

SyFi User Manual

April 15, 2011

Martin Alnæs and Kent-Andre Mardal

www.fenics.org

Visit <http://www.fenics.org/> for the latest version of this manual.
Send comments and suggestions to syfi-dev@fenics.org.

Contents

1	Introduction	5
2	Software	9
2.1	License	9
2.2	Installation	10
2.3	Python Support	11
2.4	Examples and Tests	11
2.5	GiNaC Tools	12
2.5.1	The symbol factory	12
2.5.2	Symbols for spatial variables	13
3	A Finite Element	15
3.1	Basic Concepts	15
3.2	Polygons	16
3.2.1	Line	17

3.2.2	Triangle	19
3.2.3	Tetrahedron	21
3.2.4	Rectangle	24
3.2.5	Box	27
3.3	Polynomial Spaces	29
3.3.1	Bernstein Polynomials	31
3.3.2	Legendre Polynomials	33
3.3.3	Homogeneous Polynomials	35
3.4	A Finite Element	36
3.5	Degrees of Freedom	39
4	Some Examples of Finite Elements	49
4.1	Finite Elements in H^1	49
4.1.1	The Lagrangian Element	49
4.1.2	The Crouizex-Raviart Element	52
4.2	Finite Elements in L^2	55
4.2.1	The P_0 Element	55
4.2.2	The Discontinuous Lagrangian Element	55
4.3	Finite Elements in $\mathbf{H}(div)$	59
4.3.1	The Raviart-Thomas Element	59
4.3.2	The Nedelec element of second kind	63

4.4	Finite Elements in $H(\text{div}, \mathbf{M})$	65
4.5	A Finite Element in Both $\mathbf{H}(\text{div})$ and H^1	65
4.6	Finite Elements in $\mathbf{H}(\text{curl})$	68
4.6.1	The Nedelec Element	68
5	Mixed Finite Elements	73
5.1	The Taylor–Hood and the $\mathbb{P}_n^d - \mathbb{P}_{n-2}$ Elements	73
5.2	The Mixed Crouizex-Raviart Element	74
5.3	The Mixed Raviart-Thomas Element	75
5.4	The Mixed Arnold-Falk-Winther element	76
6	Computing Element Matrices	77
6.1	A Poisson Problem	78
6.2	A Poisson Problem on Mixed Form	82
6.3	A Stokes Problem	83
6.4	A Nonlinear Convection Diffusion Problem	85
6.5	Expression Simplification	87
7	Python Support	89
8	Code Generation	93
8.1	Basic Tools	94
8.2	Debugging	98

9 Using the SyFi Form Compiler	101
9.1 Quickstart	102
9.2 Defining Form Arguments	103
9.2.1 Defining Finite Elements	103
9.2.2 Defining Basisfunctions	104
9.2.3 Defining Coefficients	104
9.3 Defining a Form	105
9.4 Defining an Integral	106
9.4.1 Argument expressions	106
9.4.2 Geometric Quantities on Cells	107
9.4.3 Symbolic Language	107
9.4.4 Examples	108
9.5 Defining forms with callback functions	110
9.6 Computing the Jacobi matrix form from a nonlinear vector form	113
9.7 Compiling a Form (Generating Code)	113
9.8 Options	114
9.9 Compiling a function	114
10 Behind the SyFi Form Compiler	117
10.1 Example of generated code	117
10.2 Data Flow During Code Generation	119

10.3 Code generation design 120

10.4 Code Formatting Utilities 123

Chapter 1

Introduction

The reader should be aware that this manual is not quite up to date. Discrepancies between this manual and the current source code is to be expected, but the general concepts should stay the same.

The finite element package SyFi is a C++ library built on top of the symbolic math library GiNaC [?]. The name SyFi stands for Symbolic Finite elements. The package provides polygonal domains, polynomial spaces, and degrees of freedom as symbolic expressions that are easily manipulated. This makes it easy to define and use finite elements.

The SyFi Form Compiler (SFC) allows the use of the symbolic expressions for finite elements from SyFi to compile efficient C++ code.

All the test examples described in this tutorial can be found in the directory `tests`. The reader is of course encouraged to run the examples along with the reading.

Before we start to describe SyFi, we need to briefly review the basic concepts in GiNaC. GiNaC is an open source C++ library for symbolic mathematics, which has a strong support for polynomials. The basic structure in GiNaC is an `ex`, which may contain either a number, a symbol, a function, a list of expressions, etc. (see `simple.cpp`):

```
ex pi = 3.14;
ex x = symbol("x");
ex f = cos(x);
ex list = lst(pi,x,f);
```

Hence, `ex` is a quite general class, and it is the cornerstone of GiNaC. It has a lot of functionality, for instance differentiation and integration (see `simple2.cpp`),

```
// initialization (f = x^2 + y^2)
ex f = x*x + y*y;

// differentiation (dfdx = df/dx = 2x)
ex dfdx = f.diff(x,1);

// integration (intf=1/3+y^2, integral of f(x,y) on x=[0,1])
ex intf = integral(x,0,1,f);
```

GiNaC also has a structure for lists of expressions, `lst`, with the function `nops()` which returns the size of the list, and operator `[int i]` or the function `op(int i)` which returns the i 'th element.

We recommend the reader to glance through the GiNaC documentation before proceeding with this tutorial. GiNaC provides all the basic tools for manipulation of polynomials, such as differentiation and integration with respect to one variable. Our goal with the SyFi package is to employ GiNaC, but also to provide higher level constructs such as differentiation with respect to several variables (e.g., ∇), integration of functions over polygonal domains, and polynomial spaces. All of which are basic ingredients in the finite element method.

Our motivation behind this project is threefold. First, we wish to make advanced finite element methods more readily available. We want to do this by implementing a variety of finite elements and functions for computing element matrices. Second, in our experience developing and debugging codes

for finite element methods is hard. Having the basis functions and the weak form as symbolic expressions, and being able to manipulate them may be extremely valuable. For instance, being able to differentiate the weak form to compute the Jacobian in the case of a nonlinear PDE, eliminates a lot of handwriting and coding. Third, having the symbolic expressions and employing GiNaCs tools for code generation, we are able to write efficient and directly compilable C++ code for the computation of element matrices etc.

To try to motivate the reader, we also show an example where we compute the element matrix for the weak form of the Poisson equation, i.e.,

$$A_{ij} = \int_T \nabla N_i \cdot \nabla N_j \, dx.$$

We remark that the following example is an attempt to make an appetizer. All concepts will be carefully described during the tutorial.

```
void compute_element_matrix(Polygon& T, int order) {
    std::map<std::pair<int,int>, ex> A;           // matrix of expression
    std::pair<int,int> index;                     // index in matrix
    LagrangeFE fe;                               // Lagrange element (any order)
    fe.set_order(order);                         // set the order
    fe.set_polygon(domain);                      // set the polygon
    fe.compute_basis_functions();                 // compute the basis functions
    for (int i=0; i< fe.nbf(); i++) {
        index.first = i;
        for (int j=0; j< fe.nbf(); j++) {
            index.second = j;
            ex nabla = inner(grad(fe.N(i)),      // compute integrands
                             grad(fe.N(j)));
            ex Aij = T.integrate(nabla);         // compute weak form
            A[index] = Aij;                     // update element matrix
        }
    }
}
```

In the above example, everything is computed symbolically. Even the polygon may be an abstract polygon, e.g., specified as a triangle with vertices \mathbf{x}_0 , \mathbf{x}_1 , and \mathbf{x}_2 , where the vertices are symbols and not concrete points. Notice also, that we usually use STL containers to store our datastructure. This leads to the somewhat unfamiliar notation `A[index]` instead of `A[i,j]`.

There are quite a few other projects that are similar in various respects to SyFi. We will not give a comprehensive description of these projects, only mention the projects such that the readers can look them up by themselves. Within the FEniCS [?] project there are two Python projects: FIAT [?] and FFC [?]. FIAT is a Python module for defining finite elements while FFC generates C++ code based on a high-level Python description of variational forms. The DSEL project [?] is a project which employs high-level C++ programming techniques such as expression templates and meta-programming for defining variational forms, performing automatic differentiation, interpolation and more. Sundance [?] is a C++ library with a powerful symbolic engine which supports automatic generation of discrete system for a given variational form. Analysa [?], GetDP [?], and FreeFem++ [?] define domain-specific languages for finite element computations.

Finally, we have to warn the reader: This project is still within its initial phase.

Chapter 2

Software

2.1 License

SyFi employs GiNaC and is therefore limited by GiNaCs license, which is the GPL-2 licence listed below.

However, SyFi is usually used to generated code. The generated code is free.

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Notice, however that SyFi is usually used to generate code. This code is free, but it comes without any warranty for fitness of any purpose.

In the case where the GNU licence does not fit your need. Contact the authors at `syfi-dev@fenics.org`.

2.2 Installation

Dependencies

SyFi is a C++ library and therefore a C++ compiler is needed. At present the library has only been tested with the GNU C++ compiler. The `configure` script is a shell script made by the tools Automake and Autoconf. Hence, you can run this script with, e.g., the GNU Bourne-again shell. Finally, SyFi relies on the C++ library GiNaC.

Configuration and Installation

As mention earlier, the configuration, build and installation scripts are all made by the Autoconf and Automake tools. Hence, to configure, build and install the package, simply execute the commands,

```
bash >./configure
bash >make
bash >make install
```

If this does not work, it is most likely because GiNaC is not properly installed on your system. Check if you have the script `ginac-config` in your path.

Reporting Bugs/Submitting Patches

In case, you want to contribute code, please create a patch with `diff`,

```
bash >diff -u -N -r SyFi SyFi-mod > SyFi-<name>-<date>.patch
```

Here `<name>` should be replaced with your name and `<date>` should be replaced with the current date.

The patch should be mailed to the SyFi mailing-list at `syfi-dev@fenics.org`.

2.3 Python Support

SyFi comes with Python support. The Python module is made by using the tool SWIG [?]. In addition, one should also install Swiginac [?], which is a Python interface to GiNaC created by using SWIG. More about the usage of the Python interface can be found in Section 7.

2.4 Examples and Tests

A series of tests are located in the subdir `tests`, these test serve as unit test and document the features of SyFi as described in this tutorial. If the tests are simple we use the function `EQUAL_OR_DIE`, an example is (see also `simple_test.cpp`)

```
symbol x("x");
ex f = x*x;
ex intf = integral(x,0,1,f);
intf = eval_integ(intf);
EQUAL_OR_DIE(integral1, "1/3");
```

When the tests or computed expressions are bigger we typically store the expressions in a GiNaC archive (`.gar` files) and compare the archive with a previously created and verified archive. The following code demonstrates how the basis functions and degrees of freedom of a first order Lagrangian element is computed, stored in an archive and then compared with the previously verified basis functions and degrees of freedom.

```
int order = 1;
ReferenceTriangle triangle;
LagrangeFE fe(triangle, order);

// regression test
archive ar;
for (int i=0; i< fe.nbf(); i++) {
    ar.archive_ex(fe.N(i) , istr("N",i).c_str());
    ar.archive_ex(fe.dof(i) , istr("D",i).c_str());
}
ofstream vfile("fe_ex1.gar.v");
vfile << ar; vfile.close();
if(!compare_archives("fe_ex1.gar.v", "fe_ex1.gar.r")) {
    cerr << "Failure!" << endl;
    return -1;
}
```

All examples described in this tutorial are also implemented as tests in the `tests` subdir.

2.5 GiNaC Tools

2.5.1 The symbol factory

In GiNaC, the identity of a symbol is not defined by its name, but by an internal number. Because of this, the code

```
ex a = symbol("x");
ex b = symbol("x");
ex c = a-b;
```


does not yield 0 in `c`, since `a` and `b` refer to different symbols. To solve this we have implemented a simple symbol factory, so we can refer to variables by name without passing the symbol objects around everywhere. The user can ask if a symbol exists, or get numbered symbols and vectors or matrices of numbered symbols in a convenient way.

```
ex x1 = get_symbol("x");
ex x2 = get_symbol("x");
assert( is_zero(x1-x2) );

ex u = get_symbolic_vector(3, "u");
ex A = get_symbolic_matrix(3, 3, "A");
ex c = isymb("c", 2);
assert( symbol_exists("c2") );
```

2.5.2 Symbols for spatial variables

Certain operations like the differential operators `diff` and `grad` needs to know certain symbols to operate correctly. The spatial variables `x,y,z` and `t` are particularly important, and because of this we have a shortcut to these variables. Operations like `grad` also need to know the number of spatial dimensions, often abbreviated `nsd` in SyFi. Therefore, a call to `initSyFi(nsd)` must be made before one can use these operators. It is safe to call `initSyFi` more than once. The spatial symbols `x,y,z,t` can also be retrieved from the symbol factory.

```
initSyFi(3);
int space_dim = SyFi::nsd;
ex x = SyFi::x;
ex y = SyFi::y;
ex z = SyFi::z;
ex t = SyFi::t;
// or:
ex x = get_symbol("x");
```


Chapter 3

A Finite Element

3.1 Basic Concepts

To keep the abstractions clear we briefly review the general definition of a finite element, see e.g., Brenner and Scott [?] or Ciarlet [?].

Definition 3.1.1 (A Finite Element) *A finite element consists of,*

1. *A polygonal domain, T .*
2. *A polynomial space, V .*
3. *A set of degrees of freedom (linear forms), $L_i : V \rightarrow \mathbb{R}$, for $i = 1, \dots, n$, where $n = \dim(V)$, that determines V uniquely.*

Furthermore, to determine a basis in V , $\{N_i\}_{i=1}^n$, we form the linear system of equations,

$$L_i(N_j) = \delta_{ij}, \tag{3.1}$$

and solve it.

Example 3.1.1 (Linear Lagrangian element on the reference triangle)

In this example we describe how the linear Lagrangian element is defined on

the reference triangle. Let T be the unit triangle with vertices $(0,0)$, $(1,0)$, and $(0,1)$. Furthermore, the polynomial space V consists of linear polynomials, i.e., polynomials on the form $N(x,y) = a + bx + cy$. The degrees of freedom for a linear Lagrangian element are simply the pointvalues at the vertices, \mathbf{x}_i , $L_i(N_j) = N_j(\mathbf{x}_i)$. The degrees of freedom and (3.1) determined a_j , b_j , and c_j for each basis function N_j . For instance N_1 , which is on the form $a_1 + b_1x + c_1y$, is determined by,

$$L_i(N_1) = N_1(\mathbf{x}_i) = \delta_{i1},$$

or written out as a system of linear equations,

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ b_1 \\ c_1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

Hence,

$$N_1 = 1 - x - y.$$


The functions N_2 and N_3 can be determined similarly.

In the next sections we will go detailed through polygons, polynomial spaces and degrees of freedom, and the corresponding software components.

3.2 Polygons

In the finite element method we need the concept of simple polygons to define integration, polynomial spaces etc. The basic polygons are lines, triangles, tetrahedra, and orthogonal rectangles and boxes. These basic components will be briefly described in this section.

Figure 3.1: A line.



3.2.1 Line

A line segment, L , between two points $\mathbf{x}_0 = [x_0, y_0, z_0]$ and $\mathbf{x}_1 = [x_1, y_1, z_1]$ in 3D is defined as, see also Figure 3.2.1,

$$x = x_0 + a t, \quad (3.2)$$

$$y = y_0 + b t, \quad (3.3)$$

$$z = z_0 + c t, \quad (3.4)$$

$$t \in [0, 1], \quad (3.5)$$

where $a = x_1 - x_0$, $b = y_1 - y_0$, and $c = z_1 - z_0$.

Integration of a function $f(x, y, z)$ along the line segment L is performed by substitution,

$$\int_L f(x, y, z) dx dy dz = \int_0^1 f(x(t), y(t), z(t)) D dt, \quad (3.6)$$

where $D = \sqrt{a^2 + b^2 + c^2}$.

Software Component: Line

The class `Line` implements a general line. It is defined as follows (see `Polygon.h`):

```
class Line : public Polygon {
ex a_;
ex b_;
public:
  Line() {}
  Line(ex x0, ex x1, // x0_ and x1_ are points
        string subscript = "");
  ~Line(){}

  virtual int no_vertices();
  virtual ex vertex(int i);
  virtual ex repr(ex t);
  virtual string str();
  virtual ex integrate(ex f);
};
```

Most of the functions in this class are self-explanatory. However, the function `repr` deserves special attention. The function `repr` returns the mathematical definition of a line. To be precise, this function returns a list of expressions (`1st`), where the items are the items in (3.2)-(3.5) (see also the example below).

The basic usage of a line is as follows (see `line_ex1.cpp`),

```
ex p0 = 1st(0.0,0.0,0.0);
ex p1 = 1st(1.0,1.0,1.0);

Line line(p0,p1);

// show usage of repr
symbol t("t");
ex l_repr = line.repr(t);
cout <<"l.repr "<<l_repr<<endl;
EQUAL_OR_DIE(l_repr, "{x==t,y==t,z==t,{t,0,1}}");
```

```

for (int i=0; i< l_repr.nops(); i++) {
    cout <<"l_repr.op(" <<i<<"):  "<<l_repr.op(i)<<endl;
}

// compute the integral of a function along the line
ex f = x*x + y*y*y + z;
ex intf = line.integrate(f);
cout <<"intf "<<intf<<endl;
EQUAL_OR_DIE(intf, "13/12");

```

The function `EQUAL_OR_DIE` compares the string representation of the expression with an expected expression represented as a character array. If the string representation of the expression and the character array are not equal the program dies, and this tells the programmer that the test faulted. The reason for the use of this function is that our test examples also serve as regression tests for the package.

3.2.2 Triangle

A triangle is defined in terms of three points \mathbf{x}_0 , \mathbf{x}_1 , and \mathbf{x}_2 . Associated with each triangle are three lines; the first line is between the points \mathbf{x}_1 and \mathbf{x}_2 , the second line is between the points \mathbf{x}_0 and \mathbf{x}_2 , and the third line is between the points \mathbf{x}_0 and \mathbf{x}_1 . This is shown in Figure 3.2. The triangle can be represented as

$$x = x_0 + ar + bs, \quad (3.7)$$

$$y = y_0 + cr + ds, \quad (3.8)$$

$$z = z_0 + er + fs, \quad (3.9)$$


$$s \in [0, 1 - r], \quad (3.10)$$

$$r \in [0, 1], \quad (3.11)$$

where $(a, c, e) = (x_1 - x_0, y_1 - y_0, z_1 - z_0)$ and $(b, d, f) = (x_2 - x_0, y_2 - y_0, z_2 - z_0)$.

