for this kind of equation under suitable assumptions on f , u and u_h (in the non-linear case) the following estimate

$$\begin{aligned}
 \|u - u_h\|_{H^1(\Omega)}^2 &\leq \sum_{S \in \mathcal{S}} C_0^2 h_S^2 \| -\nabla \cdot A \nabla u_h + f(\cdot, u_h, \nabla u_h) \|_{L^2(S)}^2 \\
 &\quad + C_1^2 \sum_{\Gamma \subset \partial\mathcal{S} \cap (\Omega \cup \Gamma_N)} h_S \| [\nu \cdot A \nabla u_h] \|_{L^2(\Gamma)}^2,
 \end{aligned}$$

where $[\cdot]$ denotes the jump of a quantity across an interior edge/face or the its value for an edge/face on the Neumann boundary.

Bänsch and Siebert [10] proved a similar L^2 error estimate for semi-linear problems, i.e. $f = f(x, u)$, namely

$$\begin{aligned} \|u - u_h\|_{L^2(\Omega)}^2 &\leq \sum_{S \in \mathcal{S}} C_0^2 h_S^4 \| -\nabla \cdot A \nabla u_h + f(\cdot, u_h) \|_{L^2(S)}^2 \\ &\quad + C_1^2 \sum_{\Gamma \subset \partial S \cap (\Omega \cup \Gamma_N)} h_S^3 \| [\nu \cdot A \nabla u_h] \|_{L^2(\Gamma)}^2. \end{aligned}$$

The following function is an implementation of the above estimators:

```
REAL ellipt_est(const DOF_REAL_VEC *, ADAPT_STAT *, REAL (*)(EL *),
               REAL (*)(EL *), int, int, REAL[3], const REAL_DD,
               REAL (*f)(const EL_INFO *, const QUAD *, int, REAL,
                        const REAL_D), FLAGS);
```

Description:

ellipt_est(uh, adapt, rw_e, rw_ec, deg, norm, C, A, f, f_flag):
 computes an error estimate of the above type for the H^1 or L^2 norm; the return value is an approximation of the estimate $\|u - u_h\|$ by quadrature.
 uh is a vector storing the coefficients of the discrete solution; if uh is a nil pointer, nothing is done, the return value is 0.0.

adapt is a pointer to an ADAPT_STAT structure; if not nil, the entries adapt->p=2, err_sum, and err_max of adapt are set by ellipt_est() (compare Section 3.13.1).

rw_e is a function for writing the local error indicator for a single element (usually to some location inside leaf_data, compare Section 3.2.12); if this function is nil, only the global estimate is computed, no local indicators are stored. rw_e(e1) returns for each leaf element e1 a pointer to a REAL for storing the square of the element indicator, which can directly be used in the adaptive method, compare the get_el_est() function pointer in the ADAPT_STAT structure (compare Section 3.13.1).

rw_ec is a function for writing the local coarsening error indicator for a single element (usually to some location inside leaf_data, compare Section 3.2.12); if this function is nil, no coarsening error indicators are computed and stored; rw_ec(e1) returns for each leaf element e1 a pointer to a REAL for storing the square of the element coarsening error indicator.

deg is the degree of the quadrature that should be used for the approximation of the norms on the elements and edges/faces; if deg is less than zero a quadrature which is exact of degree 2*uh->fe_space->bas_fcts->degree is used.

norm can be either H1_NORM or L2_NORM (which are defined as symbolic constants in alberta.h) to indicate that the H^1 or L^2 error estimate has to be calculated.

C[0], C[1], C[2] are the constants in front of the element residual, edge/face residual, and coarsening term respectively. If C is nil, then all constants are set to 1.0.

A is the constant matrix of the second order term.

`f` is a pointer to a function for the evaluation of the lower order terms at all quadrature nodes, i.e. $f(x(\lambda), u(\lambda), \nabla u(\lambda))$; if `f` is a `nil` pointer, $f \equiv 0$ is assumed;

`f(el_info, quad, iq, uh_iq, grd_uh_iq)` returns the value of the lower order terms of the element residual on element `el_info->el` at the quadrature node `quad->lambda[iq]`, where `uh_iq` is the value and `grd_uh_iq` the gradient (with respect to the world coordinates) of the discrete solution at that quadrature node.

`f_flag` specifies whether the function `f()` actually needs values of `uh_iq` or `grd_uh_iq`. This flag may hold zero, the predefined values `INIT_UH` or `INIT_GRD_UH`, or their composition `INIT_UH|INIT_GRD_UH`; the arguments `uh_iq` and `grd_uh_iq` of `f()` only hold valid information, if the flags `INIT_UH` respectively `INIT_GRD_UH` are set.

The estimate is computed by a mesh traversal of all leaf elements of `uh->fe_space->mesh`, using the quadrature for the approximation of the residuals and storing the square of the element indicators on the elements (if `rw_e` and `rw_ec` are not `nil`).

Example 3.34 (Linear problem). Consider the linear model problem (1.7) with constant coefficients A , b , and c :

$$\begin{aligned} -\nabla \cdot A \nabla u + b \cdot \nabla u + c u &= r && \text{in } \Omega, \\ u &= 0 && \text{on } \partial\Omega. \end{aligned}$$

Let A be a `REAL_DD` matrix storing A , which is then the eighth argument of `ellipt_est()`. Assume that `const REAL *b(const REAL_D)` is a function returning a pointer to a vector storing b , `REAL c(REAL_D)` returns the value of c and `REAL r(const REAL_D)` returns the value of the right hand side r of (1.7) at some point in world coordinates. The implementation of the function `f` is:

```
static REAL f(const EL_INFO *el_info, const QUAD *quad, int iq,
              REAL uh_iq, const REAL_D grd_uh_iq)
{
    FUNCNAME("f");
    const REAL *bx, *x;
    extern const REAL b(const REAL_D);
    extern REAL      c(const REAL_D), r(const REAL_D);

    x = coord_to_world(el_info, quad->lambda[iq], nil);
    bx = b(x);

    return(SCP_DOW(bx, grd_uh_iq) + c(x)*uh_iq - r(x));
}
```

As both `uh_iq` and `grd_uh_iq` are used, the estimator parameter `f_flag` must be given as `INIT_UH|INIT_GRD_UH`.

3.14.2 Error estimator for parabolic problems

Similar to the stationary case, the ALBERTA library provides an error estimator for the non-linear parabolic problem

$$\begin{aligned} \partial_t u - \nabla \cdot A \nabla u(x) + f(x, t, u(x), \nabla u(x)) &= 0 & x \in \Omega, t > 0, \\ u(x, t) &= 0 & x \in \Gamma_D, t > 0, \\ \nu \cdot A \nabla u(x, t) &= 0 & x \in \Gamma_N, t > 0, \\ u(x, 0) &= u_0 & x \in \Omega, \end{aligned}$$

where $A \in \mathbb{R}^{d \times d}$ is a positive definite matrix and $\partial\Omega = \Gamma_D \cup \Gamma_N$. The estimator is split in several parts, where the initial error

$$\eta_0 = \|u_0 - U_0\|_{L^2(\Omega)}$$

can be approximated by the function `L2_err()`, e.g. (compare Section 3.11).

For the estimation of the spatial discretization error, the coarsening error, and the time discretization error, the ALBERTA estimator is given by: The local error indicator on element $S \in \mathcal{S}$

$$\begin{aligned} \eta_S^2 &= C_0^2 h_S^4 \left\| \frac{U_{n+1} - I_{n+1}U_n}{\tau_{n+1}} - \nabla \cdot A \nabla U_{n+1} + f(\cdot, t_{n+1}, U_{n+1}, \nabla U_{n+1}) \right\|_{L^2(S)}^2 \\ &\quad + C_1^2 \sum_{\Gamma \subset \partial S \cap (\Omega \cup \Gamma_N)} h_S^3 \| \llbracket \nu \cdot A \nabla U_{n+1} \rrbracket \|_{L^2(\Gamma)}^2, \end{aligned}$$

the local coarsening error indicator for $S \in \mathcal{S}$

$$\eta_{S,c}^2 = C_2^2 h_S^3 \| \llbracket \nabla U_n \rrbracket \|_{L^2(\Gamma_c)}^2$$

and the estimator for the time error

$$\eta_\tau = C_3 \|U_{n+1} - I_{n+1}U_n\|_{L^2(\Omega)}.$$

The coarsening indicator is motivated by the fact that, for piecewise linear Lagrange finite element functions, it holds $\|U_n - I_{n+1}U_n\|_{L^2(S)}^2 = \eta_{S,c}^2$ with $C_2 = C_2(d)$ and Γ_c the element vertex/edge/face which would be removed during a coarsening operation.

The implementation is done by the function

```
REAL heat_est(const DOF_REAL_VEC *, ADAPT_INSTAT *, REAL (*)(EL *),
              REAL (*)(EL *), int, REAL[4], const DOF_REAL_VEC *,
              const REAL_DD,
              REAL (*)(const EL_INFO *, const QUAD *, int, REAL,
                      REAL, const REAL_D), FLAGS);
```

Description:

heat_est(**uh**, **adapt**, **rw_e**, **rw_ec**, **deg**, **C**, **uh_old**, **A**, **f**, **f_flag**):
 computes an error estimate of the above type, the local and global space discretization estimators are stored in **adapt**->**adapt_space** and via the **rw_e** and **rw_ec** pointers; the return value is the time discretization estimate η_τ .
uh is a vector storing the coefficients of the discrete solution U_{n+1} ; if **uh** is a **nil** pointer, nothing is done, the return value is 0.0.

adapt is an optional pointer to an **ADAPT_INSTAT** structure; if it is not a **nil** pointer, then the entries **adapt_space**->**p**=2, **adapt_space**->**err_sum** and **adapt_space**->**err_max** of **adapt** are set by **heat_est**() (compare Section 3.13.1).

rw_e is a function for writing the local error indicator η_S^2 for a single element (usually to some location inside **leaf_data**, compare Section 3.2.12); if this function is **nil**, only the global estimate is computed, no local indicators are stored. **rw_e**(**e1**) returns for each leaf element **e1** a pointer to a **REAL** for storing the square of the element indicator, which can directly be used in the adaptive method, compare the **get_el_est**() function pointer in the **ADAPT_STAT** structure (compare Section 3.13.1).

rw_ec is a function for writing the local coarsening error indicator $\eta_{S,c}^2$ for a single element (usually to some location inside **leaf_data**, compare Section 3.2.12); if this function is **nil**, no coarsening error indicators are computed and stored; **rw_ec**(**e1**) returns for each leaf element **e1** a pointer to a **REAL** for storing the square of the element coarsening error indicator.

degree is the degree of the quadrature used for the approximation of the norms on the elements and edges/faces; if **degree** is less than zero a default quadrature which is exact of degree $2 * \text{uh} \rightarrow \text{fe_space} \rightarrow \text{bas_fcts} \rightarrow \text{degree}$ is used.

C[0], **C[1]**, **C[2]**, **C[3]** are the constants in front of the element residual, edge/face residual, coarsening term, and time residual, respectively. If **C** is **nil**, then all constants are set to 1.0.

uh_old is a vector storing the coefficients of the discrete solution U_n from previous time step; if **uh_old** is a **nil** pointer, nothing is done, the return value is 0.0.

A is the constant matrix of the second order term.

f is a pointer to a function for the evaluation of the lower order terms at all quadrature nodes, i.e. $f(x(\lambda), t, u(\lambda), \nabla u(\lambda))$; if **f** is a **nil** pointer, $f \equiv 0$ is assumed;

f(**el_info**, **quad**, **iq**, **t**, **uh_iq**, **grd_uh_iq**) returns the value of the lower order terms of the element residual on element **el_info**->**e1** at the quadrature node **quad**->**lambda**[**iq**], where **uh_iq** is the value and **grd_uh_iq** the gradient (with respect to the world coordinates) of the discrete solution at that quadrature node.

f_flag specifies whether the function **f**() actually needs values of **uh_iq** or **grd_uh_iq**. This flag may hold zero, the predefined values **INIT_UH** or

INIT_GRD_UH, or their composition INIT_UH|INIT_GRD_UH; the arguments `uh_iq` and `grd_uh_iq` of `f()` only hold valid information, if the flags INIT_UH respectively INIT_GRD_UH are set.

The estimate is computed by a mesh traversal of all leaf elements of `uh->fe_space->mesh`, using the quadrature for the approximation of the residuals and storing the square of the element indicators on the elements (if `rw_e` and `rw_ec` are not `nil`).

3.15 Solver for linear and nonlinear systems

ALBERTA uses a library for solving general linear and nonlinear systems. The solvers use `REAL` vectors for storing coefficients. They do not know about ALBERTA data structures especially they do not know DOF vectors and matrices used in ALBERTA. The linear solvers only need a subroutine for the matrix–vector multiplication, and in the case that a preconditioner is used, a function for preconditioning. The nonlinear solvers need subroutines for assemblage and solution of a linearized system.

In the subsequent sections we describe the basic data structures for the OEM (Orthogonal Error Methods) library, an ALBERTA interface for solving systems involving a `DOF_MATRIX` and `DOF_REAL[_D]` vectors, and the access of functions for matrix–vector multiplication and preconditioning for a direct use of the OEM solvers. Then we describe the basic data structures for multi-grid solvers and for the available solvers of nonlinear equations. Most of the implemented methods (and more) are described for example in [47, 58].

3.15.1 General linear solvers

Highly efficient solvers for linear systems are (preconditioned) Krylov-space solvers (or Orthogonal Error Methods). The OEM library provides such solvers for the solution of general linear systems

$$Ax = b$$

with $A \in \mathbb{R}^{N \times N}$ and $x, b \in \mathbb{R}^N$. The library solvers work on vectors and do not need to know the storage of the system matrix, or the matrix used for preconditioning. Matrix–vector multiplication and preconditioning is done by application dependent routines.

The OEM solvers are *not* part of ALBERTA. For the access of the basic data structure and prototypes of solvers, the header file

```
#include <oem.h>
```

has to be included in each file using a solver from the library. Most of the implemented solvers are a C-translation from the solvers of FORTRAN OFM

library (Orthogonale Fehler Methoden), by Dörfler [29]. All solvers allow a *left* preconditioning and some also a *right* preconditioning.

