



Nektar++: Spectral/hp Element Framework

Version 5.0.2

User Guide

August 18, 2021

Department of Aeronautics, Imperial College London, UK
Scientific Computing and Imaging Institute, University of Utah, USA

Contents

Introduction	xi
I Getting Started	1
1 Installation	2
1.1 Installing Debian/Ubuntu Packages	2
1.2 Installing CentOS/Fedora Packages	3
1.3 Installing from Source	4
1.3.1 Obtaining the source code	4
1.3.2 Linux	5
1.3.3 OS X	7
1.3.4 Windows	11
1.3.5 CMake Option Reference	13
2 Mathematical Formulation	18
2.1 Background	18
2.2 Methods overview	19
2.2.1 The finite element method (FEM)	19
2.2.2 High-order finite element methods	19
2.2.3 The Galerkin formulation	21
3 XML Session File	23
3.1 Geometry	24
3.1.1 Vertices	25
3.1.2 Edges	26
3.1.3 Faces	26
3.1.4 Element	26
3.1.5 Curved Edges and Faces	27
3.1.6 Composites	27
3.1.7 Domain	28

3.2	Expansions	28
3.3	Conditions	29
3.3.1	Parameters	29
3.3.2	Solver Information	29
3.3.3	Variables	30
3.3.4	Global System Solution Algorithm	31
3.3.5	Boundary Regions and Conditions	34
3.3.6	Functions	38
3.3.7	Quasi-3D approach	39
3.4	Filters	41
3.4.1	Phase sampling	41
3.4.2	Aerodynamic forces	42
3.4.3	Benchmark	43
3.4.4	Cell history points	44
3.4.5	Checkpoint cell model	44
3.4.6	Checkpoint fields	45
3.4.7	Electrogram	45
3.4.8	FieldConvert checkpoints	46
3.4.9	History points	47
3.4.10	Kinetic energy and enstrophy	47
3.4.11	Modal energy	48
3.4.12	Moving body	49
3.4.13	Moving average of fields	49
3.4.14	One-dimensional energy	50
3.4.15	Reynolds stresses	50
3.4.16	Time-averaged fields	51
3.4.17	ThresholdMax	51
3.4.18	ThresholdMin value	52
3.5	Forcing	52
3.5.1	Absorption	52
3.5.2	Body	53
3.5.3	MovingReferenceFrame	53
3.5.4	Programmatic	53
3.5.5	Noise	54
3.6	Coupling	54
3.6.1	File	55
3.6.2	Cwipi	56
3.7	Expressions	57
3.7.1	Variables and coordinate systems	58
3.7.2	Performance considerations	62

II Preprocessing & Postprocessing 64

4 NekMesh 65

4.1	Exporting a mesh from Gmsh	66
4.2	Defining physical surfaces and volumes	66
4.3	Converting the MSH to Nektar++ format	67
4.4	NekMesh modules	68
4.4.1	Input modules	69
4.4.2	Output modules	71
4.4.3	Extract surfaces from a mesh	72
4.4.4	Negative Jacobian detection	73
4.4.5	Spherigon patches	73
4.4.6	Periodic boundary condition alignment	74
4.4.7	Boundary layer splitting	76
4.4.8	High-order cylinder generation	77
4.4.9	Linearisation	77
4.4.10	Extracting interface between tetrahedra and prismatic elements	78
4.4.11	Boundary identification	78
4.4.12	Scalar function curvature	79
4.4.13	Link Checking	79
4.4.14	2D mesh extrusion	79
4.4.15	Variational Optimisation	79
4.4.16	Mesh projection	80
4.5	Mesh generation	80
4.5.1	Methodology	81
4.5.2	Mesh generation manual	83
5	FieldConvert	87
5.1	Basic usage	87
5.1.1	Input formats	88
5.2	Convert .fld / .chk files into Paraview, VisIt or Tecplot format	88
5.3	Convert field files between XML and HDF5 format	89
5.4	Range option <i>-r</i>	89
5.5	FieldConvert modules <i>-m</i>	90
5.5.1	Smooth the data: <i>C0Projection</i> module	92
5.5.2	Calculate Q-Criterion: <i>QCriterion</i> module	92
5.5.3	Calculate λ_2 : <i>L2Criterion</i> module	92
5.5.4	Add composite ID: <i>addcompositeid</i> module	93
5.5.5	Add new field: <i>fieldfromstring</i> module	93
5.5.6	Sum two .fld files: <i>addFld</i> module	93
5.5.7	Combine two .fld files containing time averages: <i>combineAvg</i> module	94
5.5.8	Concatenate two files: <i>concatenate</i> module	94
5.5.9	Count the number of DOF: <i>dof</i> module	94
5.5.10	Equi-spaced output of data: <i>equispacedoutput</i> module	94
5.5.11	Extract a boundary region: <i>extract</i> module	95
5.5.12	Compute the gradient of a field: <i>gradient</i> module	95
5.5.13	Convert HalfMode expansion to SingleMode for further processing: <i>halfmodetofourier</i> module	95

5.5.14	Extract a plane from 3DH1D expansion: <i>homplane</i> module	96
5.5.15	Stretch a 3DH1D expansion: <i>homstretch</i> module	96
5.5.16	Inner Product of a single or series of fields with respect to a single or series of fields: <i>innerproduct</i> module	96
5.5.17	Interpolate one field to another: <i>interpfield</i> module	97
5.5.18	Interpolate scattered point data to a field: <i>interpdatapointtofld</i> module	98
5.5.19	Interpolate a field to a series of points: <i>interppts</i> module	99
5.5.20	Interpolate a set of points to another: <i>interppts</i> module	100
5.5.21	Isocontour extraction: <i>iscontour</i> module	101
5.5.22	Show high frequency energy of the Jacobian: <i>jacobianenergy</i> module	102
5.5.23	Calculate mesh quality: <i>qualitymetric</i> module	102
5.5.24	Evaluate the mean of variables on the domain: <i>mean</i> module	103
5.5.25	Extract mean mode of 3DH1D expansion: <i>meanmode</i> module	103
5.5.26	Project point data to a field: <i>pointdatapointtofld</i> module	103
5.5.27	Print L2 and LInf norms: <i>printfldnorms</i> module	104
5.5.28	Removes one or more fields from .fld files: <i>removefield</i> module . .	104
5.5.29	Computes the scalar gradient: <i>scalargrad</i> module	104
5.5.30	Scale a given .fld: <i>scaleinputfld</i> module	105
5.5.31	Time-averaged shear stress metrics: <i>shear</i> module	105
5.5.32	Stream function of a 2D incompressible flow: <i>streamfunction</i> module	106
5.5.33	Boundary layer height calculation: <i>surfdistance</i> module	106
5.5.34	Calculate vorticity: <i>vorticity</i> module	106
5.5.35	Computing the wall shear stress: <i>wss</i> module	107
5.5.36	Manipulating meshes with FieldConvert	107
5.6	FieldConvert in parallel	108
5.7	Processing large files in serial	109
5.7.1	Using the <i>part-only</i> and <i>part-only-overlapping</i> options	109
5.7.2	Using the <i>nparts</i> options	109
5.7.3	Running in parallel with the <i>nparts</i> option	110

IIISolver Applications 111

6	Acoustic Solver	112
6.1	Synopsis	112
6.1.1	Linearized Euler Equations	112
6.1.2	Acoustic Perturbation Equations	113
6.2	Usage	114
6.3	Session file configuration	114
6.3.1	Solver Info	115
6.3.2	Variables	115
6.3.3	Functions	116
6.3.4	Boundary Conditions	116
6.4	Examples	117

6.4.1	Wave Propagation in a Sheared Base Flow	117
7	Advection-Diffusion-Reaction Solver	120
7.1	Synopsis	120
7.2	Usage	121
7.3	Session file configuration	121
7.3.1	Solver Info	121
7.3.2	Parameters	122
7.3.3	Functions	122
7.4	Examples	122
7.4.1	1D Advection equation	123
7.4.2	2D Helmholtz Problem	124
7.4.3	Advection dominated mass transport in a pipe	127
7.4.4	Unsteady reaction-diffusion systems	130
8	Cardiac Electrophysiology Solver	133
8.1	Synopsis	133
8.1.1	Bidomain Model	133
8.1.2	Monodomain Model	133
8.1.3	Cell Models	134
8.2	Usage	134
8.3	Session file configuration	134
8.3.1	Solver Info	134
8.3.2	Parameters	135
8.3.3	Functions	136
8.3.4	Filters	136
8.3.5	Stimuli	136
9	Compressible Flow Solver	139
9.1	Synopsis	139
9.1.1	Euler equations	139
9.1.2	Compressible Navier-Stokes equations	140
9.1.3	Numerical discretisation	140
9.2	Usage	141
9.3	Session file configuration	141
9.4	Examples	147
9.4.1	Shock capturing	147
9.4.2	Variable polynomial order	149
9.4.3	De-Aliasing Techniques	149
10	Dummy Solver	151
10.1	Synopsis	151
11	Incompressible Navier-Stokes Solver	152
11.1	Synopsis	152

11.1.1	Velocity Correction Scheme	152
11.1.2	Direct solver (coupled approach)	159
11.1.3	Linear Stability Analysis	160
11.1.4	Steady-state solver using Selective Frequency Damping	163
11.2	Usage	164
11.3	Session file configuration	164
11.3.1	Solver Info	164
11.3.2	Parameters	168
11.3.3	Womersley Boundary Condition	168
11.3.4	Forcing	170
11.3.5	Filters	172
11.4	Session file configuration: Linear stability analysis	172
11.4.1	Solver Info	173
11.4.2	Parameters	174
11.4.3	Functions	175
11.5	Session file configuration: Steady-state solver	175
11.5.1	Execution of the classical steady-state solver	175
11.5.2	Execution of the adaptive steady-state solver	176
11.6	Session file configuration: Coordinate transformations	177
11.6.1	Solver Info	177
11.6.2	Parameters	177
11.6.3	Mapping	178
11.6.4	Functions	178
11.6.5	Boundary conditions	178
11.7	Session file configuration: Adaptive polynomial order	179
11.7.1	Solver Info	180
11.7.2	Parameters	180
11.7.3	Functions	180
11.7.4	Restarting the simulation	181
11.8	Advecting extra passive scalar fields	181
11.9	Imposing a constant flowrate	182
11.10	Examples	183
11.10.1	Kovaszny Flow 2D	183
11.10.2	Kovaszny Flow 2D using high-order outflow boundary conditions	186
11.10.3	Steady Kovaszny Oseen Flow using the direct solver	188
11.10.4	Laminar Channel Flow 2D	189
11.10.5	Laminar Channel Flow 3D	190
11.10.6	Laminar Channel Flow Quasi-3D	192
11.10.7	Turbulent Channel Flow	193
11.10.8	Turbulent Pipe Flow	196
11.10.9	Aortic Blood Flow	197
11.10.10	Infinite-strip modeling of flow past flexible cables	201
11.10.11	2D direct stability analysis of the channel flow	204
11.10.12	2D adjoint stability analysis of the channel flow	207
11.10.13	2D Transient Growth analysis of a flow past a backward-facing step	210

11.10.1 BiGlobal Floquet analysis of a of flow past a cylinder	214
12 Linear elasticity solver	219
12.1 Synopsis	219
12.1.1 The linear elasticity equations	219
12.2 Usage	220
12.3 Session file configuration	220
12.3.1 Solver Info	220
12.3.2 Parameters	221
12.4 Examples	221
12.4.1 L-shaped domain	221
12.4.2 Boundary layer deformation	222
13 Pulse Wave Solver	224
13.1 Synopsis	224
13.2 Usage	225
13.3 Session file configuration	225
13.3.1 Pulse Wave Solver mesh connectivity	225
13.3.2 Session Info	227
13.3.3 Parameters	227
13.3.4 Boundary conditions	228
13.3.5 Functions	229
13.4 Examples	229
13.4.1 Human Vascular Network	229
13.4.2 Stented Artery	234
13.4.3 Stented Artery	234
13.5 Further Information	239
13.6 Future Development	239
14 Shallow Water Solver	240
14.1 Synopsis	240
14.1.1 The Shallow Water Equations	240
14.2 Usage	241
14.3 Session file configuration	241
14.3.1 Solver Info	241
14.3.2 Parameters	241
14.3.3 Functions	241
14.4 Examples	241
14.4.1 Rossby modon case	241
IV Reference	244
15 Optimisation	245
15.1 Operator evaluation strategies	245

15.1.1	Selecting an operator strategy	246
15.1.2	XML syntax	246
15.1.3	Selecting different operator strategies	247
15.2	Collections	247
15.2.1	Default implementation	248
15.2.2	Auto-tuning	248
15.2.3	Manual selection	248
15.2.4	Collection size	249
16	Command-line Options	250
17	Frequently Asked Questions	252
17.1	Compilation and Testing	252
17.2	Usage	254
	Bibliography	256

Introduction

Nektar++ [8] is a tensor product based finite element package designed to allow one to construct efficient classical low polynomial order h -type solvers (where h is the size of the finite element) as well as higher p -order piecewise polynomial order solvers. The framework currently has the following capabilities:

- Representation of one, two and three-dimensional fields as a collection of piecewise continuous or discontinuous polynomial domains.
- Segment, plane and volume domains are permissible, as well as domains representing curves and surfaces (dimensionally-embedded domains).
- Hybrid shaped elements, i.e triangles and quadrilaterals or tetrahedra, prisms and hexahedra.
- Both hierarchical and nodal expansion bases.
- Continuous or discontinuous Galerkin operators.
- Cross platform support for Linux, Mac OS X and Windows.

The framework comes with a number of solvers and also allows one to construct a variety of new solvers.

Our current goals are to develop:

- Automatic auto-tuning of optimal operator implementations based upon not only h and p but also hardware considerations and mesh connectivity.
- Temporal and spatial adaption.
- Features enabling evaluation of high-order meshing techniques.

For further information and to download the software, visit the Nektar++ website at <http://www.nektar.info>.

Part I

Getting Started

Installation

Nektar++ is available in both a source-code distribution and as pre-compiled binary packages for a number of operating systems. We recommend using the pre-compiled packages if you wish to use the existing Nektar++ solvers for simulation and do not need to perform additional code development.

1.1 Installing Debian/Ubuntu Packages

Binary packages are available for current Debian/Ubuntu based Linux distributions. These can be installed through the use of standard system package management utilities, such as Apt, if administrative access is available.

1. Create a configuration file `/etc/apt/sources.list.d/nektar.list` containing the appropriate line from the following table.

Distribution	Repository
Debian 9.0 (stretch)	<code>deb http://www.nektar.info/debian-stretch stretch contrib</code>
Debian 10.0 (buster)	<code>deb http://www.nektar.info/debian-buster buster contrib</code>
Debian 11.0 (testing/bullseye)	<code>deb http://www.nektar.info/debian-bullseye bullseye contrib</code>
Debian (unstable)	<code>deb http://www.nektar.info/debian-unstable unstable contrib</code>
Ubuntu 16.04 (xenial xerus)	<code>deb http://www.nektar.info/ubuntu-xenial xenial contrib</code>
Ubuntu 18.04 (bionic beaver)	<code>deb http://www.nektar.info/ubuntu-bionic bionic contrib</code>
Ubuntu 20.04 (focal fossa)	<code>deb http://www.nektar.info/ubuntu-focal focal contrib</code>

2. Update the main package list `/etc/apt/sources.list` to include the non-free component, by appending (if not already present) the word `non-free` after `main`.

3. Install the Nektar++ repository GPG key:

```
wget -qO- https://www.nektar.info/nektar-apt.gpg | sudo apt-key add -
```

4. Update the package lists

```
apt update
```

5. Install specific Nektar++ packages as required, or install the complete suite with:

```
apt install nektar++
```

Any additional dependencies required by Nektar++ will be automatically installed.

Tip



Nektar++ is split into multiple packages for the different components of the software. A list of available Nektar++ packages can be found using:

```
apt search nektar++
```

1.2 Installing CentOS/Fedora Packages

1. Add a file `nektar.repo` to the directory `/etc/yum.repos.d/` with the following contents

```
[Nektar]
name=nektar
baseurl=<baseurl>
```

substituting `<baseurl>` for the appropriate line from the table below.

Distribution	<baseurl>
CentOS	<code>http://www.nektar.info/centos/\$releasever/\$basearch</code>
Fedora	<code>http://www.nektar.info/fedora/\$releasever/\$basearch</code>

Note



The `$releasever` and `$basearch` variables are automatically replaced by Yum with the OS version and architecture of your system.

2. Download and install the GPG key used to sign the packages:

```
wget https://www.nektar.info/nektar-yum.gpg
sudo rpm --import nektar-yum.gpg
```

3. Now install the Nektar++ packages as required. For example,

```
yum install nektar++-openmpi-incnavierstokes-solver
```

Any additional dependencies required by Nektar++ will be automatically installed.

Tip



Nektar++ is split into multiple packages for the different components of the software. A list of available Nektar++ packages can be found using:

```
yum search nektar++
```

1.3 Installing from Source

This section explains how to build Nektar++ from the source-code package.

Nektar++ uses a number of third-party libraries. Some of these are required, others are optional. It is generally more straightforward to use versions of these libraries supplied pre-packaged for your operating system, but if you run into difficulties with compilation errors or failing regression tests, the Nektar++ build system can automatically build tried-and-tested versions of these libraries for you. This requires enabling the relevant options in the CMake configuration.

1.3.1 Obtaining the source code

There are two ways to obtain the source code for *Nektar++*:

- Download the latest source-code archive from the [Nektar++ downloads page](#).
- Clone the git repository
 - Using anonymous access. This does not require credentials but any changes to the code cannot be pushed to the public repository. Use this initially if you would like to try using Nektar++ or make local changes to the code.

```
git clone https://gitlab.nektar.info/nektar/nektar.git nektar++
```

- Using authenticated access. This will allow you to directly contribute back into the code.

```
git clone git@gitlab.nektar.info:nektar/nektar.git nektar++
```


**Tip**

You can easily switch to using the authenticated access from anonymous access at a later date by running

```
git remote set-url origin git@gitlab.nektar.info:nektar/nektar.git
```

1.3.2 Linux**1.3.2.1 Prerequisites**

Nektar++ uses a number of external programs and libraries for some or all of its functionality. Some of these are *required* and must be installed prior to compiling *Nektar++*, most of which are available as pre-built *system* packages on most Linux distributions or can be installed manually by a *user*. Typically, the development packages, with a *-dev* or *-devel* suffix, are required to compile codes against these libraries. Others are optional and required only for specific features, or can be downloaded and compiled for use with *Nektar++* *automatically* (but not installed system-wide).

Package	Req.	Installation			Note
		Sys.	User	Auto.	
C++ compiler	✓	✓			gcc, icc, etc, supporting C++11
CMake > 2.8.11	✓	✓	✓		Ncurses GUI optional
BLAS	✓	✓	✓	✓	Or MKL, ACML, OpenBLAS
LAPACK	✓	✓	✓	✓	
Boost >= 1.56	✓	✓	✓	✓	Compile with iostreams
TinyXML	✓	✓	✓	✓	For reading XML input files
Scotch	✓	✓	✓	✓	Required for multi-level static condensation, highly recommended
METIS		✓	✓	✓	Alternative mesh partitioning
FFTW > 3.0		✓	✓	✓	For high-performance FFTs
ARPACK > 2.0		✓	✓	✓	For arnoldi algorithms
MPI		✓	✓		For parallel execution (OpenMPI, MPICH, Intel MPI, etc)
GSMPi				✓	For parallel execution
HDF5		✓	✓	✓	For large-scale parallel I/O (requires CMake >3.1)
OpenCascade CE		✓	✓	✓	For mesh generation and optimisation
PETSc		✓	✓	✓	Alternative linear solvers
PT-Scotch		✓	✓	✓	Required when MPI enabled
Tetgen		✓	✓	✓	For 3D mesh generation
Triangle		✓	✓	✓	For 2D mesh generation
VTK > 5.8		✓	✓		Not required to convert field output files to VTK, only mesh files

1.3.2.2 Quick Start

Open a terminal.

If you have downloaded the tarball, first unpack it:

```
tar -zxvf nektar++-5.0.2.tar.gz
```

Change into the `nektar++` source code directory

```
mkdir -p build && cd build
ccmake ../
make install
```

1.3.2.3 Detailed instructions

From a terminal:

1. If you have downloaded the tarball, first unpack it

```
tar -zxvf nektar++-5.0.2.tar.gz
```

2. Change into the source-code directory, create a `build` subdirectory and enter it

```
mkdir -p build && cd build
```

3. Run the CMake GUI and configure the build by pressing `c`

```
ccmake ../
```

- Select the components of Nektar++ (prefixed with `NEKTAR_BUILD_`) you would like to build. Disabling solvers which you do not require will speed up the build process.
- Select the optional libraries you would like to use (prefixed with `NEKTAR_USE_`) for additional functionality.
- Select the libraries not already available on your system which you wish to be compiled automatically (prefixed with `THIRDPARTY_BUILD_`). Some of these will be automatically enabled if not found on your system.
- Choose the installation location by adjusting `CMAKE_INSTALL_PREFIX`. By default, this will be a `dist` subdirectory within the `build` directory, which is satisfactory for most users initially.

A full list of configuration options can be found in Section 1.3.5.

Note



Selecting `THIRDPARTY_BUILD_` options will request CMake to automatically download thirdparty libraries and compile them within the *Nektar++* directory. If you have administrative access to your machine, it is recommended to install the libraries system-wide through your package-management system.

If you have installed additional system packages since running CMake, you may need to wipe your build directory and rerun CMake for them to be detected.

4. Press `c` to configure the build. If errors arise relating to missing libraries, review the `THIRDPARTY_BUILD_` selections in the configuration step above or install the missing libraries manually or from system packages.
5. When configuration completes without errors, press `c` again until the option `g` to generate build files appears. Press `g` to generate the build files and exit CMake.
6. Compile the code

```
make install
```

During the build, missing third-party libraries will be automatically downloaded, configured and built in the *Nektar++* `build` directory.

Tip



If you have multiple processors/cores on your system, compilation can be significantly increased by adding the `-jX` option to make, where X is the number of simultaneous jobs to spawn. For example, use

```
make -j4 install
```

on a quad-core system.

7. Test the build by running unit and regression tests.

```
ctest
```

1.3.3 OS X

1.3.3.1 Prerequisites

Nektar++ uses a number of external programs and libraries for some or all of its functionality. Some of these are *required* and must be installed prior to compiling

Nektar++, most of which are available on *MacPorts* (www.macports.org) or can be installed manually by a *user*. Others are optional and required only for specific features, or can be downloaded and compiled for use with Nektar++ *automatically* (but not installed system-wide).

Note



To compile *Nektar++* on OS X, Apple's Xcode Developer Tools must be installed. They can be installed either from the App Store (only on Mac OS 10.7 and above) or downloaded directly from <http://connect.apple.com/> (you are required to have an Apple Developer Connection account). Xcode includes Apple implementations of BLAS and LAPACK (called the Accelerate Framework).

Package	Req.	Installation			Note
		MacPorts	User	Auto.	
Xcode	✓				Provides developer tools
CMake > 2.8.11	✓	cmake	✓		Ncurses GUI optional
BLAS	✓				Part of Xcode
LAPACK	✓				Part of Xcode
Boost >= 1.56	✓	boost	✓	✓	Compile with iostreams
TinyXML	✓	tinyxml	✓	✓	
Scotch	✓	scotch	✓	✓	Required for multi-level static condensation, highly recommended
METIS		metis	✓	✓	Alternative mesh partitioning
FFTW > 3.0		fftw-3	✓	✓	For high-performance FFTs
ARPACK > 2.0		arpack	✓		For arnoldi algorithms
OpenMPI		openmpi			For parallel execution
GSMPi				✓	For parallel execution
HDF5			✓	✓	For large-scale parallel I/O (requires CMake >3.1)
OpenCascade CE			✓	✓	For mesh generation and optimisation
PETSc		petsc	✓	✓	Alternative linear solvers
PT-Scotch			✓	✓	Required when MPI enabled
Tetgen			✓	✓	For 3D mesh generation
Triangle			✓	✓	For 2D mesh generation
VTK > 5.8		vtk	✓		Not required to convert field output files to VTK, only mesh files

Tip

CMake, and some other software, is available from MacPorts (<http://macports.org>) and can be installed using, for example,

```
sudo port install cmake
```

Package names are given in the table above. Similar packages also exist in other package managers such as Homebrew.

1.3.3.2 Quick Start

Open a terminal (Applications->Utilities->Terminal).

If you have downloaded the tarball, first unpack it:

```
tar -zxvf nektar++-5.0.2.tar.gz
```

Change into the `nektar++` source code directory

```
mkdir -p build && cd build
ccmake ../
make install
```

1.3.3.3 Detailed instructions

From a terminal (Applications->Utilities->Terminal):

1. If you have downloaded the tarball, first unpack it

```
tar -zxvf nektar++-5.0.2.tar.gz
```

2. Change into the source-code directory, create a `build` subdirectory and enter it

```
mkdir -p build && cd build
```

3. Run the CMake GUI and configure the build

```
ccmake ../
```

Use the arrow keys to navigate the options and `ENTER` to select/edit an option.

- Select the components of Nektar++ (prefixed with `NEKTAR_BUILD_`) you would like to build. Disabling solvers which you do not require will speed up the build process.

- Select the optional libraries you would like to use (prefixed with `NEKTAR_USE_`) for additional functionality.
- Select the libraries not already available on your system which you wish to be compiled automatically (prefixed with `THIRDPARTY_BUILD_`)
- Choose the installation location by adjusting `CMAKE_INSTALL_PREFIX`. By default, this will be a `dist` subdirectory within the `build` directory, which is satisfactory for most users initially.

A full list of configuration options can be found in Section 1.3.5.

Note



Selecting `THIRDPARTY_BUILD_` options will request CMake to automatically download thirdparty libraries and compile them within the *Nektar++* directory. If you have administrative access to your machine, it is recommended to install the libraries system-wide through MacPorts.

4. Press `c` to configure the build. If errors arise relating to missing libraries (variables set to `NOTFOUND`), review the `THIRDPARTY_BUILD_` selections in the previous step or install the missing libraries manually or through MacPorts.
5. When configuration completes without errors, press `c` again until the option `g` to generate build files appears. Press `g` to generate the build files and exit CMake.
6. Compile the code

```
make install
```

During the build, missing third-party libraries will be automatically downloaded, configured and built in the *Nektar++* `build` directory.

Tip



If you have multiple processors/cores on your system, compilation can be significantly increased by adding the `-jX` option to make, where X is the number of simultaneous jobs to spawn. For example,

```
make -j4 install
```

7. Test the build by running unit and regression tests.

```
ctest
```

1.3.4 Windows

Windows compilation is supported, but the build process is somewhat convoluted at present. As such, only serial execution is supported with a minimal amount of additional build packages. These can either be installed by the user, or automatically in the build process.

Package	Req.	Installation		Note
		User	Auto.	
MS Visual Studio	✓	✓		2012, 2013 and 2015 known working
CMake ≥ 3.0	✓	✓		
BLAS	✓	✓	✓	
LAPACK	✓	✓	✓	Compile with iostreams
Boost ≥ 1.56	✓	✓	✓	

1.3.4.1 Detailed instructions

1. Install Microsoft Visual Studio 2015 (preferred), 2013 or 2012 (known to work). This can be obtained from Microsoft free of charge by using their Community developer tools from <https://www.visualstudio.com/en-us/products/visual-studio-community-vs.aspx>.
2. Install CMake 3.0+ from <http://www.cmake.org/download/>. When prompted, select the option to add CMake to the system PATH.
3. (Optional) Install Git from <http://git-scm.com/download/win> to use the development versions of *Nektar++*. When prompted, select the option to add Git to the system PATH. You do not need to select the option to add Unix tools to the PATH.
4. (Optional) If you do not wish to build boost during the compilation process (which can take some time), then boost binaries can be found at <http://sourceforge.net/projects/boost/files/boost-binaries/1.61.0/>. By default these install into `C:\local\boost_1_61_0`. If you use these libraries, you will need to:
 - Rename `libs-msvc14.0` to `lib`
 - Inside the `lib` directory, create duplicates of `boost_zlib.dll` and `boost_bzip2.dll` called `zlib.dll` and `libbz2.dll`
 - Add a `BOOST_HOME` environment variable. To do so, navigate to *Control Panel > System and Security > System*, select *Advanced System Settings*, and in the *Advanced* tab click the *Environment Variables*. In the *System Variables* box, click *New*. In the *New System Variable* window, type `BOOST_HOME` next to *Variable name* and `C:\local\boost_1_61_0` next to *Variable value*.

5. Unpack `nektar++-5.0.2.zip`.

Note



Some Windows versions do not recognise the path of a folder which has `++` in the name. If you think that your Windows version can not handle path containing special characters, you should rename `nektar++-5.0.2` to `nektar-5.0.2`.

6. Create a `builds` directory within the `nektar++-5.0.2` subdirectory.
7. Open a Visual Studio terminal. From the start menu, this can be found in *All Programs > Visual Studio 2015 > Visual Studio Tools > Developer Command Prompt for VS2015*.
8. Change directory into the `builds` directory and run the CMake graphical utility:

```
cd C:\path\to\nektar\builds
cmake-gui ..
```

9. Select the build system you want to generate build scripts for. Note that *Visual Studio 2015* is listed as *Visual Studio 14* in the drop-down list. If you have a 64-bit installation of Windows, you should select the *Win64* variant, otherwise 32-bit executables will be generated. Select the option to use the native compilers.
10. Click the *Configure* button, then the *Generate* button.
11. Return to the command line and issue the command:

```
msbuild INSTALL.vcxproj /p:Configuration=Release
```

To build in parallel with, for example, 12 processors, issue:

```
msbuild INSTALL.vcxproj /p:Configuration=Release /m:12
```

12. After the installation process is completed, the executables will be available in `builds\dist\bin`.
13. To use these executables, you need to modify your system `PATH` to include the library directories where DLLs are stored. To do this, navigate to *Control Panel > System and Security > System*, select *Advanced System Settings*, and in the *Advanced* tab click the *Environment Variables*. In the *System Variables* box, select *Path* and click *Edit*. To the end of this list, add the **full paths** to directories:
 - `builds\dist\lib\nektar++-5.0.2`
 - `builds\dist\bin`

- Optionally, if you installed Boost from the binary packages, `C:\local\boost_1_61_0 \lib`
14. To run the test suite, open a **new** command line window, change to the `builds` directory, and then issue the command

```
ctest -C Release
```

1.3.5 CMake Option Reference

This section describes the main configuration options which can be set when building Nektar++. The default options should work on almost all systems, but additional features (such as parallelisation and specialist libraries) can be enabled if needed.

1.3.5.1 Components

The first set of options specify the components of the Nektar++ toolkit to compile. Some options are dependent on others being enabled, so the available options may change.

Components of the *Nektar++* package can be selected using the following options:

- `NEKTAR_BUILD_DEMOS` (Recommended)

Compiles the demonstration programs. These are primarily used by the regression testing suite to verify the *Nektar++* library, but also provide an example of the basic usage of the framework.

- `NEKTAR_BUILD_DOC`

Compiles the Doxygen documentation for the code. This will be put in

```
$BUILDDIR/doxygen/html
```

- `NEKTAR_BUILD_LIBRARY` (Required)

Compiles the Nektar++ framework libraries. This is required for all other options.

- `NEKTAR_BUILD_PYTHON`

Installs the Python wrapper to *Nektar++*. Requires running the following command after installing *Nektar++* in order to install the Python package for the current user:

```
make nekpy-install-user
```

Alternatively, the Python package can be installed for all users by running the following command with appropriate privileges:

```
make nekpy-install-system
```

- `NEKTAR_BUILD_SOLVERS` (Recommended)

Compiles the solvers distributed with the *Nektar++* framework.

If enabling `NEKTAR_BUILD_SOLVERS`, individual solvers can be enabled or disabled. See Part III for the list of available solvers. You can disable solvers which are not required to reduce compilation time. See the `NEKTAR_SOLVER_X` option.

- `NEKTAR_BUILD_TESTS` (Recommended)

Compiles the testing program used to verify the *Nektar++* framework.

- `NEKTAR_BUILD_TIMINGS`

Compiles programs used for timing *Nektar++* operations.

- `NEKTAR_BUILD_UNIT_TESTS`

Compiles tests for checking the core library functions.

- `NEKTAR_BUILD_UTILITIES`

Compiles utilities for pre- and post-processing simulation data, including the mesh conversion and generation tool `NekMesh` and the `FieldConvert` post-processing utility.

- `NEKTAR_SOLVER_X`

Enable compilation of the 'X' solver.

- `NEKTAR_UTILITY_X`

Enable compilation of the 'X' utility.

A number of ThirdParty libraries are required by *Nektar++*. There are also optional libraries which provide additional functionality. These can be selected using the following options:

- `NEKTAR_USE_ARPACK`

Build *Nektar++* with support for ARPACK. This provides routines used for linear stability analyses. Alternative Arnoldi algorithms are also implemented directly in *Nektar++*.

- `NEKTAR_USE_CCM`

Use the ccmio library provided with the Star-CCM package for reading ccm files. This option is required as part of NekMesh if you wish to convert a Star-CCM mesh into the Nektar format. It is possible to read a Tecplot plt file from Star-CCM but this output currently needs to be converted to ascii format using the Tecplot package.

- `NEKTAR_USE_CWIPI`

Use the CWIPI library for enabling inter-process communication between two solvers. Solvers may also interface with third-party solvers using this package.

- `NEKTAR_USE_FFTW`

Build *Nektar++* with support for FFTW for performing Fast Fourier Transforms (FFTs). This is used only when using domains with homogeneous coordinate directions.

- `NEKTAR_USE_HDF5`

Build *Nektar++* with support for HDF5. This enables input/output in the HDF5 parallel file format, which can be very efficient for large numbers of processes. HDF5 output can be enabled by using a command-line option or in the `SOLVERINFO` section of the XML file. This option requires that *Nektar++* be built with MPI support with `NEKTAR_USE_MPI` enabled and that HDF5 is compiled with MPI support.

- `NEKTAR_USE_MESHGEN`

Build the NekMesh utility with support for generating meshes from CAD geometries. This enables use of the OpenCascade Community Edition library, as well as Triangle and Tetgen.

- `NEKTAR_USE_METIS`

Build *Nektar++* with support for the METIS graph partitioning library. This provides both an alternative mesh partitioning algorithm to SCOTCH for parallel simulations.

- `NEKTAR_USE_MPI` (Recommended)

Build *Nektar++* with MPI parallelisation. This allows solvers to be run in serial or parallel.

- `NEKTAR_USE_PETSC`

Build *Nektar++* with support for the PETSc package for solving linear systems.

- `NEKTAR_USE_PYTHON3` (Requires `NEKTAR_BUILD_PYTHON`)

Enables the generation of Python3 interfaces.

- `NEKTAR_USE_SCOTCH` (Recommended)

Build *Nektar++* with support for the SCOTCH graph partitioning library. This provides both a mesh partitioning algorithm for parallel simulations and enabled support for multi-level static condensation, so is highly recommended and enabled by default. However for systems that do not support SCOTCH build requirements (e.g. Windows), this can be disabled.

- `NEKTAR_USE_SYSTEM_BLAS_LAPACK` (Recommended)

On Linux systems, use the default BLAS and LAPACK library on the system. This may not be the implementation offering the highest performance for your architecture, but it is the most likely to work without problem.

- `NEKTAR_USE_VTK`

Build *Nektar++* with support for VTK libraries. This is only needed for specialist utilities and is not needed for general use.



Note

The VTK libraries are not needed for converting the output of simulations to VTK format for visualization as this is handled internally.

The `THIRDPARTY_BUILD_X` options select which third-party libraries are automatically built during the *Nektar++* build process. Below are the choices of X:

- `ARPACK`

Library of iterative Arnoldi algorithms.

- `BLAS_LAPACK`

Library of linear algebra routines.

- `BOOST`

The *Boost* libraries are frequently provided by the operating system, so automatic compilation is not enabled by default. If you do not have Boost on your system, you can enable this to have Boost configured automatically.

- `CCMIO`

I/O library for the Star-CCM+ format.

- `CWIPI`

Library for inter-process exchange of data between different solvers.

- `FFTW`

Fast-Fourier transform library.

- `GSMPI`

(MPI-only) Parallel communication library.

- `HDF5`

Hierarchical Data Format v5 library for structured data storage.

- **METIS**
A graph partitioning library used for mesh partitioning when *Nektar++* is run in parallel.
- **OCE**
OpenCascade Community Edition 3D modelling library.
- **PETSC**
A package for the parallel solution of linear algebra systems.
- **SCOTCH**
A graph partitioning library used for mesh partitioning when *Nektar++* is run in parallel, and reordering routines that are used in multi-level static condensation.
- **TETGEN**
3D tetrahedral meshing library.
- **TINYXML**
Library for reading and writing XML files.
- **TRIANGLE**
2D triangular meshing library.