Integration is performed by substitution,

$$\int_T f(x, y, z) dx dy dz = \int_0^1 \int_0^{1-r} f(x, y, z) D ds dr,$$



where $D = \sqrt{(cf - de)^2 + (af - be)^2 + (ad - bc)^2}$.

Software Component: Triangle

The class `Triangle` implements a general triangle. It is defined as follows (see `Polygon.h`):

```
class Triangle : public Polygon {
public:
    Triangle(ex x0, ex x1, ex x1, string subscript = "");
    Triangle() {}
    ~Triangle(){}

    virtual int no_vertices();
    virtual ex vertex(int i);
    virtual Line line(int i);
    virtual ex repr();
};
```



```
virtual string str();  
virtual ex integrate(ex f);  
};
```

Here the function `repr` returns a list with the items (3.7)-(3.11). In addition to the functions also found in `Line`, `Triangle` has a function `line(int i)`, returning a line.

The basic usage of a triangle is as follows (see `triangle.ex1.cpp`),

```
ex p0 = lst(0.0,0.0,1.0);  
ex p1 = lst(1.0,0.0,1.0);  
ex p2 = lst(0.0,1.0,1.0);  
  
Triangle triangle(p0,p1,p2);  
  
ex repr = triangle.repr();  
cout <<"t.repr " <<repr<<endl;  
EQUAL_OR_DIE(repr, "{x==r,y==s,z==1.0,{r,0,1},{s,0,1-r}}");  
  
ex f = x*y*z;  
ex intf = triangle.integrate(f);  
cout <<"intf " <<intf<<endl;  
EQUAL_OR_DIE(intf, "1/24");
```

3.2.3 Tetrahedron

A tetrahedron is defined by four points \mathbf{x}_0 , \mathbf{x}_1 , \mathbf{x}_2 , and \mathbf{x}_3 . Associated with a tetrahedron are four triangles and six lines. The convention used so far is that

- the first line connects \mathbf{x}_0 and \mathbf{x}_1 ,
- the second line connects \mathbf{x}_0 and \mathbf{x}_2 ,

- the third line connects \mathbf{x}_0 and \mathbf{x}_3 ,
- the fourth line connects \mathbf{x}_1 and \mathbf{x}_2 ,
- the fifth line connects \mathbf{x}_1 and \mathbf{x}_3 ,
- the sixth line connects \mathbf{x}_2 and \mathbf{x}_3 .

The i 'th triangle contains all vertices except the i 'th vertex. The tetrahedron can be represented as, see also Figure 3.3,

$$x = x_0 + ar + bs + ct, \quad (3.12)$$

$$y = y_0 + dr + es + ft, \quad (3.13)$$

$$z = z_0 + gr + hs + kt, \quad (3.14)$$

$$t \in [0, 1 - r - s], \quad (3.15)$$

$$s \in [0, 1 - r], \quad (3.16)$$

$$r \in [0, 1], \quad (3.17)$$

where $(a, d, g) = (x_1 - x_0, y_1 - y_0, z_1 - z_0)$, $(b, e, h) = (x_2 - x_0, y_2 - y_0, z_2 - z_0)$, and $(c, f, k) = (x_3 - x_0, y_3 - y_0, z_3 - z_0)$.

As earlier, integration is performed with substitution,

$$\begin{aligned} \int_T f(x, y, z) dx dy dz = \\ \int_0^1 \int_0^{1-r} \int_0^{1-r-s} f(x(r, s, t), y(r, s, t), z(r, s, t)) D dt ds dr, \end{aligned}$$


where D is the determinant of,

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & k \end{bmatrix}.$$

Software Component: Tetrahedron

The class `Tetrahedron` implements a general tetrahedron. It is defined as follows (see `Polygon.h`):

Figure 3.3: A tetrahedron.



```

class Tetrahedron : public Polygon {
public:
    Tetrahedron(string subscript) {}
    Tetrahedron(ex x0, ex x1, ex x1, ex x2, string s= "");
    ~Tetrahedron(){}

    virtual int no_vertices();
    virtual ex vertex(int i);
    virtual Line line(int i);
    virtual Triangle triangle(int i);
    virtual ex repr();
    virtual string str();
    virtual ex integrate(ex f);
};

```

The function `repr` returns a list representing (3.12) –(3.17). In addition to the usual functions it has the functions `line(int i)` and `triangle(int i)` for

returning the i 'th line and the i 'th triangle, respectively.

Its basic usage is as follows (see `tetrahedron.ex1.cpp`),

```
ex p0 = lst(0.0,0.0,0.0);
ex p1 = lst(1.0,0.0,0.0);
ex p2 = lst(0.0,1.0,0.0);
ex p3 = lst(0.0,0.0,1.0);

Tetrahedron tetrahedron(p0,p1,p2,p3);

ex repr = tetrahedron.repr();
cout <<"t.repr "<<repr<<endl;
EQUAL_OR_DIE(repr, "{x==r,y==s,z==t,{r,0,1},
                  {s,0,1-r},{t,0,1-s-r}}");

ex f = x*y*z;
ex intf = tetrahedron.integrate(f);
EQUAL_OR_DIE(intf, "1/720");
```

3.2.4 Rectangle

The rectangles currently supported by SyFi are orthogonal. Such a rectangle is defined in terms of two points \mathbf{x}_0 and \mathbf{x}_1 , as shown in Figure 3.4.

The rectangle can be represented as

$$x = x_0 + ar, \quad (3.18)$$

$$y = y_0 + bs, \quad (3.19)$$


$$z = z_0 + ct, \quad (3.20)$$

$$r \in [0, 1], \quad (3.21)$$

$$s \in [0, 1], \quad (3.22)$$

$$t \in [0, 1], \quad (3.23)$$

Figure 3.4: A rectangle.



where $a = x_1 - x_0$, $b = y_1 - y_0$, and $c = z_1 - z_0$. Notice that either a , b , or c needs to be zero, or else (3.18)-(3.23) defines a box (which will be described later).

As earlier, integration is performed with substitution,

$$\int_R f(x, y, z) dx dy dz = \int_0^1 \int_0^1 \int_0^1 f(x(r, s, t), y(r, s, t), z(r, s, t)) D dt ds dr,$$

where $D = ab$ if $c = 0$, $D = bc$, if $a = 0$, and $D = ac$, if $b = 0$.

Software Component: Rectangle

The class `Rectangle` implements a general orthogonal rectangle. It is defined as follows (see `Polygon.h`):

```
class Rectangle : public Polygon {
public:
```

```
Rectangle(GiNaC::ex p0, GiNaC::ex p1, string s = "");
Rectangle() {}
virtual ~Rectangle(){}

virtual int no_vertices();
virtual GiNaC::ex vertex(int i);
virtual Line line(int i);
virtual GiNaC::ex repr(Repr_format format = SUBS_PERFORMED);
virtual string str();
virtual GiNaC::ex integrate(GiNaC::ex f);
};
```

As described with the previous polygons, the function `repr` returns a list with the items (3.18)-(3.23). The basic usage of the rectangle is as follows (see `rectangle_ex1.cpp`),

```
ex f = x*y;

ex p0 = lst(0.0,0.0);
ex p1 = lst(1.0,1.0);

Rectangle rectangle(p0,p1);

ex repr = rectangle.repr();
cout <<"s.repr " <<repr<<endl;

ex intf = rectangle.integrate(f);
cout <<"intf " <<intf<<endl;

ex f2 = (x+1)*y*z;
p0 = lst(0.0,0.0,1.0);
p1 = lst(0.0,1.0,0.0);

Rectangle rectangle2(p0,p1);

ex repr2 = rectangle2.repr();
```

```
cout <<"s2.repr "<<repr2<<endl;

ex intf2 = rectangle2.integrate(f2);
cout <<"intf2 "<<intf2<<endl;
```

3.2.5 Box

Currently, SyFi only supports orthogonal boxes (as was also the case with rectangles). Such a box is defined in terms of two points \mathbf{x}_0 and \mathbf{x}_1 , as can be seen in Figure 3.5. The box can be represented as

$$x = x_0 + ar, \quad (3.24)$$

$$y = y_0 + bs, \quad (3.25)$$

$$z = z_0 + ct, \quad (3.26)$$

$$r \in [0, 1], \quad (3.27)$$

$$s \in [0, 1], \quad (3.28)$$

$$t \in [0, 1], \quad (3.29)$$

where $a = x_1 - x_0$, $b = y_1 - y_0$, and $c = z_1 - z_0$.

As earlier, integration is performed with substitution,


$$\begin{aligned} \int_R f(x, y, z) dx dy dz = \\ \int_0^1 \int_0^1 \int_0^1 f(x(r, s, t), y(r, s, t), z(r, s, t)) D dt ds dr, \end{aligned}$$

where $D = abc$.

Software Component: Box

The class `Box` implements a general orthogonal box. It is defined as follows (see `Polygon.h`):

Figure 3.5: A Box.



```
class Box: public Polygon {
public:
    Box(GiNaC::ex p0, GiNaC::ex p1, string subscript = "");
    Box(){}
    virtual ~Box(){}

    virtual int no_vertices();
    virtual GiNaC::ex vertex(int i);
    virtual Line line(int i);
    virtual GiNaC::ex repr(Repr_format format = SUBS_PERFORMED);
    virtual string str();
    virtual GiNaC::ex integrate(GiNaC::ex f);
};
```

The `repr` function returns a list of the definition of a orthogonal box in (3.24)-(3.29). A box can be used as follows (see `box_ex1.cpp`),


```

ex p0 = lst(-1.0,-1.0,-1.0);
ex p1 = lst( 1.0, 1.0, 1.0);

Box box(p0,p1);

ex repr = box.repr();
cout <<"b.repr "<<repr<<endl;

ex intf = box.integrate(f);
cout <<"intf "<<intf<<endl;

```

Finally, we also mention that in addition to the above mentioned classes, `Line`, `Triangle`, `Tetrahedron`, `Rectangle`, and `Box`, we have implemented the corresponding reference geometries in the subclasses `ReferenceLine`, `ReferenceTriangle`, `ReferenceTetrahedron`, `ReferenceRectangle`, and `ReferenceBox`.

3.3 Polynomial Spaces

The space of polynomials of degree less or equal to n , \mathbb{P}_n , plays a fundamental role in the construction of finite elements. There are many ways to represent this polynomial space. The perhaps visually nicest representation is having it spanned by the basis (in 1D) $1, x, x^2, \dots, x^n$. This representation is not suitable for polynomials of high degree¹.

In 1D, \mathbb{P}_n is spanned by functions on the form

$$v = a_0 + a_1x + \dots a_nx^n = \sum_{i=0}^n a_ix^i \quad (3.30)$$

In 2D on triangles, \mathbb{P}_n is spanned by functions on the form:

$$v = \sum_{i,j=0}^{i+j \leq n} a_{ij}x^iy^j \quad (3.31)$$

¹In that case, one should use the Bernstein or Legendre polynomials.

In 2D on quadrilaterals, \mathbb{P}_n is spanned by functions on the form:

$$v = \sum_{i,j=0}^{i,j \leq n} a_{ij} x^i y^j \quad (3.32)$$

The corresponding polynomial spaces in 3D are completely analogous.

Software Component: Polynomial Space

The following functions generate symbolic expressions for the above polynomial spaces (3.30), (3.31), and (3.32), their corresponding polynomial spaces in 3D and their vector counterparts.

```
// generates a polynomial of any order on a line,
// a triangle, or a tetrahedron
ex pol(int order, int nsd, const string a);

// generates a vector polynomial of any order on a line,
// a triangle or a tetrahedron
lst polv(int order, int nsd, const string a);

// generates a polynomial of any order on a square or a box
ex polb(int order, int nsd, const string a);

// generates a vector polynomial of any order
// on a square or a box
lst polbv(int order, int nsd, const string a);
```

The function `pol` returns a list with the following 3 items,

1. The polynomial, e.g., $a_0 + a_1x + \dots + a_nx^n$ in 1D.
2. A list of variables, e.g., a_0, a_1, \dots, a_n in 1D.
3. A list containing the basis, e.g., $1, x, \dots, x^n$ in 1D.

The functions `polb`, `polv`, and `polbv` return lists that are completely analogous.

These abstract polynomials (or polynomial spaces) can be easily manipulated, e.g., (see also `pol.cpp`),

```
int order = 2;
int nsd    = 2;

ex polynom_space = pol(order,nsd, "a");
cout <<"polynom_space "<<polynom_space<<endl;

ex p = polynom_space.op(0);
cout <<"polynom p = "<<p<<endl;
EQUAL_OR_DIE(p, "y^2*a5+x^2*a3+a2*y+y*x*a4+a0+a1*x");

ex dpdx = diff(p,x);
cout <<"dpdx = "<<dpdx<<endl;
EQUAL_OR_DIE(dpdx, "y*a4+a1+2*x*a3");

Triangle triangle(lst(0,0), lst(1,0), lst(0,1));
ex intp = triangle.integrate(p);
cout <<"integral of p on reference triangle="<<intp<<endl;
EQUAL_OR_DIE(intp, "1/6*a2+1/6*a1+1/12*a5
                  +1/2*a0+1/24*a4+1/12*a3");
```

3.3.1 Bernstein Polynomials

Another basis for \mathbb{P}_n is the Bernstein polynomials. This basis is much better suited for polynomials of high degree. Moreover, these polynomials can be easily expressed in barycentric coordinates, which makes them easy to adapt to, e.g., faces of polygons² etc.

²This will be used in the definition of the Raviart-Thomas element.

In 1D, the polynomial basis is on the form,

$$B_{i,n} = \binom{i}{n} x^i (1-x)^{n-i}, \quad i = 0, \dots, n.$$

And with this basis, \mathbb{P}_n can be spanned by functions on the form,

$$v = a_0 B_{0,n} + a_1 B_{1,n} + \dots a_n B_{n,n}.$$

One reason for the widespread use of these polynomials is that they adapt easily to general triangles and tetrahedra, by using barycentric coordinates. Let b_1 , b_2 , and b_3 be the barycentric coordinates for the triangle shown in Figure 3.2. Then the basis is on the form,

$$B_{i,j,k,n} = \frac{n!}{i!j!k!} b_1^i b_2^j b_3^k, \quad \text{for } i + j + k = n.$$

and \mathbb{P}_n is spanned by functions on the form,

$$v = \sum_{i+j+k=n} a_{i,j,k} B_{i,j,k,n}.$$

The Bernstein polynomials in 3D are completely analogous.

Software Components: Bernstein polynomials

The following functions generate symbolic expressions for \mathbb{P}^n using the Bernstein basis,

```
// polynom of arbitrary order on a line, a triangle,
// or a tetrahedron using the Bernstein basis
ex bernstein(int order, Polygon& p, const string a);

// vector polynom of arbitrary order on a line, a triangle,
// or a tetrahedron using the Bernstein basis
lst bernsteinv(int order, Polygon& p, const string a);
```

These functions return lists that are analogous to the lists made by the functions `pol` and `polv` described on page 30.

As described earlier, GiNaC has the tools for manipulating these polynomial spaces, (see also `pol.cpp`)

```
ex polynom_space2 = bernstein(order,triangle, "a");
ex p2 = polynom_space2.op(0);
cout <<"polynom p2 = "<<p2<<endl;
EQUAL_OR_DIE(p2, "y^2*a0+2*(1-y-x)*x*a4+2*(1-y-x)*a3*y
                +(1-y-x)^2*a5+2*a1*y*x+a2*x^2");

ex dp2dx = diff(p2,x);
cout <<"dp2dx = "<<dp2dx<<endl;
EQUAL_OR_DIE(dp2dx, "2*a1*y+2*(-1+y+x)*a5+2*a2*x
                  +2*(1-y-x)*a4-2*a3*y-2*x*a4");

ex intp2 = triangle.integrate(p2);
cout <<"integral of p2 on reference triangle="<<intp2<<endl;
EQUAL_OR_DIE(intp2, "1/12*a3+1/12*a2+1/12*a1+1/12*a5
                  +1/12*a0+1/12*a4");
```

3.3.2 Legendre Polynomials

A popular polynomial basis for polygons that are either rectangles or boxes are the Legendre polynomials. This polynomial basis is also usable to represent polynomials of high degree. The basis is defined on the interval $[-1, 1]$, as

$$L_k(x) = \frac{1}{2^k k!} \frac{d^k}{dx^k} (x^2 - 1), \quad k = 0, 1, \dots,$$

A nice feature with these polynomials is that they are orthogonal with respect to the L_2 inner product, i.e.,

$$\int_{-1}^1 L_k(x) L_l(x) dx = \begin{cases} \frac{2}{2k+1}, & k = l, \\ 0, & k \neq l, \end{cases}$$

The Legendre polynomials are extended to 2D and 3D simply by taking the tensor product,

$$L_{k,l,m}(x, y, z) = L_k(x)L_l(y)L_m(z).$$

and \mathbb{P}^n is defined by functions on the form (in 3D),

$$v = \sum_{k,l,m=0}^{k,l,m \leq n} a_{k,l,m} L_{k,l,m}.$$

Software Components: Legendre polynomials

The following functions generate symbolic expressions for \mathbb{P}^n using the Legendre basis,

```
// generates a Legendre polynom of arbitrary order
GiNaC::ex legendre(int order, int nsd, const string a);
// generates a Legendre vector polynom of arbitrary order
GiNaC::lst legendrev(int no_fields, int order,
                     int nsd, const string a);
```

These functions return lists that are analogous to the lists made by the functions `pol` and `polv` described on page 30.