The data structure (defined in `oem.h`) for passing information about matrix–vector multiplication, preconditioning and tolerances, etc. to the solvers is

```
typedef struct oem_data OEM_DATA;
struct oem_data
{
    int      (*mat_vec)(void *, int, const REAL *, REAL *);
    void     *mat_vec_data;
    int      (*mat_vec_T)(void *, int, const REAL *, REAL *);
    void     *mat_vec_T_data;
    void     (*left_precon)(void *, int, REAL *);
    void     *left_precon_data;
    void     (*right_precon)(void *, int, REAL *);
    void     *right_precon_data;

    REAL     (*scp)(void *, int, const REAL *, const REAL *);
    void     *scp_data;

    WORKSPACE *ws;

    REAL     tolerance;
    int      restart;
    int      max_iter;
    int      info;

    REAL     initial_residual;
    REAL     residual;
};
```

Description:

mat_vec: pointer to a function for the matrix–vector multiplication with the system matrix;

mat_vec(mat_vec_data, dim, u, b) applies the system matrix to the input vector *u* and stores the product in *b*; *dim* is the dimension of the linear system, **mat_vec_data** a pointer to user data.

mat_vec_data: pointer to user data for the matrix–vector multiplication, first argument to **mat_vec()**.

mat_vec_T: pointer to a function for the matrix–vector multiplication with the transposed system matrix;

mat_vec_T(mat_vec_data, dim, u, b) applies the transposed system matrix to the input vector *u* and stores the product in *b*; *dim* is the dimension of the linear system, **mat_vec_T_data** a pointer to user data.

mat_vec_T_data: pointer to user data for the matrix–vector multiplication with the transposed system matrix, first argument to **mat_vec_T()**.

left_precon: pointer to function for left preconditioning; it may be a `nil` pointer; in this case no left preconditioning is done;
left_precon(left_precon_data, dim, r) is the implementation of the left preconditioner; **r** is input and output vector of length **dim** and **left_precon_data** a pointer to user data.

left_precon_data: pointer to user data for the left preconditioning, first argument to **left_precon()**.

right_precon: pointer to function for right preconditioning; it may be a `nil` pointer; in this case no right preconditioning is done;
right_precon(right_precon_data, dim, r) is the implementation of the right preconditioner; **r** is input and output vector of length **dim** and **right_precon_data** a pointer to user data.

right_precon_data: pointer to user data for the right preconditioning, first argument to **right_precon()**.

scp: pointer to a function for computing a problem dependent scalar product; it may be a `nil` pointer; in this case the Euclidian scalar product is used;
scp(scp_data, dim, x, y) computes a problem dependent scalar product of two vectors **x** and **y** of length **dim**; **scp_data** is a pointer to user data.

scp_data: pointer to user data for computing the scalar product, first argument to **scp()**.

ws: a pointer to a **WORKSPACE** structure for storing additional vectors used by a solver; if the space is not sufficient, the used solver will enlarge this workspace; if **ws** is `nil`, then the used solver allocates memory, which is freed before exit.

tolerance: tolerance for the residual; if the norm of the residual is less than or equal to **tolerance**, the solver returns the actual iterate as the solution of the system.

restart: restart for the linear solver; used only by **oem_gmres()** at the moment.

max_iter: maximal number of iterations to be performed although the tolerance may not be reached.

info: the level of information produced by the solver; 0 is the lowest level of information (no information is printed) and 10 the highest level.

initial_residual: stores the norm of the initial residual on exit.

residual: stores the norm of the last residual on exit.

The following linear solvers are currently implemented. Table 3.19 gives an overview over the implemented solvers, the matrix types they apply to, and the cost of one iteration.

Table 3.19. OEM methods with applicable matrix types, numbers of operations per iteration (MV matrix-vector products, V vector operations), and storage requirements (N number of unknowns, k GMRES subspace dimension)

Method	Matrix	Operations	Storage
BiCGstab	symmetric	2 MV + 12 V	$5N$
CG	symmetric positive definite	1 MV + 5 V	$3N$
GMRES	regular	k MV + ...	$(k + 2)N + k(k + 4)$
ODir	symmetric positive	1 MV + 11 V	$5N$
ORes	symmetric	1 MV + 12 V	$7N$

```

int oem_bicgstab(OEM_DATA *, int, const REAL *, REAL *);
int oem_cg(OEM_DATA *, int, const REAL *, REAL *);
int oem_gmres(OEM_DATA *, int, const REAL *, REAL *);
int oem_gmres_k(OEM_DATA *, int, const REAL *, REAL *);
int oem_odir(OEM_DATA *, int, const REAL *, REAL *);
int oem_ores(OEM_DATA *, int, const REAL *, REAL *);

```

Description:

`oem_bicgstab(oem_data, dim, b, x0)`: solves a linear system by a stabilized BiCG method and can be used for symmetric system matrices; `oem_data` stores information about matrix vector multiplication, preconditioning, tolerances, etc. `dim` is the dimension of the linear system, `b` the right hand side vector, and `x0` the initial guess on input and the solution on output; `oem_bicgstab()` needs a workspace for storing $5 \cdot \text{dim}$ additional REALs; the return value is the number of iterations; `oem_bicgstab()` only uses left preconditioning.

`oem_cg(oem_data, dim, b, x0)`: solves a linear system by the conjugate gradient method and can be used for symmetric positive definite system matrices; `oem_data` stores information about matrix vector multiplication, preconditioning, tolerances, etc. `dim` is the dimension of the linear system, `b` the right hand side vector, and `x0` the initial guess on input and the solution on output; `oem_cg()` needs a workspace for storing $3 \cdot \text{dim}$ additional REALs; the return value is the number of iterations; `oem_cg()` only uses left preconditioning.

`oem_gmres(oem_data, dim, b, x0)`: solves a linear system by the GM-Res method with restart and can be used for regular system matrices; `oem_data` stores information about matrix vector multiplication, preconditioning, tolerances, etc. `dim` is the dimension of the linear system, `b` the right hand side vector, and `x0` the initial guess on input and the solution on output; `oem_data->restart` is the dimension of the Krylov-space for the minimizing procedure; `oem_data->restart` must be bigger than 0 and less or equal `dim`, otherwise `restart=10` will be used; `oem_gmres()` needs a workspace for storing $(\text{oem_data->restart}+2) \cdot \text{dim} + \text{oem_data->restart} \cdot (\text{oem_data->restart}+4)$ additional REALs.

`oem_gmres_k(oem_data, dim, b, x0)`: a single restart step (minimization on a k -dimensional Krylov subspace) of the GMRES method. This routine can be used as subroutine in other solvers. For example, `oem_gmres()` just iterates this until the tolerance is met. Other applications are nonlinear GMRES solvers, where a new linearization is done after each linear GMRES restart step.

`oem_odir(oem_data, dim, b, x0)`: solves a linear system by the method of orthogonal directions and can be used for symmetric, positive system matrices; `oem_data` stores information about matrix vector multiplication, preconditioning, tolerances, etc. `dim` is the dimension of the linear system, `b` the right hand side vector, and `x0` the initial guess on input and the solution on output; `oem_dir()` needs a workspace for storing $5 \times \text{dim}$ additional REALs; the return value is the number of iterations; `oem_odir()` only uses left preconditioning.

`oem_ores(oem_data, dim, b, x0)`: solves a linear system by the method of orthogonal residuals and can be used for symmetric system matrices; `oem_data` stores information about matrix vector multiplication, preconditioning, tolerances, etc. `dim` is the dimension of the linear system, `b` the right hand side vector, and `x0` the initial guess on input and the solution on output; `oem_res()` needs a workspace for storing $7 \times \text{dim}$ additional REALs; the return value is the number of iterations; `oem_ores()` only uses left preconditioning.

3.15.2 Linear solvers for DOF matrices and vectors

OEM solvers

The following functions are an interface to the above described solvers for DOF matrices and vectors. The function `oem_solve_s` is used for scalar valued problems, i.e.

$$Ax = b$$

with $A \in \mathbb{R}^{N \times N}$ and $x, b \in \mathbb{R}^N$, and `oem_solve_d` for decoupled vector valued problems of the form

$$\begin{bmatrix} A & 0 & \dots & 0 \\ 0 & A & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & A \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_d \end{bmatrix} = \begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_d \end{bmatrix}$$

with $A \in \mathbb{R}^{N \times N}$ and $u_i, f_i \in \mathbb{R}^N$, $i = 1, \dots, d$, where $d = \text{DIM_OF_WORLD}$. The vectors (u_1, \dots, u_d) and (f_1, \dots, f_d) are stored in `DOF_REAL_D_VECs`, whereas the matrix is stored as a single `DOF_MATRIX`.

For the solver identification we use the following type

```
typedef enum {NoSolver, BiCGStab, CG, GMRes, ODir, ORes} OEM_SOLVER;
```

This type will be changed at the time when additional solvers will be available.

The following functions can directly be used for solving a linear system involving a `DOF_MATRIX` and `DOF_REAL[_D]_VECS`.

```
int oem_solve_s(const DOF_MATRIX *, const DOF_REAL_VEC *,
               DOF_REAL_VEC *, OEM_SOLVER, REAL, int, int, int,
               int);
int oem_solve_d(const DOF_MATRIX *, const DOF_REAL_D_VEC *,
               DOF_REAL_D_VEC *, OEM_SOLVER, REAL, int, int, int,
               int);
```

Description:

`oem_solve_s[d](A, f, u, isol, tol, icon, restart, miter, info):`
solves the linear system for a scalar or decoupled vector valued problem in ALBERTA by an OEM solver; the return value is the number of used iterations;

A: pointer to a `DOF_MATRIX` storing the system matrix;

f: pointer to a `DOF_REAL[_D]_VEC` storing the right hand side of the linear system.

u: pointer to a `DOF_REAL[_D]` storing the initial guess on input and the calculated solution on output;

the values for **u** and **f** have to be the same at all Dirichlet DOFs (compare Section 3.12.5).

isol: use solver **isol** from the OEM library for solving the linear system; may be one of `BiCGStab`, `CG`, `GMRes`, `ODir`, or `ORes`; the meaning of **isol** is more or less self explaining.

tol: tolerance for the residual; if the norm of the residual is less or equal **tol**, `oem_solve_s[d]()` returns the actual iterate as the solution of the system.

icon: parameter for performing standard preconditioning, if argument **precon** is `nil`:

- 0: no preconditioning,
- 1: diagonal preconditioning,
- 2: hierarchical basis preconditioning,
- 3: BPX preconditioning.

miter: maximal number of iterations to be performed by the linear solver although the tolerance may not be reached.

info: is the level of information of the linear solver; 0 is the lowest level of information (no information is printed) and 10 the highest level.

The function initializes a data structure `oem_data` with information about the matrix–vector multiplication, preconditioning, tolerances, etc. and the additional memory needed by the linear solver is allocated automatically. The linear system is then solved by the chosen OEM solver.

SOR solvers

The SOR and SSOR methods are implemented directly for a linear system involving a `DOF_MATRIX` and `DOF_REAL_[D_]` VECs.

```
int sor_s(DOF_MATRIX *, const DOF_REAL_VEC *, const DOF_SCHAR_VEC *,
          DOF_REAL_VEC *, REAL, REAL, int, int);
int sor_d(DOF_MATRIX *, const DOF_REAL_D_VEC *,
          const DOF_SCHAR_VEC *, DOF_REAL_D_VEC *,
          REAL, REAL, int, int);
int ssor_s(DOF_MATRIX *, const DOF_REAL_VEC *, const DOF_SCHAR_VEC *,
          DOF_REAL_VEC *, REAL, REAL, int, int);
int ssor_d(DOF_MATRIX *, const DOF_REAL_D_VEC *,
          const DOF_SCHAR_VEC *, DOF_REAL_D_VEC *,
          REAL, REAL, int, int);
```

`[s]sor_s[d](matrix, f, bound, u, omega, tol, miter, info)`: solves the linear system for a scalar or decoupled vector valued problem in ALBERTA by the [Symmetric] Successive Over Relaxation method; the return value is the number of used iterations to reach the prescribed tolerance;

`matrix`: pointer to a DOF matrix storing the system matrix;

`f`: pointer to a DOF vector storing the right hand side of the system;

`bound`: optional pointer to a DOF vector giving Dirichlet boundary information;

`u`: pointer to a DOF vector storing the initial guess on input and the calculated solution on output;

`omega`: the relaxation parameter and must be in the interval $(0, 2]$; if it is not in this interval then `omega=1.0` is used;

`tol`: tolerance for the maximum norm of the correction; if this norm is less than or equal to `tol`, then `sor_s[d]()` returns the actual iterate as the solution of the system;

`miter`: maximal number of iterations to be performed by `sor_s[d]()` although the tolerance may not be reached;

`info`: level of information of `sor_s[d]()`; 0 is the lowest level of information (no information is printed) and 6 the highest level.

3.15.3 Access of functions for matrix–vector multiplication

The general `oem_...()` solvers all need pointers to matrix–vector multiplication routines which do not work with `DOF_REAL_VEC`s and a `DOF_MATRIX` but directly on `REAL` vectors. For the application to a scalar or vector–valued linear system described by a `DOF_MATRIX` (and an optional `DOF_SCHAR_VEC` holding boundary information), the following routines are provided.

```

void *init_mat_vec_s(MatrixTranspose, const DOF_MATRIX *,
                    const DOF_SCHAR_VEC *);
int mat_vec_s(void *, int, const REAL *, REAL *);

void *init_mat_vec_d(MatrixTranspose, const DOF_MATRIX *,
                    const DOF_SCHAR_VEC *);
int mat_vec_d(void *, int, const REAL *, REAL *);

```

Description:

init_mat_vec_s[d](transpose, matrix, bound): initialization routine for the general matrix–vector multiplication routine **mat_vec_s[d]()** for a system matrix from a scalar or decoupled vector valued problem stored in a DOF matrix **matrix**;

the original matrix is used if **transpose == NoTranspose (= 0)** and the transposed matrix if **transpose == Transpose (= 1)**;

if **bound** is not **nil**, the **DOF_SCHAR_VEC** provides information about boundary types, and the matrix–vector multiplication will be the identity for all Dirichlet boundary DOFs; this also enforced in the case that Dirichlet boundary values are not assembled into the system matrix **matrix** (compare Sections 3.12.1 and 3.12.2); if Dirichlet boundary conditions are assembled into the system matrix or no Dirichlet boundary conditions are prescribed, **bound** may be a pointer to **nil**;

the return value is a pointer to data used by the routine **mat_vec_s[d]()** as first argument.

mat_vec_s[d](data, dim, x, b): applies the system matrix, with data initialized by **init_mat_vec_s[d]()**, to the input vector **x** and stores the product in **b**; **dim** is the dimension of the linear system, **data** is a pointer to data used for the matrix–vector multiplication and is the return value of **init_mat_vec_s[d]()**;

mat_vec_s[d] can be used as the entry **mat_vec** or **mat_vec.T** in an **OEM_DATA** structure together with the return value of **init_mat_vec_s[d]()** as the corresponding pointer **mat_vec_data** respectively **mat_vec.T_data**.

3.15.4 Access of functions for preconditioning

While a preconditioner can be selected in **oem_solve_s[d]()** just by an integer parameter, a direct access to the functions is needed for a more general application, which directly calls one of the **oem_...()** routines.

A preconditioner may need some initialization phase, which depends on the matrix of the linear system, but is independent of the actual application of the preconditioner to a vector. Thus, a preconditioner is described by three functions for initialization, application, and a final exit routine which may free memory which was allocated during initialization, e.g. All three functions are collected in the structure

```

typedef struct precon PRECON;
struct precon
{
    void    *precon_data;

    int      (*init_precon)(void *);
    void      (*precon)(void *, int, REAL *);
    void      (*exit_precon)(void *);
};

```

Description:

precon_data: data for the preconditioner; always the first argument to the functions **init_precon()**, **precon()**, and **exit_precon()**.

init_precon(precon_data): pointer to a function for initializing the preconditioning method; the return value is **false** if initialization fails, otherwise **true**.

precon(precon_data): pointer to a function for executing the preconditioning method;

precon can be used as the entry **left_precon** or **right_precon** in an **OEM_DATA** structure together with **precon_data** as the corresponding pointer **left_precon_data** respectively **right_precon_data**.

exit_precon(precon_data): frees all data that was used by the preconditioning method.

Currently, a diagonal preconditioner and two hierarchical basis preconditioners (classical Yserentant [74] and Bramble-Pasciak-Xu [22] types) are implemented. Access to the corresponding routines for scalar (**...s()**) and vector valued problems (**...d()**) are given via the following functions, which return a pointer to such a **PRECON** structure.

```

const PRECON *get_diag_precon_s(const DOF_MATRIX *,
                                const DOF_SCHAR_VEC *);
const PRECON *get_diag_precon_d(const DOF_MATRIX *,
                                const DOF_SCHAR_VEC *);
const PRECON *get_HB_precon_s(const FE_SPACE *,
                               const DOF_SCHAR_VEC *, int, int);
const PRECON *get_HB_precon_d(const FE_SPACE *,
                               const DOF_SCHAR_VEC *, int, int);
const PRECON *get_BPX_precon_s(const FE_SPACE *,
                                const DOF_SCHAR_VEC *, int, int);
const PRECON *get_BPX_precon_d(const FE_SPACE *,
                                const DOF_SCHAR_VEC *, int, int);

```

Description:

get_diag_precon_s[d](matrix, bound): returns a pointer to a **PRECON** data structure for diagonal preconditioning of scalar or decoupled vector valued problems;

`matrix` is a pointer to a DOF matrix holding information about the system matrix and `bound` a pointer to a DOF vector holding information about boundary types of DOFs. `bound` may be `nil`, if boundary information is assembled into the system matrix (compare Sections 3.12.1 and 3.12.2)) or no Dirichlet boundary conditions are prescribed.