There are also a number of additional options to fine-tune the build:

- **NEKTAR_DISABLE_BACKUPS**
By default, Nektar++ solvers and the FieldConvert utility will not overwrite any generated field files or output files they find an existing file with the same name. Instead, the existing file will be either moved to a backup file or you will be prompted to overwrite them. If you do not want this behaviour, then enabling this option will cause all pre-existing output to be overwritten silently.
- **NEKTAR_TEST_ALL**
Enables an extra set of more substantial and long-running tests.
- **NEKTAR_TEST_USE_HOSTFILE**
By default, MPI tests are run directly with the `mpiexec` command together with the number of cores. If your MPI installation requires a hostfile, enabling this option adds the command line argument `-hostfile hostfile` to the command line arguments when tests are run with `ctest` or the `Tester` executable.

Mathematical Formulation

2.1 Background

The spectral/hp element method combines the geometric flexibility of classical h -type finite element techniques with the desirable resolution properties of spectral methods. In this approach a polynomial expansion of order P is applied to every elemental domain of a coarse finite element type mesh. These techniques have been applied in many fundamental studies of fluid mechanics [41] and more recently have gained greater popularity in the modelling of wave-based phenomena such as computational electromagnetics [18] and shallow water problems [5] - particularly when applied within a Discontinuous Galerkin formulation.

There are at least two major challenges which arise in developing an efficient implementation of a spectral/hp element discretisation:

- implementing the mathematical structure of the technique in a digestible, generic and coherent manner, and
- designing and implementing the numerical methods and data structures in a manner so that both high- and low-order discretisations can be efficiently applied.

In order to design algorithms which are efficient for both low- and high-order spectral/hp discretisations, it is important clearly define what we mean with low- and high-order. The spectral/hp element method can be considered as bridging the gap between the high-order end of the traditional finite element method and low-order end of conventional spectral methods. However, the concept of high- and low-order discretisations can mean very different things to these different communities. For example, the seminal works by Zienkiewicz & Taylor [47] and Hughes list examples of elemental expansions only up to third or possibly fourth-order, implying that these orders are considered to be high-order for the traditional h -type finite element community. In contrast the text books of the spectral/hp element community typically show examples of problems ranging from a

low-order bound of minimally fourth-order up to anything ranging from 10^{th} -order to 15^{th} -order polynomial expansions. On the other end of the spectrum, practitioners of global (Fourier-based) spectral methods [16] would probably consider a 16^{th} -order global expansion to be relatively low-order approximation.

One could wonder whether these different definitions of low- and high-order are just inherent to the tradition and lore of each of the communities or whether there are more practical reasons for this distinct interpretation. Proponents of lower-order methods might highlight that some problems of practical interest are so geometrically complex that one cannot computationally afford to use high-order techniques on the massive meshes required to capture the geometry. Alternatively, proponents of high-order methods highlight that if the problem of interest can be captured on a computational domain at reasonable cost then using high-order approximations for sufficiently *smooth* solutions will provide a higher accuracy for a given computational cost. If one however probes even further it also becomes evident that the different communities choose to implement their algorithms in different manners. For example the standard h -type finite element community will typically uses techniques such as sparse matrix storage formats (where only the non-zero entries of a global matrix are stored) to represent a global operator. In contrast the spectral/hp element community acknowledges that for higher polynomial expansions more closely coupled computational work takes place at the individual elemental level and this leads to the use of elemental operators rather than global matrix operators. In addition the global spectral method community often make use of the tensor-product approximations where products of one-dimensional rules for integration and differentiation can be applied.

2.2 Methods overview

Here a review of some terminology in order to situate the spectral/hp element method within the field of the finite element methods.

2.2.1 The finite element method (FEM)

Nowadays, the finite element method is one of the most popular numerical methods in the field of both solid and fluid mechanics. It is a discretisation technique used to solve (a set of) partial differential equations in its equivalent variational form. The classical approach of the finite element method is to partition the computational domain into a mesh of many small subdomains and to approximate the unknown solution by piecewise linear interpolation functions, each with local support. The FEM has been widely discussed in literature and for a complete review of the method, the reader is also directed to the seminal work of Zienkiewicz and Taylor [47].

2.2.2 High-order finite element methods

While in the classical finite element method the solution is expanded in a series of linear basis functions, high-order FEMs employ higher-order polynomials to approximate the

solution. For the high-order FEM, the solution is locally expanded into a set of $P + 1$ linearly independent polynomials which span the polynomial space of order P . Confusion may arise about the use of the term *order*. While the order, or *degree*, of the expansion basis corresponds to the maximal polynomial degree of the basis functions, the order of the method essentially refers to the accuracy of the approximation. More specifically, it depends on the convergence rate of the approximation with respect to mesh-refinement. It has been shown by Babuska and Suri [3], that for a sufficiently smooth exact solution $u \in H^k(\Omega)$, the error of the FEM approximation u^δ can be bounded by:

$$\|u - u^\delta\|_E \leq Ch^P \|u\|_k.$$

This implies that when decreasing the mesh-size h , the error of the approximation algebraically scales with the P^{th} power of h . This can be formulated as:

$$\|u - u^\delta\|_E = O(h^P).$$

If this holds, one generally classifies the method as a P^{th} -order FEM. However, for non-smooth problems, i.e. $k < P + 1$, the order of the approximation will in general be lower than P , the order of the expansion.

2.2.2.1 h-version FEM

A finite element method is said to be of h -type when the degree P of the piecewise polynomial basis functions is fixed and when any change of discretisation to enhance accuracy is done by means of a mesh refinement, that is, a reduction in h . Dependent on the problem, local refinement rather than global refinement may be desired. The h -version of the classical FEM employing linear basis functions can be classified as a first-order method when resolving smooth solutions.

2.2.2.2 p-version FEM

In contrast with the h -version FEM, finite element methods are said to be of p -type when the partitioning of domain is kept fixed and any change of discretisation is introduced through a modification in polynomial degree P . Again here, the polynomial degree may vary per element, particularly when the complexity of the problem requires local enrichment. However, sometimes the term p -type FEM is merely used to indicate that a polynomial degree of $P > 1$ is used.

2.2.2.3 hp-version FEM

In the hp -version of the FEM, both the ideas of mesh refinement and degree enhancement are combined.

2.2.2.4 The spectral method

As opposed to the finite element methods which builds a solution from a sequence of local elemental approximations, spectral methods approximate the solution by a truncated series of global basis functions. Modern spectral methods, first presented by Gottlieb and Orzag [16], involve the expansion of the solution into high-order orthogonal expansion, typically by employing Fourier, Chebyshev or Legendre series.

2.2.2.5 The spectral element method

Patera [34] combined the high accuracy of the spectral methods with the geometric flexibility of the finite element method to form the spectral element method. The multi-elemental nature makes the spectral element method conceptually similar to the above mentioned high-order finite element. However, historically the term spectral element method has been used to refer to the high-order finite element method using a specific nodal expansion basis. The class of nodal higher-order finite elements which have become known as spectral elements, use the Lagrange polynomials through the zeros of the Gauss-Lobatto(-Legendre) polynomials.

2.2.2.6 The spectral/hp element method

The spectral/hp element method, as its name suggests, incorporates both the multi-domain spectral methods as well as the more general high-order finite element methods. One can say that it encompasses all methods mentioned above. However, note that the term spectral/hp element method is mainly used in the field of fluid dynamics, while the terminology p and hp -FEM originates from the area of structural mechanics.

2.2.3 The Galerkin formulation

Finite element methods typically use the Galerkin formulation to derive the weak form of the partial differential equation to be solved. We will primarily adopt the classical Galerkin formulation in combination with globally C^0 continuous spectral/hp element discretisations.

To describe the Galerkin method, consider a steady linear differential equation in a domain Ω denoted by

$$L(u) = f,$$

subject to appropriate boundary conditions. In the Galerkin method, the weak form of this equation can be derived by pre-multiplying this equation with a test function v and integrating the result over the entire domain Ω to arrive at: Find $u \in \mathcal{U}$ such that

$$\int_{\Omega} v L(u) d\mathbf{x} = \int_{\Omega} v f d\mathbf{x}, \quad \forall v \in \mathcal{V},$$

where \mathcal{U} and \mathcal{V} respectively are a suitably chosen trial and test space (in the traditional Galerkin method, one typically takes $\mathcal{U} = \mathcal{V}$). In case the inner product of v and $\mathbb{L}(u)$

can be rewritten into a bi-linear form $a(v, u)$, this problem is often formulated more concisely as: Find $u \in \mathcal{U}$ such that

$$a(v, u) = (v, f), \quad \forall v \in \mathcal{V},$$

where (v, f) denotes the inner product of v and f . The next step in the classical Galerkin finite element method is the discretisation: rather than looking for the solution u in the infinite dimensional function space \mathcal{U} , one is going to look for an approximate solution u^δ in the reduced finite dimensional function space $\mathcal{U}^\delta \subset \mathcal{U}$. Therefore we represent the approximate solution as a linear combination of basis functions Φ_n that span the space \mathcal{U}^δ , i.e.

$$u^\delta = \sum_{n \in \mathcal{N}} \Phi_n \hat{u}_n.$$

Adopting a similar discretisation for the test functions v , the discrete problem to be solved is given as: Find \hat{u}_n ($n \in \mathcal{N}$) such that

$$\sum_{n \in \mathcal{N}} a(\Phi_m, \Phi_n) \hat{u}_n = (\Phi_m, f), \quad \forall m \in \mathcal{N}.$$

It is customary to describe this set of equations in matrix form as

$$\mathbf{A} \hat{\mathbf{u}} = \hat{\mathbf{f}},$$

where $\hat{\mathbf{u}}$ is the vector of coefficients \hat{u}_n , \mathbf{A} is the system matrix with elements

$$\mathbf{A}[m][n] = a(\Phi_m, \Phi_n) = \int_{\Omega} \Phi_m L(\Phi_n) d\mathbf{x},$$

and the vector $\hat{\mathbf{f}}$ is given by

$$\hat{\mathbf{f}}[m] = (\Phi_m, f) = \int_{\Omega} \Phi_m f d\mathbf{x}.$$

XML Session File

The Nektar++ native file format is compliant with XML version 1.0. The root element is NEKTAR which contains a number of other elements which describe configuration for different aspects of the simulation. The required elements are shown below:

```
1 <NEKTAR>
2   <GEOMETRY>
3     ...
4   </GEOMETRY>
5   <EXPANSIONS>
6     ...
7   </EXPANSIONS>
8   <CONDITIONS>
9     ...
10  </CONDITIONS>
11  ...
12 </NEKTAR>
```

The different sub-elements can be split across multiple files, however each file must have a top-level NEKTAR tag. For example, one might store the geometry information separate from the remaining configuration in two separate files as illustrated below:

geometry.xml

```
1 <NEKTAR>
2   <GEOMETRY>
3     ...
4   </GEOMETRY>
5 </NEKTAR>
```

conditions.xml

```
1 <NEKTAR>
2   <CONDITIONS>
3     ...
4   </CONDITIONS>
5   <EXPANSIONS>
```

```

6   ...
7   </EXPANSIONS>
8   ...
9   </NEKTAR>

```

Note

When specifying multiple files, repeated first-level XML sub-elements are not merged. The sub-elements from files appearing later in the list will, in general, override those elements from earlier files.



For example, the `NekMesh` utility will produce a default `EXPANSIONS` element and blank `CONDITIONS` element. Specifying a custom-written XML file containing these sections *after* the file produced by `NekMesh` will override these defaults.

The exception to this rule is when an empty XML sub-element would override a non-empty XML sub-element. In this case the empty XML sub-element will be ignored. If the custom-written XML file containing `CONDITIONS` were specified before the file produced by `NekMesh`, the empty `CONDITIONS` tag in the latter file would be ignored.

3.1 Geometry

This section defines the mesh. It specifies a list of vertices, edges (in two or three dimensions) and faces (in three dimensions) and how they connect to create the elemental decomposition of the domain. It also defines a list of composites which are used in the Expansions and Conditions sections of the file to describe the polynomial expansions and impose boundary conditions.

The GEOMETRY section is structured as

```

1 <GEOMETRY DIM="2" SPACE="2">
2   <VERTEX> ... </VERTEX>
3   <EDGE> ... </EDGE>
4   <FACE> ... </FACE>
5   <ELEMENT> ... </ELEMENT>
6   <CURVED> ... </CURVED>
7   <COMPOSITE> ... </COMPOSITE>
8   <DOMAIN> ... </DOMAIN>
9 </GEOMETRY>

```

It has two (required) attributes:

- `DIM` specifies the dimension of the expansion elements.
- `SPACE` specifies the dimension of the space in which the elements exist.

These attributes allow, for example, a two-dimensional surface to be embedded in a three-dimensional space.



Note

The attribute `PARTITION` may also appear in a partitioned mesh. However, this attribute should not be explicitly specified by the user.

The contents of each of the `VERTEX`, `EDGE`, `FACE`, `ELEMENT` and `CURVED` sections may optionally be compressed and stored in base64-encoded gzipped binary form, using either little-endian or big-endian ordering, as specified by the `COMPRESSED` attribute to these sections. Currently supported values are:

- `B64Z-LittleEndian`: Base64 Gzip compressed using little-endian ordering.
- `B64Z-BigEndian`: Base64 Gzip compressed using big-endian ordering.

When generating mesh input files for *Nektar++* using `NekMesh`, the binary compressed form will be used by default. To convert a compressed XML file into human-readable ASCII format use, for example:

```
NekMesh file.msh newfile.xml:xml:uncompress
```



Note

The description in the remainder of this section explains how the `GEOMETRY` section is laid out in uncompressed ASCII format.

3.1.1 Vertices

Vertices have three coordinates. Each has a unique vertex ID. In uncompressed form, they are defined within `VERTEX` subsection as follows:

```
1 <V ID="0"> 0.0 0.0 0.0 </V> ...
```

The `VERTEX` subsection has optional attributes which can be used to apply a transformation to the mesh:

`XSCALE`, `YSCALE`, `ZSCALE`, `XMOVE`, `YMOVE`, `ZMOVE`

They specify scaling factors (centred at the origin) and translations to the vertex coordinates. For example, the following snippet

```
1 <VERTEX XSCALE="5">
2   <V ID="0"> 0.0 0.0 0.0 </V>
3   <V ID="1"> 1.0 2.0 0.0 </V>
4 </VERTEX>
```

defines two vertices with coordinates (0.0, 0.0, 0.0), (1.0, 2.0, 0.0).

All of these attributes can be arbitrary analytic expressions depending on pre-defined constants and parameters defined in the XML file and mathematical operations/functions of the latter. If omitted, default scaling factors 1.0, and translations of 0.0, are assumed.

3.1.2 Edges



Tip

The `EDGES` section is only necessary when `DIM=2` or `DIM=3` in the parent `GEOMETRY` element and may be omitted for one-dimensional meshes.

Edges are defined by two vertices. Each edge has a unique edge ID. In uncompressed form, they are defined in the file with a line of the form

```
1 <E ID="0"> 0 1 </E>
```

3.1.3 Faces



Tip

The `FACES` section is only necessary when `DIM=3` in the parent `GEOMETRY` element and may otherwise be omitted.

Faces are defined by three or more edges. Each face has a unique face ID. They are defined in the file with a line of the form

```
1 <T ID="0"> 0 1 2 </T>
2 <Q ID="1"> 3 4 5 6 </Q>
```

The choice of tag specified (T or Q), and thus the number of edges specified depends on the geometry of the face (triangle or quadrilateral).

3.1.4 Element

Elements define the top-level geometric entities in the mesh. Their definition depends upon the dimension of the expansion. For two-dimensional expansions, an element is defined by a sequence of three or four edges. For three-dimensional expansions, the element is defined by a list of faces. Elements are defined in the file with a line of the form

```
1 <T ID="0"> 0 1 2 </T>
2 <H ID="1"> 3 4 5 6 7 8 </H>
```

Again, the choice of tag specified depends upon the geometry of the element. The element tags are:

- **S** Segment
- **T** Triangle
- **Q** Quadrilateral
- **A** Tetrahedron
- **P** Pyramid
- **R** Prism
- **H** Hexahedron

3.1.5 Curved Edges and Faces



Tip

The **CURVED** section is only necessary if curved edges or faces are present in the mesh and may otherwise be omitted.

For mesh elements with curved edges and/or curved faces, a separate entry is used to describe the control points for the curve. Each curve has a unique curve ID and is associated with a predefined edge or face. The total number of points in the curve (including end points) and their distribution is also included as attributes. The control points are listed in order, each specified by three coordinates. Curved edges are defined in the file with a line of the form

```
1 <E ID="3" EDGEID="7" TYPE="PolyEvenlySpaced" NUMPOINTS="3">
2   0.0 0.0 0.0    0.5 0.5 0.0    1.0 0.0 0.0
3 </E>
```

Note



In the compressed form, this section contains different sub-elements to efficiently encode the high-order curvature data. This is not described further in this document.

3.1.6 Composites

Composites define collections of elements, faces or edges. Each has a unique composite ID associated with it. All components of a composite entry must be of the same type. The syntax allows components to be listed individually, using ranges, or a mixture of the two. Examples include

```
1 <C ID="0"> T[0-862] </C>
2 <C ID="1"> E[61-67,69,70,72-74] </C>
```

3.1.7 Domain

This tag specifies composites which describe the entire problem domain. It has the form of

```
1 <DOMAIN> C[0] </DOMAIN>
```

3.2 Expansions

This section defines the polynomial expansions used on each of the defined geometric composites and variables. Expansion entries specify the number of modes and the expansion type, or a full list of data of basis type, number of modes, points type and number of points. The short-hand version has the following form

```
1 <E COMPOSITE="C[0]" NUMMODES="5" FIELDS="u" TYPE="MODIFIED" />
```

or, if we have more then one variable we can apply the same basis to all using

```
1 <E COMPOSITE="C[0]" NUMMODES="5" FIELDS="u,v,p" TYPE="MODIFIED" />
```

The expansion basis can also be specified in detail as a combination of one-dimensional bases, and thus the user is able to, for example, increase the quadrature order. For tet elements this takes the form:

```
1 <E COMPOSITE="C[0]"
2   BASISTYPE="Modified_A,Modified_B,Modified_C"
3   NUMMODES="3,3,3"
4   POINTSTYPE="GaussLobattoLegendre,GaussRadauMAlpha1Beta0,GaussRadauMAlpha2Beta0"
5   NUMPOINTS="4,3,3"
6   FIELDS="u" />
```

and for prism elements:

```
1 <E COMPOSITE="C[1]"
2   BASISTYPE="Modified_A,Modified_A,Modified_B"
3   NUMMODES="3,3,3"
4   POINTSTYPE="GaussLobattoLegendre,GaussLobattoLegendre,GaussRadauMAlpha1Beta0"
5   NUMPOINTS="4,4,3"
6   FIELDS="u" />
```

The expansions can be defined with a list of `<E>` elements (e.g., to represent different polynomial orders for different variables or to address different composites). The user can define a default expansion field by entering `<E>` tags without the `FIELDS` attribute. The default expansion is used to define any variables not explicitly listed in the `<E>` entries. In the following example, the default expansion is used to define the expansions for the composites C[0], C[1] and C[2]:

```
1 <E COMPOSITE="C[0-2]" NUMMODES="5" TYPE="MODIFIED" />
2 <E COMPOSITE="C[3]" NUMMODES="4" TYPE="MODIFIED" FIELDS="u,v"/>
3 <E COMPOSITE="C[3]" NUMMODES="3" TYPE="MODIFIED" FIELDS="p"/>
```

The expansions of each field should be defined only once for each composite.

3.3 Conditions

This section of the file defines parameters and boundary conditions which define the nature of the problem to be solved. These are enclosed in the `CONDITIONS` tag.

3.3.1 Parameters

Numerical parameters may be required by a particular solver (for instance time-integration or physical parameters), or may be arbitrary and only used for the purpose of simplifying the problem specification in the session file (e.g. parameters which would otherwise be repeated in the definition of an initial condition and boundary conditions). All parameters are enclosed in the `PARAMETERS` XML element.

```
1 <PARAMETERS>
2   ...
3 </PARAMETERS>
```

A parameter may be of integer or real type and may reference other parameters defined previous to it. It is expressed in the file as

```
1 <P> [PARAMETER NAME] = [PARAMETER VALUE] </P>
```

For example,

```
1 <P> NumSteps = 1000 </P>
2 <P> TimeStep = 0.01 </P>
3 <P> FinTime = NumSteps*TimeStep </P>
```

A number of pre-defined constants may also be used in parameter expressions, for example `PI`. A full list of supported constants is provided in Section 3.7.1.2.

3.3.2 Solver Information

These specify properties to define the actions specific to solvers, typically including the equation to solve, the projection type and the method of time integration. The property/value pairs are specified as XML attributes. For example,

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE" VALUE="UnsteadyAdvection" />
3   <I PROPERTY="Projection" VALUE="Continuous" />
4   <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4" />
5 </SOLVERINFO>
```

Boolean-valued solver properties are specified using `True` or `False`. The list of available solvers in Nektar++ can be found in Part III.

3.3.2.1 Drivers

Drivers are defined under the `CONDITIONS` section as properties of the `SOLVERINFO` XML element. The role of a driver is to manage the solver execution from an upper level.

The default driver is called `Standard` and executes the following steps:

1. Prints out on screen a summary of all the conditions defined in the input file.
2. Sets up the initial and boundary conditions.
3. Calls the solver defined by `SolverType` in the `SOLVERINFO` XML element.
4. Writes the results in the output (.fld) file.

In the following example, the driver `Standard` is used to manage the execution of the incompressible Navier-Stokes equations:

```

1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE"           VALUE="UnsteadyNavierStokes" />
3   <I PROPERTY="SolverType"        VALUE="VelocityCorrectionScheme" />
4   <I PROPERTY="Projection"        VALUE="Galerkin" />
5   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2" />
6   <I PROPERTY="Driver"            VALUE="Standard" />
7 </SOLVERINFO>

```

If no driver is specified in the session file, the driver `Standard` is called by default. Other drivers can be used and are typically focused on specific applications. As described in Sec. 11.3.1 and 11.4.1, the other possibilities are:

- `ModifiedArnoldi` - computes of the leading eigenvalues and eigenmodes using modified Arnoldi method.
- `Arpack` - computes of eigenvalues/eigenmodes using Implicitly Restarted Arnoldi Method (ARPACK).
- `SteadyState` - uses the Selective Frequency Damping method (see Sec. 11.1.4) to obtain a steady-state solution of the Navier-Stokes equations (compressible or incompressible).

3.3.3 Variables

These define the number (and name) of solution variables. Each variable is prescribed a unique ID. For example a two-dimensional flow simulation may define the velocity variables using

```

1 <VARIABLES>
2   <V ID="0"> u </V>
3   <V ID="1"> v </V>
4 </VARIABLES>

```

3.3.4 Global System Solution Algorithm

Many *Nektar++* solvers use an implicit formulation of their equations to, for instance, improve timestep restrictions. This means that a large matrix system must be constructed and a global system set up to solve for the unknown coefficients. There are several approaches in the spectral/*hp* element method that can be used in order to improve efficiency in these methods, as well as considerations as to whether the simulation is run in parallel or serial. *Nektar++* opts for ‘sensible’ default choices, but these may or may not be optimal depending on the problem under consideration.

This section of the XML file therefore allows the user to specify the global system solution parameters, which dictates the type of solver to be used for any implicit systems that are constructed. This section is particularly useful when using a multi-variable solver such as the incompressible Navier-Stokes solver, as it allows us to select different preconditioning and residual convergence options for each variable. As an example, consider the block defined by:

```

1 <GLOBALSYSSOLNINFO>
2   <V VAR="u,v,w">
3     <I PROPERTY="GlobalSysSoln"           VALUE="IterativeStaticCond" />
4     <I PROPERTY="Preconditioner"          VALUE="LowEnergyBlock"/>
5     <I PROPERTY="IterativeSolverTolerance" VALUE="1e-8"/>
6   </V>
7   <V VAR="p">
8     <I PROPERTY="GlobalSysSoln"           VALUE="IterativeStaticCond" />
9     <I PROPERTY="Preconditioner"          VALUE="FullLinearSpaceWithLowEnergyBlock"/>
10    <I PROPERTY="IterativeSolverTolerance" VALUE="1e-6"/>
11  </V>
12 </GLOBALSYSSOLNINFO>

```

The above section specifies that the variables *u,v,w* should use the *IterativeStaticCond* global solver alongside the *LowEnergyBlock* preconditioner and an iterative tolerance of 10^{-8} on the residuals. However the pressure variable *p* is generally stiffer: we therefore opt for a more expensive *FullLinearSpaceWithLowEnergyBlock* preconditioner and a larger residual of 10^{-6} . We now outline the choices that one can use for each of these parameters and give a brief description of what they mean.

Defaults for all fields can be defined by setting the equivalent property in the *SOLVERINFO* section. Parameters defined in this section will override any options specified there.

3.3.4.1 GlobalSysSoln options

Nektar++ presently implements four methods of solving a global system:

- **Direct** solvers construct the full global matrix and directly invert it using an appropriate matrix technique, such as Cholesky factorisation, depending on the properties of the matrix. Direct solvers **only** run in serial.

- **Iterative** solvers instead apply matrix-vector multiplications repeatedly, using the conjugate gradient method, to converge to a solution to the system. For smaller problems, this is typically slower than a direct solve. However, for larger problems it can be used to solve the system in parallel execution.
- **Xxt** solvers use the XX^T library to perform a parallel direct solve. This option is only available if the `NEKTAR_USE_MPI` option is enabled in the CMake configuration.
- **PETSc** solvers use the PETSc library, giving access to a wide range of solvers and preconditioners. See section 3.3.4.4 below for some additional information on how to use the PETSc solvers. This option is only available if the `NEKTAR_USE_PETSC` option is enabled in the CMake configuration.

Warning



Both the **Xxt** and **PETSc** solvers are considered advanced and are under development – either the direct or iterative solvers are recommended in most scenarios.

These solvers can be run in one of three approaches:

- The **Full** approach constructs the global system based on all of the degrees of freedom contained within an element. For most of the *Nektar++* solvers, this technique is not recommended.
- The **StaticCond** approach applies a technique called *static condensation* to instead construct the system using only the degrees of freedom on the boundaries of the elements, which reduces the system size considerably. This is the **default option in parallel**.
- **MultiLevelStaticCond** methods apply the static condensation technique repeatedly to further reduce the system size, which can improve performance by 25-30% over the normal static condensation method. It is therefore the **default option in serial**. Note that whilst parallel execution technically works, this is under development and is likely to be slower than single-level static condensation: this is therefore not recommended.

The `GlobalSysSoln` option is formed by combining the method of solution with the approach: for example `IterativeStaticCond` or `PETScMultiLevelStaticCond`.

3.3.4.2 Preconditioner options

Preconditioners can be used in the iterative and PETSc solvers to reduce the number of iterations needed to converge to the solution. There are a number of preconditioner

choices, the default being a simple Jacobi (or diagonal) preconditioner, which is enabled by default. There are a number of choices that can be enabled through this parameter, which are all generally discretisation and dimension-dependent:

Name	Dimensions	Discretisations
<code>Null</code>	All	All
<code>Diagonal</code>	All	All
<code>FullLinearSpace</code>	2/3D	CG
<code>LowEnergyBlock</code>	3D	CG
<code>Block</code>	2/3D	All
<code>FullLinearSpaceWithDiagonal</code>	All	CG
<code>FullLinearSpaceWithLowEnergyBlock</code>	2/3D	CG
<code>FullLinearSpaceWithBlock</code>	2/3D	CG

For a detailed discussion of the mathematical formulation of these options, see the developer guide.

3.3.4.3 SuccessiveRHS options

The `SuccessiveRHS` option can be used in the iterative solver only, to attempt to reduce the number of iterations taken to converge to a solution. It stores a number of previous solutions, dictated by the setting of the `SuccessiveRHS` option, to give a better initial guess for the iterative process.

3.3.4.4 PETSc options and configuration

The PETSc solvers, although currently experimental, are operational both in serial and parallel. PETSc gives access to a wide range of alternative solver options such as GMRES, as well as any packages that PETSc can link against, such as the direct multi-frontal solver MUMPS.

Configuration of PETSc options using its command-line interface dictates what matrix storage, solver type and preconditioner should be used. This should be specified in a `.petsrc` file inside your working directory, as command line options are not currently passed through to PETSc to avoid conflict with *Nektar++* options. As an example, to select a GMRES solver using an algebraic multigrid preconditioner, and view the residual convergence, one can use the configuration:

```
-ksp_monitor
-ksp_view
-ksp_type gmres
-pc_type gamg
```

Or to use MUMPS, one could use the options:

```
-ksp_type preonly
-pc_type lu
-pc_factor_mat_solver_package mumps
-mat_mumps_icntl_7 2
```

A final choice that can be specified is whether to use a *shell* approach. By default, *Nektar++* will construct a PETSc sparse matrix (or whatever matrix is specified on the command line). This may, however, prove suboptimal for higher order discretisations. In this case, you may choose to use the *Nektar++* matrix-vector operators, which by default use an assembly approach that can prove faster, by setting the PETScMatMult SOLVERINFO option to Shell:

```
1 <I PROPERTY="PETScMatMult" VALUE="Shell" />
```

The downside to this approach is that you are now constrained to using one of the *Nektar++* preconditioners. However, this does give access to a wider range of Krylov methods than are available inside *Nektar++* for more advanced users.

3.3.5 Boundary Regions and Conditions

Boundary conditions are defined by two XML elements. The first defines the boundary regions in the domain in terms of composite entities from the **GEOMETRY** section of the file. Each boundary region has a unique ID and are defined as,

```
1 <BOUNDARYREGIONS>
2   <B ID=[id]> [composite-list] </B>
3   ...
4 </BOUNDARYREGIONS>
```

For example,

```
1 <BOUNDARYREGIONS>
2   <B ID="0"> C[2] </B>
3   <B ID="1"> C[3] </B>
4 </BOUNDARYREGIONS>
```

The second XML element defines, for each variable, the condition to impose on each boundary region, and has the form,

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="[regionID]">
3     <[type1] VAR="[variable1]" VALUE="[expression1]" />
4     ...
5     <[typeN] VAR="[variableN]" VALUE="[expressionN]" />
6   </REGION>
7   ...
8 </BOUNDARYCONDITIONS>
```

There should be precisely one `REGION` entry for each `B` entry defined in the `BOUNDARYREGION` section above. For example, to impose a Dirichlet condition on both variables for a domain with a single region,

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="sin(PI*x)*cos(PI*y)" />
4     <D VAR="v" VALUE="sin(PI*x)*cos(PI*y)" />
5   </REGION>
6 </BOUNDARYCONDITIONS>
```

Boundary condition specifications may refer to any parameters defined in the session file. The `REF` attribute corresponds to a defined boundary region. The tag used for each variable specifies the type of boundary condition to enforce.

3.3.5.1 Dirichlet (essential) condition

Dirichlet conditions are specified with the `D` tag.

Projection	Homogeneous support	Time-dependent support	Dimensions
CG	Yes	Yes	1D, 2D and 3D
DG	Yes	Yes	1D, 2D and 3D
HDG	Yes	Yes	1D, 2D and 3D

Example:

```
1 <!-- homogeneous condition -->
2 <D VAR="u" VALUE="0" />
3 <!-- inhomogeneous condition -->
4 <D VAR="u" VALUE="x^2+y^2+z^2" />
5 <!-- time-dependent condition -->
6 <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="x+t" />
```

3.3.5.2 Neumann (natural) condition

Neumann conditions are specified with the `N` tag.

Projection	Homogeneous support	Time-dependent support	Dimensions
CG	Yes	Yes	1D, 2D and 3D
DG	No	No	1D, 2D and 3D
HDG	?	?	?

Example:

```
1 <!-- homogeneous condition -->
2 <N VAR="u" VALUE="0" />
3 <!-- inhomogeneous condition -->
4 <N VAR="u" VALUE="x^2+y^2+z^2" />
5 <!-- time-dependent condition -->
6 <N VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="x+t" />
```

```

7 <!-- high-order pressure boundary condition (for IncNavierStokesSolver) -->
8 <N VAR="u" USERDEFINEDTYPE="H" VALUE="0" />

```

3.3.5.3 Periodic condition

Periodic conditions are specified with the **P** tag.

Projection	Homogeneous support	Dimensions
CG	Yes	1D, 2D and 3D
DG	No	2D and 3D

Example:

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[2] </B>
4 </BOUNDARYREGIONS>
5
6 <BOUNDARYCONDITIONS>
7   <REGION REF="0">
8     <P VAR="u" VALUE="[1]" />
9   </REGION>
10  <REGION REF="1">
11    <P VAR="u" VALUE="[0]" />
12  </REGION>
13 </BOUNDARYCONDITIONS>

```

Periodic boundary conditions are specified in a significantly different form to other conditions. The **VALUE** property is used to specify which **BOUNDARYREGION** is periodic with the current region in square brackets.

Caveats:

- A periodic condition must be set for "both" boundary regions; simply specifying a condition for region 0 or 1 in the above example is not enough.
- The order of the elements inside the composites defining periodic boundaries is important. For example, if 'C[0]' above is defined as edge IDs '0,5,4,3' and 'C[1]' as '7,12,2,1' then edge 0 is periodic with edge 7, 5 with 12, and so on.
- For the above reason, the composites must also therefore be of the same size.
- In three dimensions, care must be taken to correctly align triangular faces which are intended to be periodic. The top (degenerate) vertex should be aligned so that, if the faces were connected, it would lie at the same point on both triangles.
- It is possible specify periodic boundaries that are related by a rotation about a cartesian axis. In three-dimensions it is necessary to specify the rotational arguments to allow the orientation of each periodic face to be determined. This is

not required in two-dimensions. An example of how two periodic boundaries are related by a rotation about the x-axis of $PI/6$ is shown below. The last number specifies an optional tolerance to which the rotation is considered as equivalent (default value is $1e-8$).

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[2] </B>
4 </BOUNDARYREGIONS>
5
6 <BOUNDARYCONDITIONS>
7   <REGION REF="0">
8     <P VAR="u" USERDEFINEDTYPE="Rotated:x:PI/6:1e-6" VALUE="[1]" />
9   </REGION>
10  <REGION REF="1">
11    <P VAR="u" USERDEFINEDTYPE="Rotated:x:-PI/6:1e-6" VALUE="[0]" />
12  </REGION>
13 </BOUNDARYCONDITIONS>

```

3.3.5.4 Time-dependent boundary conditions

Time-dependent boundary conditions may be specified through setting the `USERDEFINEDTYPE` attribute and using the parameter `t` where the current time is required. For example,

```

1 <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="sin(PI*(x-t))" />

```

3.3.5.5 Boundary conditions from file

Boundary conditions can also be loaded from file. The following example is from the Incompressible Navier-Stokes solver,

```

1 <REGION REF="1">
2   <D VAR="u" FILE="Test_ChanFlow2D_bcsfromfiles_u_1.bc" />
3   <D VAR="v" VALUE="0" />
4   <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
5 </REGION>

```

Boundary conditions can also be loaded simultaneously from a file and from an expression (currently only implemented in 3D). For example, in the scenario where a spatial boundary condition is read from a file, but needs to be modulated by a time-dependent expression:

```

1 <REGION REF="1">
2   <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="sin(PI*(x-t))"
3     FILE="bcsfromfiles_u_1.bc" />
4 </REGION>

```

In the case where both `VALUE` and `FILE` are specified, the values are multiplied together to give the final value for the boundary condition.

3.3.6 Functions

Finally, multi-variable functions such as initial conditions and analytic solutions may be specified for use in, or comparison with, simulations. These may be specified using expressions (`<E>`) or imported from a file (`<F>`) using the Nektar++ FLD file format

```
1 <FUNCTION NAME="ExactSolution">
2   <E VAR="u" VALUE="sin(PI*x-advx*t))*cos(PI*(y-advy*t))" />
3 </FUNCTION>
4 <FUNCTION NAME="InitialConditions">
5   <F VAR="u" FILE="session.rst" />
6 </FUNCTION>
```

A restart file is a solution file (in other words an .fld renamed as .rst) where the field data is specified. The expansion order used to generate the .rst file must be the same as that for the simulation. .pts files contain scattered point data which needs to be interpolated to the field. For further information on the file format and the different interpolation schemes, see section 5.5.18. All filenames must be specified relative to the location of the .xml file.

With the additional argument `TIMEDEPENDENT="1"`, different files can be loaded for each timestep. The filenames are defined using `boost::format` syntax where the step time is used as variable. For example, the function `Baseflow` would load the files `UOV0_1.00000000E-05.fld`, `UOV0_2.00000000E-05.fld` and so on.