The following code demonstrates the use of the Legendre polynomials, and (when runned) that the basis is orthogonal (see also `legendre.cpp`).

```
int order = 2;
int nsd   = 2;

ex polynom_space = legendre(order, nsd, "a");
cout << "polynom_space " << polynom_space << endl;

ex p = polynom_space.op(0);
cout << "polynom p = " << p << endl;
```

```

ex dpdx = diff(p,x);
cout <<"dpdx = "<<dpdx<<endl;

ex p0 = lst(-1,-1);
ex p1 = lst(1,1);

Rectangle rectangle(p0,p1) ;
ex basis = polynom_space.op(2);
for (int i=0; i< basis.nops(); i++) {
    cout <<"i "<<i<<endl;
    ex integrand = p*basis.op(i);
    ex ai = rectangle.integrate(integrand);
    cout <<"ai "<<ai<<endl;
}

```

3.3.3 Homogeneous Polynomials

Another set of polynomials which sometimes are useful are the set of homogeneous polynomials \mathbb{H}^n . These are polynomials where all terms have the same degree. \mathbb{H}^n is in 2D spanned by polynomials on the form:

$$v = \sum_{\substack{i,j, \\ i+j=n}} a_{i,j,k} x^i y^j$$

Software Components: Homogeneous polynomials

The following functions generate symbolic expressions for \mathbb{H}^n ,

```

// generates a homogeneous polynom of arbitrary order
GiNaC::ex homogenous_pol(int order, int nsd, const string a);
// generates a homogenous vector polynom of arbitrary order

```

```
GiNaC::lst homogenous_polv(int no_fields, int order,  
                           int nsd, const string a);
```

The use of these polynomials are similar to the other polynomials described earlier.

3.4 A Finite Element

Before we start describing how to construct a finite element based on the Definition 3.1.1, we will concentrate on the *usage* of a finite element. A finite element has only two interesting components, the basis functions N_i and the corresponding degrees of freedom L_i . The basis functions (and their derivatives) are used to compute the element matrices and the element vectors, while the degrees of freedom are used to define the mapping between the element matrices/vectors and the global matrix/vector. As we see in the following, the observation that only these two components are needed leads us to a minimalistic definition of a finite element in our software tools.

Software Component: Finite Element

Due to the powerful expression class in GiNaC, `ex`, our base class for the finite elements can be very small. Both the basis function N_i and the corresponding degree of freedom L_i can be well represented as an `ex`. Hence, the following definition of a finite element is suitable,

```
class FE {  
    public:  
    FE() {}  
    ~FE() {}  
  
    virtual void set_polygon(Polygon& p); // Set domain  
    virtual Polygon& get_polygon(); // Get polygonal domain
```



```
virtual ex N(int i);           // The i'th basis function
virtual ex dof(int i);        // The i'th degree of freedom
virtual int nbf();            // The number of basis
                              // functions/degrees of
                              // freedom
};
```

The usage of a finite element is as follows (see `fe_ex1.cpp` where Lagrangian elements are used),

```
ex Ni;
ex gradNi;
ex dofi;
for (int i=0; i< fe.nbf(); i++) {
    Ni = fe.N(i);
    gradNi = grad(Ni);
    dofi = fe.dof(i);
    cout <<"The basis function, N("<<i<<"="<<Ni<<endl;
    cout <<"The gradient of N("<<i<<"="<<gradNi<<endl;
    cout <<"The corresponding dof, L("<<i<<"="<<dofi<<endl;
}
```

When you run `fe_ex1`, it produces the following output:

```
The basis function, N(1)=2*y^2-y
The gradient of N(1)=[[0],[-1+4*y]]
The corresponding dof, L(1)={0,1}
The basis function, N(2)=4*y*x
The gradient of N(2)=[[4*y],[4*x]]
The corresponding dof, L(2)={1/2,1/2}
The basis function, N(3)=2*x^2-x
The gradient of N(3)=[[-1+4*x],[0]]
The corresponding dof, L(3)={1,0}
The basis function, N(4)=-4*y*x-4*y^2+4*y
The gradient of N(4)=[[-4*y],[4-8*y-4*x]]
```

```

The corresponding dof, L(4)={0,1/2}
The basis function, N(5)=-4*y*x-4*x^2+4*x
The gradient of N(5)=[[4-4*y-8*x],[-4*x]]
The corresponding dof, L(5)={1/2,0}
The basis function, N(6)=1+4*y*x+2*x^2+2*y^2-3*y-3*x
The gradient of N(6)=[[-3+4*y+4*x],[-3+4*y+4*x]]
The corresponding dof, L(6)={0,0}

```

The computation of the element matrix for a Poisson problem is as follows (see `fe_ex2.cpp`),

```

Triangle T(1st(0,0), 1st(1,0), 1st(0,1), "t");
int order = 2;

std::map<std::pair<int,int>, ex> A;
std::pair<int,int> index;
LagrangeFE fe;
fe.set_order(order);
fe.set_polygon(T);
fe.compute_basis_functions();
for (int i=0; i< fe.nbf(); i++) {
    index.first = i;
    for (int j=0; j< fe.nbf(); j++) {
        index.second = j;
        ex nabla = inner(grad(fe.N(i)), grad(fe.N(j)));
        ex Aij = T.integrate(nabla);
        A[index] = Aij;
    }
}

```

Here, we have used the class `LagrangeFE`, which is a subclass of `FE`, that implements Lagrangian elements of arbitrary order. The construction of this element is described later in Section [4.1.1](#).

3.5 Degrees of Freedom

As we have seen earlier, for each element e , we have a local set of degrees of freedom L_i^e , which in general are linear forms on the polynomial space. Degrees of freedom and linear forms are quite general concepts, but the reader not familiar with this general definition can think of them for instance as nodal values at vertices, i.e.,

$$L_i(v) = v(\mathbf{x}_i).$$

Another example is the integral of v over an edge (or a face), e_i , of the polygon,

$$L_i(v) = \int_{e_i} v \, ds.$$

The most important thing with the degrees of freedom, besides defining a basis for the polynomial space, is that they provide a mapping from the local degree of freedom, L_i^e , on a given element, e , to the global degree of freedom, L_j . This mapping does in turn provide the mapping between the element matrices/vectors and the global matrix/vector. Hence, we have the following mapping,

$$(e, i) \rightarrow L_i^e \rightarrow L_j \rightarrow j. \quad (3.33)$$

Here e , i , and j are integers, while L_i^e and L_j are degrees of freedom (or linear forms). Additionally, given a global degree of freedom we have a mapping to the local degrees of freedom,

$$j \rightarrow L_j \rightarrow L_{i(e)}^e_{e \in E(j)} \rightarrow (e, i(e))_{e \in E(j)}. \quad (3.34)$$

Here $E(j)$ is the set of elements sharing the degree of freedom L_j .

Software Component: Degrees of Freedom Handler

A degree of freedom, local or global, is well represented as an `ex` (in fact `ex` is more general than a linear form). Hence, to implement proper tools for handling the degrees of freedom, we only need to provide the mappings (3.33) and (3.34). We have implemented a class `Dof` which provides these mappings,

```

class Dof {
protected:
    int counter;
    // the structures loc2dof, dof2index, and doc2loc
    // are completely dynamic. They are all initialized
    // and updated by insert_dof(int e, int d, ex dof)

    // (int e, int i) -> ex Li
    map<pair<int,int>, ex>          loc2dof;
    // (ex Lj)      -> int j
    map<ex,int,ex_is_less>        dof2index;
    // (int j)      -> ex Lj
    map<int,ex>                   index2dof;
    // (ex Lj)      -> vector< pair<e1, i1>, .. pair<en, in> >
    map <ex, vector<pair<int,int> >,ex_is_less >  dof2loc;

public:
    Dof() { counter = 0; }
    ~Dof() {}
    int insert_dof(int e, int j, ex Lj); // update internal
                                         // structures

    // Helper functions to be used when the dofs have been set.
    // These do not modify the internal structure
    int glob_dof(int e, int j);
    int glob_dof(ex Lj);
    ex  glob_dof(int j);
    int size();
    vector<pair<int, int> > glob2loc(int j);
    void clear();
};

```

Here, the function `int insert_dof(int e, int i, ex Li)` creates the various mappings between the local dof L_i^e , in element e , and the global dof L_j . This is the only function for initializing the mappings. After the mappings have been initialized, they can be used as follows,

- `int glob_dof(int e, int i)` is the mapping $(e, i) \rightarrow j$,
- `int glob_dof(ex Lj)` is the mapping $L_j \rightarrow j$,
- `ex glob_dof(int j)` is the mapping $j \rightarrow L_j$,
- `vector<pair<int, int> > glob2loc(int j)` is the mapping $j \rightarrow (e, i(e))$.

The following code shows how to make two Lagrangian elements, implemented by the class `LagrangeFE` (The description of `LagrangeFE` is postponed until Section `sec:fem:examples`), assign their local degrees of freedom to the global set of degrees of freedom in `Dof`, and print out the local degrees of freedom associated with each global degree of freedom (see also `dof_ex.cpp`):

```
Dof dof;

Triangle t1(lst(0,0), lst(1,0), lst(0,1));
Triangle t2(lst(1,1), lst(1,0), lst(0,1));

// Create a finite element and corresponding
// degrees of freedom on the first triangle
int order = 2;
LagrangeFE fe;
fe.set_order(order);
fe.set_polygon(t1);
fe.compute_basis_functions();
for (int i=0; i< fe.nbf(); i++) {
    cout <<"fe.dof("<<i<<"= "<<fe.dof(i)<<endl;
    // insert local dof in global set of dofs
    dof.insert_dof(1,i, fe.dof(i));
}

// Create a finite element and corresponding
// degrees of freedom on the second triangle
fe.set_polygon(t2);
fe.compute_basis_functions();
for (int i=0; i< fe.nbf(); i++) {
    cout <<"fe.dof("<<i<<"= "<<fe.dof(i)<<endl;
```

```

    // insert local dof in global set of dofs
    dof.insert_dof(2,i, fe.dof(i));
}

// Print out the global degrees of freedom an their
// corresponding local degrees of freedom
vector<pair<int,int> > vec;
pair<int,int> index;
ex exdof;
for (int i=1; i<= dof.size(); i++) {
    exdof = dof.glob_dof(i);
    vec = dof.glob2loc(i);
    cout <<"global dof " <<i<<" dof "<<exdof<<endl;
    for (int j=0; j<vec.size(); j++) {
        index = vec[j];
        cout <<"    element "<<index.first<<
            "    local dof "<<index.second<<endl;
    }
}
}

```

In the previous example, the reader that also runs the companion code will notice that the degrees of freedom in `LagrangeFE` are not linear forms on polynomial spaces, i.e.,

$$L_i(v) = v(\mathbf{x}_i).$$

They are instead represented as points, \mathbf{x}_i , which is the usual way to represent these degrees of freedom in finite element software (because of their obvious simplicity compared to linear forms on polynomial spaces). Hence, the degrees of freedom in `LagrangeFE` are actually implemented in the standard fashion. However, the tools we have described are far more general than conventional finite element codes. Still the tools are equally simple to use, due to the powerful expression class `ex` in GiNaC.

Our next example concerns degrees of freedom which are line integrals over the edges of triangles. Let T be a triangle with the edges e_i , $i \in [1, 3]$. The degree of freedom associated with e_i is then simply,

$$L_i(v) = \int_{e_i} v \, ds.$$

As our next example shows, such degrees of freedom can be implemented equally easy as the point values shown in the previous example (see `dof_ex2.cpp`):

```
Dof dof;

// create two triangles
Triangle t1(lst(0,0), lst(1,0), lst(0,1));
Triangle t2(lst(1,1), lst(1,0), lst(0,1));

// create the polynomial space
ex Nj = pol(1,2,"a");
cout <<"Nj " <<Nj<<endl;
Line line;
ex dofi;

// dofs on first triangle
for (int i=1; i<= 3; i++) {
    line = t1.line(i);          // pick out the i'th line
    dofi = line.integrate(Nj);  // create the dof which is
                                // a line integral
    dof.insert_dof(1,i, dofi); // insert local dof in
                                // global set of dofs
}

// dofs on second triangle
for (int i=1; i<= 3; i++) {
    line = t2.line(i);          // pick out the i'th line
    dofi = line.integrate(Nj);  // create the dof which is
                                // a line integral
    dof.insert_dof(2,i, dofi); // insert local dof in
                                // global set of dofs
}
```

Software Component: Degrees of Freedom Handler Template

We will also describe an equally general degree of freedom handler which is not based on GiNaC, but which employs templates instead. This template class relies on two classes, the degree of freedom `D` and a comparison function. The rest is basically identical to the previously described `Dof`, except that we have added two boolean variables which can be used to turn off the computation of the global to local mapping in (3.34) and the $j \rightarrow N_j$ mapping. This class can be found in the header file `DofT.h`:

```
template <class D, class C>
class DofT {
protected:
    bool create_index2dof, create_dof2loc;
    int counter;
    // the structures loc2dof, dof2index, and doc2loc are
    // completely dynamic. They are all initialized and
    // updated by insert_dof(int e, int i, ex Li).

    // (int e, int i) -> int j
    map<pair<int,int>, int>          loc2dof;
    // (ex Lj)      -> int j
    map<D,int,C>                   dof2index;
    typename map<D,int,C>::iterator iter;

    // (int j)      -> ex Lj
    map<int,D>                  index2dof;
    // (ex j)      -> vector< pair<e1, i1>, .. pair<en, in> >
    map <int, vector<pair<int,int> > > dof2loc;

public:
    DofT(bool create_index2dof_ = false,
         bool create_dof2loc_   = false)
    {
        counter = -1;
        create_index2dof = create_index2dof_;
        create_dof2loc   = create_dof2loc_;
    }
};
```



```
}
~DofT() {}
int insert_dof(int e, int i, D Li); // update internal
                                   // structures

// Helper functions to be used when the dofs have been set.
// These do not modify the internal structure.
int glob_dof(int e, int i);
int glob_dof(D Lj);
D glob_dof(int j);
int size();
vector<pair<int, int> > glob2loc(int j);
void clear();
};
```

The typical way to represent most common degrees of freedom is as points. Hence, we have implemented a simple point class `Ptv` and its comparison function. The header file (see also `Ptv.h`) is as follows:

```
class Ptv {
private:
    int dim;
    double* v;
    static double tol;

public:
    Ptv(int size_);
    Ptv(int size_, double* v_);
    Ptv(const Ptv& p);
    Ptv();

    virtual ~Ptv();

    const int size() const;

    const double& operator [] (int i) const;
```

```

double& operator [] (int i);
Ptv& operator = (const Ptv& p);

bool is_less(const Ptv& p) const;

};

struct Ptv_is_less :
    public std::binary_function<Ptv, Ptv, bool> {
    bool operator() (const Ptv &lh, const Ptv &rh) const {
        return lh.is_less(rh);    }
};

std::ostream & operator<< ( std::ostream& os, const Ptv& p);

```

The `Ptv` class simply contain an array of doubles with variable size. The comparison function should check whether a point $x \in \mathbb{R}^n$ is less than $y \in \mathbb{R}^m$, which is not necessarily obvious how to do. For instance, which is the smallest of $x_1 = (1, 0) \in \mathbb{R}^2$, $x_2 = (0, 1) \in \mathbb{R}^2$ and $x_3 = (0, 0, 1) \in \mathbb{R}^3$? There are many possible ways to compare points. The convention we have chosen so far is to first check the size of the points. Hence, $x < y$, where $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$, if $n < m$. If $n = m$, then $x < y$ if $x_0 < y_0$. If $x_0 = y_0$, then $x < y$ if $x_1 < y_1$ and we continue in this fashion, if $x_j = y_j, 0 \leq j < i$ then $x < y$ if $x_i < y_j$. Notice that this comparison operator only affects the ordering of the degrees of freedom internally in the STL `map` structure. But it might be that other comparison conventions will speed up the search and insert routines in `map`.

Finally, we remark that the `Ptv` class and `DofT` can be used also for degrees of freedom associated with lines, edges, faces or general polygons. For instance the edge of a 2D triangle, between the points $\mathbf{x}_0 = (x_0, y_0)$ and $\mathbf{x}_1 = (x_1, y_1)$ can be represented as a point in \mathbb{R}^4 , e.g., (x_0, y_0, x_1, y_1) if $\mathbf{x}_0 < \mathbf{x}_1$ and (x_1, y_1, x_0, y_0) otherwise. Another simpler approach is to represent an edge by its midpoint.

For degrees of freedom that are not well represented as points we have created the structure `OrderedPtvSet`. This class contains a ordered set of points

Ptv, where the ordering is determined by the `Ptv::less` function. The class declaraton is as follows (see also `OrderedPtvSet.h`):

```
class OrderedPtvSet
{
    vector<Ptv> PtvS;

public:
    OrderedPtvSet();
    OrderedPtvSet(const Ptv& p0, const Ptv& p1);
    OrderedPtvSet(const Ptv& p0, const Ptv& p1, const Ptv& p2);
    OrderedPtvSet(const Ptv& p0, const Ptv& p1, const Ptv& p2,
                  const Ptv& p3);
    virtual ~OrderedPtvSet();

    void append(const Ptv& p);
    int size() const;
    const Ptv& operator [] (int i) const;
    Ptv& operator [] (int i);
    OrderedPtvSet& operator = (const OrderedPtvSet& p);
    bool less(const OrderedPtvSet& s) const;
};
```


Chapter 4

Some Examples of Finite Elements

Earlier in Section 3.4, we described the usage of a general finite element. In this section we will show how various finite elements are constructed/implemented in SyFi.

4.1 Finite Elements in H^1

4.1.1 The Lagrangian Element

We will describe the construction of a Lagrangian element on a 2D triangle. The actual implementation of the element in both 1D, 2D and 3D can be found in the class `LagrangeFE`.

As we saw in Section 3.3, the polynomial space \mathbb{P}_n in 2D can be written on the form

$$N = \sum_{i,j=0}^{i+j \leq n} a_{ij} x^i y^j.$$

Hence, to determine the basis functions N_k we simply represented them in

abstract form,

$$N_k = \sum_{i,j=0}^{i+j \leq n} a_{ij}^k x^i y^j.$$

Then the coefficients a_{ij}^k are to be determined by the $(n+1)(n+2)/2$ degrees of freedom that are the nodal values at the the points \mathbf{x}_i , i.e.,

$$L_i(N_k) = N_k(\mathbf{x}_i).$$

Hence, we need a set of $(n+1)(n+2)/2$ nodal points to determine the coefficients a_{ij}^k for each basis function. We have chosen to use the Bezier ordinates. When this is done, it is simply a matter of solving the linear system

$$L_i(N_k) = N_k(\mathbf{x}_i) = \delta_{ik},$$

for each basis function N_k .