`get_HB_precon_s[d](matrix, bound, use_get_bound, info)`: returns a pointer to a new PRECON data structure for hierarchical basis preconditioning of scalar or vector valued problems; the preconditioner needs information about Dirichlet DOFs;

`fe_space` is a pointer to the used finite element space; `bound` is a pointer to a DOF vector holding information about boundary types of DOFs; if `bound` is not `nil`, the argument `use_get_bound` is ignored; if `bound` is `nil` and `use_get_bound` is `true`, information about boundary DOFs is generated by `fe_space->bas_fcts->get_bound()`; if `bound` is `nil` and `use_get_bound` is `false` it is assumed that no Dirichlet boundary conditions are prescribed; the last argument `info` is the level of information given during initialization.

`get_BPX_precon_s[d](matrix, bound, use_get_bound, info)`: returns a pointer to a new PRECON data structure for the Bramble-Pasciak-Xu type preconditioning of scalar or vector valued problems; the preconditioner needs information about Dirichlet DOFs;

the arguments to `get_BPX_precon_s[d]()` are the same as to the function `get_HB_precon_s[d]()`, described above.

3.15.5 Multigrid solvers

A abstract framework for multigrid solvers is available. The main data structure for the multigrid solver `MG()` is

```
typedef struct multi_grid_info MULTI_GRID_INFO;
struct multi_grid_info
{
    REAL    tolerance;                /* tol. for resid      */
    REAL    exact_tolerance;          /* tol. for exact_solver */

    int     cycle;                    /* 1=V-cycle, 2=W-cycle */
    int     n_pre_smooth, n_in_smooth; /* no of smoothing loops */
    int     n_post_smooth;            /* no of smoothing loops */
    int     mg_levels;                /* current no. of levels */
    int     exact_level;              /* level for exact_solver */
    int     max_iter;                 /* max. no of MG iter's */
    int     info;

    int     (*init_multi_grid)(MULTI_GRID_INFO *mg_info);
    void     (*pre_smooth)(MULTI_GRID_INFO *mg_info, int level, int n);
    void     (*in_smooth)(MULTI_GRID_INFO *mg_info, int level, int n);
    void     (*post_smooth)(MULTI_GRID_INFO *mg_info, int level, int n);
}
```

```

void    (*mg_restrict)(MULTI_GRID_INFO *mg_info, int level);
void    (*mg_prolongate)(MULTI_GRID_INFO *mg_info, int level);
void    (*exact_solver)(MULTI_GRID_INFO *mg_info, int level);
REAL    (*mg_resid)(MULTI_GRID_INFO *mg_info, int level);
void    (*exit_multi_grid)(MULTI_GRID_INFO *mg_info);

void    *data;                                /* application dep. data */
};

```

The entries yield following information:

tolerance: tolerance for norm of residual;

exact_tolerance: tolerance for “exact solver” on coarsest level;

cycle: selection of multigrid cycle type: 1 =V-cycle, 2 =W-cycle, ...

n_pre_smooth: number of smoothing steps on each level before (first) coarse level correction;

n_in_smooth: number of smoothing steps on each level between coarse level corrections (for $\text{cycle} \geq 2$);

n_post_smooth: number of smoothing steps on each level after (last) coarse level correction;

mg_levels: number of levels;

exact_level: selection of grid level where the “exact” solver is used (and no further coarse grid correction), usually **exact_level=0**;

max_iter: maximal number of multigrid iterations;

info: level of information produced by the multigrid method;

init_multi_grid: pointer to a function for initializing the multigrid method; may be nil;

if not nil, **init_multi_grid(mg_info)** initializes data needed by the multigrid method, returns **true** if an error occurs;

pre_smooth: pointer to a function for performing the smoothing step before coarse grid corrections;

pre_smooth(mg_info, level, n) performs **n** smoothing iterations on grid level;

in_smooth: pointer to a function for performing the smoothing step between coarse grid corrections;

in_smooth(mg_info, level, n) performs **n** smoothing iterations on grid level;

post_smooth: pointer to a function for performing the smoothing step after coarse grid corrections;

post_smooth(mg_info, level, n) performs **n** smoothing iterations on grid level;

mg_restrict: pointer to a function for computing and restricting the residual to a coarser level;
mg_restrict(mg_info, level) computes and restricts the residual from grid level to next coarser grid (level-1);

mg_prolongate: pointer to a function for prolongating and adding coarse grid corrections to the fine grid solution;
mg_prolongate(mg_info, level) prolongates and adds the coarse grid (level-1) correction to the fine grid solution on grid level;

exact_solver: pointer to a function for the “exact” solver;
exact_solver(mg_info, level) computes the “exact” solution of the problem on grid level with tolerance mg_info->exact_tolerance;

mg_resid: pointer to a function for computing the norm of the actual residual;
mg_resid(mg_info, level) returns the norm of residual on grid level;

exit_multi_grid: a pointer to a cleanup routine, may be nil;
 if not nil **exit_multi_grid**(mg_info) is called after termination of the multigrid method for freeing used data;

data: pointer to application dependent data, holding information on or about different grid levels, e.g.

The abstract multigrid solver is implemented in the routine

```
int MG(MULTI_GRID_INFO *)
```

Description:

MG(mg_info): based upon information given in the data structure **mg_info**, the subroutine **MG()** iterates until the prescribed tolerance is met or the prescribed number of multigrid cycles is performed.

Main parts of the **MG()** routine are:

```
{
    int iter;
    REAL resid;

    if (mg_info->init_multi_grid)
        if (mg_info->init_multi_grid(mg_info))
            return(-1);

    resid = mg_info->resid(mg_info, mg_info->mg_levels-1);
    if (resid <= mg_info->tolerance)
        return(0);

    for (iter = 0; iter < mg_info->max_iter; iter++)
    {
        recursive_MG_iteration(mg_info, mg_info->mg_levels-1);
        resid = mg_info->resid(mg_info, mg_info->mg_levels-1);
        if (resid <= mg_info->tolerance)
            break;
    }
}
```

```

    }
    if (mg_info->exit_multi_grid)
        mg_info->exit_multi_grid(mg_info);

    return(iter+1);
}

```

The subroutine `recursive_MG_iteration()` performs smoothing, restriction of the residual and prolongation of the coarse grid correction:

```

void recursive_MG_iteration(MULTI_GRID_INFO *mg_info, int level)
{
    int cycle;

    if (level <= mg_info->exact_level)
    {
        mg_info->exact_solver(mg_info, level);
    }
    else
    {
        if (mg_info->pre_smooth)
            mg_info->pre_smooth(mg_info, level, mg_info->n_pre_smooth);

        for (cycle = 0; cycle < mg_info->cycle; cycle++) {
            if ((cycle > 0) && mg_info->in_smooth)
                mg_info->in_smooth(mg_info, level, mg_info->n_in_smooth);

            mg_info->mg_restrict(mg_info, level);
            recursive_MG_iteration(mg_info, level-1);
            mg_info->prolongate(mg_info, level);
        }

        if (mg_info->post_smooth)
            mg_info->post_smooth(mg_info, level, mg_info->n_post_smooth);
    }
}

```

For multigrid solution of a scalar linear system

$$Au = f$$

given by a `DOF_MATRIX` `A` and a `DOF_REAL_VEC` `f`, the following subroutine is available:

```

int mg_s(DOF_MATRIX *, DOF_REAL_VEC *, const DOF_REAL_VEC *,
        const DOF_SCHAR_VEC *, REAL, int, int, char *);

```

Description:

`mg_s(matrix, u, f, bound, tol, miter, info, prefix)`: solves a linear system for a scalar valued problem by a multigrid method; the return value is the number of performed iterations;

matrix is a pointer to a DOF matrix storing the system matrix, **u** is a pointer to a DOF vector for the solution, holding an initial guess on input; **f** is a pointer to a DOF vector storing the right hand side and **bound** a pointer to a DOF vector with information about boundary DOFs; **bound** must not be **nil** if Dirichlet DOFs are used;

tol is the tolerance for multigrid solver, **miter** the maximal number of multigrid iterations and **info** gives the level of information for the solver;

prefix is a parameter key prefix which is used during the initialization of additional data via `GET_PARAMETER`, see Table 3.20, may be **nil**; an SOR

Table 3.20. Parameters read by `mg_s()` and `mg_s_init()`

member	default	key
<code>mg_info->cycle</code>	1	<code>prefix->cycle</code>
<code>mg_info->n_pre_smooth</code>	1	<code>prefix->n_pre_smooth</code>
<code>mg_info->n_in_smooth</code>	1	<code>prefix->n_in_smooth</code>
<code>mg_info->n_post_smooth</code>	1	<code>prefix->n_post_smooth</code>
<code>mg_info->exact_level</code>	0	<code>prefix->exact_level</code>
<code>mg_info->info</code>	<code>info</code>	<code>prefix->info</code>
<code>mg_s_info->smoother</code>	1	<code>prefix->smoother</code>
<code>mg_s_info->smooth_omega</code>	1.0	<code>prefix->smooth_omega</code>
<code>mg_s_info->exact_solver</code>	1	<code>prefix->exact_solver</code>
<code>mg_s_info->exact_omega</code>	1.0	<code>prefix->exact_omega</code>

`smoother` (`mg_s_info->smoother=1`) and an SSOR smoother (`smoother=2`) are available; an under- or over relaxation parameter can be specified by `mg_s_info->smooth_omega`. These SOR/SSOR smoothers are used for `exact_solver`, too.

For applications, where several systems with the same matrix have to be solved, computing time can be saved by doing all initializations like setup of grid levels and restriction of matrices only once. For such cases, three subroutines are available:

```
MG_S_INFO *mg_s_init(DOF_MATRIX *, const DOF_SCHAR_VEC *, int,
                    char *);
int mg_s_solve(MG_S_INFO *, DOF_REAL_VEC *, const DOF_REAL_VEC *,
              REAL, int);
void mg_s_exit(MG_S_INFO *);
```

Description:

`mg_s_init(matrix, bound, info, prefix)`: initializes a standard multigrid method for solving a scalar valued problem by `mg_s_solve()`; the return value is a pointer to data used by `mg_s_solve()` and is the first argument to this function; the structure `MG_S_INFO` contains matrices and vectors for linear problems on all used grid levels.

matrix is a pointer to a DOF matrix storing the system matrix, **bound** a pointer to a DOF vector with information about boundary DOFs; **bound** must not be **nil** if Dirichlet DOFs are used;

info gives the level of information for **mg_s_solve()**; **prefix** is a parameter key prefix for the initialization of additional data via **GET_PARAMETER**, see Table 3.20, may be **nil**.

mg_s_solve(mg_s_info, u, f, tol, miter): solves the linear system for a scalar valued problem by a multigrid method; the routine has to be initialize by **mg_s_init()** and the return value **mg_s_info** of **mg_s_init()** is the first argument; the return value of **mg_s_solve()** is the number of performed iterations;

u is a pointer to a DOF vector for the solution, holding an initial guess on input; **f** is a pointer to a DOF vector storing the right hand side; **tol** is the tolerance for multigrid solver, **miter** the maximal number of multigrid iterations;

the function may be called several times with different right hand sides **f**.

mg_s_exit(mg_s_info): frees data needed for the multigrid method and which is allocated by **mg_s_init()**.

Remark 3.35. The multigrid solver is currently available only for Lagrange finite elements of first order (**lagrange1**). An implementation for higher order elements is future work.

3.15.6 Nonlinear solvers

For the solution of a nonlinear equation

$$u \in \mathbb{R}^N : \quad F(u) = 0 \quad \text{in } \mathbb{R}^N \quad (3.1)$$

several Newton methods are provided. For testing the convergence a (problem dependent) norm of either the correction d_k in the k th step, i.e.

$$\|d_k\| = \|u_{k+1} - u_k\|,$$

or the residual, i.e.

$$\|F(u_{k+1})\|,$$

is used.

These solvers are *not* part of **ALBERTA** and need for the access of the basic data structure and prototypes of solvers the header file

```
#include <nls.h>
```

The data structure (defined in **nls.h**) for passing information about assembling and solving a linearized equation, tolerances, etc. to the solvers is

```

typedef struct nls_data NLS_DATA;
struct nls_data
{
    void      (*update)(void *, int, const REAL *, int, REAL *);
    void      *update_data;
    int       (*solve)(void *, int, const REAL *, REAL *);
    void      *solve_data;
    REAL      (*norm)(void *, int, const REAL *);
    void      *norm_data;

    WORKSPACE *ws;

    REAL      tolerance;
    int       restart;
    int       max_iter;
    int       info;

    REAL      initial_residual;
    REAL      residual;
};

```

Description:

update: subroutine for computing a linearized system;

update(update_data, dim, uk, update_matrix, F) computes the system matrix of a linearization, if **update_matrix** is not zero, and the right hand side vector **F**, if **F** is not **nil**, for the actual iterate **uk**; **dim** is the dimension of the nonlinear system, and **update_data** a pointer to user data.

update_data: pointer to user data for the update of a linearized equation, first argument to **update()**.

solve: function for solving a linearized system for the new correction; the return value is the number of iterations used by an iterative solver or zero; this number is printed, if information about the solution process should be produced;

solve(solve_data, dim, F, d) solves the linearized equation of dimension **dim** with right hand side **F** for a correction **d** of the actual iterate; **d** is initialized with zeros and **update_data** is a pointer to user data.

solve_data: pointer to user data for solution of the linearized equation, first argument to **solve()**;

the nonlinear solver does not know how the system matrix is stored; such information can be passed from **update()** to **solve()** by using pointers to the same DOF matrix in both **update_data** and **solve_data**, e.g.

norm: function for computing a problem dependent norm $\|\cdot\|$; if **norm** is **nil**, the Euclidian norm is used;

norm(norm_data, dim, x) returns the norm of the vector **x**; **dim** is the dimension of the nonlinear system, and **norm_data** pointer to user data.

norm_data: pointer to user data for the calculation of the problem dependent norm, first argument to **norm()**.

ws: a pointer to a **WORKSPACE** structure for storing additional vectors used by a solver; if the space is not sufficient, the used solver will enlarge this workspace; if **ws** is **nil**, then the used solver allocates memory, which is freed before exit.

tolerance: tolerance for the nonlinear solver; if the norm of the correction/residual is less or equal **tolerance**, the solver returns the actual iterate as the solution of the nonlinear system.

restart: restart for the nonlinear solver.

max_iter: is a maximal number of iterations to be performed, even if the tolerance may not be reached.

info: the level of information produced by the solver; 0 is the lowest level of information (no information is printed) and 4 the highest level.

initial_residual: stores the norm of the initial correction/residual on exit.

residual: stores the norm of the last correction/residual on exit.