```
1 <FUNCTION NAME="Baseflow">
2   <F VAR="U0,V0" TIMEDEPENDENT="1" FILE="UOV0_%14.8E.fld"/>
3 </FUNCTION>
```

For .pts files, the time consuming computation of interpolation weights is only performed for the first timestep. The weights are stored and reused in all subsequent steps, which is why all consecutive .pts files must use the same ordering, number and location of data points.

Other examples of this input feature can be the insertion of a forcing term,

```
1 <FUNCTION NAME="BodyForce">
2   <E VAR="u" VALUE="0" />
3   <E VAR="v" VALUE="0" />
4 </FUNCTION>
5 <FUNCTION NAME="Forcing">
6   <E VAR="u" VALUE="-(Lambda + 2*PI*PI)*sin(PI*x)*sin(PI*y)" />
7 </FUNCTION>
```

or of a linear advection term

```
1 <FUNCTION NAME="AdvectionVelocity">
2   <E VAR="Vx" VALUE="1.0" />
3   <E VAR="Vy" VALUE="1.0" />
4   <E VAR="Vz" VALUE="1.0" />
5 </FUNCTION>
```

3.3.6.1 Remapping variable names

Note that it is sometimes the case that the variables being used in the solver do not match those saved in the FLD file. For example, if one runs a three-dimensional incompressible Navier-Stokes simulation, this produces an FLD file with the variables `u`, `v`, `w` and `p`. If we wanted to use this velocity field as input for an advection velocity, the advection-diffusion-reaction solver expects the variables `Vx`, `Vy` and `Vz`. We can manually specify this mapping by adding a colon to the filename, indicating the variable names in the target file that align with the desired function variable names. This gives a definition such as:

```
1 <FUNCTION NAME="AdvectionVelocity">
2   <F VAR="Vx,Vy,Vz" FILE="file.fld:u,v,w" />
3 </FUNCTION>
```

There are some caveats with this syntax:

- The same number of fields must be defined for both the `VAR` attribute and in the comma-separated list after the colon. For example, the following is not valid:

```
1 <FUNCTION NAME="AdvectionVelocity">
2   <F VAR="Vx,Vy,Vz" FILE="file.fld:u" />
3 </FUNCTION>
```

- This syntax is not valid with the wildcard operator `*`, so one cannot write for example:

```
1 <FUNCTION NAME="AdvectionVelocity">
2   <F VAR="*" FILE="file.fld:u,v,w" />
3 </FUNCTION>
```

3.3.6.2 Time-dependent file-based functions

With the additional argument `TIMEDEPENDENT="1"`, different files can be loaded for each timestep. The filenames are defined using `boost::format` syntax where the step time is used as variable. For example, the function `Baseflow` would load the files `UOV0_1.00000000E-05.fld`, `UOV0_2.00000000E-05.fld` and so on.

```
1 <FUNCTION NAME="Baseflow">
2   <F VAR="U0,V0" TIMEDEPENDENT="1" FILE="UOV0_%14.8R.fld" />
3 </FUNCTION>
```

Section 3.7 provides the list of acceptable mathematical functions and other related technical details.

3.3.7 Quasi-3D approach

To generate a Quasi-3D approach with Nektar++ we only need to create a 2D or a 1D mesh, as reported above, and then specify the parameters to extend the problem to a

3D case. For a 2D spectral/hp element problem, we have a 2D mesh and along with the parameters we need to define the problem (i.e. equation type, boundary conditions, etc.). The only thing we need to do, to extend it to a Quasi-3D approach, is to specify some additional parameters which characterise the harmonic expansion in the third direction. First we need to specify in the solver information section that that the problem will be extended to have one homogeneous dimension; here an example

```

1 <SOLVERINFO>
2   ...
3   <I PROPERTY="HOMOGENEOUS"           VALUE="1D"           />
4 </SOLVERINFO>

```

then we need to specify the parameters which define the 1D harmonic expansion along the z-axis, namely the homogeneous length (LZ) and the number of modes in the homogeneous direction (HomModesZ). HomModesZ corresponds also to the number of quadrature points in the homogeneous direction, hence on the number of 2D planes discretized with a spectral/hp element method.

```

1 <PARAMETERS>
2   ...
3   <P> HomModesZ      = 4      </P>
4   <P> LZ             = 1.0    </P>
5 </PARAMETERS>

```

In case we want to create a Quasi-3D approach starting from a 1D spectral/hp element mesh, the procedure is the same, but we need to specify the parameters for two harmonic directions (in Y and Z direction). For Example,

```

1 <SOLVERINFO>
2   ...
3   <I PROPERTY="HOMOGENEOUS"           VALUE="2D"           />
4 </SOLVERINFO>
5 <PARAMETERS>
6   ...
7   <P> HomModesY      = 10     </P>
8   <P> LY             = 6.5    </P>
9   <P> HomModesZ      = 6      </P>
10  <P> LZ             = 2.0    </P>
11 </PARAMETERS>

```

By default the operations associated with the harmonic expansions are performed with the Matrix-Vector-Multiplication (MVM) defined inside the code. The Fast Fourier Transform (FFT) can be used to speed up the operations (if the FFTW library has been compiled in ThirdParty, see the compilation instructions). To use the FFT routines we need just to insert a flag in the solver information as below:

```

1 <SOLVERINFO>
2   ...
3   <I PROPERTY="HOMOGENEOUS"           VALUE="2D"           />
4   <I PROPERTY="USEFFT"                VALUE="FFTW"         />
5 </SOLVERINFO>

```

The number of homogeneous modes has to be even. The Quasi-3D approach can be created starting from a 2D mesh and adding one homogenous expansion or starting from a 1D mesh and adding two homogeneous expansions. Not other options available. In case of a 1D homogeneous extension, the homogeneous direction will be the z-axis. In case of a 2D homogeneous extension, the homogeneous directions will be the y-axis and the z-axis.

3.4 Filters

Filters are a method for calculating a variety of useful quantities from the field variables as the solution evolves in time, such as time-averaged fields and extracting the field variables at certain points inside the domain. Each filter is defined in a `FILTER` tag inside a `FILTERS` block which lies in the main `NEKTAR` tag. In this section we give an overview of the modules currently available and how to set up these filters in the session file.

Here is an example `FILTER`:

```
1 <FILTER TYPE="FilterName">
2   <PARAM NAME="Param1"> Value1 </PARAM>
3   <PARAM NAME="Param2"> Value2 </PARAM>
4 </FILTER>
```

A filter has a name – in this case, `FilterName` – together with parameters which are set to user-defined values. Each filter expects different parameter inputs, although where functionality is similar, the same parameter names are shared between filter types for consistency. Numerical filter parameters may be expressions and so may include session parameters defined in the `PARAMETERS` section.

Some filters may perform a large number of operations, potentially taking up a significant percentage of the total simulation time. For this purpose, the parameter `IO_FiltersInfoSteps` is used to set the number of steps between successive total filter CPU time stats are printed. By default it is set to 10 times `IO_InfoSteps`. If the solver is run with the verbose `-v` flag, further information is printed, detailing the CPU time of each individual filter and percentage of time integration.

In the following we document the filters implemented. Note that some filters are solver-specific and will therefore only work for a given subset of the available solvers.

3.4.1 Phase sampling



Note

This feature is currently only supported for filters derived from the FieldConvert filter: AverageFields, MovingAverage, ReynoldsStresses.

When analysing certain time-dependent problems, it might be of interest to activate a filter in a specific physical phase and with a certain period (for instance, to carry out phase averaging). The simulation time can be written as $t = m\mathcal{T} + n_{\mathcal{T}}\mathcal{T}$, where m is an integer representing the number of periods \mathcal{T} elapsed, and $0 \leq n_{\mathcal{T}} \leq 1$ is the phase. This feature is not a filter in itself and it is activated by adding the parameters below to the filter of interest:

Option name	Required	Default	Description
<code>PhaseAverage</code>	✓		Feature activation
<code>PhaseAveragePeriod</code>	✓		Period \mathcal{T}
<code>PhaseAveragePhase</code>	✓		Phase $n_{\mathcal{T}}$.

For instance, to activate phase averaging with a period of $\mathcal{T} = 10$ at phase $n_{\mathcal{T}} = 0.5$:

```

1 <FILTER TYPE="FilterName">
2   <PARAM NAME="Param1"> Value1 </PARAM>
3   <PARAM NAME="Param2"> Value2 </PARAM>
4   <PARAM NAME="PhaseAverage"> True </PARAM>
5   <PARAM NAME="PhaseAveragePeriod"> 10 </PARAM>
6   <PARAM NAME="PhaseAveragePhase"> 0.5 </PARAM>
7 </FILTER>

```

Since this feature monitors $n_{\mathcal{T}}$ every `SampleFrequency`, for best results it is recommended to set `SampleFrequency` = 1.

The maximum error in sampling phase is $n_{\mathcal{T},\text{tol}} = \frac{\Delta t}{2\mathcal{T}} \cdot \text{SampleFrequency}$, which is displayed at the beginning of the simulation if the solver is run with the verbose `-v` option.

The number of periods elapsed is calculated based on simulation time. Caution is therefore recommended when modifying time information in the restart field, because if the new time does not correspond to the same phase, the feature will produce erroneous results.

3.4.2 Aerodynamic forces



Note

This filter is only supported for the incompressible Navier-Stokes solver.

This filter evaluates the aerodynamic forces along a specific surface. The forces are projected along the Cartesian axes and the pressure and viscous contributions are computed in each direction.

The following parameters are supported:

Option name	Required	Default	Description
OutputFile	✗	session	Prefix of the output filename to which the forces are written.
Frequency	✗	1	Number of timesteps after which output is written.
Boundary	✓	-	Boundary surfaces on which the forces are to be evaluated.

An example is given below:

```

1 <FILTER TYPE="AeroForces">
2   <PARAM NAME="OutputFile">DragLift</PARAM>
3   <PARAM NAME="OutputFrequency">10</PARAM>
4   <PARAM NAME="Boundary"> B[1,2] </PARAM>
5 </FILTER>

```

During the execution a file named `DragLift.fce` will be created and the value of the aerodynamic forces on boundaries 1 and 2, defined in the `GEOMETRY` section, will be output every 10 time steps.

3.4.3 Benchmark



Note

This filter is only supported for the Cardiac Electrophysiology Solver.

Filter `Benchmark` records spatially distributed event times for activation and repolarisation (recovert) during a simulation, for undertaking benchmark test problems.

```

1 <FILTER TYPE="Benchmark">
2   <PARAM NAME="ThresholdValue"> -40.0 </PARAM>
3   <PARAM NAME="InitialValue"> 0.0 </PARAM>
4   <PARAM NAME="OutputFile"> benchmark </PARAM>
5   <PARAM NAME="StartTime"> 0.0 </PARAM>
6 </FILTER>

```

- `ThresholdValue` specifies the value above which tissue is considered to be depolarised and below which is considered repolarised.
- `InitialValue` specifies the initial value of the activation or repolarisation time map.
- `OutputFile` specifies the base filename of activation and repolarisation maps output from the filter. This name is appended with the index of the event and the suffix `‘.fid’`.

- `StartTime` (optional) specifies the simulation time at which to start detecting events.

3.4.4 Cell history points



Note

This filter is only supported for the Cardiac Electrophysiology Solver.

Filter `CellHistoryPoints` writes all cell model states over time at fixed points. Can be used along with the `HistoryPoints` filter to record all variables at specific points during a simulation.

```

1 <FILTER TYPE="CellHistoryPoints">
2   <PARAM NAME="OutputFile">crn.his</PARAM>
3   <PARAM NAME="OutputFrequency">1</PARAM>
4   <PARAM NAME="Points">
5     0.00 0.0 0.0
6   </PARAM>
7 </FILTER>

```

- `OutputFile` specifies the filename to write history data to.
- `OutputFrequency` specifies the number of steps between successive outputs.
- `Points` lists coordinates at which history data is to be recorded.

3.4.5 Checkpoint cell model



Note

This filter is only supported for the Cardiac Electrophysiology Solver.

Filter `CheckpointCellModel` checkpoints the cell model. Can be used along with the `Checkpoint` filter to record complete simulation state and regular intervals.

```

1 <FILTER TYPE="CheckpointCellModel">
2   <PARAM NAME="OutputFile"> session </PARAM>
3   <PARAM NAME="OutputFrequency"> 1 </PARAM>
4 </FILTER>

```

- `OutputFile` (optional) specifies the base filename to use. If not specified, the session name is used. Checkpoint files are suffixed with the process ID and the extension `‘.chk’`.

- `OutputFrequency` specifies the number of timesteps between checkpoints.

3.4.6 Checkpoint fields

The checkpoint filter writes a checkpoint file, containing the instantaneous state of the solution fields at a given timestep. This can subsequently be used for restarting the simulation or examining time-dependent behaviour. This produces a sequence of files, by default named `session_*.chk`, where `*` is replaced by a counter. The initial condition is written to `session_0.chk`. Existing files are not overwritten, but renamed to e.g. `session_0.bak0.chk`. In case this file already exists, too, the `chk`-file is renamed to `session_0.bak*.chk` and so on.



Note

This functionality is equivalent to setting the `IO_CheckSteps` parameter in the session file.

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	✗	session	Prefix of the output filename to which the checkpoints are written.
<code>OutputFrequency</code>	✓	-	Number of timesteps after which output is written.

For example, to output the fields every 100 timesteps we can specify:

```

1 <FILTER TYPE="Checkpoint">
2   <PARAM NAME="OutputFile">IntermediateFields</PARAM>
3   <PARAM NAME="OutputFrequency">100</PARAM>
4 </FILTER>

```

3.4.7 Electrogram



Note

This filter is only supported for the Cardiac Electrophysiology Solver.

Filter `Electrogram` computes virtual unipolar electrograms at a prescribed set of points.

```

1 <FILTER TYPE="Electrogram">
2   <PARAM NAME="OutputFile"> session </PARAM>
3   <PARAM NAME="OutputFrequency"> 1 </PARAM>

```

```

4 <PARAM NAME="Points">
5     0.0  0.5  0.7
6     1.0  0.5  0.7
7     2.0  0.5  0.7
8 </PARAM>
9 </FILTER>

```

- `OutputFile` (optional) specifies the base filename to use. If not specified, the session name is used. The extension ‘.ecg’ is appended if not already specified.
- `OutputFrequency` specifies the number of resolution of the electrogram data.
- `Points` specifies a list of coordinates at which electrograms are desired. *They must not lie within the domain.*

3.4.8 FieldConvert checkpoints

This filter applies a sequence of FieldConvert modules to the solution, writing an output file. An output is produced at the end of the simulation into `session_fc.fld`, or alternatively every M timesteps as defined by the user, into a sequence of files `session_*_fc.fld`, where `*` is replaced by a counter.

Module options are specified as a colon-separated list, following the same syntax as the `FieldConvert` command-line utility (see Section 5).

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	X	<code>session.fld</code>	Output filename. If no extension is provided, it is assumed as .fld
<code>OutputFrequency</code>	X	<code>NumSteps</code>	Number of timesteps after which output is written, M .
<code>Modules</code>	X		FieldConvert modules to run, separated by a white space.

As an example, consider:

```

1 <FILTER TYPE="FieldConvert">
2   <PARAM NAME="OutputFile">MyFile.vtu</PARAM>
3   <PARAM NAME="OutputFrequency">100</PARAM>
4   <PARAM NAME="Modules"> vorticity isocontour:fieldid=0:fieldvalue=0.1 </PARAM>
5 </FILTER>

```

This will create a sequence of files named `MyFile_*_fc.vtu` containing isocontours. The result will be output every 100 time steps.

3.4.9 History points

The history points filter can be used to evaluate the value of the fields in specific points of the domain as the solution evolves in time. By default this produces a file called `session.his`. For each timestep, and then each history point, a line is output containing the current solution time, followed by the value of each of the field variables. Commented lines are created at the top of the file containing the location of the history points and the order of the variables.

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	✗	<code>session</code>	Prefix of the output filename to which the checkpoints are written.
<code>OutputFrequency</code>	✗	1	Number of timesteps after which output is written.
<code>OutputPlane</code>	✗	0	If the simulation is homogeneous, the plane on which to evaluate the history point. (No Fourier interpolation is currently implemented.)
<code>Points</code>	✓	-	A list of the history points. These should always be given in three dimensions.

For example, to output the value of the solution fields at three points (1, 0.5, 0), (2, 0.5, 0) and (3, 0.5, 0) into a file `TimeValues.his` every 10 timesteps, we use the syntax:

```

1 <FILTER TYPE="HistoryPoints">
2   <PARAM NAME="OutputFile">TimeValues</PARAM>
3   <PARAM NAME="OutputFrequency">10</PARAM>
4   <PARAM NAME="Points">
5     1 0.5 0
6     2 0.5 0
7     3 0.5 0
8   </PARAM>
9 </FILTER>

```

3.4.10 Kinetic energy and enstrophy



Note

This filter is only supported for the incompressible and compressible Navier-Stokes solvers **in three dimensions**.

The purpose of this filter is to calculate the kinetic energy and enstrophy

$$E_k = \frac{1}{2\mu(\Omega)} \int_{\Omega} \|\mathbf{u}\|^2 dx, \quad \mathcal{E} = \frac{1}{2\mu(\Omega)} \int_{\Omega} \|\omega\|^2 dx$$

where $\mu(\Omega)$ is the volume of the domain Ω . This produces a file containing the time-evolution of the kinetic energy and enstrophy fields. By default this file is called `session.eny` where `session` is the session name.

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	✗	<code>session.eny</code>	Output file name to which the energy and enstrophy are written.
<code>OutputFrequency</code>	✓	-	Number of timesteps at which output is written.

To enable the filter, add the following to the `FILTERS` tag:

```
1 <FILTER TYPE="Energy">
2   <PARAM NAME="OutputFrequency"> 1 </PARAM>
3 </FILTER>
```

3.4.11 Modal energy



Note

This filter is only supported for the incompressible Navier-Stokes solver.

This filter calculates the time-evolution of the kinetic energy. In the case of a two- or three-dimensional simulation this is defined as

$$E_k(t) = \frac{1}{2} \int_{\Omega} \|\mathbf{u}\|^2 dx$$

However if the simulation is written as a one-dimensional homogeneous expansion so that

$$\mathbf{u}(\mathbf{x}, t) = \sum_{k=0}^N \hat{\mathbf{u}}_k(t) e^{2\pi i k \mathbf{x}}$$

then we instead calculate the energy spectrum

$$E_k(t) = \frac{1}{2} \int_{\Omega} \|\hat{\mathbf{u}}_k\|^2 dx.$$

Note that in this case, each component of $\hat{\mathbf{u}}_k$ is a complex number and therefore $N = \text{HomModesZ}/2$ lines are output for each timestep. This is a particularly useful tool in examining turbulent and transitional flows which use the homogeneous extension. In either case, the resulting output is written into a file called `session.mdl` by default.

The following parameters are supported:

Option name	Required	Default	Description
OutputFile	X	session	Prefix of the output filename to which the energy spectrum is written.
OutputFrequency	X	1	Number of timesteps after which output is written.

An example syntax is given below:

```

1 <FILTER TYPE="ModalEnergy">
2   <PARAM NAME="OutputFile">EnergyFile</PARAM>
3   <PARAM NAME="OutputFrequency">10</PARAM>
4 </FILTER>

```

3.4.12 Moving body



Note

This filter is only supported for the Quasi-3D incompressible Navier-Stokes solver, in conjunction with the MovingBody forcing.

This filter `MovingBody` is encapsulated in the forcing module to evaluate the aerodynamic forces along the moving body surface. It is described in detail in section [11.3.4.1](#)

3.4.13 Moving average of fields

This filter computes the exponential moving average (in time) of fields for each variable defined in the session file. The moving average is defined as:

$$\bar{u}_n = \alpha u_n + (1 - \alpha) \bar{u}_{n-1}$$

with $0 < \alpha < 1$ and $\bar{u}_1 = u_1$.

The same parameters of the time-average filter are supported, with the output file in the form `session_*_movAvg.fld`. In addition, either α or the time-constant τ must be defined. They are related by:

$$\alpha = \frac{t_s}{\tau + t_s}$$

where t_s is the time interval between consecutive samples.

As an example, consider:

```

1 <FILTER TYPE="MovingAverage">
2   <PARAM NAME="OutputFile">MyMovingAverage</PARAM>
3   <PARAM NAME="OutputFrequency">100</PARAM>
4   <PARAM NAME="SampleFrequency"> 10 </PARAM>
5   <PARAM NAME="tau"> 0.1 </PARAM>
6 </FILTER>

```

This will create a file named `MyMovingAverage_movAvg.fld` with a moving average sampled every 10 time steps. The averaged field is however only output every 100 time steps.

3.4.14 One-dimensional energy

This filter is designed to output the energy spectrum of one-dimensional elements. It transforms the solution field at each timestep into a orthogonal basis defined by the functions

$$\psi_p(\xi) = L_p(\xi)$$

where L_p is the p -th Legendre polynomial. This can be used to show the presence of, for example, oscillations in the underlying field due to numerical instability. The resulting output is written into a file called `session.eny` by default. The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	X	<code>session</code>	Prefix of the output filename to which the energy spectrum is written.
<code>OutputFrequency</code>	X	1	Number of timesteps after which output is written.

An example syntax is given below:

```

1 <FILTER TYPE="Energy1D">
2   <PARAM NAME="OutputFile">EnergyFile</PARAM>
3   <PARAM NAME="OutputFrequency">10</PARAM>
4 </FILTER>

```

3.4.15 Reynolds stresses



Note

This filter is only supported for the incompressible Navier-Stokes solver.

This filter is an extended version of the time-average fields filter (see Section 3.4.16). It outputs not only the time-average of the fields, but also the Reynolds stresses. The same parameters supported in the time-average case can be used, for example:

```

1 <FILTER TYPE="ReynoldsStresses">
2   <PARAM NAME="OutputFile">MyAverageField</PARAM>
3   <PARAM NAME="RestartFile">MyAverageRst.fld</PARAM>
4   <PARAM NAME="OutputFrequency">100</PARAM>
5   <PARAM NAME="SampleFrequency"> 10 </PARAM>
6 </FILTER>

```

By default, this filter uses a simple average. Optionally, an exponential moving average can be used, in which case the output contains the moving averages and the Reynolds stresses calculated based on them. For example:

```
1 <FILTER TYPE="ReynoldsStresses">
2   <PARAM NAME="OutputFile">MyAverageField</PARAM>
3   <PARAM NAME="MovingAverage">true</PARAM>
4   <PARAM NAME="OutputFrequency">100</PARAM>
5   <PARAM NAME="SampleFrequency"> 10 </PARAM>
6   <PARAM NAME="alpha"> 0.01 </PARAM>
7 </FILTER>
```

3.4.16 Time-averaged fields

This filter computes time-averaged fields for each variable defined in the session file. Time averages are computed by either taking a snapshot of the field every timestep, or alternatively at a user-defined number of timesteps N . An output is produced at the end of the simulation into `session_avg.fld`, or alternatively every M timesteps as defined by the user, into a sequence of files `session*_avg.fld`, where `*` is replaced by a counter. This latter option can be useful to observe statistical convergence rates of the averaged variables.

This filter is derived from FieldConvert filter, and therefore support all parameters available in that case. The following additional parameter is supported:

Option name	Required	Default	Description
<code>SampleFrequency</code>	X	1	Number of timesteps at which the average is calculated, N .
<code>RestartFile</code>	X		Restart file used as initial average. If no extension is provided, it is assumed as <code>.fld</code>

As an example, consider:

```
1 <FILTER TYPE="AverageFields">
2   <PARAM NAME="OutputFile">MyAverageField</PARAM>
3   <PARAM NAME="RestartFile">MyRestartAvg.fld</PARAM>
4   <PARAM NAME="OutputFrequency">100</PARAM>
5   <PARAM NAME="SampleFrequency"> 10 </PARAM>
6 </FILTER>
```

This will create a file named `MyAverageField.fld` averaging the instantaneous fields every 10 time steps. The averaged field is however only output every 100 time steps.

3.4.17 ThresholdMax

The threshold value filter writes a field output containing a variable m , defined by the time at which the selected variable first exceeds a specified threshold value. The default

name of the output file is the name of the session with the suffix `_max.fld`. Thresholding is applied based on the first variable listed in the session by default.

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	X	<code>session_max.fld</code>	Output filename to which the threshold times are written.
<code>ThresholdVar</code>	X	<i>first variable name</i>	Specifies the variable on which the threshold will be applied.
<code>ThresholdValue</code>	✓	-	Specifies the threshold value.
<code>InitialValue</code>	✓	-	Specifies the initial time.
<code>StartTime</code>	X	0.0	Specifies the time at which to start recording.

An example is given below:

```

1 <FILTER TYPE="ThresholdMax">
2   <PARAM NAME="OutputFile"> threshold_max.fld </PARAM>
3   <PARAM NAME="ThresholdVar"> u </PARAM>
4   <PARAM NAME="ThresholdValue"> 0.1 </PARAM>
5   <PARAM NAME="InitialValue"> 0.4 </PARAM>
6 </FILTER>

```

which produces a field file `threshold_max.fld`.

3.4.18 ThresholdMin value

Performs the same function as the `ThresholdMax` filter (see Section ??) but records the time at which the threshold variable drops below a prescribed value.

3.5 Forcing

An optional section of the file allows forcing functions to be defined. These are enclosed in the `FORCING` tag. The forcing type is enclosed within the `FORCE` tag and expressed in the file as:

```

1 <FORCE TYPE="[NAME]">
2   ...
3 </FORCE>

```

The force type can be any one of the following.

3.5.1 Absorption

This force type allows the user to apply an absorption layer (essentially a porous region) anywhere in the domain. The user may also specify a velocity profile to be imposed

at the start of this layer, and in the event of a time-dependent simulation, this profile can be modulated with a time-dependent function. These velocity functions and the function defining the region in which to apply the absorption layer are expressed in the `CONDITIONS` section, however the name of these functions are defined here by the `COEFF` tag for the layer, the `REFFLOW` tag for the velocity profile, and the `REFFLOWTIME` for the time-dependent function.

```
1 <FORCE TYPE="Absorption">
2   <COEFF> [FUNCTION NAME] <COEFF/>
3   <REFFLOW> [FUNCTION NAME] <REFFLOW/>
4   <REFFLOWTIME> [FUNCTION NAME] <REFFLOWTIME/>
5   <BOUNDARYREGIONS> 1,4 <BOUNDARYREGIONS/>
6 </FORCE>
```

If a list of `BOUNDARYREGIONS` is specified, the distance to these regions is available as additional variable `r` in the definition of the `COEFF` function:

```
1 <FUNCTION NAME="AbsorptionCoefficient">
2   <E VAR="p" EVARS="r" VALUE="-5000 * exp(-0.5 * (3*r / 0.4)^2)" />
3   <E VAR="u" EVARS="r" VALUE="-5000 * exp(-0.5 * (3*r / 0.4)^2)" />
4   <E VAR="v" EVARS="r" VALUE="-5000 * exp(-0.5 * (3*r / 0.4)^2)" />
5 </FUNCTION>
```

3.5.2 Body

This force type specifies the name of a body forcing function expressed in the `CONDITIONS` section.

```
1 <FORCE TYPE="Body">
2   <BODYFORCE> [FUNCTION NAME] <BODYFORCE/>
3 </FORCE>
```

3.5.3 MovingReferenceFrame

This force type specifies the name of a moving frame function expressed in the `CONDITIONS` section.

```
1 <FORCE TYPE="MovingReferenceFrame">
2   <FRAMEVELOCITY> [FUNCTION NAME] <FRAMEVELOCITY/>
3 </FORCE>
```

The frame velocity functions defines the constant velocity of the reference frame.

```
1 <FUNCTION NAME="MovingReferenceFrame">
2   <E VAR="u" VALUE="1" />
3   <E VAR="v" VALUE="0" />
4 </FUNCTION>
```

3.5.4 Programmatic

This force type allows a forcing function to be applied directly within the code, thus it has no associated function.

```

1 <FORCE TYPE="Programmatic">
2 </FORCE>

```

3.5.5 Noise

This force type allows the user to specify the magnitude of a white noise force. Optional arguments can also be used to define the frequency in time steps to recompute the noise (default is never) and the number of time steps to apply the noise (default is the entire simulation).

```

1 <FORCE TYPE="Noise">
2   <WHITENOISE> [VALUE] <WHITENOISE/>
3   <!-- Optional arguments -->
4   <UPDATEFREQ> [VALUE] <UPDATEFREQ/>
5   <NSTEPS> [VALUE] <NSTEPS/>
6 </FORCE>

```

3.6 Coupling

Nektar++ Solvers can be run in parallel with third party applications and other Nektar++ solvers, where run-time data exchange is enabled by the coupling interface. The interface is configured in the `COUPLING` tag as

```

1 <COUPLING TYPE="[type]" NAME="[name]">
2   <I PROPERTY="SendSteps" VALUE="1" />
3   <I PROPERTY="SendVariables" VALUE="u0S,v0S" />
4   <I PROPERTY="ReceiveSteps" VALUE="1" />
5   <I PROPERTY="ReceiveVariables" VALUE="u0R,v0R" />
6   ...
7 </COUPLING>

```

The coupling type may be any of the types described later in this section, while the name can be chosen arbitrarily. Inside each coupling block, the send and receive frequencies are defined by the `SendSteps` and `ReceiveSteps` parameters, respectively. Which variables are to be sent or received is specified by the `SendVariables` and `ReceiveVariables`. By default, the send and receive frequencies is set to zero, which disables the corresponding exchange in this coupling. An empty `SendVariables` or `ReceiveVariables` list has the same effect.

Option name	Required	Default	Description
<code>SendSteps</code>	✗	0	Frequency (in steps) at which fields are sent. Sending is disabled if set to zero.
<code>SendVariables</code>	✗	<empty>	Comma-separated list of sent variables. Sending is disabled if the list is empty.
<code>ReceiveSteps</code>	✗	0	Frequency (in steps) at which fields are received. Receiving is disabled if set to zero.
<code>ReceiveVariables</code>	✗	<empty>	Comma-separated list of received variables. Receiving is disabled if the list is empty.

3.6.1 File

This coupling type allows the user to exchange fields at run time by reading from and writing to files. Besides the basic parameters which define the exchanged variables and the exchange frequency, the file coupling type requires the `SendFileName` and `ReceiveFunction` parameters to be set. The Coupling name is not used for this type and can be ignored.

```

1 <COUPLING NAME="coupling1" TYPE="File">
2   <I PROPERTY="SendSteps"      VALUE="1" />
3   <I PROPERTY="SendVariables"  VALUE="uOS,vOS" />
4   <I PROPERTY="SendFileName"   VALUE="DummyOut_%14.8E.pts" />
5   <I PROPERTY="ReceiveSteps"   VALUE="1" />
6   <I PROPERTY="ReceiveVariables" VALUE="uOR,vOR" />
7   <I PROPERTY="ReceiveFunction" VALUE="CouplingIn" />
8 </COUPLING>

```

`SendFileName` specifies a file name template to write the field data to. Currently, only `.pts` files are supported and the file is only created once fully written, avoiding race conditions between sender and receiver. Receiving is implemented by evaluating a session function specified in the `ReceiveFunction` parameter. The coupling waits for the file given in the receive function to appear.

Option name	Required	Default	Description
<code>SendFileName</code>	(✓)	-	File name where the sent fields should be written to. Required if sending is enabled. Time dependent file names are supported.
<code>ReceiveFunction</code>	(✓)	-	Function to evaluate to obtain the received fields. Required if receiving is enabled.

3.6.2 Cwipi



Note

The Cwipi coupling is only available when Nektar++ is compiled with OpenMPI and CWIPI

The Cwipi coupling uses CWIPI¹ to facilitate real time data exchange over MPI. See [24] for details. All data transfers are non-blocking to minimize the computational overhead. The interface must be enabled with the command line option `-cwipi` and a unique application name, e.g:

```
DummySolver --cwipi 'Dummy1' Dummy_3DCubeCwipi_1.xml
```

CWIPI uses the names of the current application and the coupling to identify two peers in cosimulation setups. The name of the remote application must be provided by the `RemoteName` parameter. Unlike the File-type coupling, a linear interpolation in time is applied to the received fields if non-unity values are set for `ReceiveSteps`.

```
1 <COUPLING NAME="coupling1" TYPE="Cwipi">
2   <I PROPERTY="RemoteName"      VALUE="Dummy1" />
3   <I PROPERTY="SendSteps"       VALUE="1" />
4   <I PROPERTY="SendVariables"   VALUE="uOS,vOS" />
5   <I PROPERTY="SendMethod"      VALUE="Evaluate" />
6   <I PROPERTY="ReceiveSteps"    VALUE="1" />
7   <I PROPERTY="ReceiveVariables" VALUE="uOR,vOR" />
8   <I PROPERTY="Oversample"      VALUE="5" />
9   <I PROPERTY="FilterWidth"     VALUE="10E-3" />
10  <I PROPERTY="NotLocMethod"    VALUE="Extrapolate" />
11 </COUPLING>
```

Additional options which define the coupling include `SendMethod`, the method used to retrieve the physical values at the locations requested by the remote application. Available options are `NearestNeighbour`, `Shepard` and the default `Evaluate`. The last option directly evaluates the expansions using a backward transform, giving superior accuracy at acceptable computational cost.

When using non-conforming domains, the current application might request values outside of the computational domain of the remote application. How to handle these not-located points is specified by the `NotLocMethod` parameter. When set to `keep`, the point value is not altered. With `Extrapolate`, the nearest neighbor value of the current application is used. Note that this can be very inefficient when using many MPI ranks.

¹<http://sites.onera.fr/cwipi/>

Option name	Required	Default	Description
<code>RemoteName</code>	✓	-	Name of the remote application.
<code>SendMethod</code>	✗	<code>Evaluate</code>	Specifies how to evaluate fields before sending. Available options are <code>NearestNeighbour</code> , <code>Shepard</code> and <code>Evaluate</code> .
<code>Oversample</code>	✗	0	Receive fields at a higher (or lower) number of quadrature points before filtering to avoid aliasing.
<code>FilterWidth</code>	✗	0	Apply a spatial filter of a given filter width to the received fields. Disabled when set to zero.
<code>NotLocMethod</code>	✗	<code>keep</code>	Specifies how not located points in non-conformal domains are handled. Possible values are <code>keep</code> and <code>Extrapolate</code> .

3.7 Expressions

This section discusses particulars related to expressions appearing in Nektar++. Expressions in Nektar++ are used to describe spatially or temporally varying properties, for example

- velocity profiles on a boundary
- some reference functions (e.g. exact solutions)

which can be retrieved in the solver code.

Expressions appear as the content of `VALUE` attribute of

- parameter values;
- boundary condition type tags within `<REGION>` subsection of `<BOUNDARYCONDITIONS>`, e.g. `<D>`, `<N>` etc;
- expression declaration tag `<E>` within `<FUNCTION>` subsection.

The tags above declare expressions as well as link them to one of the field variables declared in `<EXPANSIONS>` section. For example, the declaration

```
1 <D VAR="u" VALUE="sin(PI*x)*cos(PI*y)" />
```

registers expression $\sin(\pi x) \cos(\pi y)$ as a Dirichlet boundary constraint associated with field variable `u`.

Enforcing the same velocity profile at multiple boundary regions and/or field variables results in repeated re-declarations of a corresponding expression. Currently one cannot directly link a boundary condition declaration with an expression uniquely specified somewhere else, e.g. in the `<FUNCTION>` subsection. However this duplication does not affect an overall computational performance.

3.7.1 Variables and coordinate systems

Declarations of expressions are formulated in terms of problem space-time coordinates. The library code makes a number of assumptions to variable names and their order of appearance in the declarations. This section describes these assumptions.

Internally, the library uses 3D global coordinate space regardless of problem dimension. Internal global coordinate system has natural basis $(1,0,0), (0,1,0), (0,0,1)$ with coordinates \bar{x} , \bar{y} and \bar{z} . In other words, variables \bar{x} , \bar{y} and \bar{z} are considered to be first, second and third coordinates of a point $(\bar{x}, \bar{y}, \bar{z})$.

Declarations of problem spatial variables do not exist in the current XML file format. Even though field variables are declarable as in the following code snippet,

```
1 <VARIABLES>
2   <V ID="0"> u </V>
3   <V ID="1"> v </V>
4 </VARIABLES>
```

there are no analogous tags for space variables. However an attribute `SPACE` of `<GEOMETRY>` section tag declares the dimension of problem space. For example,

```
1 <GEOMETRY DIM="1" SPACE="2"> ...
2 </GEOMETRY>
```

specifies 1D flow within 2D problem space. The number of spatial variables presented in expression declaration should match space dimension declared via `<GEOMETRY>` section tag.

The library assumes the problem space also has natural basis and spatial coordinates have names \bar{x} , \bar{y} and \bar{z} .

Problem space is naturally embedded into the global coordinate space: each point of

- 1D problem space with coordinate x is represented by 3D point $(x,0,0)$ in the global coordinate system;
- 2D problem space with coordinates (x,y) is represented by 3D point $(x,y,0)$ in the global coordinate system;
- 3D problem space with coordinates (x,y,z) has the same coordinates in the global space coordinates.

Currently, there is no way to describe rotations and translations of problem space relative to the global coordinate system.

The list of variables allowed in expressions depends on the problem dimension:

- For 1D problems, expressions must make use of variable \mathbf{x} only;
- For 2D problems, expressions should make use of variables \mathbf{x} and \mathbf{y} only;
- For 3D problems, expressions may use any of variables \mathbf{x} , \mathbf{y} and \mathbf{z} .

Violation of these constraints yields unpredictable results of expression evaluation. The current implementation assigns magic value -9999 to each dimensionally excessive spacial variable appearing in expressions. For example, the following declaration

```

1  <GEOMETRY DIM="2" SPACE="2"> ...
2  </GEOMETRY> ...
3  <CONDITIONS> ...
4    <BOUNDARYCONDITIONS>
5      <REGION REF="0">
6        <D VAR="u" VALUE="x+y+z" /> <D VAR="v" VALUE="sin(PI*x)*cos(PI*y)" />
7      </REGION>
8    </BOUNDARYCONDITIONS>
9    ...
10 </CONDITIONS>

```

results in expression $x + y + z$ being evaluated at spatial points $(x_i, y_i, -9999)$ where x_i and y_i are the spacial coordinates of boundary degrees of freedom. However, the library behaviour under this constraint violation may change at later stages of development (e.g., magic constant 0 may be chosen) and should be considered unpredictable.

Another example of unpredictable behaviour corresponds to wrong ordering of variables:

```

1  <GEOMETRY DIM="1" SPACE="1"> ...
2  </GEOMETRY> ...
3  <CONDITIONS> ...
4    <BOUNDARYCONDITIONS>
5      <REGION REF="0">
6        <D VAR="u" VALUE="sin(y)" />
7      </REGION>
8    </BOUNDARYCONDITIONS>
9    ...
10 </CONDITIONS>

```

Here one declares 1D problem, so Nektar++ library assumes spacial variable is \mathbf{x} . At the same time, an expression $\sin(y)$ is perfectly valid on its own, but since it does not depend on \mathbf{x} , it will be evaluated to constant $\sin(-9999)$ regardless of degree of freedom under consideration.

3.7.1.1 Time dependence

Variable `t` represents time dependence within expressions. The boundary condition declarations need to add an additional property `USERDEFINEDTYPE="TimeDependent"` in order to flag time dependency to the library.

3.7.1.2 Syntax of expressions

Analytic expressions are formed of

- brackets `()`. Bracketing structure must be balanced.
- real numbers: every representation is allowed that is correct for `boost::lexical_cast<double>()`, e.g.