Software Component: The Lagrangian Element

The Lagrangian element is implemented as a subclass of `StandardFE`. The class definition is:

```
class LagrangeFE : public StandardFE {
public:
    LagrangeFE() {}
    virtual ~LagrangeFE() {}

    virtual void set_order(int order);
    virtual void set_polygon(Polygon& p);
    virtual void compute_basis_functions();
    virtual int nbf();
    virtual GiNaC::ex N(int i);
    virtual GiNaC::ex dof(int i);
};
```

The Construction of the Lagrangian Element

The Lagrangian element of arbitrary order in 1D, 2D, and 3D, is implemented in `LagrangeFE.cpp`. The following code is taken from `fe_ex3.cpp`.

```
Triangle t(lst(0,0), lst(1,0), lst(0,1));
int order = 2;                      //second order elements
ex polynom;
lst variables;

// the polynomial spaces on the form:
// first item, the polynom:
//      a0 + a1*x + a2*y + a3*x^2 + a4*x*y ...
// second item, the coefficients:
//      a0, a1, a2, ...
// third item, the basis:
//      1, x, y, x^2
// Could also do:
// GiNaC::ex polynom_space = bernstein(order, t, "a");
ex polynom_space = pol(order, 2, "a");
ex polynom = polynom_space.op(0);

// the variables a0,a1,a2 ..
variables = ex_to<lst>(polynom_space.op(1));

ex Nj;
// The Bezier ordinates in which the
// basis function should be either 0 or 1
lst points = bezier_ordinates(t,order);

// Loop over all basis functions Nj and all points.
// Each basis function Nj is determined
// by a set of linear equations:
//      Nj(xi) = dirac(i,j)
// This system of equations is then solved by lsolve
for (int j=1; j <= points.nops(); j++) {
    lst equations;
```

```

int i=0;
for (int i=1; i<= points.nops() ; i++ ) {
    // The point xi
    ex point = points.op(i-1);
    // The equation Nj(x) = dirac(i,j)
    ex eq = polynom == dirac(i,j);
    // Substitute x = xi and y = yi and
    // appended the equation to the list of equations
    // to the list of equations
    equations.append(eq.subs(lst(x == point.op(0) ,
                                y == point.op(1))));
}

// We solve the linear system
ex subs = lsolve(equations, variables);
// Substitute to get the Nj
Nj = polynom.subs(subs);
cout <<"Nj " <<Nj<<endl;
}

```

In this example the degrees of freedom are very simple. It is only a matter of evaluating the function v_k in the point \mathbf{x}_i (which in GiNaC is performed by substitution). Later we will see that more advanced degrees of freedom are readily available since we have stored the degrees of freedom as a set of `exes`.

4.1.2 The Crouizex-Raviart Element

The Crouizex-Raviart element [?] is the nonconforming equivalent of linear continuous Lagrangian elements. The degrees of freedom are the values at the midpoint of the sides, i.e.,

$$L_i(v) = v(x_{m(e_i)}),$$

where $x_m(e_i)$ is the midpoint on the edge, e_i . An equivalent definition of the degrees of freedom is,

$$L_i(v) = \int_{e_i} v \, ds,$$

This is the definition we will use.

Software Component: The Crouzeix-Raviart Element

The Crouzeix-Raviart class definition is similar to class defined for the Lagrangian element:

```
class CrouzeixRaviart : public StandardFE {
public:
  CrouzeixRaviart();
  virtual ~CrouzeixRaviart() {}

  void set_order(int order);
  void set_polygon(Polygon& p);
  void compute_basis_functions();
  virtual int nbf();
  virtual GiNaC::ex N(int i);
  virtual GiNaC::ex dof(int i);
};
```

The Construction of the Crouzeix-Raviart Element

The following code, which is from the file `CrouzeixRaviart.cpp`, shows how this element can be defined in 2D. The definition of the element in 3D can also be found in this file.

```
Triangle triangle;

// create the polynomial space
```

```
ex polynom_space = bernstein(1, triangle, "a");
ex polynom = polynom_space.op(0);
ex variables = polynom_space.op(1);
ex basis = polynom_space.op(2);

// create the dofs
int counter = 0;
symbol t("t");
for (int i=1; i<= 3; i++) {
    Line line = triangle.line(i);
    ex dofi = line.integrate(polynom);
    dofs.insert(dofs.end(),dofi);
}

// solve the linear system to compute
// each of the basis functions
for (int i=1; i<= 3; i++) {
    lst equations;
    for (int j=1; j<= 3; j++) {
        equations.append(dofs[j-1] == dirac(i,j));
    }
    ex sub = lsolve(equations, variables);
    ex Ni = polynom.subs(sub);
    Ns.insert(Ns.end(),Ni);
}
```

This element can be used in a standard fashion, (see also `crouzeixraviart_ex.cpp`),

```
CrouzeixRaviart fe;
fe.set_order(1);
fe.set_polygon(p);
fe.compute_basis_functions();
for (int i=0; i< fe.nbf(); i++) {
    cout <<"fe.N("<i<<")="<<fe.N(i)<<endl;
}
```

See also the Python implementation of this element in [Section 7](#).

4.2 Finite Elements in L^2

4.2.1 The P_0 Element

The P_0 element consists of piecewise constants, i.e.,

$$v|_T = 1,$$

where T is the polygon. This element is discontinuous across elements.


Software Component: The P_0 Element

The P_0 element is implemented in the class `P0`. The implementation is straightforward.

4.2.2 The Discontinuous Lagrangian Element

The discontinuous Lagrangian elements are similar to the continuous Lagrangian elements except for the fact that they are discontinuous. Hence, locally on the polygon T , the basis functions are the same. The difference is that discontinuous Lagrangian elements are not continuous between elements.

To exemplify this we consider the continuous and the discontinuous linear Lagrangian elements in 2D. In [Figure 4.1](#) we see that the triangles $1, \dots, 5$ all share the common vertex V . For continuous Lagrangian elements, this means that there will be only one degree of freedom associated with V . On the other hand, for discontinuous Lagrangian elements, there will be one degree of freedom associated with V per triangle. Hence, in the concrete case depicted in [Figure 4.1](#), there will be 5 degrees of freedom associated

Figure 4.1: Some triangles with the common vertex V .

with V . Each degree of freedom is associated with a basis function which is 1 in V , 0 in the other vertices, and zero outside the triangle.

Software Component: The discontinuous Lagrangian element

The implementation of the discontinuous Lagrangian element is really easy because this element is identical to the continuous Lagrangian element locally. Hence, the basis functions on each element is the same. We only need to modify the degrees of freedom.

The degrees of freedom for the discontinuous Lagrangian elements are such that for each element, each degree of freedom is new. Hence, none of degrees of freedom are shared among elements. It is fairly easy to implement this. Assume that the polygons in the mesh or the elements in the finite element space are numbered. Then the degree of freedom can be represented by both the vertex x_i and the element number e associated with the polygon T_e ,

$$L_i^e(v) = v|_{T_e}(x_i),$$

where $v|_{T_e}$ means the restriction of v to the polygon T_e . It is important to take the restriction to T_e since v is in general discontinuous in x_i .

We have implemented the discontinuous Lagrangian element as a subclass

of the continuous Lagrangian element, with an additional integer parameter `element` which is the element number. The class declaration is as follows,

```
class DiscontinuousLagrangeFE : public LagrangeFE {
    int element;
public:
    DiscontinuousLagrangeFE();
    ~DiscontinuousLagrangeFE() {}

    virtual void set_order(int order);
    virtual void set_element_number(int element);
    virtual void set_polygon(Polygon& p);
    virtual void compute_basis_functions();
    virtual int nbf();
    virtual GiNaC::ex N(int i);
    virtual GiNaC::ex dof(int i);
};
```

Earlier, the degrees of freedom for continuous Lagrangian elements were represented as vertices or points (instead of linear forms), as is usual in finite element codes. We do the same simplification here, and store the degrees of freedom as (x_i, e) , where x_i is the vertex/point and e is the element number associated with T_e . This is implemented in the functions `compute_basis_functions`:

```
void DiscontinuousLagrangeFE:: compute_basis_functions() {
    LagrangeFE:: compute_basis_functions();
    for (int i=0; i< dofs.size(); i++) {
        dofs[i] = lst(dofs[i], element);
    }
}
```

The usage is standard (see `disconlagrange_ex.cpp`),

```
Dof dof;
// create two triangles
```

```
Triangle t1(lst(0,0), lst(1,0), lst(0,1));
Triangle t2(lst(1,1), lst(1,0), lst(0,1));

int order = 2;

DiscontinuousLagrangeFE fe;
fe.set_order(order);
fe.set_polygon(t1);
fe.set_element_number(1);
fe.compute_basis_functions();
usage(fe);
for (int i=0; i< fe.nbf(); i++) {
    dof.insert_dof(1,i,fe.dof(i));
}

fe.set_polygon(t2);
fe.set_element_number(2);
fe.compute_basis_functions();
usage(fe);
for (int i=0; i< fe.nbf(); i++) {
    dof.insert_dof(2,i,fe.dof(i));
}

// Print out the global degrees of freedom and their
// corresponding local degrees of freedom
vector<pair<int,int> > vec;
pair<int,int> index;
ex exdof;
for (int i=1; i<= dof.size(); i++) {
    exdof = dof.glob_dof(i);
    vec = dof.glob2loc(i);
    cout <<"global dof " <<i<<" dof "<<exdof<<endl;
    for (int j=0; j<vec.size(); j++) {
        index = vec[j];
        cout <<"  element " <<index.first<<
            " local dof " <<index.second<<endl;
    }
}
```

}

When this program (`disconlagrange`) runs, it prints out 12 degrees of freedom in contrast to 9 which it would be for continuous Lagrangian elements.

4.3 Finite Elements in $H(\text{div})$

4.3.1 The Raviart-Thomas Element

The family of Raviart-Thomas elements [?] is popular when considering the mixed formulation of elliptic problems. In this case the polynomial space is not \mathbb{P}_n^d , but

$$\mathbb{P}_n^d + \mathbf{x}\mathbb{P}_n. \quad (4.1)$$

And the degrees of freedom are,

$$\int_{e_i} \mathbf{v} \cdot \mathbf{n} p_k ds, \quad \forall p_k \in \mathbb{P}_k(e_i), \quad (4.2)$$

$$\int_T \mathbf{v} \cdot \mathbf{p}_{k-1} dx, \quad \forall p_{k-1} \in \mathbb{P}_{k-1}^d(T), \quad (4.3)$$

where T is the polygon domain and e_i is its edges (in 2D) or faces (in 3D). Degrees of freedom which are integrals have been dealt with already for the Crouizex-Raviart element in Section 4.1.2. Hence, there are mainly two new concepts we need to deal with to implement this element. It is the polynomial space, which is on the form (4.1), and the polynomial spaces on faces or edges of the polygon, as in (4.2). Both concepts will be dealt with below.

Software Component: The Raviart-Thomas Element

Notice that for the previously defined Lagrangian and Crouizex-Raviart elements, the basis functions were scalar functions. The basis functions of the Raviart-Thomas elements are vector functions, but still, thanks to the general `ex` class, the Raviart-Thomas element class can be defined in the same way as earlier. The class definition is:

```

class RaviartThomas : public StandardFE {
public:
    RaviartThomas() {}
    virtual ~RaviartThomas() {}

    virtual void set_order(int order);
    virtual void set_polygon(Polygon& p);
    virtual void compute_basis_functions();
    virtual int nbf();
    virtual GiNaC::ex N(int i);
    virtual GiNaC::ex dof(int i);
};

```

The Construction of the Raviart-Thomas Element

First, we described how to make the polynomial space (4.1). The polynomial spaces, $\mathbb{P}_n(T)$ and $\mathbb{P}_n^d(T)$ on a polygonal domain, can be made by the functions `bernstein` and `bernsteinv`, respectively. However, we can not just add the spaces $\mathbb{P}_n^d(T)$ and $\mathbf{x}\mathbb{P}_n(T)$ together. Because, some of the basis functions are the same in both space, while others are not. Consider for instance $\mathbb{P}_1^d(T)$, which has the basis functions,

$$(0, 1)^T, (1, 0)^T, (x, 0)^T, (0, x)^T, (y, 0)^T, (0, y)^T$$

while $\mathbf{x}\mathbb{P}_1(K)$ has the following basis functions

$$(x, 0)^T, (x^2, 0)^T, (xy)^T, (0, y)^T, (0, y^2)^T, (0, xy)^T.$$

Hence $(x, 0)^T$ and $(0, y)^T$ are common.

The way we solve this problem is that we create the two spaces $\mathbb{P}_n^d(T)$ and $\mathbf{x}\mathbb{P}_n(T)$ independently. We then have two polynomial spaces, each with two independent sets of variables (or degrees of freedom). The variables associated with a basis in $\mathbf{x}\mathbb{P}_n(T)$ which is also a basis in $\mathbb{P}_n^d(T)$ is then removed. This is done by removing all variables associated with basis functions that have degree less than $n - 1$ in \mathbb{P}_n from $\mathbf{x}\mathbb{P}_n(T)$. This is done as follows

in 2D (both 2D and 3D elements of arbitrary order are implemented in `RaviartThomas.cpp`),

```
Triangle& triangle = (Triangle&)(*p);
lst equations;
lst variables;
ex polynom_space1 = bernstein(order-1, triangle, "a");
ex polynom1 = polynom_space1.op(0);
ex polynom1_vars = polynom_space1.op(1);
ex polynom1_basis = polynom_space1.op(2);

lst polynom_space2 = bernsteinv(order-1, triangle, "b");
ex polynom2 = polynom_space2.op(0).op(0);
ex polynom3 = polynom_space2.op(0).op(1);

lst pspace = lst( polynom2 + polynom1*x,
                  polynom3 + polynom1*y);

// remove multiple dofs
if ( order >= 2 ) {
  ex expanded_pol = expand(polynom1);
  for (int c1=0; c1<= order-2;c1++) {
    for (int c2=0; c2<= order-2;c2++) {
      for (int c3=0; c3<= order-2;c3++) {
        if ( c1 + c2 + c3 <= order -2 ) {
          ex eq = expanded_pol.coeff(x,c1)
                    .coeff(y,c2).coeff(z,c3);
          if ( eq != numeric(0) ) {
            equations.append(eq == 0);
          }
        }
      }
    }
  }
}
```

Second, we described how to implement the degrees of freedom (4.2)-(4.3). The degrees of freedom associated with the edges,

$$\int_{e_i} \mathbf{v} \cdot \mathbf{n} p_k ds, \forall p_k \in \mathbb{P}_k(e_i),$$

are implemented as follows (Notice that the polynomial space on the edges of the triangle is made by creating Bernstein polynomials in standard fashion).

```
ex bernstein_pol;

int counter = 0;
symbol t("t");
ex dofi;
// loop over all edges
for (int i=1; i<= 3; i++) {
    Line line = triangle.line(i);
    lst normal_vec = normal(triangle, i);
    bernstein_pol = bernstein(order-1, line, istr("a",i));
    ex basis_space = bernstein_pol.op(2);
    ex pspace_n = inner(pspace, normal_vec);

    // loop over all basis functions on current edge
    ex basis;
    for (int i=0; i< basis_space.nops(); i++) {
        counter++;
        basis = basis_space.op(i);
        ex integrand = pspace_n*basis;
        dofi = line.integrate(integrand);
        dofs.insert(dofs.end(), dofi);
        ex eq = dofi == numeric(0);
        equations.append(eq);
    }
}
```

The degrees of freedom associated with the whole triangle,

$$\int_T \mathbf{v} \cdot \mathbf{p}_{k-1} dx, \forall p_{k-1} \in \mathbb{P}_{k-1}^d(T),$$

is implemented as

```
// dofs related to the whole triangle
lst bernstein_polv;
if ( order > 1) {
  counter++;
  bernstein_polv = bernsteinv(order-2, triangle, "a");
  ex basis_space = bernstein_polv.op(2);
  for (int i=0; i< basis_space.nops(); i++) {
    lst basis = ex_to<lst>(basis_space.op(i));
    ex integrand = inner(pspace, basis);
    dofi = triangle.integrate(integrand);
    dofs.insert(dofs.end(), dofi);
    ex eq = dofi == numeric(0);
    equations.append(eq);
  }
}
```

In the above code we have formed the linear system,

$$L_i(v) = 0$$

To compute the different v_j we then produce different right hand sides corresponding to δ_{ij} and solve the system. How this is done can be seen in the `RaviartThomas.cpp`.

4.3.2 The Nedelec element of second kind

The Nedelec $\mathbf{H}(\text{div})$ element introduced in [?], is very similar to the Raviart-Thomas element, except that the polynomial space is \mathbb{P}_n^d instead of $\mathbb{P}_n^d + \mathbf{x}\mathbb{P}_n$. Hence, it is the \mathbb{R}^3 analog of the Brezzi-Douglas-Marini element [?]. The degrees of freedom in the Nedelec $\mathbf{H}(\text{div})$ element are,

$$\begin{aligned} \int_f (\mathbf{p} \cdot \mathbf{n}) ds, \quad \forall \mathbf{p} \in \mathbb{P}_k(f), \\ \int_K (\mathbf{p} \cdot \mathbf{q}) dx, \quad \forall \mathbf{q} \in \mathbb{R}_{k-1}. \end{aligned}$$

Here

$$\begin{aligned}\mathbb{R}_k &= (\mathbb{P}_{k-1}^3) \oplus \mathbb{S}^k, \\ \mathbb{S}^k &= \{\mathbf{p} \in \mathbb{H}_k^3 \mid (\mathbf{r} \cdot \mathbf{p}) = 0\},\end{aligned}$$

where $\mathbf{r} = (x, y, z)$.

Software Component: The Nedelec $\mathbf{H}(\text{div})$ Element

The Nedelec $\mathbf{H}(\text{div})$ element class definition is similar to the previous element definitions.

```
class Nedelec2Hdiv : public StandardFE {
public:
    Nedelec2Hdiv() {}
    virtual ~Nedelec2Hdiv() {}

    virtual void set_order(int order);
    virtual void set_polygon(Polygon& p);
    virtual void compute_basis_functions();
    virtual int nbf();
    virtual GiNaC::ex N(int i);
    virtual GiNaC::ex dof(int i);
};
```

The Construction of the Nedelec $\mathbf{H}(\text{div})$ Element

The construction of this element is very similar to the construction of the Raviart–Thomas element. We will therefore not discuss this here.

4.4 Finite Elements in $H(\text{div}, \mathbb{M})$

The Arnold, Falk and Winther element [?] for mixed elasticity problems in 3D with weak symmetry, has recently been added to SyFi. This element consists of basis functions which take values in \mathbb{M} , which is the space of 3×3 matrices. Each row is either a null row or the Nedgelec $H(\text{div})$ element of second kind as described in Section 4.3.2. The implementation is straightforward, since it is essentially a loop where each row is created as the basis functions a Nedgelec element. We therefore do not comment the implementation details here. The finite element is defined in `ArnoldFalkWintherWeakSym.h`.