The following Newton methods for solving (3.1) are currently implemented:

```
int nls_newton(NLS_DATA *, int, REAL *);
int nls_newton_ds(NLS_DATA *, int, REAL *);
int nls_newton_fs(NLS_DATA *, int, REAL *);
int nls_newton_br(NLS_DATA *, REAL, int, REAL *);
```

Description:

nls_newton(nls_data, dim, u0): solves a nonlinear system by the classical Newton method; the return value is the number of iterations;

nls_data stores information about functions for the assemblage and solution of $DF(u_k)$, $F(u_k)$, calculation of a norm, tolerances, etc. **dim** is the dimension of the nonlinear system, and **u0** the initial guess on input and the solution on output; **nls_newton()** stops if the norm of the **correction** is less or equal **nls_data->tolerance**; it needs a workspace for storing $2 \cdot \text{dim}$ additional REALs.

nls_newton_ds(nls_data, dim, u0): solves a nonlinear system by a Newton method with step size control; the return value is the number of iterations;

nls_data stores information about functions for the assembling and solving of $DF(u_k)$, $F(u_k)$, calculation of a norm, tolerances, etc. **dim** is the dimension of the nonlinear system, and **u0** the initial guess on input and the solution on output; **nls_newton_ds()** stops if the norm of the **correction** is less or equal **nls_data->tolerance**; in each iteration at most **nls_data->restart** steps for controlling the step size τ are performed; the aim is to choose τ such that

$$\|DF(u_k)^{-1}F(u_k + \tau d_k)\| \leq (1 - \frac{1}{2}\tau)\|d_k\|$$

holds, where $\|\cdot\|$ is the problem dependent norm, if `nls_data->norm` is not `nil`, otherwise the Euclidian norm; each step needs the update of F , the solution of one linearized problem (the system matrix for the linearized system does not change during step size control) and the calculation of a norm; `nls_newton_ds()` needs a workspace for storing $4*\text{dim}$ additional REALs.

`nls_newton_fs(nls_data, dim, u0)`: solves a nonlinear system by a Newton method with step size control; the return value is the number of iterations; `nls_data` stores information about functions for the assembling and solving of $DF(u_k)$, $F(u_k)$, calculation of a norm, tolerances, etc. `dim` is the dimension of the nonlinear system, and `u0` the initial guess on input and the solution on output; `nls_newton_fs()` stops if the norm of the **residual** is less or equal `nls_data->tolerance`; in each iteration at most `nls_data->restart` steps for controlling the step size τ are performed; the aim is to choose τ such that

$$\|F(u_k + \tau d_k)\| \leq (1 - \frac{1}{2}\tau)\|F(u_k)\|$$

holds, where $\|\cdot\|$ is the problem dependent norm, if `nls_data->norm` is not `nil`, otherwise the Euclidian norm; the step size control is not expensive, since in each step only an update of F and the calculation of $\|F\|$ are involved; `nls_newton_fs()` needs a workspace for storing $3*\text{dim}$ additional REALs.

`nls_newton_br(nls_data, delta, dim, u0)`: solves a nonlinear system by a global Newton method by Bank and Rose [6]; the return value is the number of iterations;

`nls_data` stores information about functions for the assembling and solving of $DF(u_k)$, $F(u_k)$, calculation of a norm, tolerances, etc. `delta` is a parameter with $\delta \in (0, 1 - \alpha_0)$, where $\alpha_0 = \|DF(u_0)u_0 + F(u_0)\|/\|F(u_0)\|$; `dim` is the dimension of the nonlinear system, and `u0` the initial guess on input and the solution on output; `nls_newton_br()` stops if the norm of the **residual** is less or equal `nls_data->tolerance`; in each iteration at most `nls_data->restart` steps for controlling the step size by the method of Bank and Rose are performed; the step size control is not expensive, since in each step only an update of F and the calculation of $\|F\|$ are involved;

`nls_newton_br()` needs a workspace for storing $3*\text{dim}$ additional REALs.

3.16 Graphics output

ALBERTA provides one and two dimensional interactive graphic subroutines built on the X-Windows and GL/OpenGL interfaces, and one, two and three dimensional interactive graphics via the gltools [38]. Additionally, an interface for post-processing data with the GRAPE visualization environment [64] is supplied.

3.16.1 One and two dimensional graphics subroutines

A set of subroutines for opening, closing of graphic output windows, and several display routines are provided, like drawing the underlying mesh, displaying scalar finite element functions as a graph in 1d, and using iso-lines or iso-colors in 2d. For vector valued functions \mathbf{v} similar routines are available, which display the modulus $|\mathbf{v}|$.

The routines use the following type definitions for window identification, standard color specification in [red, green, blue] coordinates, with $0 \leq \text{red, green, blue} \leq 1$, and standard colors

```
typedef void * GRAPH_WINDOW;

typedef float  GRAPH_RGBCOLOR[3];

extern const GRAPH_RGBCOLOR rgb_black;
extern const GRAPH_RGBCOLOR rgb_white;
extern const GRAPH_RGBCOLOR rgb_red;
extern const GRAPH_RGBCOLOR rgb_green;
extern const GRAPH_RGBCOLOR rgb_blue;
extern const GRAPH_RGBCOLOR rgb_yellow;
extern const GRAPH_RGBCOLOR rgb_magenta;
extern const GRAPH_RGBCOLOR rgb_cyan;
extern const GRAPH_RGBCOLOR rgb_grey50;

extern const GRAPH_RGBCOLOR rgb_albert;
extern const GRAPH_RGBCOLOR rgb_alberta;
```

The last two colors correspond to the two different colors in the ALBERTA logo.

The following graphic routines are available for one and two dimensions:

```
GRAPH_WINDOW graph_open_window(const char *, const char *, REAL *,
                               MESH *);
void graph_close_window(GRAPH_WINDOW);
void graph_clear_window(GRAPH_WINDOW, const GRAPH_RGBCOLOR);
void graph_mesh(GRAPH_WINDOW, MESH *, const GRAPH_RGBCOLOR, FLAGS);
void graph_drv(GRAPH_WINDOW, const DOF_REAL_VEC *, REAL, REAL, int);
void graph_drv_d(GRAPH_WINDOW, const DOF_REAL_D_VEC *, REAL, REAL,
                 int);
void graph_el_est(GRAPH_WINDOW, MESH *, REAL (*)(EL *), REAL, REAL);
void graph_fvalues(GRAPH_WINDOW, MESH *, REAL (*)(const EL_INFO *,
            const REAL *), FLAGS, REAL, REAL, int);
void graph_line(GRAPH_WINDOW, , const REAL [2], const REAL [2],
               const GRAPH_RGBCOLOR, REAL);
void graph_point(GRAPH_WINDOW, const REAL [2], const GRAPH_RGBCOLOR,
                REAL);
```


Description:

graph_open_window(title, geometry, world, mesh): returns a pointer to a `GRAPH_WINDOW` which is opened for display; if the window could not be opened, the return value is `nil`; in 1d the y -direction of the graphic window is used for displaying the graphs of functions;

title is an optional string holding a window title, if **title** is `nil`, a default title is used; **geometry** is an optional string holding the window geometry in X11 format “WxH” or “WxH+X+Y”, if `nil`, a default geometry is used;

world is an optional pointer to an array of *world coordinates* (`xmin`, `xmax`, `ymin`, `ymax`) to specify which part of a triangulation is displayed in this window, if **world** is `nil` and **mesh** is not `nil`, **mesh->diam** is used to select a range of world coordinates; in 1d, the range of the y -direction is set to $[-1, 1]$; if both **world** and **mesh** are `nil`, the unit square $[0, 1] \times [0, 1]$ is displayed in 1d and 2d.

graph_close_window(win): closes the graphic window **win** which has been previously opened by the function **graph_open_window()**.

graph_clear_window(win, c): clears the graphic window **win** and sets the background color **c**; if **c** is `nil`, white is used as background color.

graph_mesh(win, mesh, c, flag): displays the underlying **mesh** in the graphic window **win**; **c** is an optional color used for drawing lines, if **c** is `nil` black as a default color is used; the last argument **flag** allows for a selection of an additional display; **flag** may be 0 or the bitwise OR of some of the following flags:

GRAPH_MESH_BOUNDARY: only boundary edges are drawn, otherwise all edges of the triangulation are drawn; **c** is the display color for all edges if not `nil`; otherwise the display color for Dirichlet boundary vertices/edges is blue and for Neumann vertices/edges the color is red;

GRAPH_MESH_ELEMENT_MARK: triangles marked for refinement are filled red, and triangles marked for coarsening are filled blue, unmarked triangles are filled white;

GRAPH_MESH_VERTEX_DOF: the *first* DOF at each vertex is written near the vertex; currently only working in 2d when the library is not using OpenGL.

GRAPH_MESH_ELEMENT_INDEX: element indices are written inside the element, only available for `EL_INDEX == 1`; currently only working in 2d when the library is not using OpenGL.

graph_drv(win, u, min, max, n_refine): displays the finite element function stored in the `DOF_REAL_VEC` **u** in the graphic window **win**; in 1d, the graph of **u** is plotted in black, in 2d an iso-color display of **u** is used; **min** and **max** specify a range of **u** which is displayed; if $\text{min} \geq \text{max}$, **min** and **max** of **u** are computed by **graph_drv()**; in 2d, coloring is adjusted to the values of **min** and **max**; the display routine always uses the linear interpolant on a simplex;

if `n_refine > 0`, each simplex is recursively bisected into $2^{\text{DIM} \times \text{n_refine}}$ sub-simplices, and the linear interpolant on these sub-simplices is displayed; for `n_refine < 0` the default value `u->admin->bas_fcts->degree-1` is used.

`graph_drv_d(win, v, min, max, n_refine)`: displays the modulus of a vector valued finite element function stored in the `DOF_REAL_D_VEC` `v` in the graphic window `win`; the other arguments are the same as for `graph_drv()`.

`graph_el_est(win, mesh, get_el_est)`: displays piecewise constant values over the triangulation `mesh`, like local error indicators, in the graphics window `win`; `get_el_est` is a pointer to a function which returns the constant value on each element; by this function the piecewise constant function is defined.

`graph_fvalues(win, mesh, f, flag, min, max, n_refine)`: displays a real-valued function `f` in the graphic window `win`; `f` is a pointer to a function for evaluating values on single elements; `f(el_info, lambda)` returns the value of the function on `el_info->el` at the barycentric coordinates `lambda`; in 1d, the graph of `f` is plotted in black, in 2d an iso-color display of `f` is used; `min` and `max` specify a range of `f` which is displayed; if `min ≥ max`, `min` and `max` of `f` are computed by `graph_fvalues()`; in 2d, coloring is adjusted to the values of `min` and `max`; the display routine always uses the linear interpolant of `f` on a simplex; if `n_refine > 0`, each simplex is recursively bisected into $2^{\text{DIM} \times \text{n_refine}}$ sub-simplices, and the linear interpolant on these sub-simplices is displayed.

`graph_line(win, p0, p1, c, lw)`: draws the line segment with start point `p0` and end point `p1` in (x, y) coordinates in the graphic window `win`; `c` is an optional argument and may specify the line color to be used; if `c` is `nil` black is used; `lw` specifies the linewidth (currently only for OpenGL graphics); if `lw ≤ 0` the default linewidth 1.0 is set.

`graph_point(win, p, c, diam)`: draws a point at the position `p` in (x, y) coordinates in the graphic window `win`; `c` is an optional argument and may specify the color to be used; if `c` is `nil` black is used; `diam` specifies the drawing diameter (currently only for OpenGL graphics); if `diam ≤ 0` the default diameter 1.0 is set.

Graphic routines for one dimension

The following routines are slight generalizations of the display routines described above for one dimension.

```
void graph1d_drv(GRAPH_WINDOW, const DOF_REAL_VEC *, REAL, REAL, int,
                const GRAPH_RGBCOLOR);
void graph1d_drv_d(GRAPH_WINDOW, const DOF_REAL_D_VEC *, REAL, REAL,
                  int, const GRAPH_RGBCOLOR);
void graph1d_fvalues(GRAPH_WINDOW, MESH *,
                    REAL(*) (const EL_INFO *, const REAL *), FLAGS,
                    REAL, REAL, int, const GRAPH_RGBCOLOR);
```

```
void graph1d_el_est(GRAPH_WINDOW, MESH *, REAL (*)(EL *), REAL, REAL,
                  const GRAPH_RGBCOLOR);
```

Description:

graph1d_drv(win, uh, min, max, n_refine, c): displays the graph of the finite element function stored in the `DOF_REAL_VEC` u in the graphic window win; the first 5 arguments are the same as for **graph_drv**() described above; but by the last optional argument c a line color for the graph can be specified; if c is nil, black is used.

graph1d_drv_d(win, uh, min, max, n_refine, c): displays the graph of the modulus of the vector valued finite element function stored in the `DOF_REAL_D_VEC` v in the graphic window win; the first 5 arguments are the same as for **graph_drv**() described above; but by the last optional argument c a line color for the graph can be specified; if c is nil, black is used.

graph1d_fvalues(win, mesh, f, flag, min, max, n_refine, c): draws a scalar valued function f in the graphic window win; the first 7 arguments are the same as for **graph_fvalues**() described above; but by the last optional argument c a line color for the graph can be specified; if c is nil, black is used.

graph1d_el_est(win, mesh, get_el_est, min, max, c): displays a piecewise constant function over the triangulation mesh in the graphics window win; the first 5 arguments are the same as for **graph_el_est**() described above; but by the last optional argument c a line color for the graph can be specified; if c is nil, black is used.

Graphic routines for two dimensions

The following routines are specialized routines for two dimensional graphic output:

```
void graph_level(GRAPH_WINDOW, const DOF_REAL_VEC *, REAL,
                const GRAPH_RGBCOLOR, int);
void graph_levels(GRAPH_WINDOW, const DOF_REAL_VEC *, int,
                 const REAL *, const GRAPH_RGBCOLOR *, int);
void graph_level_d(GRAPH_WINDOW, const DOF_REAL_D_VEC *, REAL,
                  const GRAPH_RGBCOLOR, int);
void graph_levels_d(GRAPH_WINDOW, const DOF_REAL_D_VEC *, int,
                   const REAL *, const GRAPH_RGBCOLOR *, int);
void graph_values(GRAPH_WINDOW, const DOF_REAL_VEC *, REAL, REAL,
                 int);
void graph_values_d(GRAPH_WINDOW, const DOF_REAL_D_VEC *, REAL, REAL,
                  int);
```

Description:

graph_level(win, v, level, c, n_refine): draws a single selected iso-line at value level of the scalar finite element function stored in the

DOF_REAL_VEC *u* in the graphic window *win*; by the argument *c* a line color for the isoline can be specified; if *c* is *nil*, black is used as line color; the display routine always uses the linear interpolant of *u* on a simplex; if *n_refine* > 0, each triangle is recursively bisected into 2^{2*n_refine} sub-triangles, and the selected isoline of the linear interpolant on these sub-triangles is displayed; for *n_refine* < 0 the default value *u->admin->bas_fcts->degree-1* is used.

graph_levels(win, u, n, levels, c, n_refine): draws *n* selected isolines at values *level*[0], ..., *level*[*n*-1] of the scalar finite element function stored in the DOF_REAL_VEC *u* in the graphic window *win*; if *level* is *nil*, *n* equally distant isolines between the minimum and maximum of *u* are selected; *c* is an optional vector of *n* color values for the *n* isolines, if *nil*, then default color values are used; the argument *n_refine* again chooses a level of refinement, where iso-lines of the piecewise linear interpolant is displayed; for *n_refine* < 0 the default value *u->admin->bas_fcts->degree-1* is used.

graph_level_d(win, v, level, c, n_refine): draws a single selected isoline at values *level* of the modulus of a vector valued finite element function stored in the DOF_REAL_D_VEC *v* in the graphic window *win*; the arguments are the same as for **graph_level()**.

graph_levels_d(win, v, n, levels, c, n_refine): draws *n* selected isolines at values *level*[0], ..., *level*[*n*-1] of the modulus of a vector valued finite element function stored in the DOF_REAL_D_VEC *v* in the graphic window *win*; the arguments are the same as for **graph_levels()**.

graph_values(win, u, n_refine): shows an iso-color display of the finite element function stored in the DOF_REAL_VEC *u* in the graphic window *win*; it is equivalent to **graph_drv(win, u, min, max, n_refine)**.

graph_values_d(win, v, n_refine): shows an iso-color display of the modulus of a vector valued finite element function stored in the DOF_REAL_D_VEC *v* in the graphic window *win*; it is equivalent to **graph_drv_d(win, u, min, max, n_refine)**.