```
1 1.2, 1.2e-5, .02
```

- mathematical constants

Identifier	Meaning	Real Value
Fundamental constants		
E	Natural Logarithm	2.71828182845904523536
PI	π	3.14159265358979323846
GAMMA	Euler Gamma	0.57721566490153286060
DEG	deg/radian	57.2957795130823208768
PHI	golden ratio	1.61803398874989484820
Derived constants		
LOG2E	$\log_2 e$	1.44269504088896340740
LOG10E	$\log_{10} e$	0.43429448190325182765
LN2	$\log_e 2$	0.69314718055994530942
PI_2	$\frac{\pi}{2}$	1.57079632679489661923
PI_4	$\frac{\pi}{4}$	0.78539816339744830962
1_PI	$\frac{1}{\pi}$	0.31830988618379067154
2_PI	$\frac{2}{\pi}$	0.63661977236758134308
2_SQRTPI	$\frac{2}{\sqrt{\pi}}$	1.12837916709551257390
SQRT2	$\sqrt{2}$	1.41421356237309504880
SQRT1_2	$\frac{1}{\sqrt{2}}$	0.70710678118654752440

- parameters: alphanumeric names with underscores, e.g. `GAMMA_123`, `GaM123_45a_`, `_gamma123` are perfectly acceptable parameter names. However parameter name cannot start with a numeral. Parameters must be defined with `<PARAMETERS>...</PARAMETERS>`. Parameters play the role of constants that may change their values in between of expression evaluations.
- variables (i.e., `x`, `y`, `z` and `t`)

- unary minus operator (e.g. `-x`)
- binary arithmetic operators `+`, `-`, `*`, `/`, `^` Powering operator allows using real exponents (it is implemented with `std::pow()` function)
- boolean comparison operations `<`, `<=`, `>`, `>=`, `==` evaluate their sub-expressions to real values 0.0 or 1.0.
- mathematical functions of one or two arguments:

Identifier	Meaning
<code>abs(x)</code>	absolute value $ x $
<code>asin(x)</code>	inverse sine $\arcsin x$
<code>acos(x)</code>	inverse cosine $\arccos x$
<code>ang(x,y)</code>	computes polar coordinate $\theta = \arctan(y/x)$ from (x, y)
<code>atan(x)</code>	inverse tangent $\arctan x$
<code>atan2(y,x)</code>	inverse tangent function (used in polar transformations)
<code>ceil(x)</code>	round up to nearest integer $\lceil x \rceil$
<code>cos(x)</code>	cosine $\cos x$
<code>cosh(x)</code>	hyperbolic cosine $\cosh x$
<code>exp(x)</code>	exponential e^x
<code>fabs(x)</code>	absolute value (equivalent to <code>abs</code>)
<code>floor(x)</code>	rounding down $\lfloor x \rfloor$
<code>log(x)</code>	logarithm base e , $\ln x = \log x$
<code>log10(x)</code>	logarithm base 10, $\log_{10} x$
<code>rad(x,y)</code>	computes polar coordinate $r = \sqrt{x^2 + y^2}$ from (x, y)
<code>sin(x)</code>	sine $\sin x$
<code>sinh(x)</code>	hyperbolic sine $\sinh x$
<code>sqrt(x)</code>	square root \sqrt{x}
<code>tan(x)</code>	tangent $\tan x$
<code>tanh(x)</code>	hyperbolic tangent $\tanh x$

These functions are implemented by means of the `cmath` library: <http://www.cplusplus.com/reference/clibrary/cmath/>. Underlying data type is `double` at each stage of expression evaluation. As consequence, complex-valued expressions (e.g. $(-2)^{0.123}$) get value `nan` (not a number). The operator `^` is implemented via call to `std::pow()` function and accepts arbitrary real exponents.

- random noise generation functions. Currently implemented is `awgn(sigma)` - Gaussian Noise generator, where σ is the variance of normal distribution with zero mean. Implemented using the `boost::mt19937` random number generator with boost variate generators (see <http://www.boost.org/libs/random>)

3.7.1.3 Examples

Some straightforward examples include

- Basic arithmetic operators: $0.5 * 0.3164 / (3000^{0.25})$
- Simple polynomial functions: $y * (1 - y)$
- Use of values defined in `PARAMETERS` section: $-2 * K_{invis} * (x - 1)$
- More complex expressions involving trigonometric functions, parameters and constants: $(\text{LAMBDA} / 2 / \text{PI}) * \exp(\text{LAMBDA} * x) * \sin(2 * \text{PI} * y)$
- Boolean operators for multi-domain functions: $(y < 0) * \sin(y) + (y \geq 0) * y$

3.7.2 Performance considerations

Processing expressions is split into two stages:

- parsing with pre-evaluation of constant sub-expressions,
- evaluation to a number.

Parsing of expressions with their partial evaluation take place at the time of setting the run up (reading an XML file). Each expression, after being pre-processed, is stored internally and quickly retrieved when it turns to evaluation at given spatial-time point(s). This allows to perform evaluation of expressions at a large number of spacial points with minimal setup costs.

3.7.2.1 Pre-evaluation details

Partial evaluation of all constant sub-expressions makes no sense in using derived constants from table above. This means, either make use of pre-defined constant $\text{LN}10^2$ or straightforward expression $\log_{10}(2)^2$ results in constant 5.3018981104783980105 being stored internally after pre-processing. The rules of pre-evaluation are as follows:

- constants, numbers and their combinations with arithmetic, analytic and comparison operators are pre-evaluated,
- appearance of a variable or parameter at any recursion level stops pre-evaluation of all upper level operations (but doesn't stop pre-evaluation of independent parallel sub-expressions).

For example, declaration

```
1 <D VAR="u" VALUE="exp(-x*sin(PI*(sqrt(2)+sqrt(3))/2)*y )" />
```

results in expression $\exp(-x * (-0.97372300937516503167) * y)$ being stored internally: sub-expression $\sin(\text{PI} * (\text{sqrt}(2) + \text{sqrt}(3)) / 2)$ is evaluated to constant but appearance of x and y variables stops further pre-evaluation.

Grouping predefined constants and numbers together helps. Its useful to put brackets to be sure your constants do not run out and become factors of some variables or parameters.

Expression evaluator does not do any clever simplifications of input expressions, which is clear from example above (there is no point in double negation). The following subsection addresses the simplification strategy.

3.7.2.2 Preparing expression

The total evaluation cost depends on the overall number of operations. Since evaluator is not making simplifications, it worth trying to minimise the total number of operations in input expressions manually.

Some operations are more computationally expensive than others. In an order of increasing complexity:

- `+, -, <, >, <=, >=, ==,`
- `*, /, abs, fabs, ceil, floor,`
- `^, sqrt, exp, log, log10, sin, cos, tan, sinh, cosh, tanh, asin, acos, atan.`

For example,

- `x*x` is faster than `x^2` — it is one double multiplication vs generic calculation of arbitrary power with real exponents.
- `(x+sin(y))^2` is faster than `(x+sin(y))*(x+sin(y))` - sine is an expensive operation. It is cheaper to square complicated expression rather than compute it twice and add one multiplication.
- An expression `exp(-41*((x+(0.3*cos(2*PI*t)))^2 + (0.3*sin(2*PI*t))^2))` makes use of 5 expensive operations (`exp`, `sin`, `cos` and power `^` twice) while an equivalent expression `exp(-41*(x*x+0.6*x*cos(2*PI*t) + 0.09))` uses only 2 expensive operations.

If any simplifying identity applies to input expression, it may worth applying it, provided it minimises the complexity of evaluation. Computer algebra systems may help.

3.7.2.3 Vectorized evaluation

Expression evaluator is able to calculate an expression for either given point (its space-time coordinates) or given array of points (arrays of their space-time coordinates, it uses SoA). Vectorized evaluation is faster then sequential due to a better data access pattern. Some expressions give measurable speedup factor 4.6. Therefore, if you are creating your own solver, it worth making vectorized calls.

Part II

Preprocessing & Postprocessing

NekMesh

NekMesh is a utility bundled with *Nektar++* which has two purposes:

- allow foreign mesh file formats to be converted into *Nektar++*'s XML format;
- aide in the generation of high-order meshes through a series of supplied processing modules.

Note



NekMesh replaces a previous utility called MeshConvert. This change is to reflect the fact that the program no longer only converts and manipulates meshes but can now also generate them from a CAD definition. This mesh generator is in an early stage of development and as such is disabled by default. For the time being those not using the mesh generator can use **NekMesh** as they would have used MeshConvert, none of the functionality or methodology has changed.

There is also some limited support for other output formats. We begin by running through a basic example to show how a mesh can be converted from the widely-used mesh-generator **Gmsh** to the XML file format.

Note



The default since January 2016 is to output the `.xml` files in a compressed form where the VERTEX, EDGES, FACES, ELEMENTS and CURVED information is compressed into binary format which is then converted into base64. This is identified for each section by the attribute `COMPRESSED="B64Z-LittleEndian"`. To output in ascii format add the module option `":xml:uncompress"` to the `.xml` file, i.e.

```
NekMesh file.msh newfile.xml:xml:uncompress
```

4.1 Exporting a mesh from Gmsh

To demonstrate how NekMesh works, we will define a simple channel-like 3D geometry. First, we must define the Gmsh geometry to be used. The Gmsh definition is given below, and is visualised in figure 4.1.

```

1 Point(1) = {-1, 0, 0, 1.0};
2 Point(2) = {-0.3, 0, 0, 1.0};
3 Line(3) = {1, 2};
4 s[] = Extrude {0, 0, 7} {
5   Line{3}; Layers{5}; Recombine;
6 };
7 v[] = Extrude {{0, 0, 1}, {0, 0, 0}, Pi} {
8   Surface{s[1]}; Layers{10}; Recombine;
9 };

```

Whilst a full tutorial on Gmsh is far beyond the scope of this document, note the use of the `Recombine` argument. This allows us to generate a structured hexahedral mesh; remove the first `Recombine` to generate a prismatic mesh and both occurrences to generate a tetrahedral mesh. Increasing the `Layers` numbers refines the mesh in the radial and azimuthal direction respectively.

4.2 Defining physical surfaces and volumes




Figure 4.1 Geometry definition in Gmsh (left) and resulting high-order mesh visualised in ParaView (right).

In order for us to use the mesh, we need to define the physical surfaces which correspond to the inflow, outflow and walls so that we can set appropriate boundary conditions. The numbering resulting from the extrusions in this case is not straightforward. In the graphical interface, select **Geometry > Physical Groups > Add > Surface**, and then hover over each of the surfaces which are shown by the dashed gray lines. The numbering will be revealed in the toolbar underneath the geometry as a ruled surface. In this case:

- **Walls:** surfaces 7, 8, 28, 29.
- **Inflow:** surface 16.

- **Outflow:** surface 24.

We also need to define the physical volumes, which can be done in a similar fashion. For this example, there is only one volume having ID 1. Adding these groups to the end of the `.geo` file is very straightforward:

```
1 Physical Volume(0) = {1};
2 Physical Surface(1)= {7,8,28,29};
3 Physical Surface(2) = {16};
4 Physical Surface(3) = {24};
```

Either choose the option **File->Save Mesh** or, assuming this is saved in a file named `test.geo`, run the command

```
gmsh -3 test.geo
```

which will produce the resulting MSH file `test.msh`. One can generate a high-order mesh by specifying the order on the command line, for example

```
gmsh -3 -order 6 test.geo
```

will generate a sixth-order mesh. Note that you will need to use a current version of `Gmsh` in order to do this, most likely from subversion.

4.3 Converting the MSH to Nektar++ format

Assuming that you have compiled *Nektar++* according to the compilation instructions, run the command

```
NekMesh test.msh test.xml
```

to generate the XML file.

Note



This file contains only the geometry definition (and a default `EXPANSIONS` definition). In order to use this mesh, a `CONDITIONS` section must be supplied detailing the solver and parameters to use.

To validate the mesh visually, we can use a utility such as Paraview or VisIt. To do this, we can use the `FieldConvert` command using:

```
FieldConvert test.xml test.vtu
```

which generates an unstructured VTK file `test.vtu`.

It is possible that, when the high-order information was inserted into the mesh by `Gmsh`, invalid elements are generated which self intersect. In this case, the Jacobian of the mapping defining the curvature will have negative regions, which will generate warnings such as:

```
Warning: Level 0 assertion violation
3D deformed Jacobian not positive (element ID = 48) (first vertex ID = 105)
```

This tells you the element ID that is invalid, and the ID of the first vertex of the element. Whilst a resulting simulation may run, the results may not be valid because of this problem, or excessively large amounts of time may be needed to solve the resulting linear system.

4.4 NekMesh modules

`NekMesh` is designed to provide a pipeline approach to mesh generation. To do this, we break up tasks into three different types. Each task is called a *module* and a chain of modules specifies the pipeline.

- **Input** modules read meshes in a variety of formats;
- **Processing** modules modify meshes to aide in generation processes;
- **Output** modules write meshes in a variety of formats.

The figure below depicts how these might be coupled together to form a pipeline: On the




Figure 4.2 Illustrative pipeline of the `NekMesh` process.

command line, we would define this as:

```
NekMesh -m process1 -m process2 input.msh output.xml
```

Process modules can also have parameters passed to them, that can take arguments, or not.

```
NekMesh -m process1:p1=123:booleanparam input.msh output.xml
```

To list all available modules use the `-l` command line argument:


```

Available classes:
Input: dat:
    Reads Tecplot polyhedron ascii format converted from Star CCM (.dat).
...

```

and then to see the options for a particular module, use the `-p` command line argument:

```

Options for module detect:
    vol: Tag identifying surface to process.

```

Note



Module names change when you use the `-p` option. Input modules should be preceded by `in:`, processing modules by `proc:` and output modules by `out:`.

4.4.1 Input modules

Input and output modules use file extension names to determine the correct module to use. Not every module is capable of reading high-order information, where it exists. The table below indicates support currently implemented.

Format	Extension	High-order	Notes
Gmsh	msh	✓	Only reads nodes, elements and physical groups (which are mapped to composites). File format versions 2.x and 4.x currently supported.
Nektar	rea	✓	Reads elements, fluid boundary conditions. Most curve types are unsupported: high-order information must be defined in an accompanying .hsf file.
Nektar++	xml	✓	Fully supported.
PLY	ply	✗	Reads only the ASCII format..
Semtex	sem	✓	Reads elements and boundary conditions. In order to read high-order information, run <code>meshpr session.sem > session.msh</code> and place in the same directory as the session file.
Star-CCM+	dat	✗	Star outputs plt file which currently needs to be converted to ascii using Tecplot. Reads mesh only, only support for quads and triangles (2D) and hexes, prisms, tetrahedra (3D).
Star-CCM+	ccm	✗	Reads start ccm format. Reads mesh only, only support for quads and triangles (2D) and hexes, prisms, tetrahedra (3D). Requires NEKTAR_USE_CCM option to be activated in cmake and then requires ccmio library to be compiled by user.
VTK	vtk	✗	Experimental support. Only ASCII triangular data is supported.

Note that you can override the module used on the command line. For example, **Semtex** session files rarely have extensions. So for a session called `pipe-3d` we can convert this using the syntax

```
NekMesh pipe-3d:sem pipe-3d.xml
```

Typically, mesh generators allow physical surfaces and volumes to contain many element types; for example a cube could be constructed from a mixture of hexes and prisms. In *Nektar++*, a composite can only contain a single element type. Whilst the converter will attempt to preserve the numbering of composites from the original mesh type, sometimes a renumbering will occur when a domain contains many element types. For example, for a domain with the tag (150) containing quadrilaterals and triangles, the Gmsh reader will print a notification along the lines of:

```
Multiple elements in composite detected; remapped:
- Tag 150 => 150 (Triangle), 151 (Quadrilateral)
```

The resulting file therefore has two composites of IDs `150` and `151` respectively, containing the triangular and quadrilateral elements of the original mesh.

4.4.2 Output modules

The following output formats are supported:

Format	Extension	High-order	Notes
Gmsh	msh	✓	High-order hexes, quads, tetrahedra and triangles are supported up to arbitrary order. Prisms supported up to order 4, pyramids up to order 1.
Nektar++	xml	✓	Most functionality supported.
HDF5	nekg	✓	Most functionality supported.
VTk	vtk	✗	Outputs mesh only, supports line segments in 1D, triangles and quadrilaterals in 2D and tetrahedra, hexahedra, prisms and pyramids in 3D. The VTK legacy format and XML format, both compressed and uncompressed, are supported. Requires <code>NEKTAR_USE_VTK</code> option to be activated in <code>cmake</code> .

Note that for `Gmsh`, it is highly likely that you will need to experiment with the source code in order to successfully generate meshes since robustness is not guaranteed.

The default for `xml` and `vtk` is into binary data which has been converted into base64. If you wish to see an ascii output you need to specify the output module option `uncompress`. For the uncompressed `xml` the user can execute:

```
NekMesh Mesh.msh output.xml:xml:uncompress
```

If the user wishes to obtain the `vtk` output in the legacy format, the output module option `legacy` should be specified by executing:

```
NekMesh Mesh.xml output.vtk:vtk:legacy
```

Finally, both the Gmsh and Nektar++ output modules support an `order` parameter, which allows you to generate a mesh of a uniform polynomial order. This is used in the same manner as the above, so that the command

```
NekMesh Mesh.msh output.msh:msh:order=7
```

will generate an order 7 Gmsh mesh. In the rest of these subsections, we discuss the various processing modules available within `NekMesh`.

4.4.2.1 HDF5 format

`NekMesh` and all solvers within Nektar++ - along with subsequent `FieldConvert` modules - also support the HDF5 format. This allows for faster loading of geometries and meshes within each solver - and is a significant improvement over the XML format. HDF5 is recommended input format for any larger cases.

Converting from XML to HDF5 is a simple task that only requires the one `NekMesh` command:

```
NekMesh XMLMesh.xml HDF5Mesh.nekg
```

This will create two files `HDF5Mesh.xml` and `HDF5Mesh.nekg` which are both needed in the same directory to run the simulation. An additional flag in the session file is required, ensuring it is placed before the expansion list being:

```
1 <GEOMETRY DIM="3" SPACE="3" HDF5FILE="HDF5Mesh.nekg" />
```

HDF5 also has the additional advantage of ensuring the mesh and session file are split - which allows for easy amending of the session file - whilst allowing for use of `FieldConvert` modules that require only 1 XML input file - rather than having to concatenate the session and mesh XML files. Solvers and any `FieldConvert` modules can be run by referencing only the session file after the `GEOMETRY` tag is included.

4.4.3 Extract surfaces from a mesh

Often one wants to visualise surfaces of a 3D mesh, or extract the values of variables on the surface and visualise them. To support this, `NekMesh` can extract two-dimensional surfaces which can be visualised using `FieldConvert` in order to extract the value of a 3D field on a given surface.

As an example, we can extract composite surfaces 2 and 3-5 from a mesh using the `extract` module:

```
NekMesh -m extract:surf=2,3-5 Mesh.xml output.xml
```

If you also wish to have the boundaries of the extracted surface detected add the `detectbnd` option

```
NekMesh -m extract:surf=2,3-5:detectbnd Mesh.xml output.xml
```

which will produce new composites for the extracted boundary.

4.4.4 Negative Jacobian detection

To detect elements with negative Jacobian determinant, use the `jac` module:

```
NekMesh -m jac Mesh.xml output.xml
```

To get a detailed list of elements which have negative Jacobians, one may use the `list` option:

```
NekMesh -m jac:list Mesh.xml output.xml
```

and to extract the elements for the purposes of visualisation within the domain, use the `extract` boolean parameter:

```
NekMesh -m jac:extract Mesh.xml MeshWithNegativeElements.xml
```

To turn off curvature associated with negative jacobians one can try to use the `linearise` module:

```
NekMesh -m linerise:invalid Mesh.xml output.xml
```

This option will remove all high order curvature on all element types with singular jacobians.

4.4.5 Spherigon patches

Where high-order information is not available (e.g. when using meshes from imaging software), various techniques can be used to apply a smoothing to the high-order element. In *NekMesh* we use *spherigons*, a kind of patch used in the computer graphics community used for efficiently smoothing polygon surfaces.

Spherigons work through the use of surface normals, where in this sense ‘surface’ refers to the underlying geometry. If we have either the exact or approximate surface normal at each given vertex, spherigon patches approximate the edges connecting two vertices by arcs of a circle. In *NekMesh* we can either approximate the surface normals from the linear elements which connect to each vertex (this is done by default), or supply a file which gives the surface normals.

To apply spherigon patches on two connected surfaces 11 and 12 use the following command:

```
NekMesh -m spherigon:surf=11,12 \
  MeshWithStraighEdges.xml MeshWithSpherigons.xml
```

If the two surfaces "11" and "12" are not connected, or connect at a sharp edge which is C^0 continuous but not C^1 smooth, use two separate instances of the spherigon module.

```
NekMesh -m spherigon:surf=11 -m spherigon:surf=12 \
  MeshWithStraighEdges.xml MeshWithSpherigons.xml
```

This is to avoid the approximated surface normals being incorrect at the edge.

If you have a high-resolution mesh of the surfaces 11 and 12 in `ply` format it can be used to improve the normal definition of the spherigons. Run:

```
NekMesh -m spherigon:surf=11,12:usenormalfile=Surf_11-12_Mesh.ply \
  MeshWithStraighEdges.xml MeshWithSpherigons.xml
```

This can be useful, for example, when meshing the Leading edge of an airfoil. Starting from a linear mesh (left figure) the spherigon patches curve the surface elements producing leading edge closer to the underlying geometry:




Figure 4.3 (a) Leading edge without spherigons, (b) Leading edge with spherigons

4.4.6 Periodic boundary condition alignment

When using periodic boundary conditions, the order of the elements within the boundary composite determines which element edges are periodic with the corresponding boundary composite.

To facilitate this alignment, **NekMesh** has a periodic alignment module which attempts to identify pairs of mutually periodic edges. Given two surfaces `surf1` and `surf2`, which for example correspond to the physical surface IDs specified in **Gmsh**, and an axis which defines the periodicity direction, the following command attempts to reorder the composites:

```
NekMesh -m peralign:surf1=11:surf2=12:dir=y \
        -m peralign:surf1=13:surf2=14:dir=z Mesh.xml Mesh_aligned.xml
```

Here the surfaces with IDs 11 and 12 will be aligned normal to the y -axis and the surfaces 13 and 14 will be aligned normal to the z -axis.

Note that this command cannot perform magic – it assumes that any given edge or face lying on the surface is periodic with another face on the opposing surface, that there are the same number of elements on both surfaces, and the corresponding edge or face is the same size and shape but translated along the appropriate axis.

When using periodic boundary conditions that are rotationally aligned the following rotational options should be applied:

```
NekMesh -m peralign:surf1=11:surf2=12:dir=x:rot=PI/6:tol=1e-6 \
        Mesh.xml Mesh_aligned.xml
```

where `rot` specifies the rotation angle in radians from `surf1` to `surf2` about the axis specified by `dir` (i.e. the “ x ” axis in this example). An optional tolerance `tol` can also be specified which is the tolerance within which the rotation is assumed to be exact. The default tolerance is $1e-8$.

In 3D, where prismatic or tetrahedral elements are connected to one or both of the surfaces, additional logic is needed to guarantee connectivity in the XML file. In this case we append the `orient` parameter:

```
NekMesh -m peralign:surf1=11:surf2=12:dir=y:orient input.dat output.xml
```

Note



One of the present shortcomings of `orient` is that it throws away all high-order information and works only on the linear element. This can be gotten around if you are just doing e.g. spherigon patches by running this `peralign` module before the `spherigon` module.

4.4.7 Boundary layer splitting

Often it is the case that one can generate a coarse boundary layer grid of a mesh. **NekMesh** has a method for splitting prismatic and hexahedral elements into finer elements based on the work presented in [31] and [32]. You must have a prismatic mesh that is *O*-type – that is, you can modify the boundary layer without modifying the rest of the mesh.

Given n layers, and a ratio r which defines the relative heights of elements in different layers, the method works by defining a geometric progression of points

$$x_k = x_{k-1} + ar^k, \quad a = \frac{2(1-r)}{1-r^{n+1}}$$

in the standard segment $[-1, 1]$. These are then projected into the coarse elements to construct a sequence of increasingly refined elements, as depicted in figure 4.4.




Figure 4.4 Splitting Ω_{st} and applying the mapping χ^e to obtain a high-order layer of prisms from the macro-element.

To split a prism boundary layer on surface 11 into 3 layers with a growth rate of 2 and 7 integration points per element use the following command:

```
NekMesh -m bl:surf=11:layers=3:r=2:nq=7 MeshWithOnePrismLayer.xml \
    MeshWith3PrismsLayers.xml
```

Note



You can also use an expression in terms of coordinates (x, y, z) for r to make the ratio spatially varying; e.g. $r=\sin(x)$. In this case the function should be sufficiently smooth to prevent the elements self-intersecting.




Figure 4.5 (a) LE with Spherigons but only one prism layer for resolving the boundary layer, (b) LE with Spherigons with 3 growing layers of prisms for better resolving the boundary layer.

4.4.8 High-order cylinder generation

Generating accurate high-order curved geometries in **Gmsh** is quite challenging. This module processes an existing linear cylindrical mesh, with axis aligned with the z -coordinate axis, to generate accurate high-order curvature information along the edges.

```
NekMesh -m cyl:surf=2:r=1.0:N=5 LinearCylinder.xml HighOrderCylinder.xml
```

The module parameters are:

- **surf**: Surface on which to apply curvature. This should be the outer surface of the cylinder.
- **r**: Radius of the cylinder.
- **N**: Number of high-order points along each element edge.

Note



The module could also be used to apply curvature along the interior of a hollow cylinder. However, there are no checks to ensure the resulting elements are not self-intersecting.

4.4.9 Linearisation

The ability to remove all the high-order information in a mesh can be useful at times when mesh generation is tricky or impossible in the presence of curvature

To do this in NekMesh use the command:

```
NekMesh -m linearise:all high-order-mesh.xml linear-mesh.xml
```

The output will contain only the linear mesh information, all curved information is removed. Alternatively

```
NekMesh -m linearise:invalid high-order-mesh.xml linear-mesh.xml
```

attempts to remove curvature from elements only where necessary. This is a simple algorithm that removes curvature from invalid elements and repeats until all elements are valid. Either `all` or `invalid` must be specified.

- `all`: remove curvature from all elements.
- `invalid`: remove curvature from invalid elements.
- `prismonly`: consider only prisms when removing curvature. This is useful in the presence of a prismatic boundary layer.

4.4.10 Extracting interface between tetrahedra and prismatic elements

When the mesh is three-dimensional and comprised of a prismatic boundary layer with tetrahedra in the interior of the domain, this module extracts the prismatic elements only, and constructs a boundary region for the interface between the tetrahedra and prisms. This is useful in, for example, the study of aortic flows, where the prismatic boundary layer can be extracted and refined to study unsteady advection-diffusion problems on a more refined grid inside the boundary layer.

To use this module you therefore use the command:

```
NekMesh -m extracttetprisminterface input.xml output.xml
```

There are no configuration options for this module, as it is highly specific to a certain class of meshes.

4.4.11 Boundary identification

Some mesh formats lack the ability to identify boundaries of the domain they discretise. NekMesh has a rudimentary boundary identification routine for conformal meshes, which will create a composite of edges (2D) or faces (3D) which are connected to precisely one element. This can be done using the `detect` module:

```
NekMesh -m detect volume.xml volumeWithBoundaryComposite.xml
```

4.4.12 Scalar function curvature

This module imposes curvature on a surface given a scalar function $z = f(x, y)$. For example, if on surface 1 we wish to apply a surface defined by a Gaussian $z = \exp[-(x^2 + y^2)]$ using 7 quadrature points in each direction, we may issue the command

```
NekMesh -m scalar:surf=1:nq=7:scalar=exp\(-x*x+y*y\) mesh.xml deformed.xml
```

Note



This module makes no attempt to apply the curvature to the interior of the domain. Elements must therefore be coarse in order to prevent self-intersection. If a boundary layer is required, one option is to use this module in combination with the splitting module described earlier.

4.4.13 Link Checking

It is quite possible that a mesh contains some sort of hanging entity or element connectivity error. The check link module is a fast check that, a) elements are correctly connected and b) the boundary entities (composites) match the interior domain:

```
NekMesh -m linkcheck mesh.xml mesh2.xml
```

This module should be added to the module chain if the user suspects there may be a mesh issue. The module will print a warning if there is a connectivity error.

4.4.14 2D mesh extrusion

This module allows a 2D mesh, quads, triangles or both, to be extruded in the z direction to make a simple 3D mesh made of prisms and hexahedra. It is also capable of extruding the high-order curvature within the 2D mesh. The module requires two parameters:

```
NekMesh -m extrude:layers=n:length=1 2D.xml 3D.xml
```

length which determines how long the z extrusion will be and layers, the number of elements in the z direction.

4.4.15 Variational Optimisation

This module can correct invalid and improve the quality of elements in high-order meshes by applying curvilinear deformation to the interiors of domains. It achieves this by solving a solid mechanics system which, using variational calculus has been cast as a non-linear energy optimisation problem. It is basis of the work in [43].

It works by considering the boundary (curved) mesh entities to be fixed and moving the interior nodes to a lower energy configuration. This new configuration in most scenarios is a higher quality mesh. The energy is evaluated depending on which functional is chosen. We find hyperleastic to be the most reliable but it can also model the mesh and a linearelastic solid as well as functionals based on the Winslow equation and the distortion method proposed by Roca et al. [14].

There are a large number of options which can be viewed using the help function but the basic usage is:

```
NekMesh -m varopti:type inital.xml optimised.xml
```

where type can be `hyperelastic`, `linearelastic`, `winslow` or `roca`.

4.4.16 Mesh projection

This module can take any linear mesh, providing that it is a close representation of the CAD and project the boundary of the mesh onto the CAD. This will curve the surface of the mesh. The method has a number of failsafes ensuring that even bad CAD or poor linear meshes should be able to be curved to some degree. If the method encounters an issue, such as the linear mesh being a large distance from the CAD, it will simply leave that element straight sided. A well made CAD model and accurate linear mesh should be curved with little issue.

The module needs to be informed of the CAD file to project the mesh to and the order at which to curve the surface:

```
NekMesh -m projectcad:file=cadfile.step:order=x inital.xml optimised.xml
```

4.5 Mesh generation

In addition to the functionality described previously, **NekMesh** is capable of generating high-order meshes directly from a CAD definition. By default this functionality is not activated, a user wishing to utilise the mesh generation capability of **NekMesh** must compile *Nektar++* with the `NEKTAR_USE_MESHGEN` option on. As well as compiling the relevant routines into **NekMesh** it will also download a number of other packages which are required.

The most critical dependancy of the mesh generation routines is **OpenCascade** which powers the CAD engine. **NekMesh** is capable of finding and using existing installations of **OpenCascade** 6.8 or OCE 0.17. If either are not present on the installation machine **NekMesh** will install OCE 0.17 from source. This is a very big installation and will take some time so it is advised that the user ensures **OpenCascade** is available on the machine.

As with all tasks within **NekMesh** the mesh generation capability exists as its own separate module which is of type Input. Due to the vast amount of code associated with the generation of high-order meshes and the comparatively small nature of modules in the **NekMesh** program a new library has been created for *Nektar++* called *NekMeshUtils*, which contains all the core routines and classes for the **NekMesh** mesh format as well as a series of classes for the generation of meshes. This library also contains the CAD API for *Nektar++* which is used to generate the meshes.

4.5.1 Methodology

This section outlines the approach taken by **NekMesh** to generate high-order meshes. To simplify the sometimes very complicated high-order mesh generation processes in other programs, **NekMesh** executes all the stages required to produce a high-order mesh in one single pipeline which once started requires no interaction from the user. In broad terms these stages are:

- Specification of the element sizes in the mesh,
- Coarse linear mesh generation of the domain,
- Generation of optimised high-order surface on the geometric boundary,

and are outlined in more detail in the following sections.

4.5.1.1 CAD Interaction

At the core of all the ideas in the **NekMesh** generator is that the final mesh is a high quality representation of the underlying geometry. As such all of the entities in the mesh must know where they are located with respect to the CAD and the system to be able to query any geometric information at any point in the domain easily and with accuracy. To handle this **NekMesh** has been interfaced with the third-party suite of CAD libraries called OpenCascade. In its normal state OpenCascade is a very large collection of libraries with tens of thousands of functions which are simply not needed for our purposes, because of this its installation is a very arduous and long process. Combine this with the fact that there are dozens of versions and types of OpenCascade, such as OpenCascade Community Edition, it is simply impossible for **NekMesh** to use already existing OpenCascade installations on a given machine. To solve these issues, when installing *Nektar++* with the mesh generator it will download pre-compiled binaries for the relevant OS and link against those, any previously installed versions of OpenCascade will not be searched for and therefore ignored. To reduce the massively complex libraries in OpenCascade down to a manageable set of functions to be used in **NekMesh** a set of interface classes have been created which act as buffer between it and *Nektar++*. These CAD classes mean that development of mesh generation routines is significantly easier and in the future *Nektar++* developers will be able to utilise CAD information in all aspects of the framework without having to learn OpenCascade. Another advantage with

this approach is that adding support for other CAD engines, as well as OpenCascade, in the future should be relatively simple and will not require the rewriting of any of the **NekMesh** code.

4.5.1.2 Automatic specification of the mesh

One of the key challenges of generating a high-order mesh is the creation of a suitable coarse linear mesh. It is quite difficult for a user to define a full set of spacings over a whole domain which will produce a good quality especially when aiming for coarseness. This is tackled in **NekMesh** with a system for automatically defining a set of smooth and coarse mesh spacings throughout the whole domain. This is achieved using an octree description of the domain. The domain is recursively subdivided into octants which each describe a small portion of the domain. The level to which the domain subdivides is based on the curvature of the geometric boundary. Higher curvature regions will subdivide to a finer level allowing for increased control on the mesh specification and smoothness. The geometric curvature is then related to a mesh sizing parameter and propagated throughout the domain ensuring a smooth mesh. For those unfamiliar with octrees, it is best to think of it as a non-conforming hexahedral mesh

4.5.1.3 Linear Mesh Generation

The first challenge mentioned in the previous section is addressed with the **NekMesh** approach to linear mesh generation. Primarily because of the difficulties in interfacing existing linear mesh generators for high-order applications the decision was made to include a bespoke linear mesh generator within the program. Compared with the mesh generators included in commercial packages this linear mesh generator takes the quite unconventional and more historic approach in building the mesh in a bottom up fashion from 0D to 3D. Using this approach means it is possible to guarantee a level of boundary conformity which direct to 3D approaches cannot at the desired level of coarseness. In this approach, first mesh nodes are placed on the vertices of the CAD model (0D), then the curves in the CAD are meshed in 1D using the vertex nodes as boundaries, then the surfaces are meshed in their 2D parameter plane using the curve meshes as boundaries and finally the 3D volume is meshed using the surface mesh as the boundary to complete the linear mesh. In **NekMesh**, to achieve greater robustness, the 2D mesh generation library **Triangle** is used and the **TetGen** library for the 3D. Both of which are highly developed Delaunay based mesh generators. As with all additional libraries in *Nektar++* these are automatically downloaded and installed if needed.

4.5.1.4 High-order Surface Generation

Addition of the high-order nodes to and the curving of the mesh is very open problem, no high-order mesh generator has solved this and while the methods used in **NekMesh** are not 100% full-proof, the system currently in place can create good quality high-order curved meshes with a reasonable robustness. This area will receive the greatest level of development in the future. The most critical part of defining the high-order mesh is the addition of high-order nodes on the geometric surface. The mesh generator must

achieve the greatest level of geometric accuracy as it can otherwise it will greatly affect the final flow solutions. If the linear surface triangulation is taken to be fixed during this process, the problem can be addressed in a element by element fashion. If the high-order nodes are placed by simply using an affine mapping to the CAD surface and back the resulting high-order triangle will inherit the same distortions as the CAD surface. To solve this **NekMesh** uses a system node location optimisation in the parameter plane of the CAD surface to ensure the high-order triangles have as little distortion as possible while remaining exactly on the geometric surface. To do this the system models the high-order edges and triangles as a network of springs with an associated spring energy which is minimised using a multidimensional Newton type optimisation procedure with a Gauss-Seidel matrix solver.

4.5.1.5 Mesh Correction

Due to the fact that, for the time being, no consideration is given to the curving of mesh interior entities explicitly in the mesh generation process, the curving the geometric surface can produce meshes with invalid elements, especially in the case of Euler type (Tetrahedra only) meshes. Three strategies exist within *Nektar++* to correct these elements. Firstly removing the curvature, by removing the curvature of invalid elements they become valid. However this has the massive downside of compromising the geometric accuracy of the mesh but is quick and effective, this can be enacted using the command:

```
NekMesh -m linearise:invalid invalidMesh.xml validMesh.xml
```

An alternative to this is to use the linear elastic solver within *Nektar++* to deform the mesh interior entities. Its use is very computationally expensive, as with all PDE solvers, and is also not particularly robust. It can be used with the set of commands outlined in the FieldConvert deform and displacement modules and the section on the Linear Elastic Solver.

The final and possibly most useful approach is to use the Variational Optimisation module to curve the interior of the domain. This is explained in [4.4.15](#).

4.5.2 Mesh generation manual

The mesh generation is executed with the command:

```
NekMesh session.mcf mesh.xml
```

where session.mcf is a mesh configuration file which contains all the options and parameters needed for mesh generation. Below is an example of a simple example which generates a 2D NACA wing.