4.5 A Finite Element in Both $H(\text{div})$ and H^1

In [?] an element for both Darcy and Stokes types of flow was introduced. The element is defined as:

$$\mathbf{V}(T) = \{\mathbf{v} \in \mathbb{P}_3^2 : \text{div } \mathbf{v} \in \mathbb{P}_0, (\mathbf{v} \cdot \mathbf{n}_e)|_e \in \mathbb{P}_1 \forall e \in E(T)\},$$

where T is a given triangle, $E(T)$ is the edges of T , \mathbf{n}_e is the normal vector on edge e , and \mathbb{P}_k is the space of polynomials of degree k and \mathbb{P}_k^d the corresponding vector space. The degrees of freedom are,

$$\begin{aligned} \int_e (\mathbf{v} \cdot \mathbf{n}) \tau^k d\tau, \quad k = 0, 1, & \quad \forall e \in E(T), \\ \int_e (\mathbf{v} \cdot \mathbf{t}) d\tau, & \quad \forall e \in E(T). \end{aligned}$$

This element is implemented as follows (see also the PARA06 proceeding `../para06/proceeding/para_proceeding.pdf`). First we create the polynomial space, which consist of cubic vector functions, \mathbb{P}_3^2

```
Triangle triangle
ex V_space = bernsteinv(2, 3, triangle, "a");
ex V_polynomial = V_space.op(0);
ex V_variables = V_space.op(1);
```

Here `v_space` is the above mentioned list, `v_polynomial` contains the polynomial, and `v_variables` contains the variables.

In the second step we first specify the constraint $\text{div } \mathbf{v} \in \mathbb{P}_0$:

```
lst equations;
ex divV = div(V);
ex_ex_map b2c = pol2basisandcoeff(divV);
ex_ex_it iter;
// div constraints:
for (iter = b2c.begin(); iter != b2c.end(); iter++) {
    ex basis = (*iter).first;
    ex coeff= (*iter).second;
    if ( coeff != 0 && ( basis.degree(x) > 0
|| basis.degree(y) > 0 ) ) {
        equations.append( coeff == 0 );
    }
}
```

Here, the divergence is computed with the `div` function. The divergence of a function in \mathbb{P}_3^2 is in \mathbb{P}_2 . Hence, it is on the form $b_0 + b_1x + b_2y + b_3xy + b_4x^2 + b_5y^2$. In the above code we find the coefficients b_i , as expressions involving the above mentioned variables a_i and the corresponding polynomial basis, with the function `pol2basisandcoeff`. Then we ensure that the only coefficient which is not zero is b_0 .

The next constraints $(\mathbf{v} \cdot \mathbf{n}_e)|_e \in \mathbb{P}_1$ are implemented in much of the same way as the divergence constraint. We create a loop over each edge e of the triangle and multiply \mathbf{v} with the normal \mathbf{n}_e . Then we substitute the expression for the edge, i.e., in mathematical notation $|_e$, into $\mathbf{v} \cdot \mathbf{n}$. After substituting the expression for these lines to get $(\mathbf{v} \cdot \mathbf{n}_e)|_e$, we check that the remaining polynomial is in \mathbb{P}_1 in the same way as we did above.

```
// constraints on edges:
for (int i=1; i<= 3; i++) {
    Line line = triangle.line(i);
```

```

symbol s("s");
lst normal_vec = normal(triangle, i);
ex Vn = inner(V, normal_vec);
Vn = Vn.subs(line.repr(s).op(0))
      .subs(line.repr(s).op(1));
b2c = pol2basisandcoeff(Vn,s);
for (iter = b2c.begin(); iter != b2c.end(); iter++){
    ex basis = (*iter).first;
    ex coeff= (*iter).second;
    if ( coeff != 0 && basis.degree(s) > 1 )
    {
        equations.append( coeff == 0 );
    }
}
}

```

In the third step we specify the degrees of freedom. First, we specify the equations coming from $\int_e (\mathbf{v} \cdot \mathbf{n}) \tau^k, k = 0, 1$ on all edges. To do this we need to create a loop over all edges, and on each edge we create the space of linear Bernstein polynomials in barycentric coordinates on e , i.e., $\mathbb{P}_1(e)$. Then we create a loop over the basis functions τ^k in $\mathbb{P}_1(e)$ and compute the integral $\int_e (\mathbf{v} \cdot \mathbf{n}) \tau^k d\tau$.

```

// dofs related to the normal on the edges
for (int i=1; i<= 3; i++) {
    Line line = triangle.line(i);
    lst normal_vec = normal(triangle, i);
    ex P1_space = bernstein(1, line, istr("a",i));
    ex P1 = P1_space.op(2);
    ex Vn = inner(V, normal_vec);

    ex basis;
    for (int j=0; j< P1.nops(); j++) {
        basis = P1.op(j);
        ex integrand = Vn*basis;
        ex dofi = line.integrate(integrand);
    }
}

```

```

        dofs.insert(dofs.end(), lst(line.vertex(0),
                                   line.vertex(1), j));
        ex eq = dofi == numeric(0);
        equations.append(eq);
    }
}

```

Finally, the degrees of freedom $\int_e (\mathbf{v} \cdot \mathbf{t}) d\tau$, can be implemented in basically the same fashion as the previously described degrees of freedom. To summarize, we have now specified 20 equations which is precisely the number of unknowns in \mathbb{P}_3^2 . Hence, the space $\mathbf{V}(T)$ is uniquely defined, what remains is simply to solve a linear system with 20 equations and 20 unknowns. The complete source code is in `Robust.cpp`.

4.6 Finite Elements in $\mathbf{H}(\text{curl})$

4.6.1 The Nedelec Element

In electromagnetic applications, [?] the family of Nedelec elements are very common. As was also the case with the Raviart-Thomas elements, \mathbb{P}^n is not the most convenient space to define the basis functions. Instead, we will use

$$\mathbb{P}_{n-1}^d + \hat{\mathbb{H}}^k, \quad (4.4)$$

where

$$\hat{\mathbb{H}}^k = \mathbf{h} \in \mathbb{H}_k^d : \mathbf{h} \cdot \mathbf{x} = 0$$

and \mathbb{H} is the space of homogenous polynomials described in Section 3.3.3. The degrees of freedom that defines the Nedelec elements are (in 2D),

$$\int_e \mathbf{t} \cdot \mathbf{u} p \, dx, \quad \forall p \in \mathbb{P}_{k-1}(e), \quad (4.5)$$

$$\int_T \mathbf{u} \cdot \mathbf{p} \, dx, \quad \forall \mathbf{p} \in \mathbb{P}_{k-2}^n(T). \quad (4.6)$$

Software Component: The Nedelec Element

The Nedelec element class definition is similar to the previous element definitions.

```
class Nedelec : public StandardFE {
public:
    Nedelec() {}
    virtual ~Nedelec() {}

    virtual void set_order(int order);
    virtual void set_polygon(Polygon& p);
    virtual void compute_basis_functions();
    virtual int nbf();
    virtual GiNaC::ex N(int i);
    virtual GiNaC::ex dof(int i);
};
```

The Construction of the Nedelec Element

The Nedelec element of arbitrary order in both 2D and 3D is implemented in `Nedelec.cpp`. Here we will for simplicity describe how the element is implemented in 2D.

We first consider the polynomial space (4.4),

```
// create r
GiNaC::ex R_k = homogenous_polv(2,k+1, 2, "a");
GiNaC::ex R_k_x = R_k.op(0).op(0);
GiNaC::ex R_k_y = R_k.op(0).op(1);

// Equations that make sure that r*x = 0
GiNaC::ex rx = (R_k_x*x + R_k_y*y).expand();
ex_ex_map pol_map = pol2basisandcoeff(rx);
```

```

ex_ex_it iter;
for (iter = pol_map.begin();
     iter != pol_map.end(); iter++) {
    if ((*iter).second != 0 ) {
        equations.append((*iter).second == 0 );
        removed_dofs++;
    }
}

```

The degree of freedom associated with the edges (4.5) are implemented as,

```

GiNaC::ex dofi;
// dofs related to edges
for (int i=1; i<= 3; i++) {
    Line line = triangle.line(i);
    GiNaC::lst tangent_vec = tangent(triangle, i);
    GiNaC::ex bernstein_pol = bernstein(order, line,
                                         istr("a",i));
    GiNaC::ex basis_space = bernstein_pol.op(2);
    GiNaC::ex pspace_t = inner(pspace, tangent_vec);

    GiNaC::ex basis;
    for (int j=0; j< basis_space.nops(); j++) {
        counter++;
        basis = basis_space.op(j);
        GiNaC::ex integrand = pspace_t*basis;
        dofi = line.integrate(integrand);
        dofs.insert(dofs.end(), dofi);
        GiNaC::ex eq = dofi == GiNaC::numeric(0);
        equations.append(eq);
    }
}

```

The degree of freedom associated with whole triangle (4.6) are implemented

as,

```
// dofs related to the whole triangle
GiNaC::lst bernstein_polv;
if ( order > 0) {
    counter++;
    bernstein_polv = bernsteinv(2,order-1, triangle, "a");
    GiNaC::ex basis_space = bernstein_polv.op(2);
    for (int i=0; i< basis_space.nops(); i++) {
        GiNaC::lst basis = GiNaC::ex_to<GiNaC::lst> (
            basis_space.op(i));

        GiNaC::ex integrand = inner(pspace, basis);
        dofi = triangle.integrate(integrand);
        dofs.insert(dofs.end(), dofi);
        GiNaC::ex eq = dofi == GiNaC::numeric(0);
        equations.append(eq);
    }
}
```


Chapter 5

Mixed Finite Elements

Mixed finite element methods typically refer to discretization methods for systems of PDEs where different finite elements are used for the different unknowns. For instance, in incompressible flow problems, one typically has (at least) two unknowns, the velocity \mathbf{v} and the pressure p . It is wellknown that the velocity elements should have higher order than the pressure elements. The reasons for this have been extensively studied the last 30 years, and we will not go into details on this here, see e.g., Brezzi and Fortin [?] and Girault and Raviart[?].

What we will do here is to describe mixed finite elements from the programmers point of view. In this setting, we simply refer to mixed elements as a collection of finite elements of different types on the same polygon. The elements themselves and their implementation were discussed in the previous section.

5.1 The Taylor–Hood and the $\mathbb{P}_n^d - \mathbb{P}_{n-2}$ Elements

The Taylor–Hood and the $\mathbb{P}_n^d - \mathbb{P}_{n-2}$ elements are mixed elements that are popular for incompressible flow. The elements for both the velocity and the pressure are of Lagrangian type, but have different order. The Taylor–Hood

element on a polygon T is,

$$\mathbf{v}(T) \in \mathbb{P}_2^d \quad \text{and} \quad p(T) \in \mathbb{P}_1.$$

The $\mathbb{P}_n - \mathbb{P}_{n-2}$ element on a polygon T is,

$$\mathbf{v}(T) \in \mathbb{P}_n^d \quad \text{and} \quad p(T) \in \mathbb{P}_{n-2}, \quad n \geq 2.$$

For $n > 2$ the pressure element is of Lagrangian type, while for $n=2$ the pressure element is piecewise constant. These elements satisfy the Babuska-Brezzi condition.

The Taylor–Hood elements can be created as follows, (see also `taylorhood.ex.cpp`)

```
VectorLagrangeFE v_fe;
v_fe.set_order(2);
v_fe.set_size(2);
v_fe.set_polygon(domain);
v_fe.compute_basis_functions();

LagrangeFE p_fe;
p_fe.set_order(1);
p_fe.set_polygon(domain);
p_fe.compute_basis_functions();
```

The $\mathbb{P}_n^d - \mathbb{P}_{n-2}$ element can be made by changing the order of the elements with the `set` function.

5.2 The Mixed Crouizex-Raviart Element

The mixed Crouizex-Raviart element is a nonconforming linear element for the velocity and piecewise constant for the pressure. The Crouizex-Raviart element was described in Section 4.1.2, while the P_0 element was described in Section 4.2.1.

These elements can be made as follows (see also `crouzeixraviart.ex2.cpp`)

```
ReferenceTriangle domain;

VectorCrouzeixRaviart v_fe;
v_fe.set_size(2);
v_fe.set_polygon(domain);
v_fe.compute_basis_functions();

P0 p_fe;
p_fe.set_polygon(domain);
p_fe.compute_basis_functions();
```

5.3 The Mixed Raviart-Thomas Element

The velocity element is the Raviart-Thomas element described in Section 4.3.1. The pressure element is discontinuous polynomials of degree n . The \mathbb{P}_0 element is described in Section 4.2.1, while the discontinuous \mathbb{P}_n element is described in Section 4.2.2.

The can be made as such (see also `raviartthomas_ex2`):

```
int order = 3;

ReferenceTriangle triangle("t");
RaviartThomas vfe;
vfe.set_polygon(triangle);
vfe.set_order(order);
vfe.compute_basis_functions();

DiscontinuousLagrangeFE pfe;
pfe.set_polygon(triangle);
pfe.set_order(order);
pfe.compute_basis_functions();
```

```

for (int i=0; i< vfe.nbf(); i++)
    cout <<"vfe.N("<<i<<"")="<<vfe.N(i)<<endl;

for (int i=0; i< pfe.nbf(); i++)
    cout <<"pfe.N("<<i<<"")="<<pfe.N(i)<<endl;

```

5.4 The Mixed Arnold-Falk-Winther element

The mixed method of Arnold, Falk and Winther [?] for mixed elasticity problems in 3D, with weakly imposed symmetry consists of three different elements $\sigma_h \in H(\text{div}, \mathbb{M})$, $\mathbf{u}_h \in L^2(\mathbb{V})$, $\mathbf{p}_h \in L^2(\mathbb{K})$, where \mathbb{M} is the space of 3×3 matrices, \mathbb{V} is the space of vectors in \mathbb{R}^3 , and \mathbb{K} is the space of 3×3 skew symmetric matrices. The σ_h element, described in Section 4.4, is implemented as `ArnoldFalkWintherWeakSymSigma`. The \mathbf{u}_h element is a discontinuous Galerkin vector element implemented as `ArnoldFalkWintherWeakSymU`. Finally, the \mathbf{p}_h element is a discontinuous Galerkin skew symmetric matrix element implemented as `ArnoldFalkWintherWeakSymP`. The implementation of both `ArnoldFalkWintherWeakSymU` and `ArnoldFalkWintherWeakSymP` is straightforward since the code is essentially a wrap around the discontinuous Lagrangian element described in Section 4.2.2 and we will not comment on the implementation details. The code can be found in `ArnoldFalkWintherWeakSym.cpp`.

Chapter 6

Computing Element Matrices

Our next task is to compute element matrices. As earlier, everything will be done symbolically. There are several reasons for doing the computations symbolically:

- Everything is exact (No floating point precision issues)!
- Differentiation of the weak form with respect to the variables is possible (Easy to compute the Jacobian for nonlinear PDEs).
- In case one uses integers and rational numbers as input (e.g., the vertices of the polygon) one gets rational numbers as output. This enables nice output.
- In case one uses symbols as input, one get symbols as output. Hence, one might actually compute an abstract element matrix, where each entry in the matrix is a function of the vertices of the polygon, $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$, which are symbols. We will consider this in more detail later.
- Every step can be checked against analytic computations. We can even, as we will see, produce output in \LaTeX format, for easy reading.
- In Section 8 we generate C++ code from the exactly computed element matrices.

6.1 A Poisson Problem

The Poisson problem is on the form,

$$\begin{aligned} -\Delta u &= f, & \text{in } \Omega, \\ u &= h, & \text{on } \partial\Omega_E, \\ \frac{\partial u}{\partial n} &= g, & \text{on } \partial\Omega_N, \end{aligned}$$

where $\partial\Omega = \partial\Omega_E \cup \partial\Omega_N$.

The weak form of the Poisson problem is (as we have already used): Find $u \in V_h$ such that

$$a(u, v) = b(v), \quad \forall v \in V_0.$$

where,

$$\begin{aligned} a(u, v) &= \int_{\Omega} \nabla u \cdot \nabla v \, dx, \\ b(v) &= \int_{\Omega} f v \, dx + \int_{\Gamma_N} g v \, ds. \end{aligned}$$

and $V_k = \{v \in H^1; v|_{\partial\Omega_E} = k, \text{ for } k = 0, h\}$.

From this weak form we obtain the element matrix, see e.g., Brenner and Scott [?], Ciarlet [?], or Langtangen [?],

$$A_{ij} = a(N_i, N_j) = \int_T \nabla N_j \cdot \nabla N_i \, dx. \quad (6.1)$$

The computation of (6.1) is implemented in the function `compute_Poisson_element_matrix` in `ElementComputations.cpp`,

```
void compute_Poisson_element_matrix(
    FE& fe,
    Dof& dof,
    std::map<std::pair<int,int>, ex>& A)
{
    std::pair<int,int> index;
```

```

    // Insert the local degrees of freedom into the global Dof
    for (int i=0; i< fe.nbf(); i++) {
        dof.insert_dof(1,i,fe.dof(i));
    }

    Polygon& domain = fe.get_polygon();

    // The term (grad u, grad v)
    for (int i=0; i< fe.nbf(); i++) {
        index.first = dof.glob_dof(fe.dof(i));    // fetch the i'th
                                                // global dof

        for (int j=0; j< fe.nbf(); j++) {
            index.second = dof.glob_dof(fe.dof(j)); // fetch the j'th
                                                    // global dof

            ex nabla = inner(grad(fe.N(i)),        // compute the
                             grad(fe.N(j)));      // integrand

            ex Aij = domain.integrate(nabla);      // compute integral
            A[index] += Aij;                       // add to matrix
        }
    }
}

```

Notice that in this example, both the degrees of freedom `dof` and the matrix `A` are global.