3.16.2 gltools interface

The following interface for using the interactive gltools graphics of WIAS Berlin [38] is implemented. The gltools are freely available under the terms of the MIT license, see

<http://www.wias-berlin.de/software/gltools/>

The ALBERTA interface is compatible with version **gltools-2-4**. It can be used for 1d, 2d, and 3d triangulation, but only when DIM equals DIM_OF_WORLD. For window identification we use the data type

```
typedef void *    GLTOOLS_WINDOW;
```

The interface provides the following functions:

```
GLTOOLS_WINDOW open_gltools_window(const char *, const char *,
                                   const REAL *, MESH *, int);
void close_gltools_window(GLTOOLS_WINDOW);

void gltools_mesh(GLTOOLS_WINDOW, MESH *, int);
void gltools_drv(GLTOOLS_WINDOW, const DOF_REAL_VEC *, REAL, REAL);
void gltools_drv_d(GLTOOLS_WINDOW, const DOF_REAL_D_VEC *, REAL,
                  REAL);
void gltools_vec(GLTOOLS_WINDOW, const DOF_REAL_D_VEC *, REAL, REAL);
void gltools_est(GLTOOLS_WINDOW, MESH *, REAL (*)(EL *), REAL, REAL);

void gltools_disp_mesh(GLTOOLS_WINDOW, MESH *, int,
                      const DOF_REAL_VEC *);
void gltools_disp_drv(GLTOOLS_WINDOW, const DOF_REAL_VEC *, REAL,
                    REAL, const DOF_REAL_VEC *);
void gltools_disp_drv_d(GLTOOLS_WINDOW, const DOF_REAL_D_VEC *, REAL,
                      REAL, const DOF_REAL_VEC *);
void gltools_disp_vec(GLTOOLS_WINDOW, const DOF_REAL_D_VEC *, REAL,
                    REAL, const DOF_REAL_VEC *);
void gltools_disp_est(GLTOOLS_WINDOW, MESH *, REAL (*)(EL *), REAL,
                    REAL, const DOF_REAL_VEC *);
```

Description:

`open_gltools_window(title, geometry, world, mesh, dialog)`: returns a `GLTOOLS_WINDOW` which is opened for display; if the window could not be opened, the return value is `nil`; `title` is an optional string holding a title for the window; if `title` is `nil`, a default is used; `geometry` is an optional string holding the window geometry in X11 format (“WxH” or “WxH+X+Y”), if `nil`, a default geometry is used; the optional argument `world` is a pointer to an array of *world coordinates* (xmin, xmax, ymin, ymax) for 2d and (xmin, xmax, ymin, ymax, zmin, zmax) for 3d, it can be used to specify which part of the mesh will be displayed in the window; if `world` is `nil`, either `mesh` or the default domain $[0, 1]^{\text{DIM}}$ is used; `mesh` is an optional pointer to a mesh to select a range of world coordinates which will be displayed in the window; if both `world` and `mesh` are `nil`, the default domain $[0, 1]^{\text{DIM}}$ is used; display is not done or is done in an interactive mode depending on whether `dialog` equals 0 or not; in interactive mode type ‘h’ to get a list of all key bindings;

`close_gltools_window(win)`: closes the window `win` which has been previously opened by `open_gltools_window()`;

`gltools_mesh(win, mesh, mark)`: displays the elements of `mesh` in the graphic window `win`; if `mark` is not zero the piecewise constant function `sign(el->mark)` is shown;

`gltools_drv(win, u, min, max)`: shows the `DOF_REAL_VEC` `u` in the graphic window `win`; for higher order elements it is possible to display the vector on a refined grid; the key ‘P’ toggles between refined and not refined mode `min`

and `max` define the range of the discrete function for display; if `min` \geq `max` this range is adjusted automatically;

`gltools_drv_d(win, ud, min, max)`: displays the modulus of the given `DOF_REAL_D_VEC` `ud` in the graphic window `win`; for higher order elements it is possible to display the vector on a refined grid; the key 'P' toggles between refined and not refined mode `min` and `max` define the range of the modulus of discrete function for display; if `min` \geq `max` this range is adjusted automatically;

`gltools_vec(win, ud, min, max)`: displays a given vector field stored in the `DOF_REAL_D_VEC` `ud` in the graphic window `win`; for higher order elements it is possible to display the vector on a refined grid; the key 'P' toggles between refined and not refined mode `min` and `max` define the range of the modulus of discrete function for display; if `min` \geq `max` this range is adjusted automatically;

`gltools_est(win, mesh, get_est, min, max)`: displays the error estimate on `mesh` as a piecewise constant function in the graphic window `win`; the local indicators are accessed by `get_est()` on each leaf element; `min` and `max` define the range for display; if `min` \geq `max` this range is adjusted automatically;

`gltools_est()` can also be used to display any piecewise constant function on the mesh, where local values are accessed by `get_el_est()`;

`gltools_disp_mesh(win, mesh, mark, disp)`: displays a `mesh` with an additional distortion of the geometry by a displacement vector field `disp`; this can be used in solid mechanics applications, e.g.; similar to the function `gltools_mesh()`;

`gltools_disp_drv(win, u, min, max, disp)`: displays a real valued finite element function on a distorted geometry given by the vector field `disp`; similar to the function `gltools_drv()`;

`gltools_disp_drv_d(win, ud, min, max, disp)`: displays the modulus of a vector valued finite element function on a distorted geometry given by the vector field `disp`; similar to the function `gltools_drv_d()`;

`gltools_disp_vec(win, ud, min, max, disp)`: displays a vector field on a distorted geometry given by the vector field `disp`; similar to the function `gltools_vec()`;

`gltools_disp_est(win, mesh, get_el_est, min, max, disp)`: depicts an error estimate as a piecewise constant function on a distorted geometry given by the vector field `disp`; similar to the function `gltools_est()`.

3.16.3 GRAPE interface

Visualization using the GRAPE library [64] is only possible as a post-processing step. Data of the actual geometry and finite element functions is written to file by `write_mesh[_xdr]()` and `write_dof_real[_d]_vec[_xdr]()` and then read by some programs, using the GRAPE mesh interface for the visualization. We recommend to use the `xdr` routines for portability of the stored binary data. The use of the GRAPE *h*-mesh and *hp*-mesh interfaces is work in progress and the description of these programs will be done in the near future. References to visualization methods used in GRAPE applying to ALBERTA can be found in [40, 56, 57].

For obtaining the GRAPE library, please see

<http://www.iam.uni-bonn.de/sfb256/grape/>

The distribution of ALBERTA contains source files with the implementation of GRAPE mesh interface to ALBERTA in the `GRAPE` subdirectory. Having access to the GRAPE library (Version 5.4.2), this interface can be compiled and linked with the ALBERTA and GRAPE library into two executables `alberta_grape` and `alberta_movi` respectively for 2d and 3d with `DIM = DIM_OF_WORLD`. The path of the GRAPE header file and library has to be specified during the installation of ALBERTA, compare Section 2.4.

The interface can be compiled in the sub-directory `GRAPE/mesh/2d` for the 2d interface, resulting in the executable `alberta_grape22` and `alberta_movi22`, and in `GRAPE/mesh/3d` for the 3d interface, resulting in the executable `alberta_grape33` and `alberta_movi33`.

The program `alberta_grape??` is mainly designed for displaying finite element data on a single grid, i. e. one or several scalar/vector-valued finite element functions on the corresponding mesh. `alberta_grape??` expects mesh data stored by `write_mesh[_xdr]()` and `write_dof_real[_xdr]()` or `write_dof_real_d[_xdr]()` defined on the same mesh.

```
alberta_grape22 -m mesh.xdr -s scalar.xdr -v vector.xdr
```

will display the 2d mesh stored in the file `mesh.xdr` together with the scalar finite element function stored in `scalar.xdr` and the vector valued finite element function stored in `vector.xdr`.

`alberta_grape?? --help` explains the full functionality of this program and displays a list of options.

The program `alberta_movi??` is designed for displaying finite element data on a sequence of grids with one or several scalar/vector-valued finite element functions. This is the standard visualization tool for post-processing data from time-dependent simulations. `alberta_movi??` expects a sequence of mesh data stored by `write_mesh[_xdr]()` and finite element data of this mesh stored by `write_dof_real[_xdr]()` or `write_dof_real_d[_xdr]()`, where the filenames for the sequence of meshes and finite element functions are generated by the function `generate_filename()`, explained in Section 3.1.6. Sec-

tion 2.3.10 shows how to write such a sequence of data in a time-dependent problem.

Similar to `alberta_grape??` the command `alberta_movi?? --help` explains the full functionality of this program and displays a list of options.

References

1. M. AINSWORTH AND J. T. ODEN, *A unified approach to a posteriori error estimation using element residual methods*, Numer. Math., 65 (1993), pp. 23–50.
2. M. AINSWORTH AND J. T. ODEN, *A posteriori error estimation in finite element analysis*, Wiley, 2000.
3. I. BABUŠKA AND W. RHEINBOLDT, *Error estimates for adaptive finite element computations*, SIAM J. Numer. Anal., 15 (1978), pp. 736–754.
4. A. BAMBERGER, E. BÄNSCH, AND K. G. SIEBERT, *Experimental and numerical investigation of edge tones*. Preprint WIAS Berlin no. 681, 2001.
5. R. E. BANK, *PLTMG: a software package for solving elliptic partial differential equations user's guide 8.0*. Software - Environments - Tools. 5. Philadelphia, PA: SIAM. xii, 1998.
6. R. E. BANK AND D. J. ROSE, *Global approximate Newton methods.*, Numer. Math., 37 (1981), pp. 279–295.
7. R. E. BANK AND A. WEISER, *Some a posteriori error estimators for elliptic partial differential equations.*, Math. Comput., 44 (1985), pp. 283–301.
8. E. BÄNSCH, *Local mesh refinement in 2 and 3 dimensions*, IMPACT Comput. Sci. Engrg., 3 (1991), pp. 181–191.
9. E. BÄNSCH, *Adaptive finite element techniques for the Navier–Stokes equations and other transient problems*, in Adaptive Finite and Boundary Elements, C. A. Brebbia and M. H. Aliabadi, eds., Computational Mechanics Publications and Elsevier, 1993, pp. 47–76.
10. E. BÄNSCH AND K. G. SIEBERT, *A posteriori error estimation for nonlinear problems by duality techniques*. Preprint 30, Universität Freiburg, 1995.
11. P. BASTIAN, K. BIRKEN, K. JOHANNSEN, S. LANG, V. REICHENBERGER, C. WIENERS, G. WITTUM, AND C. WROBEL, *Parallel solution of partial differential equations with adaptive multigrid methods on unstructured grids*, in High performance computing in science and engineering '99, E. Krause and et al., eds., Berlin: Springer, 2000, pp. 496–508. Transactions of the High Performance Computing Center Stuttgart (HLRS). 2nd workshop, Stuttgart, Germany, October 4–6, 1999.
12. R. BECK, B. ERDMANN, AND R. ROITZSCH, *An object-oriented adaptive finite element code: Design issues and applications in hyperthermia treatment planning*, in Modern software tools for scientific computing, E. Arge and et al., eds.,

- Boston: Birkhaeuser, 1997, pp. 105–124. International workshop, Oslo, Norway, September 16–18, 1996.
13. R. BECKER AND R. RANNACHER, *A feed-back approach to error control in finite element methods: Basic analysis and examples.*, East-West J. Numer. Math., 4 (1996), pp. 237–264.
 14. J. BEY, *Tetrahedral grid refinement*, Computing, 55 (1995), pp. 355–378.
 15. J. BEY, *Simplicial grid refinement: On Freudenthal’s algorithm and the optimal number of congruence classes*, Numer. Math., 85 (2000), pp. 1–29.
 16. P. BINEV, W. DAHMEN, AND R. DEVORE, *Adaptive finite element methods with convergence rates*. IGPM Report No. 219, RWTH Aachen, 2002.
 17. M. BÖHM, A. SCHMIDT, AND M. WOLFF, *Adaptive finite element simulation of a model for transformation induced plasticity in steel*. Report ZeTeM Bremen, 2003.
 18. F. A. BORNEMAN, *An adaptive multilevel approach to parabolic equations I*, IMPACT Comput. Sci. Engrg., 2 (1990), pp. 279–317.
 19. F. A. BORNEMAN, *An adaptive multilevel approach to parabolic equations II*, IMPACT Comput. Sci. Engrg., 3 (1990), pp. 93–122.
 20. F. A. BORNEMAN, *An adaptive multilevel approach to parabolic equations III*, IMPACT Comput. Sci. Engrg., 4 (1992), pp. 1–45.
 21. F. A. BORNEMANN, B. ERDMANN, AND R. KORNUBER, *A posteriori error estimates for elliptic problems in two and three space dimensions.*, SIAM J. Numer. Anal., 33 (1996), pp. 1188–1204.
 22. J. H. BRAMBLE, J. E. PASCIAK, AND J. XU, *Parallel multilevel preconditioners*, Math. Comput., 55 (1990), pp. 1–22.
 23. S. C. BRENNER AND L. SCOTT, *The mathematical theory of finite element methods. 2nd ed.*, Springer, 2002.
 24. Z. CHEN AND R. H. NOCHETTO, *Residual type a posteriori error estimates for elliptic obstacle problems.*, Numer. Math., 84 (2000), pp. 527–548.
 25. P. G. CIARLET, *The finite element methods for elliptic problems. Repr., unabridged republ. of the orig. 1978.*, SIAM, 2002.
 26. R. COOLS AND P. RABINOWITZ, *Monomial cubature rules since “Stroud”: a compilation*, J. Comput. Appl. Math., 48 (1993), pp. 309–326.
 27. L. DEMKOWICZ, J. T. ODEN, W. RACHOWICZ, AND O. HARDY, *Toward a universal h-p adaptive finite element strategy, Part 1 – Part 3*, Comp. Methods Appl. Mech. Engrg., 77 (1989), pp. 79–212.
 28. J. DONGARRA, J. DUCROZ, S. HAMMARLING, AND R. HANSON, *An extended set of Fortran Basic Linear Algebra Subprograms*, ACM Trans. Math. Softw., 14 (1988), pp. 1–32.
 29. W. DÖRFLER, *FORTTRAN-Bibliothek der Orthogonalen Fehler-Methoden*, Manual, Mathematische Fakultät Freiburg, 1995.
 30. W. DÖRFLER, *A robust adaptive strategy for the nonlinear poisson equation*, Computing, 55 (1995), pp. 289–304.
 31. W. DÖRFLER, *A convergent adaptive algorithm for Poisson’s equation*, SIAM J. Numer. Anal., 33 (1996), pp. 1106–1124.
 32. W. DÖRFLER, *A time- and spaceadaptive algorithm for the linear time-dependent Schrödinger equation*, Numer. Math., 73 (1996), pp. 419–448.
 33. W. DÖRFLER AND K. G. SIEBERT, *An adaptive finite element method for minimal surfaces*, in Geometric Analysis and Nonlinear Partial Differential Equations, S. Hildebrandt and H. Karcher, eds., Springer, 2003, pp. 146–175.