```
1 <NEKTAR>
2   <MESHING>
3
```

```

4      <INFORMATION>
5          <I PROPERTY="CADFile"    VALUE="6412"    />
6          <I PROPERTY="MeshType"   VALUE="2D"      />
7      </INFORMATION>
8
9      <PARAMETERS>
10         <P PARAM="MinDelta"       VALUE="0.01"   />
11         <P PARAM="MaxDelta"       VALUE="1.0"     />
12         <P PARAM="EPS"            VALUE="0.1"     />
13         <P PARAM="Order"          VALUE="4"       />
14
15         <!-- 2D Domain !-->
16         <P PARAM="Xmin"           VALUE="-1.0"    />
17         <P PARAM="Ymin"           VALUE="-2.0"    />
18         <P PARAM="Xmax"           VALUE="3.0"     />
19         <P PARAM="Ymax"           VALUE="2.0"     />
20         <P PARAM="AOA"            VALUE="15.0"    />
21     </PARAMETERS>
22
23     </MESHING>
24 </NEKTAR>

```

In all cases the mesh generator needs two pieces of information and four parameters. It firstly needs to know the CAD file with which to work. In the example above this is listed as a 4 digit number, this is because the mesh generator is equipped with a NACA wing generator. In all other cases this parameter would be the name of a CAD file (in either STEP or GEO format). Secondly, what type of mesh to make, the options are `EULER` and `BndLayer` for 3D meshes and `2D` and `2DBndLayer` for 2D meshes. In the case of `EULER` the mesh will be made with only tetrahedra. For `BndLayer` the mesh generator will attempt to insert a single macro prism layer onto the geometry surface. This option requires additional parameters. This is similar for the 2D scenarios. The automatic mesh specification system requires three parameters to build the specification of a smooth, curvature refined mesh. Firstly `MinDelta` which is the size of the smallest element to be found in the final mesh. Secondly `MaxDelta` which is the maximum size of an element in the mesh and lastly `EPS` which is a sensitivity to curvature parameter with a range $1 \geq \varepsilon > 0$ which heuristically controls the size of the elements for a given degree of curvature on the geometric surface. `Order` is the polynomial order of the mesh to be generated. When generating a boundary layer mesh a few additional parameters must be given. An example is shown.

```

1      <NEKTAR>
2          <MESHING>
3
4              <INFORMATION>
5                  <I PROPERTY="CADFile"          VALUE="6412"          />
6                  <I PROPERTY="MeshType"         VALUE="2DBndLayer"       />
7              </INFORMATION>
8
9              <PARAMETERS>
10                 <P PARAM="MinDelta"             VALUE="0.01"           />
11                 <P PARAM="MaxDelta"             VALUE="1.0"           />

```



```

12      <P PARAM="EPS"                VALUE="0.1"        />
13      <P PARAM="Order"              VALUE="4"          />
14
15      <!-- Boundary layer !-->
16      <P PARAM="BndLayerSurfaces"    VALUE="5,6"        />
17      <P PARAM="BndLayerThickness"   VALUE="0.03"       />
18      <P PARAM="BndLayerLayers"      VALUE="4"          />
19      <P PARAM="BndLayerProgression" VALUE="2.0"        />
20
21      <!-- 2D Domain !-->
22      <P PARAM="Xmin"                VALUE="-1.0"       />
23      <P PARAM="Ymin"                VALUE="-2.0"       />
24      <P PARAM="Xmax"                VALUE="3.0"        />
25      <P PARAM="Ymax"                VALUE="2.0"        />
26      <P PARAM="AOA"                 VALUE="15.0"       />
27      </PARAMETERS>
28
29      </MESHING>
30      </NEKTAR>

```

A list of the CAD surfaces which will have a prism generated on is described by `BndLayerSurfaces` and the thickness of the boundary to aim for is `BndLayerThickness`. The mesh generator has been created with a range of error messages to aid in debugging. If you encounter an error and the mesh generator fails, run `NekMesh` with the verbose `-v` flag and send the stdout with the .mcf and .step files to m.turner14@imperial.ac.uk. Without the feedback this functionality cannot improve.

4.5.2.1 GEO format

Recent developments have been made to facilitate the generation of meshes from simple 2D geometries. The GEO file format, used by Gmsh, is a popular option that allows the user to script geometrical and meshing operations without the need of a GUI. A simplified reader has been implemented in `NekMesh` for 2D geometries. Although very basic this reader may be extended in the future to cover a wider range of geometrical features.

For a full description of the GEO format the user should refer to Gmsh's documentation. The following commands are currently supported:

- `//` (i.e. comments)
- `Point`
- `Line`
- `Spline` (through points)
- `BSpline` (i.e. a Bézier curve)

- **Ellipse** (arc): as defined in Gmsh's OpenCASCADE kernel, the first point defines the start of the arc, the second point the centre and the fourth point the end. The third point is not used. The start point along with the centre point form the major axis and the minor axis is then computed so that the end point falls onto the arc. The major axis must always be greater or equal to the minor axis.
- **Circle** (arc): the circle is a special case of the ellipse where the third point is skipped. The distances between the start and end points and the centre point must be equal or an error will be thrown.
- **Line Loop**
- **Plane Surface**

At the present time, NekMesh does not support the full scripting capabilities of the GEO format, but the evaluation of simple variables is supported. The used GEO files should be a straightforward succession of entity creations (see list above). This should however allow for the creation of quite a wide range of 2D and 3D geometries by transformation of arbitrary curves into generic splines and arcs.

FieldConvert

FieldConvert is a utility embedded in *Nektar++* with the primary aim of allowing the user to convert the *Nektar++* output binary files (`.chk` and `.fld`) into formats which can be read by common visualisation and post-processing software, primarily Paraview/VisIt (in unstructured VTK `.vtu` format) or Tecplot/VisIt (in ASCII `.dat` or binary `.plt` formats). FieldConvert also allows the user to manipulate the *Nektar++* output binary files by using some additional modules which can be called with the option `-m` which stands for `m`odule. Note that another flag, `-r` (which stand for `r`ange) allows the user to specify a sub-range of the domain on which the conversion or manipulation of the *Nektar++* output binary files will be performed.

Almost all of the FieldConvert functionalities can be run in parallel if *Nektar++* is compiled using MPI (see the installation documentation for additional info on how to implement *Nektar++* using MPI).¹

5.1 Basic usage

FieldConvert expects at least one input specification (such as a session file and its corresponding field file) and one output specification. These are specified on the command line as

```
FieldConvert in1.xml in2.fld out.dat
```

These can be combined with a processing module by adding the `-m` command line option. There can be more than one module specified, and they can appear anywhere in the command line arguments, although the order of execution is inferred from their order in the command line. For example, the command

```
FieldConvert in1.xml -m module1 in2.fld -m module2 out.dat
```

¹Modules that do not have parallel support will be specified in the appropriate section.

causes `in1.xml` and `in2.fld` to be read, followed by the `module1` processing module, the `module2` processing module, and finally output to the `out.dat` Tecplot file.

5.1.1 Input formats

FieldConvert supports XML and FLD-format files as produced by *Nektar++*. It also supports the reading of data files from two external spectral element codes: *Semtex*² and *Nek5000*³. These files can be directly converted to *Nektar++* format files by using the command

```
FieldConvert input.fld output.fld
```

Note that even though the `.fld` extension is typically associated with *Nektar++* files, FieldConvert can automatically identify *Semtex* and *Nek5000* input field files.

To use these files in a simulation, or to post-process the results of a simulation, an appropriate mesh must also be defined in the *Nektar++* XML format. *NekMesh* can be used to convert these input files to XML, as outlined in section 4.

5.2 Convert .fld / .chk files into Paraview, VisIt or Tecplot format

To convert the *Nektar++* output binary files (`.chk` and `.fld`) into a format which can be read by two common visualisation softwares: Paraview (`.vtu` format), VisIt (`.vtu` format) or Tecplot (`.dat` or `.plt` format) the user can run the following commands:

- Paraview or VisIt (`.vtu` format)

```
FieldConvert test.xml test.fld test.vtu
```

- Tecplot (`.dat` format)

```
FieldConvert test.xml test.fld test.dat
```

- Tecplot or VisIt(`.plt` format)

```
FieldConvert test.xml test.fld test.plt
```

where `FieldConvert` is the executable associated to the utility FieldConvert, `test.xml` is the session file and `test.vtu`, `test.dat`, `test.plt` are the desired format outputs, either Paraview, VisIt or Tecplot formats.

²<http://users.monash.edu.au/~bburn/semtex.html>

³<https://nek5000.mcs.anl.gov>

When converting to `.dat` or `.plt` format, it is possible to enable output with double precision, which is more accurate but requires larger disk space. For example, double precision output in `plt.` format can be produced with the command:

```
FieldConvert test.xml test.fld test.plt:plt:double
```



Tip

Note that the session file is also supported in its compressed format `test.xml.gz`.

5.3 Convert field files between XML and HDF5 format

When *Nektar++* is compiled with HDF5 support, solvers can select the format used for output of `.fld` files. FieldConvert can be used to convert between these formats using an option on the `.fld` output module. For example, if `in.fld` is stored in the default XML format, it can be converted to HDF5 format by issuing the command

```
FieldConvert in.fld out.fld:fld:format=Hdf5
```

5.4 Range option `-r`

The Fieldconvert range option `-r` allows the user to specify a sub-range of the mesh (computational domain) by using an additional flag, `-r` (which stands for `r`ange and either convert or manipulate the *Nektar++* output binary files. Taking as an example the conversion of the *Nektar++* binary files (`.chk` or `.fld`) shown before and wanting to convert just the 2D sub-range defined by $-2 \leq x \leq 3$, $-1 \leq y \leq 2$ the additional flag `-r` can be used as follows:

- Paraview or VisIt (`.vtu` format)

```
FieldConvert -r -2,3,-1,2 test.xml test.fld test.vtu
```

- Tecplot (`.dat` format)

```
FieldConvert -r 2,3,-1,2 test.xml test.fld test.dat
```

where `-r` defines the range option of the FieldConvert utility, the two first numbers define the range in x direction and the the third and fourth number specify the y range. A sub-range of a 3D domain can also be specified. For doing so, a third set of numbers has to be provided to define the z range.

5.5 FieldConvert modules *-m*

FieldConvert allows the user to manipulate the *Nektar++* output binary files (.chk and .fld) by using the flag `-m` (which stands for `m`odule). Specifically, FieldConvert has these additional functionalities

1. `C0Projection`: Computes the C0 projection of a given output file;
2. `QCriterion`: Computes the Q-Criterion for a given output file;
3. `L2Criterion`: Computes the Lambda 2 Criterion for a given output file;
4. `addcompositeid`: Adds the composite ID of an element as an additional field;
5. `fieldfromstring`: Modifies or adds a new field from an expression involving the existing fields;
6. `addFld`: Sum two .fld files;
7. `combineAvg`: Combine two *Nektar++* binary output (.chk or .fld) field file containing averages of fields (and possibly also Reynolds stresses) into single file;
8. `concatenate`: Concatenate a *Nektar++* binary output (.chk or .fld) field file into single file (deprecated);
9. `dof`: Count the total number of DOF;
10. `equispacedoutput`: Write data as equi-spaced output using simplices to represent the data for connecting points;
11. `extract`: Extract a boundary field;
12. `gradient`: Computes gradient of fields;
13. `halfmodetofourier`: Convert `HalfMode` expansion to `SingleMode` for further processing;
14. `homplane`: Extract a plane from 3DH1D expansions;
15. `homstretch`: Stretch a 3DH1D expansion by an integer factor;
16. `innerproduct`: take the inner product between one or a series of fields with another field (or series of fields).
17. `interpfield`: Interpolates one field to another, requires fromxml, fromfld to be defined;
18. `interpdatapointtofld`: Interpolates given discrete data using a finite difference approximation to a fld file given an xml file;

19. `interppts`: Interpolates a field to a set of points. Requires `fromfld`, `fromxml` to be defined, and a `topts`, line, plane or box of target points;
20. `interpptstopts`: Interpolates a set of points to another. Requires a `topts`, line, plane or box of target points;
21. `isocontour`: Extract an isocontour of “fieldid” variable and at value “fieldvalue”. Optionally “fieldstr” can be specified for a string definition or “smooth” for smoothing;
22. `jacobianenergy`: Shows high frequency energy of Jacobian;
23. `qualitymetric`: Evaluate a quality metric of the underlying mesh to show mesh quality;
24. `mean`: Evaluate the mean of variables on the domain;
25. `meanmode`: Extract mean mode (plane zero) of 3DH1D expansions;
26. `pointdatatofld`: Given discrete data at quadrature points project them onto an expansion basis and output fld file;
27. `printfldnorms`: Print L2 and LInf norms to stdout;
28. `removefield`: Removes one or more fields from .fld files;
29. `scalargrad`: Computes scalar gradient field;
30. `scaleinputfld`: Rescale input field by a constant factor;
31. `shear`: Computes time-averaged shear stress metrics: TAWSS, OSI, transWSS, TAAFI, TACFI, WSSG;
32. `streamfunction`: Calculates stream function of a 2D incompressible flow.
33. `surfdistance`: Computes height of a prismatic boundary layer mesh and projects onto the surface (for e.g. y^+ calculation).
34. `vorticity`: Computes the vorticity field.
35. `wss`: Computes wall shear stress field.

The module list above can be seen by running the command

```
FieldConvert -l
```

In the following we will detail the usage of each module.

5.5.1 Smooth the data: *C0Projection* module

To smooth the data of a given .fld file one can use the `C0Projection` module of FieldConvert

```
FieldConvert -m C0Projection test.xml test.fld test-C0Proj.fld
```

where the file `test-C0Proj.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

The option `localtoglobmap` will do a global gather of the coefficients and then scatter them back to the local elements. This will replace the coefficients shared between two elements with the coefficients of one of the elements (most likely the one with the highest id). Although not a formal projection it does not require any matrix inverse and so is very cheap to perform.

The option `usexmlbcs` will enforce the boundary conditions specified in the input xml file.

The option `helmsmoothing=L` will perform a Helmholtz smoothing projection of the form

$$\left(\nabla^2 + \left(\frac{2\pi}{L}\right)^2\right) \hat{u}^{new} = \left(\frac{2\pi}{L}\right)^2 \hat{u}^{orig}$$

which can be interpreted in a Fourier sense as smoothing the original coefficients using a low pass filter of the form

$$\hat{u}_k^{new} = \frac{1}{(1 + k^2/K_0^2)} \hat{u}_k^{orig} \text{ where } K_0 = \frac{2\pi}{L}$$

and so L is the length scale below which the coefficients values are halved or more. Since this form of the Helmholtz operator is not positive definite, currently a direct solver is necessary and so this smoother is mainly of use in two-dimensions.

5.5.2 Calculate Q-Criterion: *QCriterion* module

To perform the Q-criterion calculation and obtain an output data containing the Q-criterion solution, the user can run

```
FieldConvert -m QCriterion test.xml test.fld test-QCrit.fld
```

where the file `test-QCrit.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.3 Calculate λ_2 : *L2Criterion* module

To perform the λ_2 vortex detection calculation and obtain an output data containing the values of the λ_2 eigenvalue, the user can run


```
FieldConvert -m L2Criterion test.xml test.fld test-L2Crit.fld
```

where the file `test-L2Crit.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.4 Add composite ID: *addcompositeid* module

When dealing with a geometry that has many surfaces, we need to identify the composites to assign boundary conditions. To assist in this, FieldConvert has a `addcompositeid` module, which adds the composite ID of every element as a new field. To use this we simply run

```
FieldConvert -m addcompositeid mesh.xml out.dat
```

In this case, we have produced a Tecplot file which contains the mesh and a variable that contains the composite ID. To assist in boundary identification, the input file `mesh.xml` should be a surface XML file that can be obtained through the NekMesh `extract` module (see section 4.4.3).

5.5.5 Add new field: *fieldfromstring* module

To modify or create a new field using an expression involving the existing fields, one can use the `fieldfromstring` module of FieldConvert

```
FieldConvert -m fieldfromstring:fieldstr="x+y+u":fieldname="result" \
file1.xml file2.fld file3.fld
```

In this case `fieldstr` is a required parameter describing a function of the coordinates and the existing variables, and `fieldname` is an optional parameter defining the name of the new or modified field (the default is `newfield`). `file3.fld` is the output containing both the original and the new fields, and can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.6 Sum two .fld files: *addFld* module

To sum two .fld files one can use the `addFld` module of FieldConvert

```
FieldConvert -m addfld:fromfld=file1.fld:scale=-1 file1.xml file2.fld \
file3.fld
```

In this case we use it in conjunction with the command `scale` which multiply the values of a given .fld file by a constant `value`. `file1.fld` is the file multiplied by `value`, `file1.xml` is the associated session file, `file2.fld` is the .fld file which is summed to `file1.fld` and finally `file3.fld` is the output which contain the sum of the two .fld

files. `file3.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.7 Combine two .fld files containing time averages: *combineAvg* module

To combine two .fld files obtained through the AverageFields or ReynoldsStresses filters, use the `combineAvg` module of FieldConvert

```
FieldConvert -m combineAvg:fromfld=file1.fld file1.xml file2.fld \
file3.fld
```

`file3.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.8 Concatenate two files: *concatenate* module

To concatenate `file1.fld` and `file2.fld` into `file-conc.fld` one can run the following command

```
FieldConvert file.xml file1.fld file2.fld file-conc.fld
```

where the file `file-conc.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt. The `concatenate` module previously used for this purpose is not required anymore, and will be removed in a future release.

5.5.9 Count the number of DOF: *dof* module

To count the number of DOF in a solution file, one can run the following command

```
FieldConvert -m dof file.xml file.fld out.stdout
```

5.5.10 Equi-spaced output of data: *equispacedoutput* module

This module interpolates the output data to a truly equispaced set of points (not equispaced along the collapsed coordinate system). Therefore a tetrahedron is represented by a tetrahedral number of pointst. This produces much smaller output files. The points are then connected together by simplices (triangles and tetrahedrons).

```
FieldConvert -m equispacedoutput test.xml test.fld test.dat
```

or

```
FieldConvert -m equispacedouttput test.xml test.fld test.vtu
```

**Note**

Currently this option is only set up for triangles, quadrilaterals, tetrahedrons and prisms.

5.5.11 Extract a boundary region: *extract* module

The boundary region of a domain can be extracted from the output data using the following command line

```
FieldConvert -m extract:bnd=2 test.xml \
    test.fld test-boundary.fld
```

The option `bnd` specifies which boundary region to extract. Note this is different to NekMesh where the parameter `surf` is specified and corresponds to composites rather boundaries. If `bnd` is not provided, all boundaries are extracted to different fields. The output will be placed in `test-boundary_b2.fld`. If more than one boundary region is specified the extension `_b0.fld`, `_b1.fld` etc will be outputted. To process this file you will need an xml file of the same region. This can be generated using the command:

```
NekMesh -m extract:surf=5 test.xml test\_b0.xml
```

The surface to be extracted in this command is the composite number and so needs to correspond to the boundary region of interest. Finally to process the surface file one can use

```
FieldConvert test\_b0.xml test\_b0.fld test\_b0.dat
```

This will obviously generate a Tecplot output if a `.dat` file is specified as last argument. A `.vtu` extension will produce a Paraview or VisIt output.

5.5.12 Compute the gradient of a field: *gradient* module

To compute the spatial gradients of all fields one can run the following command

```
FieldConvert -m gradient test.xml test.fld test-grad.fld
```

where the file `file-grad.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.13 Convert HalfMode expansion to SingleMode for further processing: *halfmodetofourier* module

To obtain full Fourier expansion form a `HalfMode` result, use the comand:

```
FieldConvert -m halfmodetofourier file.xml half_mode_file.fld
             single_mode_file.fld
```

5.5.14 Extract a plane from 3DH1D expansion: *homplane* module

To obtain a 2D expansion containing one of the planes of a 3DH1D field file, use the command:

```
FieldConvert -m homplane:planeid=value file.xml file.fld file-plane.fld
```

If the option `wavespace` is used, the Fourier coefficients corresponding to `planeid` are obtained. The command in this case is:

```
FieldConvert -m homplane:wavespace:planeid=value file.xml \
             file.fld file-plane.fld
```

The output file `file-plane.fld` can be processed in a similar way as described in section 5.2 to visualise it either in Tecplot or in Paraview.

5.5.15 Stretch a 3DH1D expansion: *homstretch* module

To stretch a 3DH1D expansion in the z-direction, use the command:

```
FieldConvert -m homstretch:factor=value file.xml file.fld file-stretch.fld
```

The number of modes in the resulting field can be chosen using the command-line parameter `--output-points-hom-z`.

The output file `file-stretch.fld` can be processed in a similar way as described in section 5.2 to visualise it either in Tecplot or in Paraview.

5.5.16 Inner Product of a single or series of fields with respect to a single or series of fields: *innerproduct* module

You can take the inner product of one field with another field using the following command:

```
FieldConvert -m innerproduct:fromfld=file1.fld file2.xml file2.fld \
             out.stdout
```

This command will load the `file1.fld` and `file2.fld` assuming they both are spatially defined by `files.xml` and determine the inner product of these fields. The input option `fromfld` must therefore be specified in this module.

Optional arguments for this module are `fields` which allow you to specify the fields that you wish to use for the inner product, i.e.

```
FieldConvert -m innerproduct:fromfld=file1.fld:fields="0,1,2" file2.xml \
file2.fld out.stdout
```

will only take the inner product between the variables 0,1 and 2 in the two fields files. The default is to take the inner product between all fields provided.

Additional options include `multifldids` and `allfromflds` which allow for a series of fields to be evaluated in the following manner:

```
FieldConvert -m innerproduct:fromfld=file1.fld:multifldids="0-3" \
file2.xml file2.fld out.stdout
```

will take the inner product between a file names field1_0.fld, field1_1.fld, field1_2.fld and field1_3.fld with respect to field2.fld.

Analogously including the options `allfromflds`, i.e.

```
FieldConvert -m innerproduct:fromfld=file1.fld:multifldids="0-3":\
allfromflds file2.xml file2.fld out.stdout
```

Will take the inner product of all the from fields, i.e. field1_0.fld,field1_1.fld,field1_2.fld and field1_3.fld with respect to each other. This option essentially ignores file2.fld. Only the unique inner products are evaluated so if four from fields are given only the related trianuglar number $4 \times 5/2 = 10$ of inner products are evaluated.

This option can be run in parallel.

5.5.17 Interpolate one field to another: *interpfield* module

To interpolate one field to another, one can use the following command:

```
FieldConvert -m interpfield:fromxml=file1.xml:fromfld=file1.fld \
file2.xml file2.fld
```

This command will interpolate the field defined by `file1.xml` and `file1.fld` to the new mesh defined in `file2.xml` and output it to `file2.fld`. The `fromxml` and `fromfld` must be specified in this module. In addition there are two optional arguments `clamptolowervalue` and `clamptouppervalue` which clamp the interpolation between these two values. Their default values are -10,000,000 and 10,000,000.

**Tip**

This module can run in parallel where the speed is increased not only due to using more cores but also, since the mesh is split into smaller sub-domains, the search method currently adopted performs faster.

5.5.18 Interpolate scattered point data to a field: *interpdatapointtofld* module

To interpolate discrete point data to a field, use the *interpdatapointtofld* module:

```
FieldConvert -m interpdatapointtofld:frompts=file1.pts file1.xml file1.fld
```

or alternatively for csv data:

```
FieldConvert -m interpdatapointtofld:frompts=file1.csv file1.xml file1.fld
```

This command will interpolate the data from `file1.pts` (`file1.csv`) to the mesh and expansions defined in `file1.xml` and output the field to `file1.fld`. The file `file.pts` must be of the form:

```
1 <?xml version="1.0" encoding="utf-8" ?>
2 <NEKTAR>
3   <POINTS DIM="1" FIELDS="a,b,c">
4     1.0000 -1.0000 1.0000 -0.7778
5     2.0000 -0.9798 0.9798 -0.7980
6     3.0000 -0.9596 0.9596 -0.8182
7     4.0000 -0.9394 0.9394 -0.8384
8   </POINTS>
9 </NEKTAR>
```

where `DIM="1" FIELDS="a,b,c"` specifies that the field is one-dimensional and contains three variables, *a*, *b*, and *c*. Each line defines a point, while the first column contains its *x*-coordinate, the second one contains the *a*-values, the third the *b*-values and so on. In case of *n*-dimensional data, the *n* coordinates are specified in the first *n* columns accordingly. An equivalent csv file is:

```
# x, a, b, c
1.0000,-1.0000,1.0000,-0.7778
2.0000,-0.9798,0.9798,-0.7980
3.0000,-0.9596,0.9596,-0.8182
4.0000,-0.9394,0.9394,-0.8384
```

In order to interpolate 1D data to a *n*D field, specify the matching coordinate in the output field using the `interpcoord` argument:

```
FieldConvert -m interppointdatatofld:frompts=1D-file1.pts:interpcoord=1 \
3D-file1.xml 3D-file1.fld
```

This will interpolate the 1D scattered point data from `1D-file1.pts` to the y -coordinate of the 3D mesh defined in `3D-file1.xml`. The resulting field will have constant values along the x and z coordinates. For 1D Interpolation, the module implements a quadratic scheme and automatically falls back to a linear method if only two data points are given. A modified inverse distance method is used for 2D and 3D interpolation. Linear and quadratic interpolation require the data points in the `.pts`-file to be sorted by their location in ascending order. The Inverse Distance implementation has no such requirement.

5.5.19 Interpolate a field to a series of points: *interppoints* module

You can interpolate one field to a series of given points using the following command:

```
FieldConvert -m interppoints:fromxml=file1.xml:fromfld=\
file1.fld:topts=file2.pts file2.dat
```

This command will interpolate the field defined by `file1.xml` and `file1.fld` to the points defined in `file2.pts` and output it to `file2.dat`. The `fromxml` and `fromfld` must be specified in this module. The format of the file `file2.pts` is of the same form as for the *interppointdatatofld* module:

```
1 <?xml version="1.0" encoding="utf-8" ?>
2 <NEKTAR>
3   <POINTS DIM="2" FIELDS="">
4     0.0 0.0
5     0.5 0.0
6     1.0 0.0
7   </POINTS>
8 </NEKTAR>
```

Similar to the *interppointdatatofld* module, the `.pts` file can be interchanged with a `.csv` file (the output can also be written to `.csv`):

```
# x, y
0.0,0.0
0.5,0.0
1.0,0.0
```

There are three optional arguments `clamptolowervalue`, `clamptouppervalue` and `defaultvalue` the first two clamp the interpolation between these two values and the third defines the default value to be used if the point is outside the domain. Their default values are -10,000,000, 10,000,000 and 0.

In addition, instead of specifying the file `file2.pts`, a module list of the form

```
FieldConvert -m interppoints:fromxml=file1.xml:fromfld= \
    file1.fld:line=npts,x0,y0,x1,y1 file2.dat
```

can be specified where `npts` is the number of equispaced points between (x_0, y_0) to (x_1, y_1) . This also works in 3D, by specifying (x_0, y_0, z_0) to (x_1, y_1, z_1) .

An extraction of a plane of points can also be specified by

```
FieldConvert -m interppoints:fromxml=file1.xml:fromfld=file1.fld:\
    plane=npts1,npts2,x0,y0,z0,x1,y1,z1,x2,y2,z2,x3,y3,z3 file2.dat
```

where `npts1,npts2` is the number of equispaced points in each direction and (x_0, y_0, z_0) , (x_1, y_1, z_1) , (x_2, y_2, z_2) and (x_3, y_3, z_3) define the plane of points specified in a clockwise or anticlockwise direction.

In addition, an extraction of a box of points can also be specified by

```
FieldConvert -m interppoints:fromxml=file1.xml:fromfld=file1.fld:\
    box=npts1,npts2,npts3,xmin,xmax,ymin,ymax,zmin,zmax file2.dat
```

where `npts1,npts2,npts3` is the number of equispaced points in each direction and $(x_{min}, y_{min}, z_{min})$ and $(x_{max}, y_{max}, z_{max})$ define the limits of the box of points.

There is also an additional optional argument `cp=p0,q` which adds to the interpolated fields the value of $c_p = (p - p_0)/q$ and $c_{p0} = (p - p_0 + 0.5u^2)/q$ where p_0 is a reference pressure and q is the free stream dynamics pressure. If the input does not contain a field “p” or a velocity field “u,v,w” then cp and $cp0$ are not evaluated accordingly



Note

This module runs in parallel for the line, plane and box extraction of points.

5.5.20 Interpolate a set of points to another: *interpptstopts* module

You can interpolate one set of points to another using the following command:

```
FieldConvert file1.pts -m interpptstopts:topts=file2.pts file2.dat
```

This command will interpolate the data in `file1.pts` to a new set of points defined in `file2.pts` and output it to `file2.dat`.

Similarly to the *interppts* module, the target point distribution can also be specified using the `line`, `plane` or `box` options. The optional arguments `clamptolowervalue`, `clamptouppervalue`, `defaultvalue` and `cp` are also supported with the same meaning as in *interppts*.

One useful application for this module is with 3DH1D expansions, for which currently the *interppts* module does not work. In this case, we can use for example

```
FieldConvert file1.xml file1.fld -m interppts:\
plane=npts1,npts2,x0,y0,z0,x1,y1,z1,x2,y2,z2,x3,y3,z3 \
file2.dat
```

With this usage, the *equispacedoutput* module will be automatically called to interpolate the field to a set of equispaced points in each element. The result is then interpolated to a plane by the *interppts* module.



Note

This module does not work in parallel.

5.5.21 Isocontour extraction: *isocontour* module

Extract an isocontour from a field file. This option automatically take the field to an equispaced distribution of points connected by linear simplicies of triangles or tetrahedrons. The linear simplicies are then inspected to extract the isocontour of interest. To specify the field `fieldid` can be provided giving the id of the field of interest and `fieldvalue` provides the value of the isocontour to be extracted.

```
FieldConvert -m isocontour:fieldid=2:fieldvalue=0.5 test.xml test.fld \
test-isocontour.dat
```

Alternatively `fieldstr="u+v"` can be specified to calculate the field $u + v$ and extract its isocontour. You can also specify `fieldname="UplusV"` to define the name of the isocontour in the .dat file, i.e.

```
FieldConvert -m isocontour:fieldstr="u+v":fieldvalue=0.5:\
fieldname="UplusV" test.xml test.fld test-isocontour.dat
```

Optionally `smooth` can be specified to smooth the isocontour with default values `smoothnegdiffusion`=0.495, `smoothnegdiffusion`=0.5 and `smoothiter`=100. This option typically should be used with the `globalcondense` option which removes multiply defined vertices from the simplex definition which arise as isocontour are generated element by element. The `smooth` option previously automatically called the `globalcondense`

option but this has been deprecated since it is now possible to read isocontour files directly and so it is useful to have these as separate options.

In addition to the `smooth` or `globalcondense` options you can specify `removesmallcontour=100` which will remove separate isocontours of less than 100 triangles.



Note

Currently this option is only set up for triangles, quadrilaterals, tetrahedrons and prisms.

5.5.22 Show high frequency energy of the Jacobian: *jacobianenergy* module

```
FieldConvert -m jacobianenergy file.xml file.fld jacenergy.fld
```

The option `topmodes` can be used to specify the number of top modes to keep.

The output file `jacenergy.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.23 Calculate mesh quality: *qualitymetric* module

The `qualitymetric` module assesses the quality of the mesh by calculating a per-element quality metric and adding an additional field to any resulting output. This does not require any field input, therefore an example usage looks like

```
FieldConvert -m qualitymetric mesh.xml mesh-with-quality.dat
```

Two quality metrics are implemented that produce scalar fields Q :

- By default a metric outlined in [14] is produced, where all straight sided elements have quality $Q = 1$ and $Q < 1$ shows the deformation between the curved element and the straight-sided element. If $Q = 0$ then the element is invalid. Note that Q varies over the volume of the element but is not guaranteed to be continuous between elements.
- Alternatively, if the `scaled` option is passed through to the module, then the scaled Jacobian

$$J_s = \frac{\min_{\xi \in \Omega_{st}} J(\xi)}{\max_{\xi \in \Omega_{st}} J(\xi)}$$

(i.e. the ratio of the minimum to maximum Jacobian of each element) is calculated. Again $Q = 1$ denotes an ideal element, but now invalid elements are shown by $Q < 0$. Any elements with Q near zero are determined to be low quality.

5.5.24 Evaluate the mean of variables on the domain: *mean* module

To evaluate the mean of variables on the domain one can use the `mean` module of FieldConvert

```
FieldConvert -m mean file1.xml file2.fld out.stdout
```

This module does not create an output file which is reinforced by the `out.stdout` option. The integral and mean for each field variable are then printed to the stdout.

5.5.25 Extract mean mode of 3DH1D expansion: *meanmode* module

To obtain a 2D expansion containing the mean mode (plane zero in Fourier space) of a 3DH1D field file, use the command:

```
FieldConvert -m meanmode file.xml file.fld file-mean.fld
```

The output file `file-mean.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot or in Paraview or VisIt.

5.5.26 Project point data to a field: *pointdatatofld* module

To project a series of points given at the same quadrature distribution as the .xml file and write out a .fld file use the `pointdatatofld` module:

```
FieldConvert -m pointdatatofld:frompts=file.pts file.xml file.fld
```

This command will read in the points provided in the `file.pts` and assume these are given at the same quadrature distribution as the mesh and expansions defined in `file.xml` and output the field to `file.fld`. If the points do not match an error will be dumped.

The file `file.pts` which is assumed to be given by an interpolation from another source is of the form:

```
1 <?xml version="1.0" encoding="utf-8" ?>
2 <NEKTAR>
3   <POINTS DIM="3" FIELDS="p">
4     1.70415 -0.4 -0.0182028 -0.106893
5     1.70415 -0.395683 -0.0182028 -0.106794
6     1.70415 -0.3875 -0.0182028 -0.106698
7     1.70415 -0.379317 -0.0182028 -0.103815
8   </POINTS>
9 </NEKTAR>
```

where `DIM="3" FIELDS="p"` specifies that the field is three-dimensional and contains one variable, p . Each line defines a point, the first, second, and third columns contains the

x, y, z -coordinate and subsequent columns contain the field values, in this case the p -value. So in the general case of n -dimensional data, the n coordinates are specified in the first n columns accordingly followed by the field data. Alternatively, the `file.pts` can be interchanged with a csv file.

The default argument is to use the equispaced (but potentially collapsed) coordinates which can be obtained from the command.