This function can be used as follows (see `fe_ex4.cpp`),

```

//matrix in terms of rational numbers
int order = 1;
Triangle triangle(lst(0,0), lst(1,0), lst(0,1));
LagrangeFE fe;
fe.set_order(order);
fe.set_polygon(triangle);
fe.compute_basis_functions();

Dof dof;

```

```
std::map<std::pair<int,int>, ex> A;  
compute_Poisson_element_matrix(fe, dof, A);
```

In the above example, the vertices were integers, therefore the entries in the matrix will be rational numbers. In the following example the vertices are symbols.

```
//matrix in terms of symbols  
symbol x0("x0"), x1("x1"), x2("x2");  
symbol y0("y0"), y1("y1"), y2("y2");  
Triangle triangle2(lst(x0,y0), lst(x1,y1), lst(x2,y2));  
  
LagrangeFE fe2;  
fe2.set_order(order);  
fe2.set_polygon(triangle2);  
fe2.compute_basis_functions();  
  
Dof dof2;  
std::map<std::pair<int,int>, ex> A2;  
compute_Poisson_element_matrix(fe2, dof2, A2);
```

In this case `A2` will contain expressions involving the vertices, (x_0, y_0) , (x_1, y_1) , (x_2, y_2) (we used a triangle above).

The GiNaC library supports many different ways to print out the output. In the example below, we turn on \LaTeX output with the command `cout << latex;` before we print out `A2`.

```
cout << "LaTeX format on output " << endl;  
cout << latex;  
print(A2);
```

This gives the following expression (compiled by `latex`) for $A[1,1]$ (code for

the other entries are also produced, but these are not shown here).

$$\begin{aligned}
 A[1,1] = & \frac{1}{2} \frac{x_0^2 |(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2} \\
 & - \frac{y_1y_0 |(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2} \\
 & + \frac{1}{2} \frac{y_0^2 |(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2} \\
 & - \frac{x_0 |(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|x_1}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2} \\
 & + \frac{1}{2} \frac{|(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|x_1^2}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2} \\
 & + \frac{1}{2} \frac{y_1^2 |(-x_0 + x_1)(y_2 - y_0) - (-x_0 + x_2)(y_1 - y_0)|}{(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2}
 \end{aligned}$$

We can also print out C code,

```
cout <<"C code format on output "<<endl;
cout <<src;
print(A2);
```

Then the following code for $A[1,1]$ is produced,

```
A[1,1]=(x0*x0)/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)
*fabs((-x0+x1)*(y2-y0)-(-x0+x2)*(y1-y0))/2.0
-y1/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)*y0
*fabs((-x0+x1)*(y2-y0)-(-x0+x2)*(y1-y0))
+1.0/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)*(y0*y0)
*fabs((-x0+x1)*(y2-y0)-(-x0+x2)*(y1-y0))/2.0
-x0/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)
*fabs((-x0+x1)*(y2-y0)-(-x0+x2)*(y1-y0))*x1
+1.0/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)
*fabs((-x0+x1)*(y2-y0)
-(-x0+x2)*(y1-y0))*(x1*x1)/2.0+(y1*y1)
```

```
/pow(-y1*x2-x0*y2+y0*x2+y2*x1+x0*y1-y0*x1,2.0)
*fabs((-x0+x1)*(y2-y0)-(-x0+x2)*(y1-y0))/2.0
```

As is clear, these expressions can be rather large. GiNaC does not, by default, try to simplify these expressions. However, the above expressions is composed of smaller expressions that appear many times and it is possible to simplify these expressions fairly easy. For instance, the expression $(-y_1x_2 - x_0y_2 + y_0x_2 + y_2x_1 + x_0y_1 - y_0x_1)^2$ appears at least six times (and this is only in $A[1, 1]$). Of course, this expression should be computed only once. It seems that GiNaC has powerful tools for expression three traversal that could enable generation of efficient code based on finding common sub-expressions, but we have not exploited these tools to a great extent yet. Some example code can be found in `check_visitor.cpp` in the sandbox.

6.2 A Poisson Problem on Mixed Form

The Poisson problem can also be written on mixed form,

$$\begin{aligned} \mathbf{u} - \nabla p &= 0, & \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= f, & \text{in } \Omega, \\ \mathbf{u} \cdot \mathbf{n} &= g, & \text{on } \partial\Omega_N, \\ p \mathbf{n} &= h \mathbf{n} & \text{on } \partial\Omega_E. \end{aligned}$$

Notice that essential boundary conditions for the Poisson problem on standard form become natural conditions for the Poisson problem on mixed form and vice versa.

The weak form of the Poisson problem on mixed form is: Find $\mathbf{u} \in \mathbf{V}_g, p \in Q$ such that

$$a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) = \mathbf{G}(\mathbf{v}), \quad \forall \mathbf{v} \in \mathbf{V}_0, \quad (6.2)$$

$$b(\mathbf{u}, q) = F(q), \quad \forall q \in Q, \quad (6.3)$$

where

$$a(\mathbf{u}, \mathbf{v}) = \int_{\Omega} \mathbf{u} \cdot \mathbf{v} \, dx, \quad (6.4)$$

$$b(\mathbf{u}, q) = \int_{\Omega} \nabla \cdot \mathbf{u} \, q \, dx, \quad (6.5)$$

$$F(q) = \int_{\Omega} f \, q \, dx, \quad (6.6)$$

$$\mathbf{G}(\mathbf{v}) = \int_{\Omega_E} h \, \mathbf{n} \cdot \mathbf{v} \, ds \quad (6.7)$$

$$\mathbf{V}_k = \mathbf{v} \in \mathbf{H}(\text{div}) : \mathbf{v} \cdot \mathbf{n}|_{\partial\Omega_N} = k, \quad k = 0, g.$$

$$\mathbf{H}(\text{div}) = \mathbf{v} \in \mathbf{L}^2 : \nabla \cdot \mathbf{v} \in L^2,$$

$$Q = \begin{cases} L_0^2 & \text{if } \partial\Omega_N = \partial\Omega, \\ L^2 & \text{else.} \end{cases}$$

The function `compute_mixed_Poisson_element_matrix` in `ElementComputations.cpp` computes the element matrix for the mixed Poisson problem. We will not comment or list the code here because it is very similar to the code described in the next section. An example of use is in `mxpoisson.ex.cpp`.

6.3 A Stokes Problem

The Stokes problem is on the form: Find \mathbf{u} and p such that

$$\begin{aligned} -\Delta \mathbf{u} + \nabla p &= \mathbf{f}, & \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= 0, & \text{in } \Omega, \\ \mathbf{u} &= \mathbf{g}, & \text{on } \partial\Omega_E, \\ \frac{\partial \mathbf{u}}{\partial \mathbf{n}} - p \mathbf{n} &= \mathbf{h}, & \text{on } \partial\Omega_N. \end{aligned}$$

The weak form for the Stokes problem is: Find $\mathbf{u} \in \mathbf{V}_g, p \in Q$ such that

$$\begin{aligned} a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) &= \mathbf{F}(\mathbf{v}), \quad \forall \mathbf{v} \in \mathbf{V}_0, \\ b(\mathbf{u}, q) &= 0, \quad \forall q \in Q, \end{aligned}$$

where

$$\begin{aligned}
 a(\mathbf{u}, \mathbf{v}) &= \int_{\Omega} \nabla \mathbf{u} : \nabla \mathbf{v} \, dx, \\
 b(\mathbf{u}, q) &= - \int_{\Omega} \nabla \cdot \mathbf{u} \, q \, dx, \\
 \mathbf{F}(q) &= \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, dx + \int_{\Omega_N} \mathbf{h} \cdot \mathbf{v} \, ds, \\
 \mathbf{V}_{\mathbf{k}} &= \mathbf{v} \in \mathbf{H}^1 : \mathbf{v}|_{\partial\Omega_E} = \mathbf{k}, \quad \mathbf{k} = \mathbf{0}, \mathbf{g}, \\
 Q &= \begin{cases} L_0^2 & \text{if } \partial\Omega_E = \partial\Omega, \\ L^2 & \text{else.} \end{cases}
 \end{aligned}$$

Notice that we have multiplied the equation for the mass conservation, $\nabla \cdot \mathbf{u} = 0$, with -1 to obtain symmetry.

The function `compute_Stokes_element_matrix` in `ElementComputations` implements the computation of an element matrix for the Stokes problem. The code is shown below.

```

void compute_Stokes_element_matrix(
    FE& v_fe,
    FE& p_fe,
    Dof& dof,
    std::map<std::pair<int,int>, ex>& A)
{
    std::pair<int,int> index;
    std::pair<int,int> index2;

    Polygon& domain = v_fe.get_polygon();

    // Insert the local degrees of freedom into the global Dof
    for (int i=0; i< v_fe.nbf(); i++) {
        dof.insert_dof(1,i,v_fe.dof(i));
    }
    for (int i=0; i< p_fe.nbf(); i++) {
        dof.insert_dof(1,v_fe.nbf()+i,p_fe.dof(i));
    }

    // The term (grad u, grad v)
    for (int i=0; i< v_fe.nbf(); i++) {
        index.first = dof.glob_dof(v_fe.dof(i));    // fetch the dof for v_i

```



```

    for (int j=0; j< v_fe.nbf(); j++) {
        index.second = dof.glob_dof(v_fe.dof(j)); // fetch the dof for v_j
        GiNaC::ex nabla = inner(grad(v_fe.N(i)),
                                grad(v_fe.N(j))); // compute the integrand
        GiNaC::ex Aij = domain.integrate(nabla); // compute the integral
        A[index] += Aij;                        // add to global matrix
    }
}

// The term (-div u, q)
for (int i=0; i< p_fe.nbf(); i++) {
    index.first = dof.glob_dof(p_fe.dof(i)); // fetch the dof for p_i
    for (int j=1; j< v_fe.nbf(); j++) {
        index.second=dof.glob_dof(v_fe.dof(j)); // fetch the dof for v_j
        ex divV= -p_fe.N(i)*div(v_fe.N(j)); // compute the integrand
        ex Aij = domain.integrate(divV); // compute the integral
        A[index] += Aij; // add to global matrix

        // Do not need to compute the term (grad(p),v), since the system is
        // symmetric. We simply set Aji = Aij
        index2.first = index.second;
        index2.second = index.first;
        A[index2] -= Aij;
    }
}
}

```

6.4 A Nonlinear Convection Diffusion Problem

Our next example concerns a nonlinear convection diffusion equation, where we compute the element matrix for the Jacobian typically arising in a Newton iteration. Let the PDE be,

$$(\mathbf{u} \cdot \nabla) \mathbf{u} - \Delta \mathbf{u} = \mathbf{f}, \quad \text{in } \Omega, \quad (6.8)$$

$$\mathbf{u} = \mathbf{g}, \quad \text{on } \partial\Omega. \quad (6.9)$$

This can be stated on weak form as: Find $\mathbf{u} \in \mathbf{V}_g$ such that

$$\mathbf{F}(\mathbf{u}, \mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathbf{V}_0,$$

where

$$\mathbf{F}(\mathbf{u}, \mathbf{v}) = \int_{\Omega} (\mathbf{u} \cdot \nabla \mathbf{u}) \cdot \mathbf{v} \, dx + \int_{\Omega} \nabla \mathbf{u} : \nabla \mathbf{v} \, dx - \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, dx$$

and

$$\mathbf{V}_{\mathbf{k}} = \mathbf{v} \in \mathbf{H}^1 : \mathbf{v}|_{\partial\Omega} = \mathbf{k}, \quad \mathbf{k} = \mathbf{0}, \mathbf{g}.$$

The Jacobian is obtained by letting $\mathbf{u} = \hat{\mathbf{u}} = \sum_j u_j \mathbf{N}_j$, $\mathbf{v} = \mathbf{N}_i$ and differentiating \mathbf{F} with respect to u_j ,

$$J_{ij} = \frac{\partial F(\hat{\mathbf{u}}, \mathbf{N}_i)}{\partial u_j}.$$

This is precisely the way it is done with SyFi, (see also `nljacobian.ex.cpp`),

```
void compute_nlconvdiff_element_matrix(
    FE& fe,
    Dof& dof,
    std::map<std::pair<int,int>, ex>& A)
{
    std::pair<int,int> index;
    Polygon& domain = fe.get_polygon();

    // insert the local dofs into the global Dof object
    for (int i=0; i< fe.nbf() ; i++) {
        dof.insert_dof(1,i,fe.dof(i));
    }

    // create the local U field: U = sum_k u_k N_k
    ex UU = matrix(2,1,1st(0,0));
    ex ujs = symbolic_matrix(1,fe.nbf(), "u");
    for (int k=0; k< fe.nbf(); k++) {
        UU += ujs.op(k)*fe.N(k);    // U += u_k N_k
    }

    //Get U represented as a matrix
    matrix U = ex_to<matrix>(UU.evalm());

    for (int i=0; i< fe.nbf() ; i++) {
        index.first = dof.glob_dof(fe.dof(i));    // fetch global dof

        // First: the diffusion term in Fi
        ex gradU = grad(U);                        // compute the gradient
    }
}
```

```

    ex Fi_diffusion = inner(gradU,          // grad(U)*grad(Ni)
                           grad(fe.N(i)));

    // Second: the convection term in Fi
    ex Ut = U.transpose();                // get the transposed of U
    ex UgradU = (Ut*gradU).evalm();       // compute U*grad(U)
    ex Fi_convection = inner(UgradU, fe.N(i), // compute U*grad(U)*Ni
                           true);

    // add together terms for convection and diffusion
    ex Fi = Fi_convection + Fi_diffusion;

    // Loop over all uj and differentiate Fi with respect
    // to uj to get the Jacobian Jij
    for (int j=0; j< fe.nbf() ; j++) {
        index.second = dof.glob_dof(fe.dof(j)); // fetch global dof
        symbol uj = ex_to<symbol>(ujs.op(j)); // cast uj to a symbol
        ex Jij = Fi.diff(uj,1);                // differentiate Fi wrt. uj
        ex Aij = domain.integrate(Jij);        // intergrate the Jacobian Jij
        A[index] += Aij;                        // update the global matrix
    }
}
}
}

```

Running the example `nljacobian_ex`, which employs second order continuous Lagrangian elements, yields the following output for $A[1, 1]$,

$$A[1, 1] = \frac{1}{2} + \frac{2}{105}u_3 + \frac{2}{105}u_7 + \frac{1}{21}u_2\frac{13}{420}u_1 - \frac{1}{280}u_{11} \quad (6.10)$$

$$- \frac{1}{21}u_6 - \frac{1}{280}u_5\frac{1}{140}u_{10} + \frac{1}{210}u_9 - \frac{1}{140}u_4. \quad (6.11)$$

We have used GiNaC to generate the \LaTeX code, as described on Page 80.

6.5 Expression Simplification

When generating expressions for complicated forms, and specially non-linear forms, there is much room for optimization of the resulting expressions to generate more efficient code. Generating optimal code for the computation

of a large symbolic expression is a very difficult problem, and the underlying symbolic engine in GiNaC has limited support for this. However, GiNaC has many of the building blocks to perform this optimization. We have implemented a basic algorithm to simplify general expressions, which generates helper variables for basic binary operations that are repeated. The algorithm is very simple, and the resulting speedup (in the tests) ranges from a factor four to slightly negative. Obviously more work is needed in this area to make this usable. The current expression simplifier can be tested by running "make simplify && ./simplify v" under tests/. A basic code example is shown below.

```
ExpressionSimplifier es;
es.add(e_symbol, e_expression);
es.add(f_symbol, f_expression);
es.simplify();

list< pair< symbol, ex > > & selist =
    es.get_output().get_symex_list();
genCodeSymbols(cout, selist);
```

Chapter 7

Python Support

SyFi comes with Python support. The SyFi Python module is created by using the tool SWIG (<http://www.swig.org>). One should also install the Python interface to GiNaC called Swiginac (<http://swiginac.berlios.de/>).

The following code shows how Swiginac can be used (see also `simple.py`),

```
from swiginac import *

x = symbol("x")
y = symbol("y")

f = sin(x)
print "f = ", f
dfdx = diff(f,x)
print "dfdx = ", dfdx
```

SyFi classes and functions can be used in Python just as they are used in C++. The following example shows how to compute the element matrix for a Poisson problem using forth order Lagrangian elements,

```
from swiginac import *
```

```

from SyFi import *

p0 = [0,0,0]; p1 = [1,0,0]; p2 = [0,1,0]
triangle = Triangle(p0, p1, p2)
fe = LagrangeFE(triangle,4)
print fe.nbf()
for i in range(0,fe.nbf()):
    for j in range(0,fe.nbf()):
        integrand = inner(grad(fe.N(i)),grad(fe.N(j)))
        Aij = triangle.integrate(integrand)
        print "A(%d,%d)=%%(i,j), Aij.eval()

```

Finally, we show a Python implementation of the Crouzeix-Raviart element (The C++ implementation can be found in the file `CrouzeixRaviart.cpp`). Notice that in this code we inherit the functions `ex N(int i)` and `ex dof(int i)` and the `exvectors` `Ns` and `dofs` from the C++ class `StandardFE`. Hence, thanks to SWIG, cross-language inheritance works, and we therefore only need to implement the function `compute.basis.functions`. The following example is implemented in `crouzeixraviart.py`.

```

from swiginac import *
from SyFi import *

initSyFi(3)

x = cvar.x; y = cvar.y; z = cvar.z # fetch some global variables

class CrouzeixRaviart:
    """
        Python implementation of the Crouzeix-Raviart element.
        The corresponding C++ implementation is in the
        file CrouzeixRaviart.cpp.
    """

```

```
def __init__(self, polygon):
    """ Constructor """
    self.Ns = []
    self.dofs = []
    self.polygon = polygon
    self.compute_basis_functions()

def compute_basis_functions(self):
    """
    Compute the basis functions and degrees of freedom
    and put them in Ns and dofs, respectively.
    """
    polspace = bernstein(1,triangle,"a")
    N = polspace[0]
    variables = polspace[1]

    for i in range(0,3):
        line = triangle.line(i)
        dofi = line.integrate(N)
        self.dofs.append(dofi)

    for i in range(0,3):
        equations = []
        for j in range(0,3):
            equations.append(relational(self.dofs[j], dirac(i,j)))
        sub = lsolve(equations, variables)
        Ni = N.subs(sub)
        self.Ns.append(Ni);

    def N(self,i): return self.Ns[i]
    def dof(self,i): return self.dofs[i]
    def nbf(self): return len(self.Ns)

p0 = [0,0,0]; p1 = [1,0,0]; p2 = [0,1,0];

triangle = Triangle(p0, p1, p2)
```

```
fe = CrouzeixRaviart(triangle)
fe.compute_basis_functions()
print fe.nbf()

for i in range(0,fe.nbf()):
    print "N(%d)          = %i,    fe.N(i).eval().printc()
    print "grad(N(%d)) = %i,    grad(fe.N(i)).eval().printc()
    print "dof(%d)      = %i,    fe.dof(i).eval().printc()
```


Chapter 8

Code Generation

In this section we will describe some matrix factories created for the PyCC project [?], which have been made by using SyFi, GiNaC and Swiginac. At present, we have written ca. 1500 lines of Python code using SyFi, Swiginac etc., which have generated roughly 60 000 lines of C++ code for the computation (of various variants) of the mass matrix, the stiffness matrix, the convection matrix and the divergence matrix using Lagrangian elements of order 1-5 in 2D and 1-3 in 3D. Furthermore, the generated C++ code is efficient, since everything except the geometry mapping can be computed exactly. Notice also that although only Lagrangian elements have been used so far, most of the Python code that generated the C++ code is completely element independent. In addition to the generated C++ code we have also written about 1500 lines of code which loops over the cells of a Dolfin mesh [?] such that global matrices are made.