34. D. DUNAVANT, *High degree efficient symmetrical Gaussian quadrature rules for the triangle*, Int. J. Numer. Methods Eng., 21 (1985), pp. 1129–114.
35. G. DZIUK, *An algorithm for evolutionary surfaces*, Numer. Math., 58 (1991), pp. 603 – 611.
36. K. ERIKSSON AND C. JOHNSON, *Adaptive finite element methods for parabolic problems I: A linear model problem*, SIAM J. Numer. Anal., 28 (1991), pp. 43–77.
37. J. FRÖHLICH, J. LANG, AND R. ROITZSCH, *Selfadaptive finite element computations with smooth time controller and anisotropic refinement*, in Numerical Methods in Engineering, J. Desideri, P. Tallec, E. Onate, J. Periaux, and E. Stein, eds., John Wiley & Sons, New York, 1996, pp. 523–527.
38. J. FUHRMANN AND H. LANGMACH, *gltools: OpenGL based online visualization*. Software: <http://www.wias-berlin.de/software/gltools/>.
39. K. GATERMANN, *The construction of symmetric cubature formulas for the square and the triangle*, Computing, 40 (1988), pp. 229–240.
40. B. HAASDONK, M. OHLBERGER, M. RUMPF, A. SCHMIDT, AND K. G. SIEBERT, *Multiresolution visualization of higher order adaptive finite element simulations*, Computing, 70 (2003), pp. 181–204.
41. H. JARAUSCH, *On an adaptive grid refining technique for finite element approximations*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 1105–1120.
42. H. KARDESTUNCER, ed., *Finite Element Handbook*, McGraw-Hill, New York, 1987.
43. R. KORNUBER AND R. ROITZSCH, *On adaptive grid refinement in the presence of internal or boundary layers*, IMPACT Comput. Sci. Engrg., 2 (1990), pp. 40–72.
44. I. KOSSACZKÝ, *A recursive approach to local mesh refinement in two and three dimensions*, J. Comput. Appl. Math., 55 (1994), pp. 275–288.
45. C. LAWSON, R. HANSON, D. KINCAID, AND F. KROUGH, *Basic Linear Algebra Subprograms for Fortran usage*, ACM Trans. Math. Softw., 5 (1979), pp. 308–325.
46. J. M. MAUBACH, *Local bisection refinement for n-simplicial grids generated by reflection*, SIAM J. Sci. Comput., 16 (1995), pp. 210–227.
47. A. MEISTER, *Numerik linearer Gleichungssysteme*, Vieweg, 1999.
48. W. F. MITCHELL, *A comparison of adaptive refinement techniques for elliptic problems*, ACM Trans. Math. Softw., 15 (1989), pp. 326–347.
49. W. F. MITCHELL, *MGGHAT: Elliptic PDE software with adaptive refinement, multigrid and high order finite elements*, in Sixth Copper Mountain Conference on Multigrid Methods, N. D. Melson, T. A. Manteuffel, and S. F. McCormick, eds., NASA, 1993, pp. 439–448.
50. P. MORIN, R. H. NOCHETTO, AND K. G. SIEBERT, *Data oscillation and convergence of adaptive FEM*, SIAM J. Numer. Anal., 38 (2000), pp. 466–488.
51. P. MORIN, R. H. NOCHETTO, AND K. G. SIEBERT, *Convergence of adaptive finite element methods*, SIAM Review, 44 (2002), pp. 631–658.
52. P. MORIN, R. H. NOCHETTO, AND K. G. SIEBERT, *Local problems on stars: A posteriori error estimators, convergence, and performance*, Math. Comp., 72 (2003), pp. 1067–1097.
53. R. H. NOCHETTO, M. PAOLINI, AND C. VERDI, *An adaptive finite element method for two-phase Stefan problems in two space dimensions. Part II: Implementation and numerical experiments*, SIAM J. Sci. Stat. Comput., 12 (1991), pp. 1207–1244.

54. R. H. NOCHETTO, A. SCHMIDT, AND C. VERDI, *A posteriori error estimation and adaptivity for degenerate parabolic problems*, Math. Comput., 69 (2000), pp. 1–24.
55. R. H. NOCHETTO, K. G. SIEBERT, AND A. VEESER, *Pointwise a posteriori error control for elliptic obstacle problems*, Numer. Math., 95 (2003), pp. 163–195.
56. M. RUMPF, A. SCHMIDT, AND K. G. SIEBERT, *On a unified visualization approach for data from advanced numerical methods*, in Visualization in Scientific Computing '95, R. Scateni, J. V. Wijk, and P. Zanarini, eds., Springer, 1995, pp. 35–44.
57. M. RUMPF, A. SCHMIDT, AND K. G. SIEBERT, *Functions defining arbitrary meshes – a flexible interface between numerical data and visualization routines*, Computer Graphics Forum, 15 (1996), pp. 129–141.
58. Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS, 1996.
59. A. SCHMIDT AND K. G. SIEBERT, *Concepts of the finite element toolbox ALBERT*. Preprint 17/98 Freiburg, 1998. To appear in Notes on Numerical Fluid Mechanics.
60. A. SCHMIDT AND K. G. SIEBERT, *Abstract data structures for a finite element package: Design principles of ALBERT*, Z. Angew. Math. Mech., 79 (1999), pp. S49–S52.
61. A. SCHMIDT AND K. G. SIEBERT, *A posteriori estimators for the h - p version of the finite element method in 1d*, Applied Numerical Mathematics, 35 (2000), pp. 43–46.
62. A. SCHMIDT AND K. G. SIEBERT, *ALBERT – Software for scientific computations and applications*, Acta Math. Univ. Comenianae, 70 (2001), pp. 105–122.
63. J. SCHOEBERL, *NETGEN: An advancing front 2D/3D-mesh generator based on abstract rules*, Comput. Vis. Sci., 1 (1997), pp. 41–52.
64. SFB 256, *GRAPE – GRAphics Programming Environment Manual, Version 5.0*, Bonn, 1995.
65. J. R. SHEWCHUK, *Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator*, in Applied Computational Geometry: Towards Geometric Engineering, M. C. Lin and D. Manocha, eds., vol. 1148 of Lecture Notes in Computer Science, Springer-Verlag, May 1996, pp. 203–222. From the First ACM Workshop on Applied Computational Geometry.
66. K. G. SIEBERT, *A posteriori error estimator for anisotropic refinement*, Numer. Math., 73 (1996), pp. 373–398.
67. R. STEVENSON, *An optimal adaptive finite element method*. Preprint No. 1271, Department of Mathematics, University of Utrecht, 2003.
68. A. H. STROUD, *Approximate calculation of multiple integrals*, Prentice-Hall, Englewood Cliffs, NJ, 1971.
69. A. VEESER, *Efficient and reliable a posteriori error estimators for elliptic obstacle problems*, SIAM J. Numer. Anal., 39 (2001), pp. 146–167.
70. A. VEESER, *Convergent adaptive finite elements for the nonlinear Laplacian*, Numer. Math., 92 (2002), pp. 743–770.
71. R. VERFÜRTH, *A posteriori error estimates for nonlinear problems: Finite element discretization of elliptic equations*, Math. Comp., 62 (1994), pp. 445–475.
72. R. VERFÜRTH, *A posteriori error estimation and adaptive mesh-refinement techniques*, J. Comp. Appl. Math., 50 (1994), pp. 67–83.
73. R. VERFÜRTH, *A Review of A Posteriori Error Estimation and Adaptive Mesh-Refinement Techniques*, Wiley-Teubner, 1996.

- 74. H. YSERENTANT, *On the multi-level splitting of finite element spaces*, Numer. Math., 49 (1986), pp. 379–412.
- 75. O. C. ZIENKIEWICZ, D. W. KELLY, J. GAGO, AND I. BABUŠKA, *Hierarchical finite element approaches, error estimates and adaptive refinement*, in The mathematics of finite elements and applications IV, J. Whiteman, ed., Academic Press, 1982, pp. 313–346.

Index

ABS(), 113
ADAPT_INSTAT, 256
adapt_mesh(), 254
adapt_method_instat(), 258, 259
adapt_method_stat(), 253
ADAPT_STAT, 250
adaptive methods, 42–53, 250–262
 ADAPT_INSTAT, 256
 adapt_mesh(), 254
 adapt_method_instat(), 258, 259
 adapt_method_stat(), 253
 ADAPT_STAT, 250
 adaptive strategy, 43
 ALBERTA marking strategies, 256
 coarsening strategies, 47
 equidistribution strategy, 44
 estimate(), 251
 get_adapt_instat(), 260
 get_adapt_stat(), 260
 get_el_est(), 251
 get_el_estc(), 251
 guaranteed error reduction strategy, 45
 marking strategies, 43
 maximum strategy, 44
 one_timestep(), 260
 stationary problems, 42
 strategies for time dependent problems, 49
 time and space adaptive strategy, 52
 time dependent problems, 49
 time step size control, 50
add_element_d_vec(), 225
add_element_matrix(), 225
add_element_vec(), 225
ADD_PARAMETER(), 124
add_parameter(), 124
ALBERTA marking strategies, 256
alberta_alloc(), 118
alberta_calloc(), 118
alberta_free(), 118
alberta_grape, 293
alberta_matrix(), 119
alberta_movi, 293
alberta_realloc(), 118
assemblage of discrete system, 38–42
 load vector, 38
 system matrix, 39
assemblage tools, 224–250
 add_element_d_vec(), 225
 add_element_matrix(), 225
 add_element_vec(), 225
 dirichlet_bound(), 247
 dirichlet_bound_d(), 247
 EL_MATRIX_INFO, 227
 EL_VEC_D_INFO, 244
 EL_VEC_INFO, 244
 fill_matrix_info(), 234
 get_q00_psi_phi(), 243
 get_q01_psi_phi(), 241
 get_q10_psi_phi(), 242
 get_q11_psi_phi(), 239
 interpol(), 249
 interpol_d(), 249
 L2scp_fct_bas(), 246
 L2scp_fct_bas_d(), 246

- OPERATOR_INFO, 231
- Q00_PSI_PHI, 243
- Q01_PSI_PHI, 240
- Q10_PSI_PHI, 241
- Q11_PSI_PHI, 238
- update_matrix(), 229
- update_real_d_vec(), 245
- update_real_vec(), 245
- ball_project(), 149
- barycentric coordinates, 26–28
 - coord_to_world(), 209
 - el_grd_lambda(), 209
 - world_to_coord(), 209
- BAS_FCT, 184
- BAS_FCTS, 185
- bisection
 - newest vertex, 12
 - procedure of Kossaczky, 12
- BOUNDARY, 130, 131
- CALL_EL_LEVEL, 153
- CALL_EVERY_EL_INORDER, 153
- CALL_EVERY_EL_POSTORDER, 153
- CALL_EVERY_EL_PREORDER, 153
- CALL_LEAF_EL, 153
- CALL_LEAF_EL_LEVEL, 153
- CALL_MG_LEVEL, 153
- change_error_out(), 117
- change_msg_out(), 117
- check_and_get_mesh(), 143
- clear_dof_matrix(), 169
- CLEAR_WORKSPACE(), 121
- clear_workspace(), 121
- close_gltools_window(), 290
- coarse_restrict(), 167, 183
- coarsen(), 182
- coarsening
 - algorithm, 18
 - atomic coarsening operation, 18
 - coarsening algorithm, 19
 - interpolation of DOF vectors, 21, 34, 183
 - restriction of DOF vectors, 21, 33, 183
- coarsening strategies, 47
- conforming triangulation, 11
- coord_to_world(), 209
- COPY_DOW(), 128
- curved boundary, 131
- D2_BAS_FCT, 184
- D2_uh_at_qp(), 219
- D2_uh_d_at_qp(), 219
- data types
 - ADAPT_INSTAT, 256
 - ADAPT_STAT, 250
 - BAS_FCT, 184
 - BAS_FCTS, 185
 - BOUNDARY, 131
 - D2_BAS_FCT, 184
 - DOF, 161
 - DOF_ADMIN, 162
 - DOF_FREE_UNIT, 162
 - DOF_INT_VEC, 164
 - DOF_MATRIX, 169
 - DOF_REAL_D_VEC, 164
 - DOF_REAL_VEC, 165
 - DOF_SCHAR_VEC, 164
 - DOF_UCHAR_VEC, 164
 - EL, 134
 - EL_INFO, 135
 - EL_MATRIX_INFO, 227
 - EL_VEC_D_INFO, 244
 - EL_VEC_INFO, 244
 - FE_SPACE, 206
 - FLAGS, 114
 - GLTOOLS_WINDOW, 290
 - GRAPH_RGBCOLOR, 286
 - GRAPH_WINDOW, 286
 - GRD_BAS_FCT, 184
 - LEAF_DATA_INFO, 139
 - MACRO_DATA, 151
 - MACRO_EL, 132
 - MATRIX_ROW, 168
 - MatrixTranspose, 173
 - MESH, 141
 - MULTI_GRID_INFO, 277
 - NLS_DATA, 283
 - OEM_DATA, 269
 - OEM_SOLVER, 273
 - OPERATOR_INFO, 231
 - PARAMETRIC, 142
 - PRECON, 276
 - Q00_PSI_PHI, 243
 - Q01_PSI_PHI, 240
 - Q10_PSI_PHI, 241
 - Q11_PSI_PHI, 238

- QUAD, 210
- QUAD_FAST, 213
- QUADRATURE, 210
- RC_LIST_EL, 140
- REAL, 114
- REAL_D, 128
- REAL_DD, 128
- S_CHAR, 113
- TRAVERSE_STACK, 158
- U_CHAR, 113
- VOID_LIST_ELEMENT, 121
- WORKSPACE, 120
- DIM, 128
- DIM_OF_WORLD, 128
- DIRICHLET, 130
- Dirichlet boundary, 131
- dirichlet_bound(), 247
- dirichlet_bound_d(), 247
- DIST_DOW(), 128
- div_uh_d_at_qp(), 219
- DOF, 161
- DOF_ADMIN, 162
- dof_asum(), 173
- dof_axpy(), 173
- dof_axpy_d(), 173
- dof_compress(), 164
- dof_copy(), 173
- dof_copy_d(), 173
- dof_dot(), 173
- dof_dot_d(), 173
- dof_gemv(), 173
- dof_gemv_d(), 173
- DOF_INT_VEC, 164
- DOF_MATRIX, 169
 - ENTRY_NOT_USED(), 168
 - ENTRY_USED(), 168
 - MATRIX_ROW, 168
 - NO_MORE_ENTRIES, 168
 - ROW_LENGTH, 168
 - UNUSED_ENTRY, 168
- dof_max(), 173
- dof_max_d(), 173
- dof_min(), 173
- dof_min_d(), 173
- dof_mv(), 173
- dof_mv_d(), 173
- dof_nrm2(), 173
- dof_nrm2_d(), 173
- DOF_REAL_D_VEC, 164
- DOF_REAL_VEC, 165
- dof_scal(), 173
- dof_scal_d(), 173
- DOF_SCHAR_VEC, 164
- dof_set(), 173
- dof_set_d(), 173
- DOF_UCHAR_VEC, 164
- dof_xpay(), 173
- dof_xpay_d(), 173
- degree of freedom (DOFs), 24
- DOFs, 24–25, 161–173
 - adding and removing of DOFs, 179–183
 - boundary type, 130
 - entries in the `el` structure, 171
 - entries in the `mesh` structure, 171
 - FOR_ALL_DOFS, 170
 - FOR_ALL_FREE_DOFS, 170
 - DOFs
 - get_dof_indices(), 173
 - init_dof_admins(), 207
 - initialization of DOFs on a mesh, 143
 - relation global and local DOFs, 29
- EL, 134
- el_det(), 209
- el_grd_lambda(), 209
- EL_INDEX, 130
- EL_INFO, 135
- EL_MATRIX_INFO, 227
- EL_TYPE(), 137
- EL_VEC_D_INFO, 244
- EL_VEC_INFO, 244
- el_volume(), 209
- element indices, 130
- ellipt_est(), 264
- enlarge_dof_lists(), 164
- equidistribution strategy, 44
- error estimators, 263–268
 - ellipt_est(), 264
 - heat_est(), 267
- ERROR(), 116
- ERROR_EXIT(), 116
- estimate(), 251
- eval_D2_uh(), 217
- eval_D2_uh_d(), 217
- eval_D2_uh_d_fast(), 218
- eval_D2_uh_fast(), 218
- eval_div_uh_d(), 217