```
FieldConvert file.xml file.dat
```

In this case the `pointdatatofld` module should be used without the `-noequispaced` option. However this can lead to problems when performing an elemental forward projection/transform since the mass matrix in a deformed element can be singular as the equispaced points do not have a sufficiently accurate quadrature rule that spans the polynomial space. Therefore it is advisable to use the set of points given by

```
FieldConvert --noequispaced file.xml file.dat
```

which produces a set of points at the gaussian collapsed coordinates.

Finally the option `setnantovalue=0` can also be used which sets any nan values in the interpolation to zero or any specified value in this option.

5.5.27 Print L2 and LInf norms: *printfldnorms* module

```
FieldConvert -m printfldnorms test.xml test.fld out.stdout
```

This module does not create an output file which is reinforced by the `out.stdout` option. The L2 and LInf norms for each field variable are then printed to the stdout.

5.5.28 Removes one or more fields from .fld files: *removefield* module

This module allows to remove one or more fields from a .fld file:

```
FieldConvert -m removefield:fieldname="u,v,p" test.xml test.fld test-removed.fld
```

where the file `test-removed.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt. The lighter resulting file speeds up the postprocessing of large files when not all fields are required.

5.5.29 Computes the scalar gradient: *scalargrad* module

The scalar gradient of a field is computed by running:

```
FieldConvert -m scalargrad:bnd=0 test.xml test.fld test-scalgrad.fld
```

The option `bnd` specifies which boundary region to extract. Note this is different to NekMesh where the parameter `surf` is specified and corresponds to composites rather boundaries. If `bnd` is not provided, all boundaries are extracted to different fields. To process this file you will need an xml file of the same region.

5.5.30 Scale a given .fld: *scaleinputfld* module

To scale a .fld file by a given scalar quantity, the user can run:

```
FieldConvert -m scaleinputfld:scale=value test.xml test.fld test-scal.fld
```

The argument `scale=value` rescales of a factor `value` `test.fld` by the factor value. The output file `file-scal.fld` can be processed in a similar way as described in section 5.2 to visualise the result either in Tecplot, Paraview or VisIt.

5.5.31 Time-averaged shear stress metrics: *shear* module

Time-dependent wall shear stress derived metrics relevant to cardiovascular fluid dynamics research can be computed using this module. They are

- TAWSS: time-averaged wall shear stress;
- OSI: oscillatory shear index;
- transWSS: transverse wall shear stress;
- TACFI: time-averaged cross-flow index;
- TAAFI: time-averaged aneurysm formation index;
- $|WSSG|$: wall shear stress gradient.

To compute these, the user can run:

```
FieldConvert -m shear:N=value:fromfld=test_id_b0.fld \
test.xml test-multishear.fld
```

The argument `N` and `fromfld` are compulsory arguments that respectively define the number of `fld` files corresponding to the number of discrete equispaced time-steps, and the first `fld` file which should have the form of `test_id_b0.fld` where the first underscore in the name marks the starting time-step file ID.

The input `.fld` files are the outputs of the *wss* module. If they do not contain the surface normals (an optional output of the *wss* module), then the *shear* module will not compute the last metric, $|\text{WSSG}|$.

5.5.32 Stream function of a 2D incompressible flow: *streamfunction* module

The streamfunction module calculates the stream function of a 2D incompressible flow, by solving the Poisson equation

$$\nabla^2 \psi = -\omega$$

where ω is the vorticity. Note that this module applies the same boundary conditions specified for the y-direction velocity component (`v`) to the stream function, what may not be the most appropriate choice.

To use this module, the user can run

```
FieldConvert -m streamfunction test.xml test.fld test-streamfunc.fld
```

where the file `test-streamfunc.fld` can be processed in a similar way as described in section 5.2.

5.5.33 Boundary layer height calculation: *surfdistance* module

The surface distance module computes the height of a boundary layer formed by quadrilaterals (in 2D) or prisms and hexahedrons (in 3D) and projects this value onto the surface of the boundary, in a similar fashion to the `extract` module. In conjunction with a mesh of the surface, which can be obtained with `NekMesh`, and a value of the average wall shear stress, one potential application of this module is to determine the distribution of y^+ grid spacings for turbulence calculations.

To compute the height of the prismatic layer connected to boundary region 3, the user can issue the command:

```
FieldConvert -m surfdistance:bnd=3 input.xml output.fld
```

Note that no `.fld` file is required, since the mesh is the only input required in order to calculate the element height. This produces a file `output_b3.fld`, which can be visualised with the appropriate surface mesh from `NekMesh`.

5.5.34 Calculate vorticity: *vorticity* module

To perform the vorticity calculation and obtain an output data containing the vorticity solution, the user can run

```
FieldConvert -m vorticity test.xml test.fld test-vort.fld
```

where the file `test-vort.fld` can be processed in a similar way as described in section 5.2.

5.5.35 Computing the wall shear stress: *wss* module

To obtain the wall shear stress vector and magnitude, the user can run:

```
FieldConvert -m wss:bnd=0:addnormals=1 test.xml test.fld test-wss.fld
```

The option `bnd` specifies which boundary region to extract. Note this is different to NekMesh where the parameter `surf` is specified and corresponds to composites rather boundaries. If `bnd` is not provided, all boundaries are extracted to different fields. The `addnormals` is an optional command argument which, when turned on, outputs the normal vector of the extracted boundary region as well as the shear stress vector and magnitude. This option is off by default. To process the output file(s) you will need an xml file of the same region.

5.5.36 Manipulating meshes with FieldConvert

FieldConvert has support for two modules that can be used in conjunction with the linear elastic solver, as shown in chapter 12. To do this, FieldConvert has an XML output module, in addition to the Tecplot and VTK formats.

The `deform` module, which takes no options, takes a displacement field and applies it to the geometry, producing a deformed mesh:

```
FieldConvert -m deform input.xml input.fld deformed.xml
```

The `displacement` module is designed to create a boundary condition field file. Its intended use is for mesh generation purposes. It can be used to calculate the displacement between the linear mesh and a high-order surface, and then produce a `fld` file, prescribing the displacement at the boundary, that can be used in the linear elasticity solver.

Presently the process is somewhat convoluted and must be used in conjunction with NekMesh to create the surface file. However the bash input below describes the procedure. Assume the high-order mesh is in a file called `mesh.xml`, the linear mesh is `mesh-linear.xml` that can be generated by removing the `CURVED` section from `mesh.xml`, and that we are interested in the surface with ID 123.

```
# Extract high order surface
NekMesh -m extract:surf=123 mesh.xml mesh-surf-curved.xml

# Use FieldConvert to calculate displacement between two surfaces
FieldConvert -m displacement:id=123:to=mesh-surf-curved.xml \
  mesh-linear.xml mesh-deformation.fld
```

```
# mesh-deformation.fld is used as a boundary condition inside the
# solver to prescribe the deformation conditions.xml contains
# appropriate Nektar++ parameters (mu, E, other BCs, ...)
LinearElasticSolver mesh-linear.xml conditions.xml

# This produces the final field mesh-linear.fld which is the
# displacement field, use FieldConvert to apply it:
FieldConvert-g -m deform mesh-linear.xml mesh-linear.fld mesh-deformed.xml
```

5.6 FieldConvert in parallel

To run FieldConvert in parallel the user needs to compile *Nektar++* with MPI support and can employ the following command

```
mpirun -np <nprocs> FieldConvert test.xml test.fld test.dat
```

```
mpirun -np <nprocs> FieldConvert test.xml test.fld test.plt
```

or

```
mpirun -np <nprocs> FieldConvert test.xml test.fld test.vtu
```

replacing `<nprocs>` with the number of processors. For the `.dat` and `.plt` outputs the current version will produce a single output file. However it is also sometimes useful to produce multiple output files, one for each partition, and this can be done by using the `writemultiplefiles` option, i.e.

```
mpirun -np <nprocs> FieldConvert test.xml test.fld \
    test.dat:dat:writemultiplefiles
```

```
mpirun -np <nprocs> FieldConvert test.xml test.fld \
    test.plt:plt:writemultiplefiles
```

For the `.vtu` format multiple files will by default be produced of the form `test_vtu/P0000000.vtu`, `test_vtu/P0000001.vtu`, `test_vtu/P0000002.vtu`. For this format an additional file called `test.pvtu` is written out which allows for parallel reading of the individual `.vtu` files.

FieldConvert functions that produce a `.fld` file output will also be created when running in parallel. In this case when producing a `.fld` file a directory called `test.fld` (or the specified output name) is created with the standard parallel field files placed within the directory.

5.7 Processing large files in serial

When processing large files, it is not always convenient to run in parallel but process each parallel partition in serial, for example when interpolating a solution field from one mesh to another or creating an output file for visualization.

5.7.1 Using the *part-only* and *part-only-overlapping* options

Loading full `file1.xml` can be expensive if the `file1.xml` is already large. So instead you can pre-partition the file using the using the `-part-only` option. So the command

```
FieldConvert --part-only 10 file.xml file.fld
```

will partition the mesh into 10 partitions and write each partition into a directory called `file_xml`. If you enter this directory you will find partitioned XML files `P0000000.xml`, `P0000001.xml`, ..., `P0000009.xml` which can then be processed individually as outlined above.

There is also a `-part-only-overlapping` option, which can be run in the same fashion.

```
FieldConvert --part-only-overlapping 10 file.xml file.fld
```

In this mode, the mesh is partitioned into 10 partitions in a similar manner, but the elements at the partition edges will now overlap, so that the intersection of each partition with its neighbours is non-empty. This is sometime helpful when, for example, producing a global isocontour which has been smoothed. Applying the smoothed isocontour extraction routine with the `-part-only` option will produce a series of isocontour where there will be a gap between partitions, as the smoother tends to shrink the isocontour within a partition. using the `-part-only-overlapping` option will still yield a shrinking isocontour, but the overlapping partitions help to overlap the partiiton boundaries.

5.7.2 Using the *nparts* options

If you have a partitioned directory either from a parallel run or using the `-part-only` option you can now run the `FieldConvert` option using the `nparts` command line option, that is

```
FieldConvert --nparts 10 file1\_xml:xml file1.fld file1.vtu
```

Note the form `file1_xml:xml` option tells the code it is a parallel partition which should be treated as an `xml` type file. the argument of `nparts` should correpond to the number of partitions used in generating the `file1_xml` directory. This will create a parallel vtu file as it processes each partition.

Another example is to interpolate `file1.fld` from one mesh `file1.xml` to another `file2.xml`. If the mesh files are large we can do this by partitioning `file2.xml` into 10 (or more) partitions to generate the `file_xml` directory and interpolating each partition one by one using the command:

```
FieldConvert --nparts 10 -m interpfield:fromxml=file1.xml:fromfld=file1.fld
\
file2_xml:xml file2.fld
```

Note that internally the routine uses the range option so that it only has to load the part of `file1.xml` that overlaps with each partition of `file2.xml`. The resulting output will lie in a directory called `file2.fld`, with each of the different parallel partitions in files with names `P0000000.fld`, `P0000001.fld`, ..., `P0000009.fld`. In previous versions of FieldConvert it was necessary to generate an updated `Info.xml` file but in the current version it should automatically be updating this file.

5.7.3 Running in parallel with the *nparts* option

The examples above will process each partition serially which may now take a while for many partitions. You can however run this option in parallel using a smaller number of cores than the *nparts*.

For the example of creating a *vtu* file above you can use 4 processor concurrently with the command line:

```
mpirun -n 4 FieldConvert --nparts 10 file1\_xml:xml file1.fld file1.vtu
```

Obviously the executable will have to have been compiled with the MPI option for this to work.

Part III

Solver Applications

Acoustic Solver

6.1 Synopsis

The aim of the AcousticSolver is to predict acoustic wave propagation. Through the application of a splitting technique, the flow-induced acoustic field is totally decoupled from the underlying hydrodynamic field.

6.1.1 Linearized Euler Equations

The Linearized Euler Equations (LEE) are obtained by linearizing the Euler Equations about a mean flow state $(\bar{\rho}, \bar{c}^2, \bar{\mathbf{u}})$. Hence, they describe the evolution of perturbations $(p^a, \rho^a, \bar{\rho}\mathbf{u}^a)$ around this state. In conservative form, the LEE are given as:

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}_1}{\partial x_1} + \frac{\partial \mathbf{F}_2}{\partial x_2} + \frac{\partial \mathbf{F}_3}{\partial x_3} + \mathbf{C}\mathbf{U} = \mathbf{W} \quad (6.1)$$

with

$$\mathbf{U} = \begin{bmatrix} p^a \\ \rho^a \\ \bar{\rho}u_1^a \\ \bar{\rho}u_2^a \\ \bar{\rho}u_3^a \end{bmatrix}, \quad (6.2)$$

$$\mathbf{F}_1 = \begin{bmatrix} \bar{\rho}u_1^a \bar{c}^2 + \bar{u}_1 p^a \\ \bar{\rho}u_1^a + \bar{u}_1 \rho^a \\ \bar{\rho}u_1^a \bar{u}_1 + p^a \\ \bar{\rho}u_2^a \bar{u}_1 \\ \bar{\rho}u_3^a \bar{u}_1 \end{bmatrix}, \quad \mathbf{F}_2 = \begin{bmatrix} \bar{\rho}u_2^a \bar{c}^2 + \bar{u}_2 p^a \\ \bar{\rho}u_2^a + \bar{u}_2 \rho^a \\ \bar{\rho}u_1^a \bar{u}_2 \\ \bar{\rho}u_2^a \bar{u}_2 + p^a \\ \bar{\rho}u_3^a \bar{u}_2 \end{bmatrix}, \quad \mathbf{F}_3 = \begin{bmatrix} \bar{\rho}u_3^a \bar{c}^2 + \bar{u}_3 p^a \\ \bar{\rho}u_3^a + \bar{u}_3 \rho^a \\ \bar{\rho}u_1^a \bar{u}_3 \\ \bar{\rho}u_2^a \bar{u}_3 \\ \bar{\rho}u_3^a \bar{u}_3 + p^a \end{bmatrix}, \quad (6.3)$$

$$\mathbf{C} = \begin{bmatrix} (\gamma - 1) \frac{\partial \bar{u}_k}{\partial x_k} & 0 & \frac{1}{\bar{\rho}} (1 - \gamma) \frac{\partial \bar{p}}{\partial x_1} & \frac{1}{\bar{\rho}} (1 - \gamma) \frac{\partial \bar{p}}{\partial x_2} & \frac{1}{\bar{\rho}} (1 - \gamma) \frac{\partial \bar{p}}{\partial x_3} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & \bar{u}_k \frac{\partial \bar{u}_1}{\partial x_k} & \frac{\partial \bar{u}_1}{\partial x_1} & \frac{\partial \bar{u}_1}{\partial x_2} & \frac{\partial \bar{u}_1}{\partial x_3} \\ 0 & \bar{u}_k \frac{\partial \bar{u}_2}{\partial x_k} & \frac{\partial \bar{u}_2}{\partial x_1} & \frac{\partial \bar{u}_2}{\partial x_2} & \frac{\partial \bar{u}_2}{\partial x_3} \\ 0 & \bar{u}_k \frac{\partial \bar{u}_3}{\partial x_k} & \frac{\partial \bar{u}_3}{\partial x_1} & \frac{\partial \bar{u}_3}{\partial x_2} & \frac{\partial \bar{u}_3}{\partial x_3} \end{bmatrix}. \quad (6.4)$$

By default, the source term vector \mathbf{W} is zero and has to be specified by an appropriate forcing.

6.1.2 Acoustic Perturbation Equations

The acoustic perturbation equations (APE-1/APE-4) proposed by Ewert and Schroeder [13] assure stable aeroacoustic simulations. These equations are similar to the LEE, but account for acoustic perturbations exclusively. The AcousticSolver implements the APE-1/4 type operator:

$$\frac{\partial p^a}{\partial t} + \bar{c}^2 \nabla \cdot \left(\bar{\rho} \mathbf{u}^a + \bar{\mathbf{u}} \frac{p^a}{\bar{c}^2} \right) = \dot{\omega}_c \quad (6.5a)$$

$$\frac{\partial \mathbf{u}^a}{\partial t} + \nabla (\bar{\mathbf{u}} \cdot \mathbf{u}^a) + \nabla \left(\frac{p^a}{\bar{\rho}} \right) = \dot{\omega}_m, \quad (6.5b)$$

where $(\bar{\mathbf{u}}, \bar{c}^2, \bar{\rho})$ represents the base flow and (\mathbf{u}^a, p^a) the acoustic perturbations. Similar to the LEE, the acoustic source terms $\dot{\omega}_c$ and $\dot{\omega}_m$ are by default zero and must be specified e.g. by an appropriate forcing. This way, e.g. the APE-1, APE-4 [13] or revised APE equations [15] can be obtained. Expressed as hyperbolic conservation law, the APE-1/4 operator reads:

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}_1}{\partial x_1} + \frac{\partial \mathbf{F}_2}{\partial x_2} + \frac{\partial \mathbf{F}_3}{\partial x_3} = \mathbf{W} \quad (6.6)$$

with

$$U = \begin{bmatrix} p^a \\ u_1^a \\ u_2^a \\ u_3^a \end{bmatrix}, \quad (6.7)$$

$$F_1 = \begin{bmatrix} \bar{\rho} c^2 u_1^a + p^a \bar{u}_1 \\ \bar{u}_j u_j^a + p^a / \bar{\rho} \\ 0 \\ 0 \end{bmatrix}, \quad F_2 = \begin{bmatrix} \bar{\rho} c^2 u_2^a + p^a \bar{u}_2 \\ 0 \\ \bar{u}_j u_j^a + p^a / \bar{\rho} \\ 0 \end{bmatrix}, \quad F_3 = \begin{bmatrix} \bar{\rho} c^2 u_3^a + p^a \bar{u}_3 \\ 0 \\ 0 \\ \bar{u}_j u_j^a + p^a / \bar{\rho} \end{bmatrix}. \quad (6.8)$$

6.2 Usage

```
AcousticSolver session.xml
```

6.3 Session file configuration

Parameters

Under this section it is possible to set the parameters of the simulation.

```
1 <PARAMETERS>
2   <P> TimeStep      = 1e-05 /P>
3   <P> NumSteps      = 1000  /P>
4   <P> FinTime       = 0.01  /P>
5   <P> IO_CheckSteps = 100    /P>
6   <P> IO_InfoSteps  = 10     /P>
7   <P> IO_CFLSteps   = 10     /P>
8 </PARAMETERS>
```

- `TimeStep` is the time-step we want to use;
- `FinTime` is the final physical time at which we want our simulation to stop;
- `NumSteps` is the equivalent of `FinTime` but instead of specifying the physical final time we specify the number of time-steps;
- `IO_CheckSteps` sets the number of steps between successive checkpoint files;
- `IO_InfoSteps` sets the number of steps between successive info stats are printed to screen;
- `IO_CFLSteps` sets the number of steps between successive Courant number stats are printed to screen;

6.3.1 Solver Info

```

1 <SOLVERINFO>
2   <I PROPERTY="EQType"           VALUE="APE"           />
3   <I PROPERTY="Projection"       VALUE="DisContinuous" />
4   <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4" />
5   <I PROPERTY="UpwindType"       VALUE="LaxFriedrichs"  />
6 </SOLVERINFO>

```

- `EQType` is the tag which specify the equations we want solve:
 - `APE` Acoustic Perturbation Equations (variables: `p,u,v,w`);
 - `LEE` Linearized Euler Equations (variables: `p,rho,rhou,rhov,rhow`).
- `Projection` is the type of projection we want to use. Currently, only `DisContinuous` is supported.
- `AdvectionType` is the advection operator. Currently, only `WeakDG` (classical DG in weak form) is supported.
- `TimeIntegrationMethod` is the time-integration scheme we want to use. Note that only an explicit discretisation is supported:
- `UpwindType` is the numerical interface flux (i.e. Riemann solver) we want to use for the advection operator (see [24] for the implemented formulations):
 - `Upwind`;
 - `LaxFriedrichs`;

6.3.2 Variables

For the APE operator, the acoustic pressure and velocity perturbations are solved, e.g.:

```

1 <VARIABLES>
2   <V ID="0"> p </V>
3   <V ID="1"> u </V>
4   <V ID="2"> v </V>
5   <V ID="3"> w </V>
6 </VARIABLES>

```

The LEE use a conservative formulation and introduce the additional density perturbation:

```

1 <VARIABLES>
2   <V ID="0"> p </V>
3   <V ID="1"> rho </V>
4   <V ID="2"> rhou </V>
5   <V ID="3"> rhov </V>
6   <V ID="4"> rhow </V>
7 </VARIABLES>

```

6.3.3 Functions

- **BaseFlow** Baseflow $(\bar{\rho}, \bar{c}^2, \bar{\mathbf{u}})$ defined by the variables `rho0, c0sq, u0, v0, w0` for APE and $(\bar{\rho}, \bar{c}^2, \bar{\mathbf{u}}, \gamma)$ defined by `rho0, c0sq, u0, v0, w0, gamma` for LEE.
- **InitialConditions**

6.3.4 Boundary Conditions

In addition to plain Dirichlet and Neumann boundary conditions, the `AcousticSolver` features a slip-wall boundary condition, a non-reflecting boundary and a white noise boundary condition.

- Rigid (Slip-) Wall Boundary Condition, e.g. for APE:

```
1 <BOUNDARYCONDITIONS>
2 <REGION REF="0">
3   <D VAR="p" USERDEFINEDTYPE="Wall" VALUE="0" />
4   <D VAR="u" USERDEFINEDTYPE="Wall" VALUE="0" />
5   <D VAR="v" USERDEFINEDTYPE="Wall" VALUE="0" />
6   <D VAR="w" USERDEFINEDTYPE="Wall" VALUE="0" />
7 </REGION>
8 </BOUNDARYCONDITIONS>
```

This BC imposes zero wall-normal perturbation velocity in a way that is more robust than using a Dirichlet boundary condition directly.

- Non-Reflecting Boundary Condition, e.g. for APE:

```
1 <BOUNDARYCONDITIONS>
2 <REGION REF="0">
3   <D VAR="p" USERDEFINEDTYPE="RiemannInvariantBC"/>
4   <D VAR="u" USERDEFINEDTYPE="RiemannInvariantBC"/>
5   <D VAR="v" USERDEFINEDTYPE="RiemannInvariantBC"/>
6   <D VAR="w" USERDEFINEDTYPE="RiemannInvariantBC"/>
7 </REGION>
8 </BOUNDARYCONDITIONS>
```

The Riemann-Invariant BC approximates a non-reflecting (r.g. Farfield) boundary condition by setting incoming invariants to zero.

- White Noise Boundary Condition, e.g. for APE:

```
1 <BOUNDARYCONDITIONS>
2 <REGION REF="0">
3   <D VAR="p" USERDEFINEDTYPE="Wall" VALUE="10" />
4   <D VAR="u" USERDEFINEDTYPE="Wall" VALUE="10" />
5   <D VAR="v" USERDEFINEDTYPE="Wall" VALUE="10" />
6   <D VAR="w" USERDEFINEDTYPE="Wall" VALUE="10" />
7 </REGION>
8 </BOUNDARYCONDITIONS>
```


The white noise BC imposes a stochastic, uniform pressure at the boundary. The implementation uses a Mersenne-Twister pseudo random number generator to generate white Gaussian noise. The standard deviation σ of the pressure is specified by the `VALUE` attribute.

6.4 Examples

6.4.1 Wave Propagation in a Sheared Base Flow

In this section we explain how to set up a simple, 2D simulation of aeroacoustics in Nektar++. We will study the propagation of an acoustic wave in the simple case of a sheared base flow, i.e. $\bar{\mathbf{u}} = [300 \tanh(20x_2), 0]^T$, $\bar{c}^2 = (341 \text{ m/s})^2$, $\bar{\rho} = 1.204 \text{ kg/m}^3$. The geometry consists of 64 quadrilateral elements.

6.4.1.1 Input file

We require a discontinuous Galerkin projection and use an explicit fourth-order Runge-Kutta time integration scheme. We therefore set the following solver information:

```
1 <SOLVERINFO>
2   <I PROPERTY="EQType"           VALUE="APE"/>
3   <I PROPERTY="Projection"       VALUE="DisContinuous"/>
4   <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4"/>
5   <I PROPERTY="UpwindType"      VALUE="LaxFriedrichs"/>
6 </SOLVERINFO>
```

To maintain numerical stability we must use a small time-step. Finally, we set the density, heat ratio and ambient pressure.

```
1 <PARAMETERS>
2   <P> TimeStep      = 1e-05      </P>
3   <P> NumSteps      = 1000       </P>
4   <P> FinTime       = TimeStep*NumSteps </P>
5   <P> IO_CheckSteps = 10         </P>
6   <P> IO_InfoSteps  = 10         </P>
7 </PARAMETERS>
```

The initial condition and the base flow field are specified by the `Baseflow` and `InitialConditions` functions, respectively:

```
1 <FUNCTION NAME="Baseflow">
2   <E VAR="u0" VALUE="300 * tanh(2*y/0.1)"/>
3   <E VAR="v0" VALUE="0"/>
4   <E VAR="c0sq" VALUE="1.4 * Pinfinitiy / Rho0"/>
5   <E VAR="rho0" VALUE="Rho0"/>
6 </FUNCTION>
7 <FUNCTION NAME="InitialConditions">
8   <E VAR="p" VALUE="0"/>
9   <E VAR="u" VALUE="0"/>
10  <E VAR="v" VALUE="0"/>
11 </FUNCTION>
```

At all four boundaries the `RiemannInvariantBC` condition is used:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="p" USERDEFINEDTYPE="RiemannInvariantBC"/>
4     <D VAR="u" USERDEFINEDTYPE="RiemannInvariantBC"/>
5     <D VAR="v" USERDEFINEDTYPE="RiemannInvariantBC"/>
6   </REGION>
7 </BOUNDARYCONDITIONS>

```

The system is excited via an acoustic source term $\hat{\omega}_c$, which is modeled by a field forcing as:

```

1 <FORCING>
2   <FORCE TYPE="Field">
3     <FIELDFORCE> Source </FIELDFORCE/>
4   </FORCE>
5 </FORCING>

```

and the corresponding function

```

1 <FUNCTION NAME="Source">
2   <E VAR="p" VALUE="100 * 2*PI*5E2 * cos(2*PI*5E2 * t) * exp(-32*(x^2+y^2))"/>
3   <E VAR="u" VALUE="0"/>
4   <E VAR="v" VALUE="0"/>
5 </FUNCTION>

```

6.4.1.2 Running the code

```
AcousticSolver Test_pulse.xml
```

6.4.1.3 Results

Fig. 6.1 shows the acoustic source term, the velocity and the acoustic pressure and velocity perturbations at a single time step.




Figure 6.1 Acoustic source term, base flow velocity, acoustic pressure and acoustic velocity perturbations.

Advection-Diffusion-Reaction Solver

7.1 Synopsis

The ADRSolver is designed to solve partial differential equations of the form:

$$\alpha \frac{\partial u}{\partial t} + \lambda u + \nu \nabla u + \epsilon \nabla \cdot (D \nabla u) = f \quad (7.1)$$

in either discontinuous or continuous projections of the solution field. For a full list of the equations which are supported, and the capabilities of each equation, see the table below.

Equation to solve	EquationType	Dimensions	Projections
$u = f$	Projection	All	Continuous/Discontinuous
$\nabla^2 u = 0$	Laplace	All	Continuous/Discontinuous
$\nabla^2 u = f$	Poisson	All	Continuous/Discontinuous
$\nabla^2 u + \lambda u = f$	Helmholtz	All	Continuous/Discontinuous
$\epsilon \nabla^2 u + \mathbf{V} \nabla u = f$	SteadyAdvectionDiffusion	2D only	Continuous/Discontinuous
$\epsilon \nabla^2 u + \lambda u = f$	SteadyDiffusionReaction	2D only	Continuous/Discontinuous
$\epsilon \nabla^2 u + \mathbf{V} \nabla u + \lambda u = f$	SteadyAdvectionDiffusionReaction	2D only	Continuous/Discontinuous
$\frac{\partial u}{\partial t} + \mathbf{V} \nabla u = f$	UnsteadyAdvection	All	Continuous/Discontinuous
$\frac{\partial u}{\partial t} = \epsilon \nabla^2 u$	UnsteadyDiffusion	All	Continuous/Discontinuous
$\frac{\partial u}{\partial t} = \epsilon \nabla^2 u + R(u)$	UnsteadyReactionDiffusion	All	Continuous
$\frac{\partial u}{\partial t} + \mathbf{V} \nabla u = \epsilon \nabla^2 u$	UnsteadyAdvectionDiffusion	All	Continuous/Discontinuous
$\frac{\partial u}{\partial t} + u \nabla u = 0$	UnsteadyInviscidBurger	1D only	Continuous/Discontinuous

Table 7.1 Equations supported by the ADRSolver with their capabilities.

7.2 Usage

```
ADRSolver session.xml
```

7.3 Session file configuration

The type of equation which is to be solved is specified through the EquationType SOLVERINFO option in the session file. This can be set as in table 7.1. At present, the Steady non-symmetric solvers cannot be used in parallel.

7.3.1 Solver Info

The solver info are listed below:

- **Eqtype:** This sets the type of equation to solve, according to the table above.
- **TimeIntegrationMethod:** The following types of time integration methods have been tested with each solver:

EqType	Explicit	Diagonally Implicit	IMEX	Implicit
UnsteadyAdvection	✓			
UnsteadyDiffusion	✓	✓		
UnsteadyReactionDiffusion			✓	
UnsteadyAdvectionDiffusion			✓	
UnsteadyInviscidBurger	✓			

- **Projection:** The Galerkin projection used may be either:
 - **Continuous** for a C0-continuous Galerkin (CG) projection.
 - **Discontinuous** for a discontinuous Galerkin (DG) projection.
- **DiffusionAdvancement:** This specifies how to treat the diffusion term. This will be restricted by the choice of time integration scheme:
 - **Explicit** Requires the use of an explicit time integration scheme.
 - **Implicit** Requires the use of a diagonally implicit, IMEX or Implicit scheme.
- **AdvectionAdvancement:** This specifies how to treat the advection term. This will be restricted by the choice of time integration scheme:
 - **Explicit** Requires the use of an explicit or IMEX time integration scheme.
 - **Implicit** Not supported at present.

- **AdvectionType:** Specifies the type of advection:
 - `NonConservative` (for CG only).
 - `WeakDG` (for DG only).
- **DiffusionType:**
 - `LDG` (The penalty term is proportional to an optional parameter `LDGc11` which is by default set to one; proportionality to polynomial order can be manually imposed by setting the parameter `LDGc11` equal to p^2).
- **UpwindType:**
 - `Upwind`.

7.3.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file:

- `epsilon`: sets the diffusion coefficient ϵ .
Can be used in: SteadyDiffusionReaction, SteadyAdvectionDiffusionReaction, UnsteadyDiffusion, UnsteadyAdvectionDiffusion.
Default value: 0.
- `d00`, `d11`, `d22`: sets the diagonal entries of the diffusion tensor D .
Can be used in: UnsteadyDiffusion
Default value: All set to 1 (i.e. identity matrix).
- `lambda`: sets the reaction coefficient λ .
Can be used in: SteadyDiffusionReaction, Helmholtz, SteadyAdvectionDiffusionReaction
Default value: 0.

7.3.3 Functions

The following functions can be specified inside the `CONDITIONS` section of the session file:

- `AdvectionVelocity`: specifies the advection velocity \mathbf{V} .
- `InitialConditions`: specifies the initial condition for unsteady problems.
- `Forcing`: specifies the forcing function f .

7.4 Examples

Example files for the ADRSolver are provided in `solvers/ADRSolver/Examples`

7.4.1 1D Advection equation

In this example, it will be demonstrated how the Advection equation can be solved on a one-dimensional domain.

7.4.1.1 Advection equation

We consider the hyperbolic partial differential equation:

$$\frac{\partial u}{\partial t} + \frac{\partial f}{\partial x} = 0, \quad (7.2)$$

where $f = au$ is the advection flux.

7.4.1.2 Input file

The input for this example is given in the example file `Advection1D.xml`

The geometry section defines a 1D domain consisting of 10 segments. On each segment an expansion consisting of 4 Lagrange polynomials on the Gauss-Lobatto-Legendre points is used as specified by

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" FIELDS="u" TYPE="GLL_LAGRANGE_SEM" NUMMODES="4"/>
3 </EXPANSIONS>
```

Since we are solving the unsteady advection problem, we must specify this in the solver information. We also choose to use a discontinuous flux-reconstruction projection and use a Runge-Kutta order 4 time-integration scheme.

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE" VALUE="UnsteadyAdvection" />
3   <I PROPERTY="Projection" VALUE="DisContinuous" />
4   <I PROPERTY="AdvectionType" VALUE="FRDG" />
5   <I PROPERTY="UpwindType" VALUE="Upwind" />
6   <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4" />
7 </SOLVERINFO>
```

We choose to advect our solution for 20 time units with a time-step of 0.01 and so provide the following parameters

```
1 <P> FinTime      = 20          </P>
2 <P> TimeStep     = 0.01       </P>
3 <P> NumSteps     = FinTime/TimeStep </P>
```

We also specify the advection velocity. We first define dummy parameters

```
1 <P> advx          = 1          </P>
2 <P> advy          = 0          </P>
```

and then define the actual advection function as

```

1 <FUNCTION NAME="AdvectionVelocity">
2   <E VAR="Vx" VALUE="advx" />
3 </FUNCTION>

```

Two boundary regions are defined, one at each end of the domain, and periodicity is enforced

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[2] </B>
4 </BOUNDARYREGIONS>
5
6 <BOUNDARYCONDITIONS>
7   <REGION REF="0">
8     <P VAR="u" VALUE="[1]" />
9   </REGION>
10  <REGION REF="1">
11    <P VAR="u" VALUE="[0]" />
12  </REGION>
13 </BOUNDARYCONDITIONS>

```

Finally, we specify the initial value of the solution on the domain

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="u" VALUE="exp(-20.0*x*x)" />
3 </FUNCTION>
4
5 <FUNCTION NAME="ExactSolution">
6   <E VAR="u" VALUE="exp(-20.0*x*x)" />
7 </FUNCTION>

```

7.4.1.3 Running the code

```
ADRSolver Advection1D.xml
```

To visualise the output, we can convert it into either TecPlot or VTK formats

```
FieldConvert Advection1D.xml Advection1D.fld Advection1D.dat
FieldConvert Advection1D.xml Advection1D.fld Advection1D.vtu
```

7.4.2 2D Helmholtz Problem

In this example, it will be demonstrated how the Helmholtz equation can be solved on a two-dimensional domain.

7.4.2.1 Helmholtz equation

We consider the elliptic partial differential equation:

$$\nabla^2 u + \lambda u = f \quad (7.3)$$

where ∇^2 is the Laplacian and λ is a real positive constant.

7.4.2.2 Input file

The input for this example is given in the example file `Helmholtz2D_modal.xml`

The geometry for this problem is a two-dimensional octagonal plane containing both triangles and quadrilaterals. Note that a mesh composite may only contain one type of element. Therefore, we define two composites for the domain, while the rest are used for enforcing boundary conditions.

```

1 <COMPOSITE>
2   <C ID="0"> Q[22-47] </C>
3   <C ID="1"> T[0-21] </C>
4   <C ID="2"> E[0-1] </C>
5   .
6   .
7   <C ID="10"> E[84,75,69,62,51,40,30,20,6] </C>
8 </COMPOSITE>
9
10 <DOMAIN> C[0-1] </DOMAIN>

```

For both the triangular and quadrilateral elements, we use the modified Legendre basis with 7 modes (maximum polynomial order is 6).

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="7" FIELDS="u" TYPE="MODIFIED" />
3   <E COMPOSITE="C[1]" NUMMODES="7" FIELDS="u" TYPE="MODIFIED" />
4 </EXPANSIONS>

```

Only one parameter is needed for this problem. In this example $\lambda = 1$ and the Continuous Galerkin Method is used as projection scheme to solve the Helmholtz equation, so we need to specify the following parameters and solver information.

```

1 <PARAMETERS>
2   <P> Lambda = 1 </P>
3 </PARAMETERS>
4
5 <SOLVERINFO>
6   <I PROPERTY="EQTYPE" VALUE="Helmholtz" />
7   <I PROPERTY="Projection" VALUE="Continuous" />
8 </SOLVERINFO>

```

All three basic boundary condition types have been used in this example: Dirichlet, Neumann and Robin boundary. The boundary regions are defined, each of which corresponds to one of the edge composites defined earlier. Each boundary region is then assigned an appropriate boundary condition.

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[2] </B>
3   .
4   .
5   <B ID="8"> C[10] </B>
6 </BOUNDARYREGIONS>
7
8 <BOUNDARYCONDITIONS>
9   <REGION REF="0">
10    <D VAR="u" VALUE="sin(PI*x)*sin(PI*y)" />
11  </REGION>
12  <REGION REF="1">
13    <R VAR="u" VALUE="sin(PI*x)*sin(PI*y)-PI*sin(PI*x)*cos(PI*y)"
14    PRIMCOEFF="1" />
15  </REGION>
16  <REGION REF="2">
17    <N VAR="u" VALUE="(5/sqrt(61))*PI*cos(PI*x)*sin(PI*y)-
18    (6/sqrt(61))*PI*sin(PI*x)*cos(PI*y)" />
19  </REGION>
20  .
21  .
22 </BOUNDARYCONDITIONS>

```

We know that for $f = -(\lambda + 2\pi^2)\sin(\pi x)\cos(\pi y)$, the exact solution of the two-dimensional Helmholtz equation is $u = \sin(\pi x)\cos(\pi y)$. These functions are defined specified to initialise the problem and verify the correct solution is obtained by evaluating the L_2 and L_{inf} errors.

```

1 <FUNCTION NAME="Forcing">
2   <E VAR="u" VALUE="-(Lambda + 2*PI*PI)*sin(PI*x)*sin(PI*y)" />
3 </FUNCTION>
4
5 <FUNCTION NAME="ExactSolution">
6   <E VAR="u" VALUE="sin(PI*x)*sin(PI*y)" />
7 </FUNCTION>

```

7.4.2.3 Running the code

```
ADRSolver Test_Helmholtz2D_modal.xml
```

This execution should print out a summary of input file, the L_2 and L_{inf} errors and the time spent on the calculation.