We have create two matrix factories. These are implemented in `MatrixFactory` and `MatrixFactory_highorder`. There are three differences between these two factories. The first difference is that `MatrixFactory` employs the numbering of degrees of freedom in the Dolfin mesh. Therefore, this `MatrixFactory` is limited to linear Lagrangian elements. On the other hand, `MatrixFactory_highorder` uses `DofT`, described in Section 3.5, which works for general elements. The second difference is that in `MatrixFactory` the integration is performed on a

global element with global basis functions, e.g., for the stiffness matrix,

$$A_{ij} = \int_T \nabla N_i \nabla N_j dx. \quad (8.1)$$

In `MatrixFactory_highorder`, the integration is performed on the reference element with a geometry tensor G (the Jacobian of the geometry mapping) and $D = \det(G)$,

$$A_{ij} = \int_{\hat{T}} (G^{-T} \nabla \hat{N}_i) \cdot (G^{-T} \nabla \hat{N}_j) D dx. \quad (8.2)$$

which is the typical way to do it in finite element codes. At present, we favor (8.2) to (8.1) simply because it produces much smaller expressions and therefore faster code. However, the large expressions in (8.1) typically involve subexpressions repeated many times. Hence, it should be possible to postprocess these expressions to create smaller expressions and faster code. However, we have not done this yet. Finally, the third difference is that `MatrixFactory_highorder` works for the `FastMatSparse` matrix in PyCC, the `Epetra` matrix in Trilinos [?] and for STL maps of type `map<pair<int,int>,double>`.

8.1 Basic Tools

We will illustrate the code generation by considering what was done for the mass matrix in `MatrixFactory_highorder`.

The entries of a mass matrix are:

$$M_{kl} = \int_T N_k N_l dx = \int_{\hat{T}} \hat{N}_i \hat{N}_j D dx, ,$$

where T is the global polygon, N_k and N_l are the k 'th and l 'th global basis functions, respectively, \hat{T} is the reference polygon, \hat{N}_i and \hat{N}_j are the i 'th and j 'th basis functions on the reference polygon corresponding to k and l , respectively, and D is the determinant of the Jacobian of the geometry mapping. The following code shows how this can be done (see also `code_gen.py`):

```
def create_A_string_mass(fe):
    A_str = "  double A[%d][%d];\n\n "%(fe.nbf(), fe.nbf())
```

```
domain = fe.get_polygon()

# loop over all N(i)
for i in range(0,fe.nbf()):

    # loop over all N(j)
    for j in range(0,fe.nbf()):

        # compute the integrand N(i)*N(j)
        integrand = fe.N(i).eval()*fe.N(j).eval()

        # integrate over the domain
        Aij = domain.integrate(toex(integrand))

        # generate C string and append the string to the rest
        A_str += "  A[%d] [%d]=(%s)*D;\n\n %"
                (i,j,Aij.eval().evalf().printc())
```

The following output is produced, when using linear element on a 2D triangle (see also `matrix_factory_mass_2D.cc`, which also contains code for higher order Lagrangian elements),

```
double A[3][3];
A[0][0]=(8.333333333333329e-02)*D;
A[0][1]=(4.166666666666664e-02)*D;
A[0][2]=(4.166666666666664e-02)*D;
A[1][0]=(4.166666666666664e-02)*D;
A[1][1]=(8.333333333333329e-02)*D;
A[1][2]=(4.166666666666664e-02)*D;
A[2][0]=(4.166666666666664e-02)*D;
A[2][1]=(4.166666666666664e-02)*D;
A[2][2]=(8.333333333333329e-02)*D;
```

Hence, this is the mass element matrix on the reference element multiplied with D . In addition to computing the element matrix we also need to com-

puted the global degrees of freedom and generate a C function. We will not go into details on this, but recommend the reader to have a look in `code_gen.py`.

The complete function for the computation of the element matrix, in the case of linear Lagrangian elements, and the insertion of the element matrix in the global matrix can be found in `matrix_factory_mass_2D.cc` is:

```
void matrix_factory_mass_2D_order1 (map<pair<int,int>,double>& matrix,
    DofT<Ptv,Ptv_is_less>& dof,
    int element, double pp0[2], double pp1[2], double pp2[2]){

    // geometry related stuff

    double x0 = pp0[0]; double y0 = pp0[1];
    double x1 = pp1[0]; double y1 = pp1[1];
    double x2 = pp2[0]; double y2 = pp2[1];

    double G00 = x1 - x0; double G01 = x2 - x0;
    double G10 = y1 - y0; double G11 = y2 - y0;

    double D = fabs(G00*G11-G01*G10);

    // inserting local dofs in the global dof handler (dof)

    int iidof[3];
    double dof1[2];
    dof1[0]=x0;  dof1[1]=y0;
    Ptv pdof1(2,dof1);
    iidof[0] = dof.insert_dof(element,1,pdof1);

    double dof2[2];
    dof2[0]=G01+x0;  dof2[1]=y0+G11;
    Ptv pdof2(2,dof2);
    iidof[1] = dof.insert_dof(element,2,pdof2);

    double dof3[2];
    dof3[0]=G00+x0;  dof3[1]=G10+y0;
    Ptv pdof3(2,dof3);
    iidof[2] = dof.insert_dof(element,3,pdof3);

    // compute the element matrix

    double A[3][3];
```

```

A[0][0]=(8.333333333333329e-02)*D;
A[0][1]=(4.166666666666664e-02)*D;
A[0][2]=(4.166666666666664e-02)*D;
A[1][0]=(4.166666666666664e-02)*D;
A[1][1]=(8.333333333333329e-02)*D;
A[1][2]=(4.166666666666664e-02)*D;
A[2][0]=(4.166666666666664e-02)*D;
A[2][1]=(4.166666666666664e-02)*D;
A[2][2]=(8.333333333333329e-02)*D;

// insert element matrix into global matrix

int nbf = 3;
pair<int,int> index;
for (int i=0; i< nbf; i++) {
    index.first = iidof[i];
    for (int j=0; j< nbf; j++) {
        index.second = iidof[j];
        matrix[index] += A[i][j];
    }
}
}
}

```

Finally, we show how the above function is used in PyCC to compute the mass matrix on a Dolfin mesh (see also `MatrixFactory_highorder.cpp`)

```

void MapMatrixFactory:: computeMassMatrix(){

    int e = -1;

    if (mesh->numSpaceDim() == 2) {

        double p0[2];
        double p1[2];
        double p2[2];

        for (CellIterator cell(*mesh); !cell.end(); ++cell) {

            e++;

            // Obtain vertices from Dolfin mesh
            Vertex& v0 = (*cell).vertex(0);
            Vertex& v1 = (*cell).vertex(1);

```

```
Vertex& v2 = (*cell).vertex(2);

// Create double arrays with the data from the vertices
p0[0] = v0.coord().x; p0[1] = v0.coord().y;
p1[0] = v1.coord().x; p1[1] = v1.coord().y;
p2[0] = v2.coord().x; p2[1] = v2.coord().y;

switch(order1) {
  case 1 :
    matrix_factory_mass_2D_order1(*matrix,*idof,e,p0,p1,p2);
    break;
  case 2 :
    matrix_factory_mass_2D_order2(*matrix,*idof,e,p0,p1,p2);
    break;
  case 3 :
    matrix_factory_mass_2D_order3(*matrix,*idof,e,p0,p1,p2);
    break;
  case 4 :
    matrix_factory_mass_2D_order4(*matrix,*idof,e,p0,p1,p2);
    break;
  case 5 :
    matrix_factory_mass_2D_order5(*matrix,*idof,e,p0,p1,p2);
    break;
}
}
```

Notice that this code works for Lagrangian elements of order 1-5 in 2D.

8.2 Debugging

Debugging finite element codes is often extremely hard, at least that is the authors' experience. This has been one of the reasons why we have chosen to employ a symbolic math engine behind the curtain in the first place.

One of the advantages of SyFi is that one obtain explicit symbolic expressions for all the basis functions (and its derivatives). Another good thing is that

one can create global finite elements, that is finite elements that are not defined on reference geometries, and perform integration and differentiation on their geometries. For instance, when we created the divergence matrix factory we initially had a mysterious bug which took us several hours to find. To locate the bug, we computed the divergence element matrix on a global element with the vertices $\mathbf{x}_0 = (0.2, 0.2)$, $\mathbf{x}_1 = (0.4, 0.2)$, and $\mathbf{x}_2 = (0.1, 0.3)$, and compared it with the divergence element matrix on the reference element with the corresponding geometry tensor. To do this, we wrote the following code (see also `main_syfi.cpp`):

```
// create global triangle
lst p0(0.2, 0.2);
lst p1(0.4, 0.2);
lst p2(0.1, 0.3);
Triangle triangle(p0,p1,p2);

// create vector element for v on the global triangle
VectorLagrangeFE v_fe;
v_fe.set_size(2);
v_fe.set_order(vorder);
v_fe.set_polygon(triangle);
v_fe.compute_basis_functions();

// create scalar element for p on the global triangle
LagrangeFE p_fe;
p_fe.set_order(1);
p_fe.set_polygon(triangle);
p_fe.compute_basis_functions();

// compute global element matrix
map<pair<int,int>, ex> A;
pair<int,int> index;
for (int i=0; i< p_fe.nbf(); i++) {
    index.first = i;
    for (int j=0; j< v_fe.nbf(); j++) {
        index.second= j;
```

```
        ex divV= p_fe.N(i)*div(v_fe.N(j));
        ex Aij = triangle.integrate(divV);
        A[index] = Aij;
    }
}
```

The element matrix created by this code was then printed out and compared with the element matrix computed by the matrix factory on the same polygon (see `dolfin.main.cpp`). By comparing each entry of the two matrices we quickly found the (uninteresting) bug. Hence, in our experience it is extremely valuable to have the concrete basis functions etc. on global element, and being able to work with them both with a pen and a paper and the computer, to reveal what is going on.

Chapter 9

Using the SyFi Form Compiler

FIXME: This chapter has not been updated for the new rewritten UFL-based SFC.

The SyFi Form Compiler – abbreviated SFC – is a Python module that compiles a symbolic description of a finite element discretization into efficient low level C++ code.

Other FEM packages can use the code generated by SFC by implementing assembly of global finite element matrices through the C++ interface called UFC [?] (Unified Form-assembly Code). Currently, the software packages PyCC [?] and (Py)DOLFIN [?] support assembly of UFC forms, and FFC [?] also supports generating UFC compliant code. We refer to the manuals of DOLFIN and PyCC for more details about the global matrix/vector assembly, and focus here on how to define a finite element discretization as input to SFC.

We recommend reading the UFC paper [?] and introductory chapters of the UFC manual [?] before proceeding, since much mathematical notation and technical details are defined there.

9.1 Quickstart

To give a quick first impression of what SFC does, consider this code:

```
from sfc import *

# Define elements
polygon = "tetrahedron"
P2 = FiniteElement("CG", polygon, 2)
T0 = TensorElement("DG", polygon, 0)

# Define form arguments
u = TrialFunction(P2)
v = TestFunction(P2)
M = Function(T0)

# Define integrand for a weighted stiffness matrix
def stiffness(v, u, M, itg):
    GinvT = itg.GinvT()
    Du = grad(u, GinvT)
    Dv = grad(v, GinvT)
    return inner(M*Du, Dv)

# Collect the pieces as a form
a = CallbackForm(basisfunctions = [v, u],
                  coefficients    = [M],
                  cell_integrals = [stiffness])

# Generate UFC code, compile and import
a_compiled = compile_form(a)

# Assemble the global system using PyDolfin
from dolfin import *
n = 10
mesh = UnitCube(n, n, n)
A = assemble(a_compiled, mesh) # FIXME: ensure working dolfin syntax here (== valid
```

The result of this code is that the variable `a_compiled` contains a compiled

C++ object that implements `ufc::form`. This class interface provides functionality to compute the element tensor for a particular problem. By passing this object to the PyDOLFIN assembler together with a mesh, we assemble a global sparse matrix as seen at the end of the code.

9.2 Defining Form Arguments

9.2.1 Defining Finite Elements

The first step of defining a discretized form is choosing which elements to use. The basic syntax is:

```
fe = FiniteElement(family, domain, order)
fe = VectorElement(family, domain, order)
fe = TensorElement(family, domain, order)
```

The argument `family` can currently be one of

- "CG" or "Lagrange" (Continuous Galerkin)
- "DG" (Discontinuous Galerkin)
- "Bubble"
- "CR" or "Crouzeix-Raviart"

The argument `domain` can be one of

- "triangle"
- "tetrahedron"
- "quadrilateral"

- "hexahedron"

For vector and tensor elements, the default size is the geometric dimension of the polygon.

Some typical examples of elements are shown below, which should be self-explaining.

```
polygon = "tetrahedron"
P0 = FiniteElement("DG", polygon, 0)
P1 = FiniteElement("CG", polygon, 1)
V2 = VectorElement("CG", polygon, 2)
```

9.2.2 Defining Basisfunctions

One basisfunction is needed for each rank of the form, defined either using `BasisFunction` or `TestFunction` and `TrialFunction`. The element can be the same or different.

```
v = TestFunction(v_element)
u = TrialFunction(u_element)
# equivalently
v = BasisFunction(v_element)
u = BasisFunction(u_element)
```

9.2.3 Defining Coefficients

Next we need to define each coefficient function as a member of a finite element space. This is done by constructing a `Function` object with an element as its first constructor argument and an optional name.

```
f = Function(element, name="f")
```

A shortcut exists for defining (piecewise) constants,

```
mu = Constant(name="mu")
```

which is equivalent to

```
element = FiniteElement("DG", polygon, 0)
mu = Function(element, name="mu")
```

Providing a name for a coefficient is optional, but may aid SFC in making the generated code more readable and self-documented.

9.3 Defining a Form

There are a few different ways to go about defining the variational form in SFC. Lets first summarize the main properties of a form in the UFC framework. The form corresponding to a rank r element tensor has r basis functions. It also has n coefficient functions. These are defined like explained above, and passed as arguments when constructing a form representation object `FormRep`.

A form can have any number of cell integrals, exterior facet integrals (boundary integrals) or interior facet integrals (not implemented). Calling the function `add_cell_integral` on a `FormRep` object adds another integral and returns its `IntegralRep` object. We will look at how to define each integral below.

The following code shows what this looks like

```
# define a form with two basisfunctions and one coefficient
a = FormRep(name          = "my_form",
             basisfunctions = [v, u],
```

```
        coefficients = [c],
        options      = {}))

# add one cell integral to the form
citg = a.add_cell_integral()

# add two boundary integrals to the form
bitg0 = a.add_exterior_facet_integral()
bitg1 = a.add_exterior_facet_integral()
```

All arguments to `FormRep` are optional with sensible default values.

9.4 Defining an Integral

An `IntegralRep` object collects all information concerning a single integral. Most importantly, it provides a member function `set_A(index, integrand)` to set the integrand for each element tensor expression (for a way to avoid manually looping over basis functions, see the section about `CallbackForm`). Additionally, it can provide symbols and expressions for geometric quantities, basis functions, and coefficients.

9.4.1 Argument expressions

The following functions are the most important ones to define a form by manually iterating over basis functions and computing each integrand.

- `itg.v_basis(i)` Returns a list with expressions for all basis functions in finite element space `i`.
- `itg.coefficient(i)`
Returns an expression for coefficient `i`.
- `itg.set_A(index, integrand)` Set integrand expression for a given (multi-)index. (Integration and scaling with the geometry mapping is

performed by SFC.)

9.4.2 Geometric Quantities on Cells

- `itg.vx(i)`
Returns a symbol for the global coordinate of vertex i .
- `itg.G()`
Returns the affine mapping G from reference element to global element.
- `itg.detG()`
Returns the determinant of the affine mapping G . (You should usually don't need this, SFC scales the integral with this factor when integrating.)
- `itg.GinvT()`
Returns the transposed inverse of the affine mapping G , must be passed to the differential operators `div`, `grad` and `curl`.
- `itg.n()`
Returns the outwards pointing normal vector on a boundary facet (exterior facet integrals only).

9.4.3 Symbolic Language

When computing with symbolic expressions, all features of `swiginac`, which is most features of `GiNaC`, are available through SFC. We recommend reading the `GiNaC` manual for more details about its features. Some operators are defined by SFC on top of the general symbolic library. These include (u , v are vectors, A , B are matrices):

- $\nabla u = \text{grad}(u, \text{GinvT})$
- $\nabla \cdot u = \text{div}(u, \text{GinvT})$
- $\nabla \times u = \text{curl}(u, \text{GinvT})$

- $AB = \sum_k A_{ik} B_{kj} = \mathbf{A} * \mathbf{B}$ (regular matrix-matrix product)
- $Av = \mathbf{A} * \mathbf{v}$
- $v \cdot A = v^T A = \text{dot}(\mathbf{v}, \mathbf{A})$
- $A : B = \sum_i \sum_j A_{ij} B_{ij} = \text{inner}(\mathbf{A}, \mathbf{B})$
- $u \cdot v = \sum_k u_k v_k = \text{dot}(\mathbf{u}, \mathbf{v})$
- `trace(A)`
- `transpose(A)`

See the file `symbolic_utils.py` for the definition of these operators.