- eval_div_uh_d_fast(), 218
- eval_grd_uh(), 217
- eval_grd_uh_d(), 217
- eval_grd_uh_d_fast(), 218
- eval_grd_uh_fast(), 218
- eval_uh(), 217
- eval_uh_d(), 217
- eval_uh_d_fast(), 218
- eval_uh_fast(), 218
- evaluation of derivatives, 30
- evaluation of finite element functions, 29
 - D2_uh_at_qp(), 219
 - D2_uh_d_at_qp(), 219
 - div_uh_d_at_qp(), 219
 - eval_D2_uh(), 217
 - eval_D2_uh_d(), 217
 - eval_D2_uh_d_fast(), 218
 - eval_D2_uh_fast(), 218
 - eval_div_uh_d(), 217
 - eval_div_uh_d_fast(), 218
 - eval_grd_uh(), 217
 - eval_grd_uh_d(), 217
 - eval_grd_uh_d_fast(), 218
 - eval_grd_uh_fast(), 218
 - eval_uh(), 217
 - eval_uh_d(), 217
 - eval_uh_d_fast(), 218
 - eval_uh_fast(), 218
 - grd_uh_at_qp(), 219
 - grd_uh_d_at_qp(), 219
 - uh_at_qp(), 219
 - uh_d_at_qp(), 219
- f_at_qp(), 211
- f_d_at_qp(), 211
- false, 113
- FE_SPACE, 206
- file organization, 111
- FILL_ANY, 154
- FILL_BOUND, 154
- FILL_COORDS, 154
- FILL_EL_TYPE, 154
- fill_elinfo(), 154
- fill_macro_info(), 154
- fill_matrix_info(), 234
- FILL_NEIGH, 154
- FILL_NOTHING, 154
- FILL_OPP_COORDS, 154
- FILL_ORIENTATION, 154
- find_el_at_pt(), 160
- finite element discretization, 35–42
- finite element spaces, 29
- FLAGS, 114
- FOR_ALL_DOFS, 170
- FOR_ALL_FREE_DOFS, 170
- free_alberta_matrix(), 119
- free_dof_dof_vec(), 166
- free_dof_int_vec(), 166
- free_dof_matrix(), 169
- free_dof_real_d_vec(), 166
- free_dof_real_vec(), 166
- free_dof_schar_vec(), 166
- free_dof_uchar_vec(), 166
- free_int_dof_vec(), 166
- free_macro_data(), 152
- free_mesh(), 144
- free_traverse_stack(), 158
- free_void_list_element(), 121
- FREE_WORKSPACE(), 121
- free_workspace(), 121
- FUNCNAME(), 114
- generate_filename(), 127
- get_adapt_instat(), 260
- get_adapt_stat(), 260
- get_bas_fcts, 190
- GET_BOUND(), 131
- get_bound()
 - entry in BAS_FCTS structure, 187
- get_bound()
 - for linear elements, 193
- get_BPX_precon_d(), 276
- get_BPX_precon_s(), 276
- get_diag_precon_d(), 276
- get_diag_precon_s(), 276
- get_dof_dof_vec(), 166
- get_dof_indices()
 - entry in BAS_FCTS structure, 187
- get_dof_indices()
 - for linear elements, 192
 - for quadratic elements, 197
- get_dof_int_vec(), 166
- get_dof_matrix(), 169
- get_dof_real_d_vec(), 166
- get_dof_real_vec(), 166
- get_dof_schar_vec(), 166
- get_dof_uchar_vec(), 166

- GET_DOF_VEC(), 166
- get_el_est(), 251
- get_el_estc(), 251
- get_face_normal(), 215
- get_fe_space(), 207
- get_HB_precon_d(), 276
- get_HB_precon_s(), 276
- get_int_dof_vec(), 166
- get_int_vec()
 - for linear elements, 189
- get_lagrange(), 206
- get_lumping_quadrature(), 211
- get_macro_data(), 152
- GET_MESH(), 143
- get_mesh(), 143
- GET_PARAMETER(), 125
- get_parameter(), 125
- get_q00_psi_phi(), 243
- get_q01_psi_phi(), 241
- get_q10_psi_phi(), 242
- get_q11_psi_phi(), 239
- get_quad_fast(), 214
- get_quadrature(), 211
- get_traverse_stack(), 158
- get_void_list_element(), 121
- GET_WORKSPACE(), 120
- get_workspace(), 120
- global_coarsen(), 182
- global_refine(), 176
- gltools graphics, 290–292
 - close_gltools_window(), 290
 - gltools_disp_drv(), 290
 - gltools_disp_drv_d(), 290
 - gltools_disp_est(), 290
 - gltools_disp_mesh(), 290
 - gltools_disp_vec(), 290
 - gltools_drv(), 290
 - gltools_drv_d(), 290
 - gltools_est(), 290
 - gltools_mesh(), 290
 - gltools_vec(), 290
 - GLTOOLS_WINDOW, 290
 - open_gltools_window(), 290
- GRAPE
 - generate_filename(), 127
- GRAPE interface, 293–294
 - alberta_grape, 293
 - alberta_movi, 293
- graph1d_drv(), 288
- graph1d_drv_d(), 288
- graph1d_el_est(), 288
- graph1d_fvalues(), 288
- graph_clear_window(), 286
- graph_close_window(), 286
- graph_drv(), 286
- graph_drv_d(), 286
- graph_el_est(), 286
- graph_fvalues(), 286
- graph_level(), 289
- graph_level_d(), 289
- graph_levels(), 289
- graph_levels_d(), 289
- graph_line(), 286
- graph_mesh(), 286
- graph_open_window(), 286
- graph_point(), 286
- GRAPH_RGBCOLOR, 286
- graph_values(), 289
- graph_values_d(), 289
- GRAPH_WINDOW, 286
- graphics routines, 285–294
 - close_gltools_window(), 290
 - gltools_disp_drv(), 290
 - gltools_disp_drv_d(), 290
 - gltools_disp_est(), 290
 - gltools_disp_mesh(), 290
 - gltools_disp_vec(), 290
 - gltools_drv(), 290
 - gltools_drv_d(), 290
 - gltools_est(), 290
 - gltools_mesh(), 290
 - gltools_vec(), 290
 - GLTOOLS_WINDOW, 290
 - graph_drv(), 288
 - graph_drv_d(), 288
 - graph1d_el_est(), 288
 - graph1d_fvalues(), 288
 - graph_clear_window(), 286
 - graph_close_window(), 286
 - graph_drv(), 286
 - graph_drv_d(), 286
 - graph_el_est(), 286
 - graph_fvalues(), 286
 - graph_level(), 289
 - graph_level_d(), 289
 - graph_levels(), 289
 - graph_levels_d(), 289
 - graph_line(), 286

- graph_mesh(), 286
- graph_open_window(), 286
- graph_point(), 286
- GRAPH_RGBCOLOR, 286
- graph_values(), 289
- graph_values_d(), 289
- GRAPH_WINDOW, 286
- NO_WINDOW, 286
- open_gltools_window(), 290
- rgb_albert, 286
- rgb_alberta, 286
- rgb_black, 286
- rgb_blue, 286
- rgb_cyan, 286
- rgb_green, 286
- rgb_grey50, 286
- rgb_magenta, 286
- rgb_red, 286
- rgb_white, 286
- rgb_yellow, 286
- grd_f_at_qp(), 211
- GRD_BAS_FCT, 184
- grd_f_d_at_qp(), 211
- grd_uh_at_qp(), 219
- grd_uh_d_at_qp(), 219
- guaranteed error reduction strategy, 45

- H1_err(), 222
- H1_err_d(), 222
- H1_NORM, 265
- H1_norm_uh(), 221
- H1_norm_uh_d(), 221
- Heat equation
 - implementation, 93
- heat_est(), 267
- hierarchical mesh, 21

- implementation of model problems,
 - 55–111
 - Heat equation, 93
 - nonlinear reaction–diffusion equation, 68
 - Poisson equation, 56
- include files
 - alberta.h, 113
 - alberta_util.h, 113
 - nls.h, 282
 - oem.h, 269
- INDEX(), 137

- INFO(), 116
- INIT_D2_PHI, 213
- init_dof_admins(), 143, 207
- init_element(), 232
- INIT_GRD_PHI, 213
- init_leaf_data(), 143
- init_mat_vec_d(), 275
- init_mat_vec_s(), 275
- init_parameters(), 123
- INIT_PHI, 213
- initialization of meshes, 143
- installation, 111
- integrate_std_simp(), 211
- INTERIOR, 130
- interior node, 131
- interpol(_d)
 - entry in BAS_FCTS structure, 187
- interpol(), 249
- interpol() for linear elements, 193
- interpol_d(), 249
- interpolation, 167
- interpolation and restriction of DOF
 - vectors, 32–35
- interpolation of DOF vectors, 21, 33, 34
- IS_DIRICHLET(), 131
- IS_INTERIOR(), 131
- IS_LEAF_EL(), 139
- IS_NEUMANN(), 131

- L2_err(), 222
- L2_err_d(), 222
- L2_NORM, 265
- L2_norm_uh(), 221
- L2_norm_uh_d(), 221
- L2scp_fct_bas(), 246
- L2scp_fct_bas_d(), 246
- LALt(), 232
- leaf data, 23
 - transformation during coarsening, 183
 - transformation during refinement, 178
- LEAF_DATA(), 139
- LEAF_DATA_INFO, 139
- linear solver
 - NLS_DATA, 283
 - OEM_DATA, 269
 - OEM_SOLVER, 273
- linear solvers, 269–282

- oem_bicgstab(), 271
- oem_cg(), 271
- oem_gmres(), 271
- oem_gmres_k(), 271
- oem_kdir(), 271
- oem_ores(), 271
- oem_solve_d(), 273
- oem_solve_s(), 273
- sor_d(), 274
- sor_s(), 274
- ssor_d(), 274
- ssor_s(), 274
- local numbering
 - edges, 132
 - faces, 132
 - neighbours, 132
 - vertices, 13
- macro triangulation, 11
 - example of a macro triangulation in 2d/3d, 147–149
 - export macro triangulations, 151
 - import macro triangulations, 151
 - macro triangulation file, 144
 - macro_data2mesh(), 152
 - macro_test(), 153
 - read_macro(), 146
 - read_macro_bin(), 150
 - read_macro_xdr(), 150
 - reading macro triangulations, 144
 - write_macro(), 150
 - write_macro_bin(), 150
 - write_macro_data(), 152
 - write_macro_data_bin(), 152
 - write_macro_data_xdr(), 152
 - write_macro_xdr(), 150
 - writing macro triangulations, 150
- MACRO_DATA, 151
- macro_data2mesh(), 152
- MACRO_EL, 132
- macro_test(), 153
- marking strategies, 43
- MAT_ALLOC(), 119
- MAT_FREE(), 119
- mat_vec_d(), 275
- mat_vec_s(), 275
- MatrixTranspose, 173, 275
- MAX(), 113
- max_err_at_qp(), 222
- max_err_at_qp_d(), 222
- max_quad_points(), 215
- maximum strategy, 44
- MEM_ALLOC(), 118
- MEM_CALLOC(), 118
- MEM_FREE(), 118
- MEM_REALLOC(), 118
- memory (de-) allocation, 118–121
 - alberta_alloc(), 118
 - alberta_calloc(), 118
 - alberta_free(), 118
 - alberta_matrix(), 119
 - alberta_realloc(), 118
 - CLEAR_WORKSPACE(), 121
 - clear_workspace(), 121
 - free_alberta_matrix(), 119
 - free_void_list_element(), 121
 - FREE_WORKSPACE(), 121
 - free_workspace(), 121
 - get_void_list_element(), 121
 - GET_WORKSPACE(), 120
 - get_workspace(), 120
 - MAT_ALLOC(), 119
 - MAT_FREE(), 119
 - MEM_ALLOC(), 118
 - MEM_CALLOC(), 118
 - MEM_FREE(), 118
 - MEM_REALLOC(), 118
 - print_mem_use(), 121
 - REALLOC_WORKSPACE(), 120
 - realloc_workspace(), 120
- MESH, 141
- mesh coarsening, 18–19, 182–183
- mesh refinement, 12–17, 176–182
- mesh refinement and coarsening, 9–21
- mesh traversal, 153–161
- MESH_COARSENEED, 182
- MESH_COARSENEED, 252, 254
- MESH_REFINED, 176
- MESH_REFINED, 252, 254
- mesh_traverse(), 155
- messages, 114
- MG(), 279
- mg_s(), 280
- mg_s_exit(), 281
- mg_s_init(), 281
- mg_s_solve(), 281
- MIN(), 113
- MSG(), 114

- msg_info, 116
- MULTI_GRID_INFO, 277
- N_EDGES, 129
- N_FACE, 129
- N_NEIGH, 129
- N_VERTICES, 129
- NEIGH(), 137
- NEIGH_IN_EL, 129
- neighbour information, 129
- NEUMANN, 130
- Neumann boundary, 131
- new_bas_fcts, 190
- nil, 113
- NLS_DATA, 283
- nls_newton(), 284
- nls_newton_br(), 284
- nls_newton_ds(), 284
- nls_newton_fs(), 284
- NO_WINDOW, 286
- nonlinear reaction–diffusion equation
 - implementation, 68
- nonlinear solvers, 282–285
 - nls_newton(), 284
 - nls_newton_br(), 284
 - nls_newton_ds(), 284
 - nls_newton_fs(), 284
- NORM_DOW(), 128
- NoTranspose, 173, 275
- numerical quadrature, 37, 38
 - D2_uh_at_qp(), 219
 - D2_uh_d_at_qp(), 219
 - div_uh_d_at_qp(), 219
 - get_lumping_quadrature(), 211
 - get_quad_fast(), 214
 - get_quadrature(), 211
 - grd_uh_at_qp(), 219
 - grd_uh_d_at_qp(), 219
 - INIT_D2_PHI, 213
 - INIT_GRD_PHI, 213
 - INIT_PHI, 213
 - integrate_std_simp(), 211
 - max_quad_points(), 215
 - QUAD, 210
 - QUAD_FAST, 213
 - QUADRATURE, 210
 - uh_at_qp(), 219
 - uh_d_at_qp(), 219
- oem_bicgstab(), 271
- oem_cg(), 271
- OEM_DATA, 269
- oem_gmres(), 271
- oem_gmres_k(), 271
- oem_oder(), 271
- oem_ores(), 271
- oem_solve_d(), 273
- oem_solve_s(), 273
- OEM_SOLVER, 273
- one_timestep(), 260
- open_error_file(), 117
- open_gltools_window(), 290
- open_msg_file(), 117
- OPERATOR_INFO, 231
- OPP_VERTEX(), 137
- param_bound(), 131, 177
 - ball_project(), 149
- parameter file, 122
- parameter handling, 122–127
 - ADD_PARAMETER(), 124
 - add_parameter(), 124
 - GET_PARAMETER(), 125
 - get_parameter(), 125
 - init_parameters(), 123
 - save_parameters(), 124
- PARAMETRIC, 142
- parametric simplex, 10
- Poisson equation
 - implementation, 56
- PRECON, 276
- preconditioner
 - BPX, 277
 - diagonal, 274, 277
 - hierarchical basis, 274, 277
- preserve_coarse_dofs, 182
- print_dof_int_vec(), 167
- print_dof_matrix(), 169
- print_dof_real_d_vec(), 167
- print_dof_real_vec(), 167
- print_dof_schar_vec(), 167
- print_dof_uchar_vec(), 167
- PRINT_INFO(), 116
- PRINT_INT_VEC(), 115
- print_mem_use(), 121
- print_msg(), 115
- PRINT_REAL_VEC(), 115

- Q00_PSI_PHI, 243
- Q01_PSI_PHI, 240
- Q10_PSI_PHI, 241
- Q11_PSI_PHI, 238
- QUAD, 210
- QUAD_FAST, 213
- QUADRATURE, 210