7.4.2.4 Post-processing

Simulation results are written in the file `Helmholtz2D_modal.fld`. We can choose to visualise the output in Gmsh

```
FieldConvert Helmholtz2D_modal.xml Helmholtz2D_modal.fld Helmholtz2D_modal.vtu
```

which generates the file `Helmholtz2D_modal.vtu` which can be visualised and is shown in Fig. 7.1




Figure 7.1 Solution of the 2D Helmholtz Problem.

7.4.3 Advection dominated mass transport in a pipe

The following example demonstrates the application of the ADRsolver for modelling advection dominated mass transport in a straight pipe. Such a transport regime is encountered frequently when modelling mass transport in arteries. This is because the diffusion coefficient of small blood borne molecules, for example oxygen or adenosine triphosphate, is very small $O(10^{-10})$.

7.4.3.1 Background

The governing equation for modelling mass transport is the unsteady advection diffusion equation:

$$\frac{\partial u}{\partial t} + v \nabla u + \epsilon \nabla^2 u = 0$$

For small diffusion coefficient, ϵ , the transport is dominated by advection and this leads to a very fine boundary layer adjacent to the surface which must be captured in order to get a realistic representation of the wall mass transfer processes. This creates problems not only from a meshing perspective, but also numerically where classical oscillations are observed in the solution due to under-resolution of the boundary layer.

The Graetz-Nusselt solution is an analytical solution of a developing mass (or heat) transfer boundary layer in a pipe. Previously this solution has been used as a benchmark for the accuracy of numerical methods to capture the fine boundary layer which develops for high Peclet number transport (the ratio of advection to diffusion). The solution

is derived based on the assumption that the velocity field within the mass transfer boundary layer is linear i.e. the Schmidt number (the relative thickness of the momentum to mass transfer boundary layer) is sufficiently large. The analytical solution for the non-dimensional mass transfer at the wall is given by:

$$Sh(z) = \frac{2^{4/3}(PeR/z)^{1/3}}{g^{1/3}\Gamma(4/3)},$$

where z is the streamwise coordinate, R the pipe radius, $\Gamma(4/3)$ an incomplete Gamma function and Pe the Peclet number given by:

$$Pe = \frac{2UR}{\epsilon}$$

In the following we will numerically solve mass transport in a pipe and compare the calculated mass transfer at the wall with the Graetz-Nusselt solution. The Peclet number of the transport regime under consideration is 750000, which is physiologically relevant.

7.4.3.2 Input file

The geometry under consideration is a pipe of radius, $R = 0.5$ and length $l = 0.5$




Figure 7.2 Pipe.

Since the mass transport boundary layer will be confined to a very small layer adjacent to the wall we do not need to mesh the interior region, hence the mesh consists of a layer of ten prismatic elements over a thickness of $0.036R$. The elements progressively grow over the thickness of domain.

In this example we utilise heterogeneous polynomial order, in which the polynomial order normal to the wall is higher so that we avoid unphysical oscillations, and hence the incorrect solution, in the mass transport boundary layer. To do this we specify explicitly the expansion type, points type and distribution in each direction as follows:

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]"
3     NUMMODES="3,5,3"
4     BASISTYPE="Modified_A,Modified_A,Modified_B"
5     NUMPOINTS="4,6,3"
6     POINTSTYPE="GaussLobattoLegendre,GaussLobattoLegendre,GaussRadauMAlpha1Beta0"
7     FIELDS="u" />
8 </EXPANSIONS>

```

The above represents a quadratic polynomial order in the azimuthal and streamwise direction and 4th order polynomial normal to the wall for a prismatic element.

We choose to use a continuous projection and an first-order implicit-explicit time-integration scheme. The `DiffusionAdvancement` and `AdvectionAdvancement` parameters specify how these terms are treated.

```

1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE"           VALUE="UnsteadyAdvectionDiffusion" />
3   <I PROPERTY="Projection"       VALUE="Continuous" />
4   <I PROPERTY="DiffusionAdvancement" VALUE="Implicit" />
5   <I PROPERTY="AdvectionAdvancement" VALUE="Explicit" />
6   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder1" />
7   <I PROPERTY="GlobalSysSoln"     VALUE="IterativeStaticCond" />
8 </SOLVERINFO>

```

We integrate for a total of 30 time units with a time-step of 0.0005, necessary to keep the simulation numerically stable.

```

1 <P> TimeStep = 0.0005 />
2 <P> FinalTime = 30 />
3 <P> NumSteps = FinalTime/TimeStep />

```

The value of the ϵ parameter is $\epsilon = 1/Pe$

```

1 <P> epsilon = 1.33333e-6 />

```

The analytical solution represents a developing mass transfer boundary layer in a pipe. In order to reproduce this numerically we assume that the inlet concentration is a uniform value and the outer wall concentration is zero; this will lead to the development of the mass transport boundary layer along the length of the pipe. Since we do not model explicitly the mass transfer in the interior region of the pipe we assume that the inner wall surface concentration is the same as the inlet concentration; this assumption is valid based on the large Peclet number meaning the concentration boundary layer is confined to the region in the immediate vicinity of the wall. The boundary conditions are specified as follows in the input file:

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[3] /> <!-- inlet -->
3   <B ID="1"> C[4] /> <!-- outlet -->
4   <B ID="2"> C[2] /> <!-- outer surface -->
5   <B ID="3"> C[5] /> <!-- inner surface -->

```

```

6 </BOUNDARYREGIONS>
7
8 <BOUNDARYCONDITIONS>
9   <REGION REF="0">
10     <D VAR="u" VALUE="1" />
11   </REGION>
12   <REGION REF="1">
13     <N VAR="u" VALUE="0" />
14   </REGION>
15   <REGION REF="2">
16     <D VAR="u" VALUE="0" />
17   </REGION>
18   <REGION REF="3">
19     <D VAR="u" VALUE="1" />
20   </REGION>
21 </BOUNDARYCONDITIONS>

```

The velocity field within the domain is fully developed pipe flow (Poiseuille flow), hence we can define this through an analytical function as follows:

```

1 <FUNCTION NAME="AdvectionVelocity">
2   <E VAR="Vx" VALUE="0" />
3   <E VAR="Vy" VALUE="0" />
4   <E VAR="Vz" VALUE="2.0*(1-(x*x+y*y)/0.25)" />
5 </FUNCTION>

```

We assume that the initial domain concentration is uniform everywhere and the same as the inlet. This is defined by,

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="u" VALUE="1" />
3 </FUNCTION>

```

7.4.3.3 Results

To compare with the analytical expression we numerically calculate the concentration gradient at the surface of the pipe. This is then plotted against the analytical solution by extracting the solution along a line in the streamwise direction, as shown in Fig. 7.3.

7.4.4 Unsteady reaction-diffusion systems

Reaction-diffusion systems are prevalent in a number of areas relating to the modelling of various physical phenomena, and are particularly prevalent in the study of chemical interactions and pattern formation. The ADRSolver supports the solution of a single-variable system

$$\frac{\partial u}{\partial t} = \epsilon \nabla^2 u + R(u)$$

where the diffusion coefficient ϵ and reaction term $R(u)$ are defined using the session file.




Figure 7.3 Concentration gradient at the surface of the pipe.

7.4.4.1 Numerical restrictions

The reaction-diffusion system is only supported in a selected configuration, which is mostly defined inside the `SOLVERINFO` block:

- use of a continuous Galerkin discretisation;
- use an implicit-explicit (IMEX) timestepping scheme, such as `IMEXOrder3`;

This naturally leads to the following `SOLVERINFO` configuration:

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE"           VALUE="UnsteadyReactionDiffusion" />
3   <I PROPERTY="Projection"       VALUE="Continuous" />
4   <I PROPERTY="DiffusionAdvancement" VALUE="Implicit" />
5   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder3" />
6 </SOLVERINFO>
```

Further to this, the reaction term $R(u)$ is imposed by the definition of a body forcing function. For example, the reaction term $R(u) = 0.1u$ may be defined using the function:

```
1 <!-- Body force to enforce reaction term -->
2 <FUNCTION NAME="BodyForce">
3   <E VAR="u" EVARS="u" VALUE="0.1*u" />
4 </FUNCTION>
```

Note in particular the use of the `EVARS` (equation variables) attribute, which permits the definition of this function in terms of the scalar variable u . This function should be used together with an appropriate `FORCING` block (as described in section 3.5):

```
1 <FORCING>
2   <FORCE TYPE="Body">
3     <BODYFORCE> BodyForce </BODYFORCE>
```

```
4 </FORCE>  
5 </FORCING>
```

An example of a simple unsteady reaction-diffusion problem can be found in the **Tests** directory in the session file `ReactionDiffusion2D.xml`.

Cardiac Electrophysiology Solver

8.1 Synopsis

The CardiacEPSolver is used to model the electrophysiology of cardiac tissue, specifically using the monodomain or bidomain model. These models are continuum models and represent an average of the electrical activity over many cells. The system is a reaction-diffusion system, with the reaction term modeling the flow of current in and out of the cells using a separate set of ODEs.

8.1.1 Bidomain Model

The Bidomain model is given by the following PDEs,

$$\begin{aligned} g_{ix} \frac{\partial^2 V_i}{\partial x^2} + g_{iy} \frac{\partial^2 V_i}{\partial y^2} &= \chi \left[C_m \frac{\partial(V_i - V_e)}{\partial t} + G_m(V_i - V_e) \right] \\ g_{ex} \frac{\partial^2 V_e}{\partial x^2} + g_{ey} \frac{\partial^2 V_e}{\partial y^2} &= -\chi \left[C_m \frac{\partial(V_i - V_e)}{\partial t} + G_m(V_i - V_e) \right]. \end{aligned}$$

However, when solving numerically, one often rewrites these equations in terms of the transmembrane potential and extracellular potential,

$$\begin{aligned} \chi \left[C_m \frac{\partial V_m}{\partial t} + J_{ion} \right] &= g_{ex} \frac{\partial^2 V_e}{\partial x^2} + g_{ey} \frac{\partial^2 V_e}{\partial y^2} \\ (g_{ix} + g_{ex}) \frac{\partial^2 V_e}{\partial x^2} + (g_{iy} + g_{ey}) \frac{\partial^2 V_e}{\partial y^2} &= -g_{ix} \frac{\partial^2 V_m}{\partial x^2} - g_{iy} \frac{\partial^2 V_m}{\partial y^2} \end{aligned}$$

8.1.2 Monodomain Model

In the case where the intracellular and extracellular conductivities are proportional, that is $g_{ix} = k g_{ex}$ for some k , then the above two PDEs can be reduced to a single PDE:

$$\chi \left[C_m \frac{\partial V_m}{\partial t} + J_{ion} \right] = \nabla \cdot (\sigma \nabla V_m)$$

8.1.3 Cell Models

The action potential of a cardiac cell can be modelled at either a biophysical level of detail, including a number of transmembrane currents, or as a phenomenological model, to reproduce the features of the action potential, with fewer variables. Each cell model will include a unique system of ODEs to represent the gating variables of that model.

A number of ionic cell models are currently supported by the solver including:

- Courtemanche, Ramirez, Nattel, 1998
- Luo, Rudy, 1991
- ten Tusscher, Panfilov, 2006 (epicardial, endocardial and mid-myocardial variants)

Phenomenological cell models are also supported:

- Aliev-Panfilov
- Fitzhugh-Nagumo

It is important to ensure that the units of the voltage and currents from the cell model are consistent with the units expected by the tissue level solver (monodomain/bidomain). We will show as an example the Courtemanche, Ramirez, Nattel, 1998 human atrial model.

The monodomain equation:

$$\chi \left[C_m \frac{\partial V_m}{\partial t} + J_{ion} \right] = \nabla \cdot (\sigma \nabla V_m)$$

8.2 Usage

```
CardiacEPSolver session.xml
```

8.3 Session file configuration

8.3.1 Solver Info

- `Eqtype` Specifies the PDE system to solve. The following values are supported:
 - `Monodomain`: solve the monodomain equation.
 - `BidomainRoth`: solve the bidomain equations using the Roth formulation.
- `CellModel` Specifies the cell model to use. Available cell models are

Value	Description	No. of Var.	Ref.
<code>AlievPanfilov</code>	Phenomological	1	[2]
<code>CourtemancheRamirezNattel98</code>	Human atrial	20	[26]
<code>FitzHughNagumo</code>			
<code>Fox02</code>			
<code>LuoRudy91</code>	Mammalian ventricular	7	[25]
<code>PanditGilesDemir03</code>			
<code>TenTusscher06</code>	Human ventricular	18	[42]
<code>Winslow99</code>			

- `Projection` Specifies the Galerkin projection type to use. Only `Continuous` has been extensively tested.
- `TimeIntegrationMethod` Specifies the time integration scheme to use for advancing the PDE system. This must be an IMEX scheme. Suitable choices are: `IMEXOrder1`, `IMEXOrder2`, `IMEXOrder3`, `IMEXdirk_3_4_3`. The cell model state variables are time advanced using Forward Euler for the ion concentrations, and Rush-Larsen for the cell model gating variables.
- `DiffusionAdvancement` Specifies whether the diffusion is handled implicitly or explicitly in the time integration scheme. The current code only supports `Implicit` integration of the diffusion term. The cell model is always integrated explicitly.

8.3.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file. Example values are taken from [12].

- `Chi` sets the surface-to-volume ratio (Units: mm^{-1}).
Example: $\chi = 140\text{mm}^{-1}$
- `Cm` sets the specific membrane capacitance (Units: $\mu\text{F mm}^{-2}$).
Example: $C_m = 0.01\mu\text{F mm}^{-2}$
- `Substeps` sets the number of substeps taken in time integrating the cell model for each PDE timestep.
Example: 4
- `d_min`, `d_max`, `o_min`, `o_max` specifies a bijective map to assign conductivity values σ to intensity values μ when using the `IsotropicConductivity` function. The intensity map is first thresholded to the range $[d_{\min}, d_{\max}]$ and then the conductivity is calculated as

$$\sigma = \frac{o_{\max} - o_{\min}}{d_{\max} - d_{\min}}(1 - \mu) + o_{\min}$$

8.3.3 Functions

The following functions can be specified inside the `CONDITIONS` section of the session file. If both are specified, the effect is multiplicative. Example values are taken from [12].

- `IsotropicConductivity` specifies the conductivity σ of the tissue.
Example: $\sigma = 0.13341 \text{ mS mm}^{-1}$, based on $\sigma = \frac{\sigma_i \sigma_e}{\sigma_i + \sigma_e}$, $\sigma_i = 0.17$, $\sigma_e = 0.62 \text{ mS mm}^{-1}$

The variable name to use is `intensity` since the conductivity may be derived from late-Gadolinium enhanced MRA imaging. Example specifications are

```
1 <E VAR="intensity" VALUE="0.13341" />
2 <F VAR="intensity" FILE="scarmap.con" />
```

where `scarmap.con` is a Nektar++ field file containing a variable `intensity` describing the conductivity across the domain.

- `AnisotropicConductivity` specifies the conductivity σ of the tissue.

8.3.4 Filters

The following filters are supported exclusively for the cardiac EP solver. Further filters from section 3.4 are also available for this solver.

- `Benchmark` (section 3.4.3)
- `CellHistoryPoints` (section 3.4.4)
- `CheckpointCellModel` (section 3.4.5)
- `Electrogram` (section 3.4.7)

8.3.5 Stimuli

Electrophysiological propagation is initiated through the stimulus current I_{ion} . The `STIMULI` section describes one or more regions of stimulus and the time-dependent protocol with which they are applied.

```
1 <STIMULI>
2 ...
3 </STIMULI>
```

A number of stimulus types are available

8.3.5.1 Stimulus types

- `StimulusRect` stimulates a cuboid-shaped region of the domain, specified by two coordinates (x_1, y_1, z_1) and (x_2, y_2, z_2) . An additional parameter specifies the

"smoothness" of the boundaries of the region; higher values produce a sharper boundary. Finally, the maximum strength of the stimulus current is specified in $\mu\text{A}/\text{mm}^3$

```

1 <STIMULUS TYPE="StimulusRect" ID="0">
2   <p_x1> -15.24 </p_x1>
3   <p_y1>  14.02 </p_y1>
4   <p_z1>   6.87 </p_z1>
5   <p_x2>  12.23 </p_x2>
6   <p_y2>  16.56 </p_y2>
7   <p_z2>   8.88 </p_z2>
8   <p_is> 100.00 </p_is>
9   <p_strength> 50.0 </p_strength>
10 </STIMULUS>

```

- **StimulusCirc** stimulates a spherical region of the domain, as specified by a centre and radius. The smoothness and strength parameters are also specified as for 'StimulusRect'.

```

1 <STIMULUS TYPE="StimulusCirc" ID="0">
2   <p_x1> -15.24 </p_x1>
3   <p_y1>  14.02 </p_y1>
4   <p_z1>   6.87 </p_z1>
5   <p_r1>  12.23 </p_r1>
6   <p_is> 100.00 </p_is>
7   <p_strength> 50.0 </p_strength>
8 </STIMULUS>

```

8.3.5.2 Protocols

A protocol specifies the time-dependent function indicating the strength of the stimulus and one such **PROTOCOL** section should be included within each **STIMULUS**. This can be expressed as one of:

- **ProtocolSingle** a single stimulus is applied at a given start time and for a given duration

```

1 <PROTOCOL TYPE="ProtocolSingle">
2   <START> 0.0 </START>
3   <DURATION> 2.0 </DURATION>
4 </PROTOCOL>

```

- **ProtocolS1** a train of pulses of fixed duration applied at a given start time and with a given cycle length.

```

1 <PROTOCOL TYPE="ProtocolS1">
2   <START> 0.0 </START>
3   <DURATION> 2.0 </DURATION>
4   <S1CYCLELENGTH> 300.0 </S1CYCLELENGTH>
5   <NUM_S1> 5 </NUM_S1>
6 </PROTOCOL>

```

- `ProtocolS1S2` same as ‘ProtocolS1’ except with an additional single pulse applied at a different cycle length at the end of the train of S1 pulses.

```
1 <PROTOCOL TYPE="ProtocolS1S2">
2   <START>    0.0   </START>
3   <DURATION> 2.0   </DURATION>
4   <S1CYCLELENGTH> 300.0 </S1CYCLELENGTH>
5   <NUM_S1>    5     </NUM_S1>
6   <S2CYCLELENGTH> 100.0 </S2CYCLELENGTH>
7 </PROTOCOL>
```

Compressible Flow Solver

9.1 Synopsis

The CompressibleFlowSolver allows us to solve the unsteady compressible Euler and Navier-Stokes equations for 1D/2D/3D problems using a discontinuous representation of the variables. In the following we describe both the compressible Euler and the Navier-Stokes equations.

9.1.1 Euler equations

The Euler equations can be expressed as a hyperbolic conservation law in the form

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}_i}{\partial x} + \frac{\partial \mathbf{g}_i}{\partial y} + \frac{\partial \mathbf{h}_i}{\partial z} = 0, \quad (9.1)$$

where \mathbf{q} is the vector of the conserved variables, $\mathbf{f}_i = \mathbf{f}_i(\mathbf{q})$, $\mathbf{g}_i = \mathbf{g}_i(\mathbf{q})$ and $\mathbf{h}_i = \mathbf{h}_i(\mathbf{q})$ are the vectors of the inviscid fluxes

$$\mathbf{q} = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ E \end{pmatrix}, \quad \mathbf{f}_i = \begin{pmatrix} \rho u \\ p + \rho u^2 \\ \rho uv \\ \rho uw \\ u(E + p) \end{pmatrix}, \quad \mathbf{g}_i = \begin{pmatrix} \rho v \\ \rho uv \\ p + \rho v^2 \\ \rho vw \\ v(E + p) \end{pmatrix}, \quad \mathbf{h}_i = \begin{pmatrix} \rho w \\ \rho uw \\ \rho vw \\ p + \rho w^2 \\ w(E + p) \end{pmatrix}, \quad (9.2)$$

where ρ is the density, u , v and w are the velocity components in x , y and z directions, p is the pressure and E is the total energy. In this work we considered a perfect gas law for which the pressure is related to the total energy by the following expression

$$E = \frac{p}{\gamma - 1} + \frac{1}{2}\rho(u^2 + v^2 + w^2), \quad (9.3)$$

where γ is the ratio of specific heats.

9.1.2 Compressible Navier-Stokes equations

The Navier-Stokes equations include the effects of fluid viscosity and heat conduction and are consequently composed by an inviscid and a viscous flux. They depend not only on the conserved variables but also, indirectly, on their gradient. The second order partial differential equations for the three-dimensional case can be written as:

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}}{\partial x} + \frac{\partial \mathbf{g}}{\partial y} + \frac{\partial \mathbf{h}}{\partial z} = 0, \quad (9.4)$$

where \mathbf{q} is the vector of the conserved variables, $\mathbf{f} = \mathbf{f}(\mathbf{q}, \nabla(\mathbf{q}))$, $\mathbf{g} = \mathbf{g}(\mathbf{q}, \nabla(\mathbf{q}))$ and $\mathbf{h} = \mathbf{h}(\mathbf{q}, \nabla(\mathbf{q}))$ are the vectors of the fluxes which can also be written as:

$$\mathbf{f} = \mathbf{f}_i - \mathbf{f}_v, \mathbf{g} = \mathbf{g}_i - \mathbf{g}_v, \mathbf{h} = \mathbf{h}_i - \mathbf{h}_v, \quad (9.5)$$

where \mathbf{f}_i , \mathbf{g}_i and \mathbf{h}_i are the inviscid fluxes of Eq. (9.2) and \mathbf{f}_v , \mathbf{g}_v and \mathbf{h}_v are the viscous fluxes which take the following form:

$$\mathbf{f}_v = \begin{pmatrix} 0 \\ \tau_{xx} \\ \tau_{xy} \\ \tau_{yx} \\ \tau_{zx} \\ u\tau_{xx} + v\tau_{yx} + w\tau_{zx} + kT_x \end{pmatrix}, \quad \mathbf{g}_v = \begin{pmatrix} 0 \\ \tau_{xy} \\ \tau_{yy} \\ \tau_{yz} \\ \tau_{zy} \\ u\tau_{xy} + v\tau_{yy} + w\tau_{zy} + kT_y \end{pmatrix}, \quad (9.6)$$

$$\mathbf{h}_v = \begin{pmatrix} 0 \\ \tau_{xz} \\ \tau_{yz} \\ \tau_{zz} \\ u\tau_{xz} + v\tau_{yz} + w\tau_{zz} + kT_z \end{pmatrix},$$

where τ_{xx} , τ_{xy} , τ_{xz} , τ_{yx} , τ_{yx} , τ_{yy} , τ_{yz} , τ_{zx} , τ_{zy} and τ_{zz} are the components of the stress tensor¹

$$\begin{aligned} \tau_{xx} &= 2\mu \left(u_x - \frac{u_x + v_y + w_z}{3} \right), & \tau_{yy} &= 2\mu \left(v_y - \frac{u_x + v_y + w_z}{3} \right), \\ \tau_{zz} &= 2\mu \left(w_z - \frac{u_x + v_y + w_z}{3} \right), & \tau_{xy} &= \tau_{yx} = \mu(v_x + u_y), \\ \tau_{yz} &= \tau_{zy} = \mu(w_y + v_z), & \tau_{zx} &= \tau_{xz} = \mu(u_z + w_x). \end{aligned} \quad (9.7)$$

where μ is the dynamic viscosity calculated using the Sutherland's law and k is the thermal conductivity.

9.1.3 Numerical discretisation

In Nektar++ the spatial discretisation of the Euler and of the Navier-Stokes equations is projected in the polynomial space via a discontinuous projection. Specifically we make use either of the discontinuous Galerkin (DG) method or the Flux Reconstruction (FR)

¹Note that we use Stokes hypothesis $\lambda = -2/3$.

approach. In both the approaches the physical domain Ω is divided into a mesh of N non-overlapping elements Ω_e and the solution is allowed to be discontinuous at the boundary between two adjacent elements. Since the Euler as well as the Navier-Stokes equations are defined locally (on each element of the computational domain), it is necessary to define a term to couple the elements of the spatial discretisation in order to allow information to propagate across the domain. This term, called numerical interface flux, naturally arises from the discontinuous Galerkin formulation as well as from the Flux Reconstruction approach.

For the advection term it is common to solve a Riemann problem at each interface of the computational domain through exact or approximated Riemann solvers. In Nektar++ there are different Riemann solvers, one exact and nine approximated. The exact Riemann solver applies an iterative procedure to satisfy conservation of mass, momentum and energy and the equation of state. The left and right states are connected either with the unknown variables through the Rankine-Hugoniot relations, in the case of shock, or the isentropic characteristic equations, in the case of rarefaction waves. Across the contact surface, conditions of continuity of pressure and velocity are employed. Using these equations the system can be reduced to a non-linear algebraic equation in one unknown (the velocity in the intermediate state) that is solved iteratively using a Newton method. Since the exact Riemann solver gives a solution with an order of accuracy that is related to the residual in the Newton method, the accuracy of the method may come at high computational cost. The approximated Riemann solvers are simplifications of the exact solver.

Concerning the diffusion term, the coupling between the elements is achieved by using a local discontinuous Galerkin (LDG) approach as well as five different FR diffusion terms.

The boundary conditions are also implemented by exploiting the numerical interface fluxes just mentioned. For a more detailed description of the above the interested reader can refer to [9] and [29].

9.2 Usage

```
CompressibleFlowSolver session.xml
```

9.3 Session file configuration

In the following we describe the session file configuration. Specifically we consider the sections under the tag `<CONDITIONS>` in the session (.xml) file.

Parameters

Under this section it is possible to set the parameters of the simulation.

```
1 <PARAMETERS>
2   <P> TimeStep           = 0.0000001   </P>
```

```

3 <P> FinTime           = 1.0                </P>
4 <P> NumSteps          = FinTime/TimeStep    </P>
5 <P> IO_CheckSteps     = 5000                </P>
6 <P> IO_InfoSteps      = 1                   </P>
7 <P> Gamma             = 1.4                 </P>
8 <P> pInf              = 101325              </P>
9 <P> rhoInf            = 1.225               </P>
10 <P> TInf              = pInf/(287.058*rhoInf) </P>
11 <P> Twall             = pInf/(287.058*rhoInf)+15.0 </P>
12 <P> uInf              = 147.4              </P>
13 <P> vInf              = 0.0                </P>
14 <P> wInf              = 0.0                </P>
15 <P> mu                = 1e-5               </P>
16 <P> Pr                = 0.72               </P>
17 <P> thermalConductivity = 0.02             </P>
18 </PARAMETERS>

```

- `TimeStep` is the time-step we want to use;
- `FinTime` is the final physical time at which we want our simulation to stop;
- `NumSteps` is the equivalent of `FinTime` but instead of specifying the physical final time we specify the number of time-steps;
- `IO_CheckSteps` sets the number of steps between successive checkpoint files;
- `IO_InfoSteps` sets the number of steps between successive info stats are printed to screen;
- `Gamma` ratio of the specific heats. Default value = 1.4;
- `pInf` farfield pressure (i.e. p_∞). Default value = 101325 *Pa*;
- `rhoInf` farfield density (i.e. ρ_∞). Default value = 1.225 *Kg/m³*;
- `TInf` farfield temperature (i.e. T_∞). Default value = 288.15 *K*;
- `Twall` temperature at the wall when isothermal boundary conditions are employed (i.e. T_w). Default value = 300.15 *K*;
- `uint` farfield *X*-component of the velocity (i.e. u_∞). Default value = 0.1 *m/s*;
- `vInf` farfield *Y*-component of the velocity (i.e. v_∞). Default value = 0.0 *m/s*;
- `wInf` farfield *Z*-component of the velocity (i.e. w_∞). Default value = 0.0 *m/s*;
- `mu` dynamic viscosity (i.e. μ_∞). Default value = 1.78e-05 *Pa·s*;
- `Pr` Prandtl number. Default value = 0.72;

- `thermalConductivity` thermal conductivity (i.e. κ_∞). This can be set as an alternative to `Pr`, in which case the Prandtl number is calculated from κ_∞ (it is only possible to set one of them). By default, this is obtained from the Prandtl number;

Solver info

Under this section it is possible to set the solver information.

```

1 <SOLVERINFO>
2 <I PROPERTY="EQType"           VALUE="NavierStokesCFE"    />
3 <I PROPERTY="Projection"       VALUE="DisContinuous"  />
4 <I PROPERTY="AdvectionType"    VALUE="WeakDG"         />
5 <I PROPERTY="DiffusionType"    VALUE="LDGNS"          />
6 <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4" />
7 <I PROPERTY="UpwindType"       VALUE="ExactToro"      />
8 <I PROPERTY="ProblemType"      VALUE="General"        />
9 <I PROPERTY="ViscosityType"    VALUE="Constant"       />
10 <I PROPERTY="EquationOfState"  VALUE="IdealGas"       />
11 </SOLVERINFO>

```

- `EQType` is the tag which specify the equations we want solve:
 - `NavierStokesCFE` (Compressible Navier-Stokes equations);
 - `EulerCFE` (Compressible Euler equations).
 - `IsentropicVortex` (Isentropic vortex test case).
 - `RinglebFlow` (Ringleb flow test case).
- `Projection` is the type of projection we want to use:
 - `DisContinuous`.
Note that the Continuous projection is not supported in the Compressible Flow Solver.
- `AdvectionType` is the advection operator we want to use.
 - `WeakDG` (classical DG in weak form);
 - `FRDG` (Flux-Reconstruction recovering nodal DG scheme);
 - `FRSD` (Flux-Reconstruction recovering a spectral difference (SD) scheme);
 - `FRHU` (Flux-Reconstruction recovering Huynh (G2) scheme);
 - `FRcmin` (Flux-Reconstruction with $c = c_{min}$);
 - `FRcinf` (Flux-Reconstruction with $c = \infty$).

Note that only `WeakDG` is fully supported, the other operators work only with quadrilateral elements ($2D$ or $2.5D$).

- `DiffusionType` is the diffusion operator we want to use for the Navier-Stokes equations:
 - `LDGNS` (LDG with primitive variables. The penalty term is inversely proportional to the element size, proportional to the local viscosity for the momentum equations and to the thermal conductivity for the energy equation, and proportional to an optional parameter `LDGNSc11` which is by default set to one; proportionality to polynomial order can be manually imposed by setting the parameter `LDGNSc11` equal to p^2);
 - `LFRDGNS` (Flux-Reconstruction recovering nodal DG scheme);
 - `LFRSDNS` (Flux-Reconstruction recovering a spectral difference (SD) scheme);
 - `LFRHUNS` (Flux-Reconstruction recovering Huynh (G2) scheme);
 - `LFRcminNS` (Flux-Reconstruction with $c = c_{min}$);
 - `LFRcinfNS` (Flux-Reconstruction with $c = \infty$).

Note that only `LDGNS` is fully supported, the other operators work only with quadrilateral elements ($2D$ or $2.5D$).

- `TimeIntegrationMethod` is the time-integration scheme we want to use. Note that only an explicit discretisation is supported:
 - `ForwardEuler`;
 - `RungeKutta2_SSP`;
 - `RungeKutta3_SSP`;
 - `ClassicalRungeKutta4`.
- `UpwindType` is the numerical interface flux (i.e. Riemann solver) we want to use for the advection operator:
 - `AUSM0`;
 - `AUSM1`;
 - `AUSM2`;
 - `AUSM3`;
 - `Average`;
 - `ExactToro`;
 - `HLL`;
 - `HLLC`;
 - `LaxFriedrichs`;
 - `Roe`.
- `ViscosityType` is the viscosity type we want to use:

- `Constant` (Constant viscosity);
- `Variable` (Variable viscosity through the Sutherland's law.);
- `EquationOfState` allows selecting an equation of state for accounting for non-ideal gas behaviour:
 - `IdealGas` (default option);
 - `VanDerWaals` (requires additional parameters `Tcrit` and `Pcrit`);
 - `RedlichKwong` (requires additional parameters `Tcrit` and `Pcrit`);
 - `PengRobinson` (requires additional parameters `Tcrit`, `Pcrit` and `AcentricFactor`);

Boundary conditions

In this section we can specify the boundary conditions for our problem. First we need to define the variables under the section `VARIABLES`. For a 1D problem we have:

```
1 <VARIABLES>
2   <V ID="0"> rho </V>
3   <V ID="1"> rhou </V>
4   <V ID="4"> E </V>
5 </VARIABLES>
```

For a 2D problem we have

```
1 <VARIABLES>
2   <V ID="0"> rho </V>
3   <V ID="1"> rhou </V>
4   <V ID="2"> rhov </V>
5   <V ID="4"> E </V>
6 </VARIABLES>
```

For a 3D problem we have:

```
1 <VARIABLES>
2   <V ID="0"> rho </V>
3   <V ID="1"> rhou </V>
4   <V ID="2"> rhov </V>
5   <V ID="3"> rhow </V>
6   <V ID="4"> E </V>
7 </VARIABLES>
```

After having defined the variables depending on the dimensions of the problem we want to solve it is necessary to specify the boundary regions on which we want to define the boundary conditions:

```
1 <BOUNDARYREGIONS>
2   <B ID="0"> C[100] </B>
3 </BOUNDARYREGIONS>
```

Finally we can specify the boundary conditions on the regions specified under `BOUNDARYREGIONS`. In the following some examples for a 2D problem:

- Slip wall boundary conditions:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="rho" USERDEFINEDTYPE="Wall" VALUE="0" />
4     <D VAR="rhov" USERDEFINEDTYPE="Wall" VALUE="0" />
5     <D VAR="rho" USERDEFINEDTYPE="Wall" VALUE="0" />
6     <D VAR="E" USERDEFINEDTYPE="Wall" VALUE="0" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

- No-slip wall boundary conditions:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="rho" USERDEFINEDTYPE="WallViscous" VALUE="0" />
4     <D VAR="rhov" USERDEFINEDTYPE="WallViscous" VALUE="0" />
5     <D VAR="rho" USERDEFINEDTYPE="WallViscous" VALUE="0" />
6     <D VAR="E" USERDEFINEDTYPE="WallViscous" VALUE="0" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

- Adiabatic wall boundary conditions:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="rho" USERDEFINEDTYPE="WallAdiabatic" VALUE="0" />
4     <D VAR="rhov" USERDEFINEDTYPE="WallAdiabatic" VALUE="0" />
5     <D VAR="rho" USERDEFINEDTYPE="WallAdiabatic" VALUE="0" />
6     <D VAR="E" USERDEFINEDTYPE="WallAdiabatic" VALUE="0" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

- Farfield boundary conditions (including inviscid characteristic boundary conditions):

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="rho" VALUE="rhoInf" />
4     <D VAR="rhov" VALUE="rhoInf*uInf" />
5     <D VAR="rho" VALUE="rhoInf*vInf" />
6     <D VAR="E"
7     VALUE="pInf/(Gamma-1)+0.5*rhoInf*(uInf*uInf+vInf*vInf+wInf*wInf)" />
8   </REGION>
9 </BOUNDARYCONDITIONS>

```

- Pressure outflow boundary conditions:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="rho" USERDEFINEDTYPE="PressureOutflow" VALUE="0" />
4     <D VAR="rhov" USERDEFINEDTYPE="PressureOutflow" VALUE="0" />
5     <D VAR="rho" USERDEFINEDTYPE="PressureOutflow" VALUE="0" />
6     <D VAR="E" USERDEFINEDTYPE="PressureOutflow" VALUE="pOut" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

where pOut is the target static pressure at the boundary.

Initial conditions and exact solution

Under the two following sections it is possible to define the initial conditions and the exact solution (if existent).

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="rho"      VALUE="rhoInf"/>
3   <E VAR="rho_u"    VALUE="rhoInf*uInf" />
4   <E VAR="rho_v"    VALUE="rhoInf*vInf" />
5   <E VAR="E"
6   VALUE="pInf/(Gamma-1)+0.5*rhoInf*(uInf*uInf+vInf*vInf+wInf*wInf)"/>
7 </FUNCTION>
8
9 <FUNCTION NAME="ExactSolution">
10  <E VAR="rho"      VALUE="rhoInf" />
11  <E VAR="rho_u"    VALUE="rhoInf*uInf" />
12  <E VAR="rho_v"    VALUE="rhoInf*vInf" />
13  <E VAR="E"
14  VALUE="pInf/(Gamma-1)+0.5*rhoInf*(uInf*uInf+vInf*vInf+wInf*wInf)"/>
15 </FUNCTION>

```

9.4 Examples

9.4.1 Shock capturing

Compressible flows can be characterised by abrupt changes in density within the flow domain often referred to as shocks. These discontinuities can lead to numerical instabilities (Gibbs phenomena). This problem is prevented by locally adding a diffusion term to the equations to damp the numerical oscillations.