9.4.4 Examples

Lets look at a more realistic example, for defining and compiling a linear elasticity form.

```
def define_linear_elasticity(itg):
    I      = Id(itg.nsd)
    GinvT = itg.GinvT()

    # get coefficients
    lamdb, mu = itg.coefficients()

    # for all trial functions u
    for i, u in enumerate(itg.basisfunctions(1)):
        Du = grad(u, GinvT)
        DuT = Du.transpose()
        E = (Du + DuT) / 2
        sigma = lamdb * Du * I + 2*mu*E

    # for all test functions v
    for j, v in enumerate(itg.basisfunctions(1)):
        Dv = grad(v, GinvT)
```



```

    integrand = inner(sigma, Dv)
    itg.set_A((i, j), integrand)

```

```

def define_linear_elasticity_traction(itg):
    I      = Id(itg.nsd)
    GinvT = itg.GinvT()

    # get coefficients
    lambd, mu = itg.coefficients()

    # for all trial functions u
    for i, u in enumerate(itg.basisfunctions(1)):
        Du = grad(u, GinvT)
        DuT = Du.transpose()
        E = (Du + DuT) / 2
        sigma = lambd * Du * I + 2*mu*E

    # for all test functions v
    for j, v in enumerate(itg.basisfunctions(1)):
        Dv = grad(v, GinvT)
        integrand = inner(sigma, Dv)
        itg.set_A((i, j), integrand)

```

The code which ties these functions to a form is shown below.

```

from sfc import *

def define_linear_elasticity(itg):
    ...

def define_linear_elasticity_traction(itg):
    ...

polygon = "tetrahedron"
element = VectorElement("CG", polygon, 1)

```

```
v = TestFunction(element)
u = TrialFunction(element)

lamdb = Constant(polygon)
mu     = Constant(polygon)

a = Form(name          = "elasticity",
          basisfunctions = [v, u],
          coefficients   = [lamdb, mu])

itg = userform.add_cell_integral()
define_linear_elasticity(itg)

itg = userform.add_exterior_facet_integral()
define_linear_elasticity_traction(itg)

a_compiled = compile_form(a)
```

FIXME: verify this code

9.5 Defining forms with callback functions

The loop over basis functions can be automated by using callback functions.

¹ The callback function should compute and return the integrand for the given set of basis functions and coefficients.

For each integrand in your form, you define a function taking one argument for each form argument plus one. In other words, the function should have $r+n+1$ arguments where r is the rank and n is the number of coefficients. The last argument is an integral context object, of the same type as was returned from `FormRep.add*_integral()`. Your integrand function will be called once for each entry in the element tensor, with the different basis functions passed as the first r arguments. The next n arguments are expressions for the coefficients.

¹A *callback function* is a function that is passed as an argument to other code. It allows a lower-level software layer to call a subroutine (or function) defined in a higher-level layer.

We will demonstrate how to do this for the three variational forms

$$c(; f) = \|f\|_1 = \int_{\Omega} f \cdot f + \nabla f \cdot \nabla f \, dx, \quad (9.1)$$

$$b(v; f) = \int_{\Omega} f \cdot v \, dx, \quad (9.2)$$

$$a(v, u; w) = \int_{\Omega} (w \cdot \nabla u) \cdot v \, dx. \quad (9.3)$$

Note the syntax $a(v, u; M)$, where v is the test function, u is the trial function, and w is a coefficient.

Callback functions for these integrands are

```
def H1(f, itg):
    GinvT = itg.GinvT()
    Df = grad(f, GinvT)
    return inner(f, f) + inner(Df, Df)

def source(v, f, itg):
    return inner(f, v)

def convection(v, u, w, itg):
    GinvT = itg.GinvT()
    Du = grad(u, GinvT)
    return inner( inner(w, Du), v )
```

The first arguments of the callback function must be a number of basis functions equal to the rank of the form. Thus for vector and matrix forms, the first argument is the test function, and for matrix forms the second argument is the trial function. Scalar forms (f.ex. energy norms) require no basis function arguments.

The next arguments should be one for each coefficient function (f and w in **H1**, **source** and **convection** above). The last argument is the `IntegralRep` object, which is briefly described elsewhere. Basically, through this object you can get symbols representing geometry variables like the geometry mapping G^{-T} and boundary normal vector \mathbf{n} .

The below code shows how to define and compile a form for each of the above callback functions.

```
from sfc import *

P1 = FiniteElement("CG", "triangle", 1)

v = TestFunction(P1)
u = TrialFunction(P1)
f = Function(P1)
w = Function(P1)

c = CallbackForm(coefficients = [f],
                  cell_integrals = [H1])
cc = compile_form(c)

b = CallbackForm(basisfunctions = [v],
                  coefficients = [f],
                  cell_integrals = [source])
bc = compile_form(b)

a = CallbackForm(basisfunctions = [v, u],
                  coefficients = [w],
                  cell_integrals = [convection])
ac = compile_form(a)
```

We first define the linear finite element P1 and types of the arguments v, u, f and w, and then tie it all together in a CallbackForm object. This object is then sent to `compile_form` which generates C++ code for the form and elements, and returns a compiled C++ object which is a subclass of `ufc:form`. This is accomplished by using Instant [?] which compiles a Python extension module using SWIG [?] and g++ [?]. Also note that the name of the form is automatically deduced from the name of the callback function.

9.6 Computing the Jacobi matrix form from a nonlinear vector form

Let $F(v; u, \dots)$ be a rank 1 form which is nonlinear in the first coefficient u . Then a matrix form representing the Jacobi of this form, $J(v, u; \dots)$, can be computed like

```
J = Jacobi(F)
```

9.7 Compiling a Form (Generating Code)

After defining a `FormRep` or `CallbackForm` in one of the ways described above, this object can be passed to `compile_form` which generates C++ code, compiles it as a Python extension module, imports it and returns an instance of the newly generated `ufc::form` subclass.

```
a_compiled = compile_form(a)
```

This object can be used with PyDOLFIN to assemble the global sparse matrix or global vector.

Alternatively, calling

```
write_ufc_code(a)
```

will generate code and write it to file in the directory `generated_code`,

9.8 Options

Several options affect the way forms are compiled, the most important being the choice of symbolic integration or quadrature and quadrature order if applicable. (Many of the other options are likely to change a bit, and are thus not explained here.)

```
options = {
    "symbolic": False,
    "quad_order": 4,
}

a = FormRep(..., options = options)

# or

a = CallbackForm(..., options = options)
```

9.9 Compiling a function

To allow userdefined functions based on symbolic computations without sacrificing performance during assembly, `sfc` provides the function `compile_function(expression, name)`. This works similar to `compile_form(...)`. An important use of this function is during validation of a solver, as a powerful tool for the method of manufactured solutions.

UFC defines a functor interface `ufc::function`, which can be passed to the function `evaluate_dof(...)` of the class `ufc::finite_element` to evaluate the degrees of freedom of the function in the local finite element space on a cell. The return value of `compile_function` is a compiled functor object of this type.

The argument `name` is used both for naming the generated functor class and to name the Python module, so function names must be valid C++ variable names. The `expression` can be a number of different types. A string is

assumed to be a valid C++ expression which is simply pasted into the generated code. A list of strings is treated as a vector function. Tensor functions are the same as vector functions, just unpack the components into a list like a matrix in C. Finally, the expression can be a swiginac expression, either scalar or matrix. Below is an example Python code for compiling a function from a swiginac matrix, and the resulting generated C++ code.

```
from sfc import compile_function, symbols, matrix, sin, cos

x, y = symbols(["x" "y"])
mat = matrix(2, 2, [x*x, cos(x*y), sin(y*x), y*y])
matfunc = compile_function(mat, "matfunc")
```

```
class matfunc: public ufc::function
{
public:
    /// Evaluate the function at the point
    /// x = (x[0], x[1], ...) in the cell
    virtual void evaluate(double* values,
                          const double* x_,
                          const ufc::cell& c) const
    {
        const double x=x_[0], y=x_[1], z=x_[2];
        values[0] = (x*x);
        values[1] = cos(y*x);
        values[2] = sin(y*x);
        values[3] = (y*y);
    }
};
```


Chapter 10

Behind the SyFi Form Compiler

FIXME: This chapter has not been updated for the new rewritten UFL-based SFC.

This chapter gives an overview of the implementation of SFC, intended for developers and technical users during debugging.

We strongly recommend reading the UFC paper (FIXME:REFERENCE) and UFC manual (FIXME:REFERENCE) before proceeding, since much mathematical notation, numerical concepts and technical details are defined there.

10.1 Example of generated code

Let us first look at an example of UFC code that is generated by SFC.

Shown below is the function `tabulate_tensor` which computes the element tensor for a stiffness matrix with a scalar conductivity. The form uses scalar linear elements for the test and trial functions v and u , and a scalar piecewise constant element (P0) for the coefficient M .

FIXME: update code

```

void cell_integral_stiffness_with_M_LagrangeFE_1_2D::
    tabulate_tensor(double* A, const double * const * w,
                    const ufc::cell& cell) const
{
    // coordinates
    double x0 = cell.coordinates[0][0];
    double y0 = cell.coordinates[0][1];
    double x1 = cell.coordinates[1][0];
    double y1 = cell.coordinates[1][1];
    double x2 = cell.coordinates[2][0];
    double y2 = cell.coordinates[2][1];

    // affine map
    double G00 = x1 - x0;
    double G01 = x2 - x0;

    double G10 = y1 - y0;
    double G11 = y2 - y0;

    double detG_tmp = G00*G11-G01*G10;
    double detG = fabs(detG_tmp);

    double Ginv00 =  G11 / detG_tmp;
    double Ginv01 = -G10 / detG_tmp;
    double Ginv10 = -G01 / detG_tmp;
    double Ginv11 =  G00 / detG_tmp;

    A[3*0 + 0] = (5.e-01*(Ginv01*Ginv01)*w[0][0]
                  +5.e-01*(Ginv00*Ginv00)*w[0][0]
                  +Ginv11*Ginv01*w[0][0]
                  +5.e-01*(Ginv11*Ginv11)*w[0][0]
                  +Ginv10*Ginv00*w[0][0]
                  +5.e-01*(Ginv10*Ginv10)*w[0][0])*detG;
    A[3*0 + 1] = (-5.e-01*(Ginv01*Ginv01)*w[0][0]
                  -5.e-01*(Ginv00*Ginv00)*w[0][0]
                  -5.e-01*Ginv11*Ginv01*w[0][0]
                  -5.e-01*Ginv10*Ginv00*w[0][0])*detG;

```

```

A[3*0 + 2] = (-5.e-01*Ginv11*Ginv01*w[0][0]
             -5.e-01*(Ginv11*Ginv11)*w[0][0]
             -5.e-01*Ginv10*Ginv00*w[0][0]
             -5.e-01*(Ginv10*Ginv10)*w[0][0])*detG;
A[3*1 + 0] = (-5.e-01*(Ginv01*Ginv01)*w[0][0]
             -5.e-01*(Ginv00*Ginv00)*w[0][0]
             -5.e-01*Ginv11*Ginv01*w[0][0]
             -5.e-01*Ginv10*Ginv00*w[0][0])*detG;
A[3*1 + 1] = (5.e-01*(Ginv01*Ginv01)*w[0][0]
             +5.e-01*(Ginv00*Ginv00)*w[0][0])*detG;
A[3*1 + 2] = (5.e-01*Ginv11*Ginv01*w[0][0]
             +5.e-01*Ginv10*Ginv00*w[0][0])*detG;
A[3*2 + 0] = (-5.e-01*Ginv11*Ginv01*w[0][0]
             -5.e-01*(Ginv11*Ginv11)*w[0][0]
             -5.e-01*Ginv10*Ginv00*w[0][0]
             -5.e-01*(Ginv10*Ginv10)*w[0][0])*detG;
A[3*2 + 1] = (5.e-01*Ginv11*Ginv01*w[0][0]
             +5.e-01*Ginv10*Ginv00*w[0][0])*detG;
A[3*2 + 2] = (5.e-01*(Ginv11*Ginv11)*w[0][0]
             +5.e-01*(Ginv10*Ginv10)*w[0][0])*detG;
}

```

(The expressions for A had to be edited manually to fit on the page.)

10.2 Data Flow During Code Generation

A summary of the data flow in a user application can be written as

- The user defines FiniteElement objects representing all finite element spaces.
- The user defines one BasisFunction object for each axis of the element tensor.
- The user defines one Function object for each coefficient.

- The user defines a FormRep object with one list of BasisFunction objects and one list of Function objects.
- The user populates the FormRep object with IntegralRep objects.
- The user fills in each IntegralRep object with integrand expressions.
- The user passes the FormRep object to `compile_form` or another code generation function.

If CallbackForm is used, the apparent data flow in the user code will be slightly different from the above, but a quick look at the CallbackForm code in SFC (it's quite short) should remove any confusion.

10.3 Code generation design

Here we describe the overall design of the code generation software, intended for developers who wish to extend SFC, and perhaps usable for advanced users during debugging.

The C++ interface is fixed, defined by the header file `ufc.h` from UFC. UFC also contains a utility Python module with format strings for generating UFC compliant code. An example of a format string is seen below.

```
cell_integral_implementation = """\
/// Constructor
%(classname)s::%(classname)s() : ufc::cell_integral()
{
%(constructor)s
}

/// Destructor
%(classname)s::~~%(classname)s()
{
%(destructor)s
}
```

```

}

/// Tabulate the tensor for the contribution from a local cell
void %(classname)s::tabulate_tensor(double* A,
                                   const double * const * w,
                                   const ufc::cell& c) const
{
  %(tabulate_tensor)s
}
"""

```

Each UFC class (`form`, `dof_map`, `finite_element`, `cell_integral`, `exterior_facet_integral`, `interior_facet_integral`) has separate format strings. In SFC, each UFC class is mirrored by a subclass of the class `CodeGenerator` (`FormCG`, `DofMapCG`, `FiniteElementCG`, `CellIntegralCG`, etc). These classes have only one function in common, called `generate_code()`. This function should return a tuple (`header_code`, `cpp_code`) when called, and each implementation of it follows the same pattern. Each function in the UFC interface has a format string variable (like “`%(tabulate_tensor)s`”) with the same name. For each format string variable `foo` there is a corresponding function `gen_foo()` in the `CodeGenerator` subclass. The result from each of these functions `gen_*` are code for function bodies (without the function signature), which is inserted in a dictionary that is combined with the appropriate format string from UFC. The three `ufc::*_integral` classes have the same single function `tabulate_tensor`, which is reflected by a common base class `IntegralCG` for their code generators.

```

class IntegralCG(CodeGenerator):
    def __init__(self, classname, header_format, implementation_format):
        self.classname = classname
        self.header_format = header_format
        self.implementation_format = implementation_format

    def generate_code(self):
        # generate code components for integral class
        vars = {

```

```

        'classname'      : self.classname,
        'constructor'     : indent(self.gen_constructor()),
        'destructor'      : indent(self.gen_destructor()),
        'members'         : indent(self.gen_members()),
        'tabulate_tensor' : indent(self.gen_tabulate_tensor())
    }

    # combine generated code components with
    # code templates defined in the ufc module
    hcode = self.header_format % vars

    cppcode = self.gen_members_implementation()
    cppcode += self.implementation_format % vars

    return hcode, cppcode

```

The actual code generation for the most complicated function, `tabulate_tensor`, is “outsourced” to a set of functions `gen_tabulate_tensor_*` found in `gen_tabulate_tensor.py`. Among the functions defined here are

- `gen_tabulate_tensor.cell_symbolic`
- `gen_tabulate_tensor.cell_quadrature`
- `gen_tabulate_tensor.exterior_facet_symbolic`
- `gen_tabulate_tensor.exterior_facet_quadrature`
- `gen_tabulate_tensor.interior_facet_symbolic`
- `gen_tabulate_tensor.interior_facet_quadrature`

Each of these functions take a single argument `itgrep`, which is an `Integral` object describing a single integral from a user form.

```

class CellIntegralCG(IntegralCG):
    def __init__(self, itgrep):

```

```
IntegralCG.__init__(self, classname = itgrep.classname,
                    header_format = ufc.cell_integral_header,
                    implementation_format = ufc.cell_integral_implementation)
self.itgrep = itgrep

def gen_tabulate_tensor(self):
    if self.itgrep.symbolic:
        code = gen_tabulate_tensor_cell_symbolic(self.itgrep)
    else:
        code = gen_tabulate_tensor_cell_quadrature(self.itgrep)
    return code
```

10.4 Code Formatting Utilities

To help ensure a consistent formatting of the generated code, it has proved useful to have some simple code generation utilities. These utilities assist in making the code readable, and provide some basic consistency checks to ensure that the generated code is valid.

At the most basic level, if you want to indent a multiline string, use the function `indent(text)`. This ensures a consistent indentation throughout the generated code, and removes.

One of the central utilities is the `CodeFormatter` class, which handles new-lines, braces and indentation for code with block constructs like `if`, `while`, `do`, `switch` and `class`. Errors in calls to the `CodeFormatter` will often be detected during code generation, reducing the chance of C++ compiler errors on generated code which can be cumbersome to debug. It also does consistency checks for nested block constructs if you use its functions `begin_*` and `end_*`. An example of its basic usage follows:

```
#!/usr/bin/env python
from sfc.CodeFormatter import CodeFormatter

# fictional choice of integers
```

```
facet_dofs = [(2, 0, 1), (5, 3, 4), (6, 7, 8)]

code = CodeFormatter()
code.begin_switch("facet")
for i, dofs in enumerate(facet_dofs):
    code.begin_case(i)
    for j, d in enumerate(dofs):
        code += "dofs[%d] = %d;" % (j, d)
    code.end_case()
code += "default:"
code.indent()
code += 'throw std::runtime_error("Invalid facet number.");'
code.outdent()
code.end_switch()
print str(code)
```

which produces the following nicely formatted C++ code

```
switch(facet)
{
case 0:
    dofs[0] = 2;
    dofs[1] = 0;
    dofs[2] = 1;
    break;
case 1:
    dofs[0] = 5;
    dofs[1] = 3;
    dofs[2] = 4;
    break;
case 2:
    dofs[0] = 6;
    dofs[1] = 7;
    dofs[2] = 8;
    break;
default:
```



```
    throw std::runtime_error("Invalid facet number.");  
}
```

During its filetime, the object `code` keeps track of indentation level and the current scope (on a stack). If you haven't closed all blocks when calling `str(code)`, an exception is raised.

Another useful set of helper functions can generate code for declarations, definitions, assignments and additions for a list of symbol,expression pairs, or *tokens*. A token list is a list of (symbol, expression) tuples. The functions are:

- `gen_token_declarations(tokens) → double s;`
- `gen_token_definitions(tokens) → double s = e;`
- `gen_token_assignments(tokens) → s = e;`
- `gen_token_additions(tokens) → s += e;`