- RC_LIST_EL, 140, 178
- read_dof_int_vec(), 175
- read_dof_int_vec_xdr(), 176
- read_dof_real_d_vec(), 175
- read_dof_real_d_vec_xdr(), 176
- read_dof_real_vec(), 175
- read_dof_real_vec_xdr(), 176
- read_dof_schar_vec(), 175
- read_dof_schar_vec_xdr(), 176
- read_dof_uchar_vec(), 175
- read_dof_uchar_vec_xdr(), 176
- read_macro(), 146
- read_macro_bin(), 150
- read_macro_xdr(), 150
- read_mesh(), 174
- read_mesh_xdr(), 176
- REAL, 114
- real_coarse_restr()
 - for linear elements, 194
- REAL_D, 128
- REAL_DD, 128
- real_refine_inter()
 - for quadratic elements, 199
- real_refine_inter()
 - for linear elements, 194
 - for quadratic elements, 197
- REALLOC_WORKSPACE(), 120
- realloc_workspace(), 120
- reference element, 136
- refine(), 176
- refine_interpol(), 167, 181
- refinement
 - algorithm, 12
 - atomic refinement operation, 14
 - bisection, 12
 - DOFs
 - handed from parent to children, 179
 - newly created, 180
 - removed on the parent, 182
 - edge, 12
 - interpolation of DOF vectors, 21, 33, 181
 - local numbering
 - edges, 132
 - faces, 132
 - neighbours, 132
 - vertices, 13
 - recursive refinement algorithm, 17
- restriction, 167
- restriction of DOF vectors, 21, 33
- rgb_albert, 286
- rgb_alberta, 286
- rgb_black, 286
- rgb_blue, 286
- rgb_cyan, 286
- rgb_green, 286
- rgb_grey50, 286
- rgb_magenta, 286
- rgb_red, 286
- rgb_white, 286
- rgb_yellow, 286

- save_parameters(), 124
- SCAL_DOW(), 128
- SCP_DOW(), 128
- SET_DOW(), 128
- simplex, 10
- sor_d(), 274
- sor_s(), 274
- sort_face_indices(), 215
- SQR(), 113
- ssor_d(), 274
- ssor_s(), 274
- standard simplex, 10
- strategies for time dependent problems, 49

- TEST(), 116
- TEST_EXIT(), 116
- time and space adaptive strategy, 52
- time step size control, 50
- Transpose, 173, 275
- traverse_first(), 158
- traverse_neighbour(), 159
- traverse_next(), 158
- TRAVERSE_STACK, 158
- triangulation, 11
- true, 113
- S_CHAR, 113

U_CHAR, 113
uh_at_qp(), 219
uh_d_at_qp(), 219
update_matrix(), 229
update_real_d_vec(), 245
update_real_vec(), 245

vector_product(), 129
VOID_LIST_ELEMENT, 121

WAIT, 127
WAIT, 117
WAIT REALLY, 117
WARNING(), 117
WORKSPACE, 120
world_to_coord(), 209
write_dof_int_vec(), 175
write_dof_int_vec_xdr(), 176

write_dof_real_d_vec(), 175
write_dof_real_d_vec_xdr(), 176
write_dof_real_vec(), 175
write_dof_real_vec_xdr(), 176
write_dof_schar_vec(), 175
write_dof_schar_vec_xdr(), 176
write_dof_uchar_vec(), 175
write_dof_uchar_vec_xdr(), 176
write_macro(), 150
write_macro_bin(), 150
write_macro_data(), 152
write_macro_data_bin(), 152
write_macro_data_xdr(), 152
write_macro_xdr(), 150
write_mesh(), 174
write_mesh_xdr(), 176

XDR, 174

Data types, symbolic constants, functions, and macros

Data types

ADAPT_INSTAT, 256
ADAPT_STAT, 250
BAS_FCTS, 185
BAS_FCT, 184
BOUNDARY, 131
D2_BAS_FCT, 184
DOF_ADMIN, 162
DOF_FREE_UNIT, 162
DOF_INT_VEC, 164
DOF_MATRIX, 169
DOF_REAL_D_VEC, 164
DOF_REAL_VEC, 164
DOF_SCHAR_VEC, 164
DOF_UCHAR_VEC, 164
DOF, 161
EL_INFO, 135
EL_MATRIX_INFO, 227
EL_VEC_D_INFO, 243
EL_VEC_INFO, 243
EL, 134
FE_SPACE, 206
FLAGS, 114
GLTOOLS_WINDOW, 290
GRAPH_RGBCOLOR, 286
GRAPH_WINDOW, 286
GRD_BAS_FCT, 184
LEAF_DATA_INFO, 139
MACRO_EL, 132
MATRIX_ROW, 168
MatrixTranspose, 173
MESH, 141
MULTI_GRID_INFO, 277

NLS_DATA, 282
OEM_DATA, 269
OEM_SOLVER, 272
OPERATOR_INFO, 230
PARAMETRIC, 142
PRECON, 275
Q00_PSI_PHI, 242
Q01_PSI_PHI, 239
Q10_PSI_PHI, 241
Q11_PSI_PHI, 237
QUAD_FAST, 213
QUAD, 210
RC_LIST_EL, 140
REAL, 114
REAL_D, 128
REAL_DD, 128
S_CHAR, 113
TRAVERSE_STACK, 157
U_CHAR, 113
WORKSPACE, 120

Symbolic constants

CALL_EL_LEVEL, 153
CALL EVERY_EL_INORDER, 153
CALL EVERY_EL_POSTORDER, 153
CALL EVERY_EL_PREORDER, 153
CALL_LEAF_EL_LEVEL, 153
CALL_LEAF_EL, 153
CALL_MG_LEVEL, 153
CENTER, 163
DIM_OF_WORLD, 128
DIM, 128
DIRICHLET, 130

EDGE, 163
 EL_INDEX, 130
 FACE, 163
 false, 113
 FILL_BOUND, 154
 FILL_COORDS, 154
 FILL_NEIGH, 154
 FILL_NOTHING, 154
 FILL_OPP_COORDS, 154
 FILL_ORIENTATION, 154
 GRAPH_MESH_BOUNDARY, 287
 GRAPH_MESH_ELEMENT_INDEX, 287
 GRAPH_MESH_ELEMENT_MARK, 287
 GRAPH_MESH_VERTEX_DOF, 287
 H1_NORM, 264
 INIT_D2_PHI, 213
 INIT_GRD_PHI, 213
 INIT_GRD_UH, 265
 INIT_UH, 265
 INTERIOR, 130
 L2_NORM, 264
 NEIGH_IN_EL, 129
 NEUMANN, 130
 nil, 113
 NO_MORE_ENTRIES, 168
 NO_WINDOW, 286
 N_EDGES, 129
 N_FACES, 129
 N_NEIGH, 129
 N_VERTICES, 129
 ROW_LENGTH, 168
 true, 113
 UNUSED_ENTRY, 168
 VERTEX, 163
 change_error_out(), 117
 change_msg_out(), 117
 check_and_get_mesh(), 143
 clear_dof_matrix(), 169
 clear_workspace(), 121
 close_gltools_window(), 291
 coarsen(), 182
 coord_to_world(), 209
 D2_uh_at_qp(), 219
 D2_uh_d_at_qp(), 219
 dirichlet_bound(), 247
 dirichlet_bound_d(), 247
 div_uh_d_at_qp(), 219
 dof_asum(), 173
 dof_axpy(), 173
 dof_axpy_d(), 173
 dof_compress(), 163
 dof_copy(), 173
 dof_copy_d(), 173
 dof_dot(), 173
 dof_dot_d(), 173
 dof_gemv(), 173
 dof_gemv_d(), 173
 dof_max(), 173
 dof_max_d(), 173
 dof_min(), 173
 dof_min_d(), 173
 dof_mv(), 173
 dof_mv_d(), 173
 dof_nrm2(), 173
 dof_nrm2_d(), 173
 dof_scal(), 173
 dof_scal_d(), 173
 dof_set(), 173
 dof_set_d(), 173
 dof_xpay(), 173
 dof_xpay_d(), 173
 el_det(), 209
 el_grd_lambda(), 209
 el_volume(), 209
 ellipt_est(), 264
 enlarge_dof_lists(), 164
 estimate(), 251
 eval_D2_uh(), 216
 eval_D2_uh_d(), 216
 eval_D2_uh_d_fast(), 218
 eval_D2_uh_fast(), 218
 eval_div_uh_d_fast(), 218
 eval_div_uh_d(), 216

Functions

adapt_mesh(), 254
 adapt_method_instat(), 258
 adapt_method_stat(), 253
 add_element_dvec(), 225
 add_element_matrix(), 225
 add_element_vec(), 225
 add_parameter(), 124
 alberta_alloc(), 118
 alberta_calloc(), 118
 alberta_free(), 118
 alberta_matrix(), 119
 alberta_realloc(), 118

```

eval_grd_uh(), 216
eval_grd_uh_d(), 216
eval_grd_uh_d_fast(), 218
eval_grd_uh_fast(), 218
eval_uh(), 216
eval_uh_d(), 216
eval_uh_d_fast(), 218
eval_uh_fast(), 218
f_at_qp(), 211
f_d_at_qp(), 211
fill_elinfo(), 154
fill_macro_info(), 154
fill_matrix_info(), 233
find_el_at_pt(), 160
free_alberta_matrix(), 119
free_dof_dof_vec(), 166
free_dof_int_vec(), 166
free_dof_matrix(), 169
free_dof_real_d_vec(), 166
free_dof_real_vec(), 166
free_dof_schar_vec(), 166
free_dof_uchar_vec(), 166
free_int_dof_vec(), 166
free_mesh(), 144
free_traverse_stack(), 158
free_workspace(), 121
get_BPX_precon_d(), 276
get_BPX_precon_s(), 276
get_HB_precon_d(), 276
get_HB_precon_s(), 276
get_adapt_instat(), 260
get_adapt_stat(), 260
get_bas_fcts(), 190
get_diag_precon_d(), 276
get_diag_precon_s(), 276
get_dof_dof_vec(), 166
get_dof_int_vec(), 166
get_dof_matrix(), 169
get_dof_real_d_vec(), 166
get_dof_real_vec(), 166
get_dof_schar_vec(), 166
get_dof_uchar_vec(), 166
get_el_estc(), 251
get_el_est(), 251
get_face_normal(), 215
get_fe_space(), 207
get_int_dof_vec(), 166
get_lagrange(), 206
get_lumping_quadrature(), 211
get_mesh(), 143
get_parameter(), 125
get_q00_psi_phi(), 243
get_q01_psi_phi(), 240
get_q10_psi_phi(), 242
get_q11_psi_phi(), 238
get_quad_fast(), 214
get_quadrature(), 211
get_traverse_stack(), 158
global_coarsen(), 182
global_refine(), 176
gltools_disp_drv(), 291
gltools_disp_drv_d(), 291
gltools_disp_est(), 291
gltools_disp_mesh(), 291
gltools_disp_vec(), 291
gltools_drv(), 291
gltools_drv_d(), 291
gltools_est(), 291
gltools_mesh(), 291
gltools_vec(), 291
graph1d_drv(), 288
graph1d_drv_d(), 288
graph1d_el_est(), 288
graph1d_fvalues(), 288
graph_clear_window(), 286
graph_close_window(), 286
graph_el_est(), 286
graph_drv(), 286
graph_drd_v(), 286
graph_fvalues(), 286
graph_level(), 289
graph_level_d(), 289
graph_levels(), 289
graph_levels_d(), 289
graph_line(), 286
graph_mesh(), 286
graph_open_window(), 286
graph_point(), 286
graph_values(), 289
graph_values_d(), 289
grd_f_at_qp(), 211
grd_f_d_at_qp(), 211
grd_uh_at_qp(), 219
grd_uh_d_at_qp(), 219
H1_err(), 222
H1_err_d(), 222
H1_norm_uh(), 221
H1_norm_uh_d(), 221

```

```

heat_est(), 266
init_dof_admin(), 207
init_mat_vec_d(), 274
init_mat_vec_s(), 274
init_parameters(), 123
integrate_std_simp(), 211
interpol(), 248
interpol_d(), 248
L2_err(), 222
L2_err_d(), 222
L2_norm_uh(), 221
L2_norm_uh_d(), 221
L2scp_fct_bas(), 245
L2scp_fct_bas_d(), 245
max_err_at_qp(), 222
max_err_at_qp_d(), 222
mesh_traverse(), 155
MG(), 279
marking(), 255
mat_vec_d(), 274
mat_vec_s(), 274
max_quad_points(), 215
mg_s(), 280
mg_s_exit(), 281
mg_s_init(), 281
mg_s_solve(), 281
new_bas_fcts(), 190
nls_newton(), 284
nls_newton_br(), 284
nls_newton_ds(), 284
nls_newton_fs(), 284
oem_bicgstab(), 270
oem_cg(), 270
oem_gmres(), 270
oem_odir(), 270
oem_ores(), 270
oem_solve_d(), 273
oem_solve_s(), 273
one_timestep(), 259
open_error_file(), 117
open_gltools_window(), 291
open_msg_file(), 117
print_dof_int_vec(), 167
print_dof_matrix(), 169
print_dof_real_d_vec(), 167
print_dof_real_vec(), 167
print_dof_schar_vec(), 167
print_dof_uchar_vec(), 167
print_mem_use(), 121
print_msg(), 115
read_dof_int_vec(), 174
read_dof_int_vec_xdr(), 176
read_dof_real_d_vec(), 174
read_dof_real_d_vec_xdr(), 176
read_dof_real_vec(), 174
read_dof_real_vec_xdr(), 176
read_dof_schar_vec(), 174
read_dof_schar_vec_xdr(), 176
read_dof_uchar_vec(), 174
read_dof_uchar_vec_xdr(), 176
read_macro(), 146
read_mesh(), 174
read_mesh_xdr(), 175
realloc_workspace(), 120
refine(), 176
save_parameters(), 124
sort_face_indices(), 215
sor_d(), 274
sor_s(), 274
ssor_d(), 274
ssor_s(), 274
traverse_first(), 158
traverse_neighbour(), 159
traverse_next(), 158
uh_at_qp(), 219
uh_d_at_qp(), 219
update_matrix(), 228
update_real_d_vec(), 245
update_real_vec(), 245
vector_product(), 129
world_to_coord(), 209
write_dof_int_vec(), 174
write_dof_int_vec_xdr(), 176
write_dof_real_d_vec(), 174
write_dof_real_d_vec_xdr(), 176
write_dof_real_vec(), 174
write_dof_real_vec_xdr(), 176
write_dof_schar_vec(), 174
write_dof_schar_vec_xdr(), 176
write_dof_uchar_vec(), 174
write_dof_uchar_vec_xdr(), 176
write_macro(), 150
write_macro_bin(), 150
write_macro_xdr(), 150
write_mesh(), 174
write_mesh_xdr(), 175

```

Macros

ABS(), 113	MAT_ALLOC(), 119
ADD_PARAMETER(), 124	MAT_FREE(), 119
DIST_DOW(), 128	MAX(), 113
EL_TYPE(), 137	MEM_ALLOC(), 118
ENTRY_NOT_USED(), 168	MEM_CALLOC(), 118
ENTRY_USED(), 168	MEM_FREE(), 118
ERROR_EXIT(), 116	MEM_REALLOC(), 118
ERROR(), 116	MIN(), 113
FOR_ALL_DOFS(), 170	MSG(), 114
FOR_ALL_FREE_DOFS(), 170	NEIGH(), 137
FUNCNAME(), 114	NORM_DOW(), 128
GET_BOUND(), 131	OPP_VERTEX(), 137
GET_DOF_VEC(), 166	PRINT_INFO(), 116
GET_MESH(), 143	REALLOC_WORKSPACE(), 120
GET_PARAMETER, 125	SCAL_DOW(), 128
INDEX(), 137	SCP_DOW(), 128
INFO(), 116	SET_DOW(), 128
IS_DIRICHLET(), 131	SQR(), 113
IS_INTERIOR(), 131	TEST_EXIT(), 116
IS_LEAF_EL(), 139	TEST(), 116
IS_NEUMANN(), 131	WAIT REALLY, 117
LEAF_DATA(), 139	WAIT, 117
	WARNING(), 117