9.4.1.1 Non-smooth artificial viscosity model

For the non-smooth artificial viscosity model the added artificial viscosity is constant in each element and discontinuous between the elements. The Euler system is augmented by an added Laplacian term on right hand side of equation 9.1 [35]. The diffusivity of the system is controlled by a variable viscosity coefficient ε . For consistency ε is proportional to the element size and inversely proportional to the polynomial order. Finally, from physical considerations ε needs to be proportional to the maximum characteristic speed of the problem. The final form of the artificial viscosity is

$$\varepsilon = \varepsilon_0 \frac{h}{p} \lambda_{max} S, \quad (9.8)$$

where S is a sensor.

As shock sensor, a modal resolution-based indicator is used

$$s_e = \log_{10} \left(\frac{\langle q - \tilde{q}, q - \tilde{q} \rangle}{\langle q, q \rangle} \right), \quad (9.9)$$

where $\langle \cdot, \cdot \rangle$ represents a L^2 inner product, q and \tilde{q} are the full and truncated expansions of a state variable (in our case density)

$$q(x) = \sum_{i=1}^{N(P)} \hat{q}_i \phi_i, \quad \tilde{q}(x) = \sum_{i=1}^{N(P-1)} \hat{q}_i \phi_i, \quad (9.10)$$

then the constant element-wise sensor is computed as follows

$$S_\varepsilon = \begin{cases} 0 & \text{if } s_e < s_0 - \kappa \\ \frac{1}{2} \left(1 + \sin \frac{\pi(s_e - s_0)}{2\kappa} \right) & \text{if } |s_e - s_0| \leq \kappa \\ 1 & \text{if } s_e > s_0 + \kappa \end{cases}, \quad (9.11)$$

where $s_0 = s_\kappa - 4.25 \log_{10}(p)$.

To enable the non-smooth viscosity model, the following line has to be added to the `SOLVERINFO` section:

```
1 <SOLVERINFO>
2   <I PROPERTY="ShockCaptureType"      VALUE="NonSmooth" />
3 </SOLVERINFO>
```

The diffusivity and the sensor can be controlled by the following parameters:

```
1 <PARAMETERS>
2 <P> Skappa           = -1.3           </P>
3 <P> Kappa            = 0.2            </P>
4 <P> mu0              = 1.0            </P>
5 </PARAMETERS>
```




Figure 9.1 (a) Steady state solution for $M = 0.8$ flow at $\alpha = 1.25^\circ$ past a NACA 0012 profile, (b) Artificial viscosity (ε) distribution

9.4.2 Variable polynomial order

A sensor based p -adaptive algorithm is implemented to optimise the computational cost and accuracy. The DG scheme allows one to use different polynomial orders since the fluxes over the elements are determined using a Riemann solver and there is now further coupling between the elements. Furthermore, the initial p -adaptive algorithm uses the same sensor as the shock capturing algorithm to identify the smoothness of the local solution so it is rather straightforward to implement both algorithms at the same time.

The polynomial order in each element can be adjusted based on the sensor value that is obtained. Initially, a converged solution is obtained after which the sensor in each element is calculated. Based on the determined sensor value and the pre-defined sensor thresholds, it is decided to increase, decrease or maintain the degree of the polynomial approximation in each element and a new converged solution is obtained.

$$p_e = \begin{cases} p_e - 1 & \text{if } s_e > s_{ds} \\ p_e + 1 & \text{if } s_{sm} < s_e < s_{ds} \\ p_e & \text{if } s_{fl} < s_e < s_{sm} \\ p_e - 1 & \text{if } s_e < s_{fl} \end{cases} \quad (9.12)$$

For now, the threshold values s_e , s_{ds} , s_{sm} and s_{fl} are determined empirically by looking at the sensor distribution in the domain. Once these values are set, two .txt files are outputted, one that has the composites called VariablePComposites.txt and one with the expansions called VariablePExpansions.txt. These values have to be copied into a new .xml file to create the adapted mesh.

9.4.3 De-Aliasing Techniques

Aliasing effects, arising as a consequence of the nonlinearity of the underlying problem, need to be addressed to stabilise the simulations. Aliasing appears when nonlinear quantities are calculated at an insufficient number of quadrature points. We can identify two types of nonlinearities:

- PDE nonlinearities, related to the nonlinear and quasi-linear fluxes.
- Geometrical nonlinearities, related to the deformed/curved meshes.

We consider two de-aliasing strategies based on the concept of consistent integration:

- Local dealiasing: It only targets the PDE-aliasing sources, applying a consistent integration of them locally.
- Global dealiasing: It targets both the PDE and the geometrical-aliasing sources. It requires a richer quadrature order to consistently integrate the nonlinear fluxes, the geometric factors, the mass matrix and the boundary term.

Since Nektar++ tackles separately the PDE and geometric aliasing during the projection and solution of the equations, to consistently integrate all the nonlinearities in the compressible NavierStokes equations, the quadrature points should be selected based on the maximum order of the nonlinearities:

$$Q_{min} = P_{exp} + \frac{\max(2P_{exp}, P_{geom})}{2} + \frac{3}{2} \quad (9.13)$$

where Q_{min} is the minimum required number of quadrature points to exactly integrate the highest-degree of nonlinearity, P_{exp} being the order of the polynomial expansion and P_{geom} being the geometric order of the mesh. Bear in mind that we are using a discontinuous discretisation, meaning that aliasing effect are not fully controlled, since the boundary terms introduce non-polynomial functions into the problem.

To enable the global de-aliasing technique, modify the number of quadrature points by:

```
1 <E COMPOSITE="[101]"
2 BASISTYPE="Modified_A,Modified_A"
3 NUMMODES="7,7"
4 POINTSTYPE="GaussLobattoLegendre,GaussLobattoLegendre"
5 NUMPOINTS="14,14"
6 FIELDS="rho,rhou,rhov,E"
7 />
```

where `NUMMODES` corresponds to $P+1$, where P is the order of the polynomial used to approximate the solution. `NUMPOINTS` specifies the number of quadrature points.

Dummy Solver

10.1 Synopsis

The Dummy solver does not solve any equation systems but only serves to exchange fields with other solvers and applications. It is intended for demonstrating and testing the coupling implementations only.

Incompressible Navier-Stokes Solver

11.1 Synopsis

A useful tool implemented in Nektar++ is the incompressible Navier Stokes solver that allows one to solve the governing equation for viscous Newtonians fluids governed by:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f} \quad (11.1a)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (11.1b)$$

where V is the velocity, p is the specific pressure (including density) and ν the kinematic viscosity.

11.1.1 Velocity Correction Scheme

The first approach uses a splitting/projection method where the velocity system and the pressure are typically decoupled. Splitting schemes are typically favoured for their numerical efficiency since the velocity and pressure are handled independently, requiring the solution of three (in two dimensions) elliptic systems of rank N (opposed to a single system of rank $3N$ solved in the Stokes problem). However, a drawback of this approach is the splitting scheme error which is introduced when decoupling the pressure and the velocity system, although this can be made consistent with the overall temporal accuracy of the scheme by appropriate discretisation of the pressure boundary conditions.

11.1.1.1 High order splitting scheme

In the original approach a stiffly-stable time integration was proposed in the work of Karniadakis, Israeli and Orszag [20]. This was then later fully analysed in the work of Guermond and Shen [17].

Briefly, high order splitting scheme was originally proposed in three steps involving explicit advection of the non-linear terms, followed by the solution of the pressure Poisson system and finally solving a Helmholtz problem to enforce the viscous terms and velocity boundary conditions. In the following however we briefly formulate this scheme as a two steps using a formulation outline by Guermond and Shen.

1. In the first step we formulate a weak pressure Poisson problem by taking the inner product over the solution domain Ω of equation (11.1a) with respect to the gradient of the test basis, ∇q , i.e.

$$\int_{\Omega} \nabla q \cdot \frac{\partial \mathbf{u}}{\partial t} + \int_{\Omega} \nabla q \cdot \mathbf{N}(\mathbf{u}) = - \int_{\Omega} \nabla q \cdot \nabla p + \int_{\Omega} \nabla q \cdot \nu \nabla^2 \mathbf{u} \quad (11.2)$$

where $N(\mathbf{u}) = \mathbf{u} \cdot \nabla \mathbf{u}$. We recall that the term $\int_{\Omega} \nabla q \cdot \nabla p$ is the weak approximation to the Laplacian operator for pressure. To decouple this term from the velocity system a few steps are necessary. Using the identity

$$\nabla^2 \mathbf{u} = -\nabla \times \nabla \times \mathbf{u} + \nabla(\nabla \cdot \mathbf{u})$$

we can enforce the divergence to be zero by setting the last term to zero. If we now integrate the 1st, 2nd and last term in equation (11.2) by parts we can obtain the weak pressure equation

$$\begin{aligned} \int_{\Omega} \nabla q \cdot \nabla p^{n+1} &= \int_{\Omega} q \nabla \cdot \left(\frac{\partial \mathbf{u}^{n+1}}{\partial t} + \mathbf{N}(\mathbf{u})^{n+1} \right) \\ &\quad - \int_{\partial\Omega} q \left(\frac{\partial \mathbf{u}^{n+1}}{\partial t} + \mathbf{N}(\mathbf{u})^{n+1} + \nu \nabla \times \nabla \times \mathbf{u}^{n+1} \right) \cdot \mathbf{n} \end{aligned} \quad (11.3)$$

where $\partial\Omega$ is the boundary of the domain and we have used the factor that $\nabla \cdot (\nabla \times \nabla \times \mathbf{u}) = 0$. To get the final form of the weak pressure Poisson equation we can use a backward approximation of the time derivative to obtain

$$\frac{\partial \mathbf{u}^{n+1}}{\partial t} \simeq \frac{\gamma_0 \tilde{\mathbf{u}}^{n+1} - \hat{\mathbf{u}}}{\Delta t} \quad (11.4)$$

where $\tilde{\mathbf{u}}^{n+1}$ is an intermediate velocity upon which to decouple the system we impose that $\nabla \cdot \tilde{\mathbf{u}}^{n+1} = 0$ and

$$\gamma_0 = \begin{cases} 1, & \text{if } J = 1 \\ \frac{3}{2}, & \text{if } J = 2 \end{cases} \quad \hat{\mathbf{u}} = \begin{cases} u^n, & \text{if } J = 1 \\ 2u^n - \frac{1}{2}u^{n-1}, & \text{if } J = 2. \end{cases}$$

Finally we introduce a consistent extrapolation for the non-linear terms and the curl of vorticity terms of the form:

$$\mathbf{N}^{*,n+1} = \begin{cases} \mathbf{N}^n, & \text{if } J = 1 \\ 2\mathbf{N}^n - \mathbf{N}^{n-1}, & \text{if } J = 2. \end{cases}$$

A similar extrapolation can be used on the curl-curl term to end up with the final weak pressure approximation

$$\begin{aligned} \int_{\Omega} \nabla q \cdot \nabla p^{n+1} &= \int_{\Omega} q \nabla \cdot \left(-\frac{\hat{\mathbf{u}}}{\Delta t} + \mathbf{N}(\mathbf{u})^{*,n+1} \right) \\ &\quad - \int_{\partial\Omega} q \left(\frac{\partial \mathbf{u}^{n+1}}{\partial t} + \mathbf{N}(\mathbf{u})^{*,n+1} + \nu(\nabla \times \nabla \times \mathbf{u})^{*,n+1} \right) \cdot \mathbf{n} \end{aligned} \quad (11.5)$$

We note this can be recast into an equivalent strong form of the pressure Poisson equation of the form

$$\nabla^2 p^{n+1} = \nabla \cdot \left(\frac{\hat{\mathbf{u}}}{\Delta t} - \mathbf{N}^{*,n+1} \right) \quad (11.6)$$

with consistent Neumann boundary conditions prescribed as

$$\frac{\partial p^{n+1}}{\partial n} = - \left[\frac{\partial \mathbf{u}^{n+1}}{\partial t} + \nu(\nabla \times \nabla \times \mathbf{u})^{*,n+1} + \mathbf{N}^{*,n+1} \right] \cdot \mathbf{n} \quad (11.7)$$

2. The second step is discretise equation (11.1a) at time level $n + 1$, use the pressure at $n + 1$ from the first step and solve for the velocity \mathbf{u}^{n+1} .

In this step now approximate the time derivative using

$$\frac{\partial \mathbf{u}^{n+1}}{\partial t} \simeq \frac{\gamma_0 \mathbf{u}^{n+1} - \hat{\mathbf{u}}}{\Delta t} \quad (11.8)$$

which leads us to the Helmholtz problem

$$\left(\Delta - \frac{\gamma_0}{\nu \Delta t} \right) \mathbf{u}^{n+1} = - \left(\frac{\gamma_0}{\nu \Delta t} \right) \hat{\mathbf{u}} + \frac{1}{\nu} \nabla p^{n+1} \quad (11.9)$$

This scheme is activated in the SolverInfo section with the `SolverType` specification:

```
1 <I PROPERTY="SolverType" VALUE="VelocityCorrectionScheme"/>
```

11.1.1.2 Velocity Correction Scheme with a Weak Pressure formulation

As presented in the previous section in the work of Guermond and Shen [17] and subsequent work they formulate the pressure in a weak rather than strong form to obtain the pressure Poisson system. Therefore if we take the inner product of equation (11.1a) with respect to the gradient of the test space, ∇q , we obtain equation (11.2)

We again make the approximation

$$\frac{\partial \mathbf{u}^{n+1}}{\partial t} \simeq \frac{\gamma_0 \tilde{\mathbf{u}}^{n+1} - \hat{\mathbf{u}}}{\Delta t}. \quad (11.10)$$

However this time we only integrate by parts the last term and do not integrate the non-linear term by parts. However we still need to enforce the condition that $\nabla \cdot \tilde{\mathbf{u}}^{n+1} = 0$

and so we also integrate just this part of the time derivate by parts to arrive at a weak pressure system of the form:

$$\begin{aligned} \int_{\Omega} \nabla q \cdot \nabla p^{n+1} + \frac{\gamma_0}{\Delta t} \int_{\partial\Omega_0} q \tilde{\mathbf{u}}^{n+1} \cdot \mathbf{n} &= \int_{\Omega} \nabla q \cdot \left(\frac{\hat{u}}{\Delta t} - \mathbf{N}^{*,n+1} \right) \\ &\quad - \int_{\partial\Omega_d \cup \partial\Omega_0} q \nu (\nabla \times \nabla \times \mathbf{u})^{*,n+1} \cdot \mathbf{n} + \frac{\gamma_0}{\Delta t} \int_{\partial\Omega_d} q \mathbf{w}^{n+1} \cdot \mathbf{n} \end{aligned} \quad (11.11)$$

where $\partial\Omega_d$ is the Dirichlet boundary conditions for the velocity and $\partial\Omega_0$ is the outflow boundary.

This scheme is activated in the SolverInfo section with the `SolverType` specification:

```
1 <I PROPERTY="SolverType" VALUE="VCSWeakPressure"/>
```

11.1.1.3 Specifying pressure boundary conditions

In order to specify the pressure boundary conditions given by equation (11.7) or for the equivalent conditions in the VCSWeakPressure scheme the `USERDEFINEDTYPE` condition “H” can be used. Therefore a zero velocity wall boundary condition on boundary region 0 in two-dimensions can be specified as

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="0" />
4     <D VAR="v" VALUE="0" />
5     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
6   </REGION>
7 </BOUNDARYCONDITIONS>
```

11.1.1.4 Outflow boundary conditions

The most straightforward outflow condition is to specify fully developed conditions of $\nabla \mathbf{u}^{n+1} \cdot \mathbf{n} = 0$ and $p = 0$ which can be specified as

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <N VAR="u" VALUE="0" />
4     <N VAR="v" VALUE="0" />
5     <D VAR="p" VALUE="0" />
6   </REGION>
7 </BOUNDARYCONDITIONS>
```

However when energetic vortices pass through an outflow region one can experience instabilities as identified by the work of Dong, Karnidakis and Chrysosostomidis [11]. In this paper they suggest to impose a pressure Dirichlet outflow condition of the form

$$p^{n+1} = \nu \mathbf{n} \cdot \nabla \mathbf{u}^{*,n+1} \cdot \mathbf{n} - \frac{1}{2} |\mathbf{u}^{*,n+1}|^2 S_o(\mathbf{n} \cdot \mathbf{u}^{*,n+1}) + \mathbf{f}_b^{n+1} \cdot \mathbf{n} \quad (11.12)$$

with a step function defined by $S_o(n \cdot \mathbf{u}) = \frac{1}{2}(1 - \tanh \frac{n \cdot \mathbf{u}}{\mathbf{u}_0 \delta})$, where \mathbf{u}_0 is the characteristic velocity scale and δ is a non-dimensional positive constant chosen to be sufficiently small. \mathbf{f}_b is the forcing term in this case the analytical conditions can be given but if these are not known explicitly, it is set to zero, i.e. $\mathbf{f}_b = 0$. (see the test `Ko-vaFlow_m8_short_HOBC.xml` for a non-zero example). Note that in the paper [11] they define this term as the negative of what is shown here so that it could be used to impose a default pressure values. This does however mean that the forcing term is imposed through the velocity components u, v by specifying the entry `VALUE` (An example can be found in `ChanFlow_m3_VCSWeakPress_ConOBC.xml`). For the velocity component one can specify

$$\nabla \mathbf{u}^{n+1} \cdot \mathbf{n} = \frac{1}{\nu} \left[p^{n+1} \mathbf{n} + \frac{1}{2} |\mathbf{u}^{*,n+1}|^2 S_o(\mathbf{n} \cdot \mathbf{u}^{*,n+1}) - \nu (\nabla \cdot \mathbf{u}^{*,n+1}) \mathbf{n} - \mathbf{f}_b^{n+1} \right] \quad (11.13)$$

This condition can be enforced using the `USERDEFINEDTYPE` “HOutflow”, i.e.

```

1  <BOUNDARYCONDITIONS>
2    <REGION REF="0">
3      <N VAR="u" USERDEFINEDTYPE="HOutflow" VALUE="0" />
4      <N VAR="v" USERDEFINEDTYPE="HOutflow" VALUE="0" />
5      <D VAR="p" USERDEFINEDTYPE="HOutflow" VALUE="0" />
6    </REGION>
7  </BOUNDARYCONDITIONS>
```

Note that in the moving body work of Bao et al. [4] some care must be made to identify when the flow over the boundary is incoming or outgoing and so a modification of the term

$$\frac{1}{2} |\mathbf{u}^{*,n+1}|^2 S_o(\mathbf{n} \cdot \mathbf{u}^{*,n+1})$$

is replaced with

$$\frac{1}{2} \left((\theta + \alpha_2) |\mathbf{u}^{*,n+1}|^2 + (1 - \theta + \alpha_1) (\mathbf{u}^{*,n+1} \cdot \mathbf{n}) |\mathbf{u}^{*,n+1}| \right) S_o(\mathbf{n} \cdot \mathbf{u}^{*,n+1})$$

where the default values are given by $\theta = 1, \alpha_1 = 0, \alpha_2 = 0$ and these values can be set through the parameters `OutflowBC_theta`, `OutflowBC_alpha1` and `OutflowBC_alpha2`.

Dong has also suggested convective like outflow conditions in [10] which can be enforced through a Robin type specification of the form

$$\frac{\partial \mathbf{u}^{n+1}}{\partial n} + \frac{\gamma_0 D_0}{\Delta t} \mathbf{u}^{n+1} = \frac{1}{\nu} \left[\mathbf{f}^{n+1} + \mathbf{E}(\mathbf{n}, \mathbf{u}^{*,n+1}) + p^{n+1} \mathbf{n} - \nu (\nabla \cdot \mathbf{u}^{*,n+1}) \mathbf{n} \right] + \frac{D_0}{\Delta t} \hat{\mathbf{u}} \quad (11.14)$$

$$\frac{\partial p^{n+1}}{\partial n} + \frac{1}{\nu D_0} p^{n+1} = - \left(-\nu (\nabla \times \nabla \times \mathbf{u})^{*,n+1} + \mathbf{N}^{*,n+1} \right) \cdot \mathbf{n} - \frac{1}{\nu D_0} \left[\mathbf{f}^{n+1} + \mathbf{E}(\mathbf{n}, \mathbf{u}^{*,n+1} + p^{n+1} \mathbf{n} - \nu (\nabla \cdot \mathbf{u}^{*,n+1}) \mathbf{n}) \right] \quad (11.15)$$

```

1 <BOUNDARYCONDITIONS>
2 <REGION REF="0">
3 <R VAR="u" USERDEFINEDTYPE="HOutflow" VALUE="0" PRIMCOEFF="D0/TimeStep"/>
4 <R VAR="v" USERDEFINEDTYPE="HOutflow" VALUE="0" PRIMCOEFF="D0/TimeStep"/>
5 <R VAR="p" USERDEFINEDTYPE="HOutflow" VALUE="0" PRIMCOEFF="1.0/(D0*Kinvis)"/>
6 </REGION>
7 </BOUNDARYCONDITIONS>

```

11.1.1.5 Substepping/subcycling the Velocity Correction Scheme

It is possible to use different time steps in the velocity correction scheme using a substepping (also known as subcycling) [38] or auxiliary semi-Lagrangian approach [46]. Originally the scheme was proposed by Maday, Patera and Ronquist who referred to as an operator-integration-factor splitting method [27]

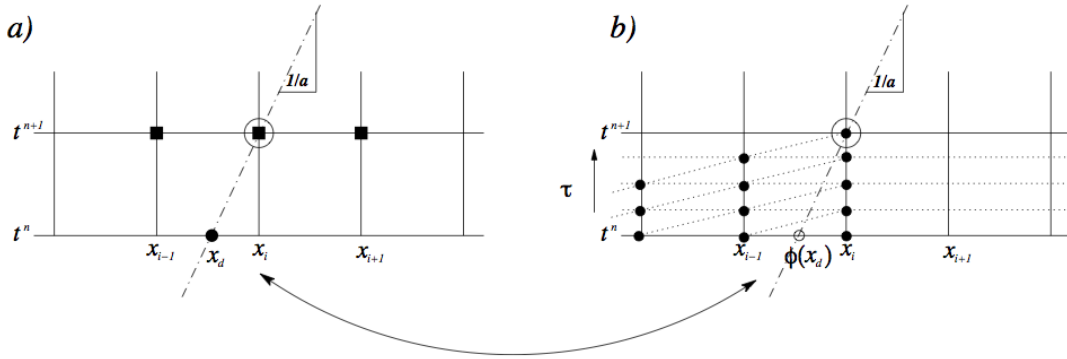


Figure 11.1 Schematic representation of the substepping approach. (a) Making an explicit time step the hyperbolic solution, travelling with a speed a , can be understood as being related to the solution at point x_d (the departure point). (b) Making smaller explicit time steps we can evaluate the solution $\phi(x_d)$ at the departure point and then use this value to make a semi-Lagrangian discretisation of the implicit components usually associated with diffusion.

A schematic of the approach can be understood from figure 11.1.1.5 where we observe that smaller time steps can be used for the explicit advection steps whilst a larger overall time step is adopted for the more expensive implicit solve for the diffusion operator. More details of the implementation can be found in [46] and [38]. In the following sections we outline the parameters that can be used to set up this scheme. Since the explicit part is advanced using a DG scheme it is necessary to use a `Mixed_CG_Discontinuous` expansion with this option.

Note

Some examples of the substepping scheme can be found in the regression tests directory under `$NEKHOME/Solver/IncNavierStokesSolver/Tests/` directory: `KovaFlow_SubStep_2order.xml`, `Hex_Kovasnay_SubStep.xml` and `Tet_Kovasnay_SubStep.xml`.

11.1.1.6 Approximation spaces for the velocity correction scheme

For well resolved simulations it appears that often using the same polynomial space for the pressure and velocity does give suitable answer but this does not satisfy the so-called LBB or inf-sup condition. Therefore, it is potentially better to specify an equivalent of the Taylor Hood approximation and use one higher polynomial order for velocity than the pressure with a continuous expansion. To specify this type of expansion you can use an expansion section of the form:

```

1  <EXPANSIONS>
2      <E COMPOSITE="C[0]" NUMMODES="8" FIELDS="u,v" TYPE="MODIFIED" />
3      <E COMPOSITE="C[0]" NUMMODES="7" FIELDS="p" TYPE="MODIFIEDQUADPLUS1" />
4  </EXPANSIONS>

```

In the above example the “u,v” fields are specified to have a polynomial order of 7 using a modified expansion. Implicitly this form of the expansion definition uses a quadrature order of 9. The above definition then also uses a modified expansion for pressure but of polynomial order 6. Since currently for this solver to run we need to use a consistent quadrature order for both the velocity and pressure fields we specify the `MODIFIEDQUADPLUS1` to tell the solver to use an additional quadrature point and therefore also use 9 quadrature points in each 1D direction for the pressure.

In other cases it is sometimes useful to run with an even higher quadrature order, for example to handle highly deformed elements where the Jacobian is represented by a polynomial expansion. This can be done by using a more detailed definition of the expansion of the form:

```

1  <EXPANSIONS>
2      <E COMPOSITE="C[0]" BASISTYPE="Modified_A,Modified_B" NUMMODES="8,8"
        POINTSTYPE="GaussLobattoLegendre,GaussRadauMAlpha1Beta0" NUMPOINTS="9,8"
        FIELDS="u,v" />
3      <E COMPOSITE="C[0]" BASISTYPE="Modified_A,Modified_B" NUMMODES="7,7"
        POINTSTYPE="GaussLobattoLegendre,GaussRadauMAlpha1Beta0" NUMPOINTS="9,8"
        FIELDS="p" />
4  </EXPANSIONS>

```

In this example we have specified an 8th order expansion for “u,v” and a 7th order expansion for “p”. The BasisType is given as “Modified_A, Modified_B” which is for a triangular expansion (note that for a quadrilateral expansion it would have been “Modified_A,Modified_A”) and so the number of quadrature points in this case is 9 in the first direction which uses Gauss-Lobatto-Legendre points but only 8 in the second

direction since this uses a Gauss-Radau formula with $\alpha = 1, \beta = 0$ weights (see [21] for details on why).

11.1.2 Direct solver (coupled approach)

The second approach consists of directly solving the matrix problem arising from the discretization of the Stokes problem. The direct solution of the Stokes system introduces the problem of appropriate spaces for the velocity and the pressure systems to satisfy the inf-sup condition and it requires the solution of the full velocity-pressure system. However, if a discontinuous pressure space is used then all but the constant mode of the pressure system can be decoupled from the velocity. When implementing this approach with a spectral/hp element discretization, the remaining velocity system may then also be statically condensed to decouple the so called interior elemental degrees of freedom, reducing the Stokes problem to a smaller system expressed on the elemental boundaries. The direct solution of the Stokes problem provides a very natural setting for the solution of the pressure system which is not easily dealt with in a splitting scheme. Further, the solution of the full coupled velocity system allows the introduction of a spatially varying viscosity, which arise for non-Newtonian flows, with only minor modifications.

Note



The coupled solver is only supported for two-dimensional or quasi-3D problems, and only using a direct solver (e.g. `DirectStaticCond`) which prevents its use in parallel.

We consider the weak form of the Stokes problem for the velocity field $\mathbf{u} = [u, v]^T$ and the pressure field p :

$$(\nabla \phi, \nu \nabla \mathbf{u}) - (\nabla \cdot \phi, p) = (\phi, \mathbf{f}) \quad (11.16a)$$

$$(q, \nabla \cdot \mathbf{u}) = 0 \quad (11.16b)$$

where the components of A, B and C are $\nabla \phi_b, \nu \nabla \mathbf{u}_b$, $\nabla \phi_b, \nu \nabla \mathbf{u}_i$ and $\nabla \phi_i, \nu \nabla \mathbf{u}_i$ and the components D_b and D_i are $q, \nabla \mathbf{u}_b$ and $q, \nabla \mathbf{u}_i$. The indices b and i refer to the degrees of freedom on the elemental boundary and interior respectively. In constructing the system we have lumped the contributions from each component of the velocity field into matrices A, B and C . However, we note that for a Newtonian fluid the contribution from each field is decoupled. Since the interior degrees of freedom of the velocity field do not overlap, the matrix C is block diagonal and to take advantage of this structure we can statically condense out the C matrix to obtain the system:

$$\begin{bmatrix} A - BC^{-1}B^T & D_b^T - BC^{-1}D_i & 0 \\ D_b - D_i^T C^{-1}B^T & -D_i^T C^{-1}D_i & 0 \\ B^T & D_i & C \end{bmatrix} \begin{bmatrix} \mathbf{u}_b \\ p \\ \mathbf{u}_i \end{bmatrix} = \begin{bmatrix} \mathbf{f}_b - BC^{-1}\mathbf{f}_i \\ -D_i^T C^{-1}\mathbf{f}_i \\ \mathbf{f}_i \end{bmatrix} \quad (11.17)$$

To extend the above Stokes solver to an unsteady Navier-Stokes solver we first introduce the unsteady term, $\partial \mathbf{u} / \partial t$, into the Stokes problem. This has the principal effect of modifying the weak Laplacian operator $\nabla \phi, \nu \nabla \mathbf{u}]$ into a weak Helmholtz operator $\nabla \phi, \nu \nabla \mathbf{u}) - \lambda(\phi, \mathbf{u}$ where λ depends on the time integration scheme. The second modification requires the explicit discretisation of the non-linear terms in a similar manner to the splitting scheme and this term is then introduced as the forcing term \mathbf{f} . For more details see [1, 39].

11.1.3 Linear Stability Analysis

Hydrodynamic stability is an important part of fluid-mechanics that has a relevant role in understanding how an unstable flow can evolve into a turbulent state of motion with chaotic three-dimensional vorticity fields and a broad spectrum of small temporal and spatial scales. The essential problems of hydrodynamic stability were recognised and formulated in 19th century, notably by Helmholtz, Kelvin, Rayleigh and Reynolds.

Conventional linear stability assumes a normal representation of the perturbation fields that can be represented as independent wave packets, meaning that the system is self-adjoint. The main aim of the global stability analysis is to evaluate the amplitude of the eigenmodes as time grows and tends to infinity. However, in most industrial applications, it is also interesting to study the behaviour at intermediate states that might affect significantly the functionality and performance of a device. The study of the transient evolution of the perturbations is seen to be strictly related to the non-normality of the linearised Navier-Stokes equations, therefore the normality assumption of the system is no longer assumed. The eigenmodes of a non-normal system do not evolve independently and their interaction is responsible for a non-negligible transient growth of the energy. Conventional stability analysis generally does not capture this behaviour, therefore other techniques should be used.

A popular approach to study the hydrodynamic stability of flows consists in performing a direct numerical simulation of the linearised Navier-Stokes equations using iterative methods for computing the solution of the associated eigenproblem. However, since linearly stable flows could show a transient increment of energy, it is necessary to extend this analysis considering the combined effect of the direct and adjoint evolution operators. This phenomenon has noteworthy importance in several engineering applications and it is known as transient growth.

In Nektar++ it is then possible to use the following tools to perform stability analysis:

- direct stability analysis;
- adjoint stability analysis;
- transient growth analysis;

11.1.3.1 Direct stability analysis

The equations that describe the evolution of an infinitesimal disturbance in the flow can be derived decomposing the solution into a basic state (\mathbf{U}, p) and a perturbed state $\mathbf{U} + \varepsilon \mathbf{u}'$ with $\varepsilon \ll 1$ that both satisfy the Navier-Stokes equations. Substituting into the Navier-Stokes equations and considering that the quadratic terms $\mathbf{u}' \cdot \nabla \mathbf{u}'$ can be neglected, we obtain the linearised Navier-Stokes equations:

$$\frac{\partial \mathbf{u}'}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{u}' + \mathbf{u}' \cdot \nabla \mathbf{U} = -\nabla p + \nu \nabla^2 \mathbf{u}' + \mathbf{f} \quad (11.18a)$$

$$\nabla \cdot \mathbf{u}' = 0 \quad (11.18b)$$

The linearised Navier-Stokes equations are identical in form to the non-linear equation, except for the non-linear advection term. Therefore the numerical techniques used for solving Navier-Stokes equations can still be applied as long as the non-linear term is substituted with the linearised one. It is possible to define the linear operator that evolved the perturbation forward in time:

$$\mathbf{u}'(\mathbf{x}, t) = \mathcal{A}(\mathbf{U})\mathbf{u}'(\mathbf{x}, 0) \quad (11.19)$$

Let us assume that the base flow \mathbf{U} is steady, then the perturbations are autonomous and we can assume that:

$$\mathbf{u}'(\mathbf{x}, t) = \mathbf{q}'(\mathbf{x}) \exp(\lambda t) \quad \text{where } \lambda = \sigma + i\omega \quad (11.20)$$

Then we obtain the associated eigenproblem:

$$\mathcal{A}(\mathbf{U})\mathbf{q}' = \lambda \mathbf{q}' \quad (11.21)$$

The dominant eigenvalue determines the behaviour of the flow. If the real part is positive then there exist exponentially growing solutions. Conversely, if all the eigenvalues have negative real part then the flow is linearly stable. If the real part of the eigenvalue is zero, it is a bifurcation point.

11.1.3.2 Adjoint Stability Analysis

The adjoint of a linear operator is one of the most important concepts in functional analysis and it plays an important role in understanding transition to turbulence. Let us write the linearised Navier-Stokes equation in a compact form:

$$\mathcal{H}\mathbf{q} = 0 \quad \text{where} \quad \mathcal{H} = \left(\frac{-\partial_t - (\mathbf{U} \cdot \nabla) + (\nabla \mathbf{U}) \cdot + \frac{1}{Re} \nabla^2}{\nabla \cdot} \middle| \frac{-\nabla}{0} \right) \quad (11.22)$$

The adjoint operator \mathcal{H}^* is defined as:

$$\langle \mathcal{H}\mathbf{q}, \mathbf{q} \rangle = \langle \mathbf{q}, \mathcal{H}^* \mathbf{q}^* \rangle \quad (11.23)$$

Integrating by parts and employing the divergence theorem, it is possible to express the adjoint equations:

$$-\frac{\partial \mathbf{u}^*}{\partial t} + (\mathbf{U} \cdot \nabla) \mathbf{u}^* + (\nabla \mathbf{U})^T \cdot \mathbf{u}^* = -\nabla p^* + \frac{1}{Re} \nabla^2 \mathbf{u} \quad (11.24a)$$

$$\nabla \cdot \mathbf{u}^* = 0 \quad (11.24b)$$

The adjoint fields are in fact related to the concept of **receptivity**. The value of the adjoint velocity at a point in the flow indicates the response that arises from an unsteady momentum source at that point. The adjoint pressure and the adjoint stream function play instead the same role for mass and vorticity sources respectively. Therefore, the adjoint modes can be seen as a powerful tool to understand where to act in order to ease/inhibit the transition.

11.1.3.3 Transient Growth Analysis

Transient growth is a phenomenon that occurs when a flow that is linearly stable, but whose perturbations exhibit a non-negligible transient response due to regions of localised convective instabilities. This situation is common in many engineering applications, for example in open flows where the geometry is complex, producing a steep variation of the base flow. Therefore, the main question to answer is if it exists a bounded solution that exhibit large growth before inevitably decaying. Let us introduce a norm to quantify the size of a perturbation. It is physically meaningful to use the total kinetic energy of a perturbation on the domain Ω . This is convenient because it is directly associated with the standard- $L2$ inner product:

$$\mathcal{A}(\tau)\mathbf{v} = \sigma\mathbf{u}, \quad \|\mathbf{u}\| = 1 \quad (11.25)$$

where $\sigma = \|\mathbf{u}'(\tau)\|$. This is no other than the singular value decomposition of $\mathcal{A}(\tau)$. The phenomenology of the transient growth can be explained considering the non-normality of the linearised Navier-Stokes evolution operator. This can be simply understood using the simple geometric example showed in figure 11.1.3.3. Let us assume a unit-length vector

\mathbf{f} represented in a non-orthogonal basis. This vector is defined as the difference of the nearly collinear vectors Φ_1 and Φ_2 . With the time progression, the component of these two vectors decrease respectively by 20% and 50%. The vector \mathbf{f} increases substantially in length before decaying to zero. Thus, the superposition of decaying non-orthogonal eigenmode can produce in short term a growth in the norm of the perturbations.




Figure 11.2 Geometric interpretation of the transient growth. Adapted from Schmid, 2007

11.1.4 Steady-state solver using Selective Frequency Damping

To compute linear stability analysis, the choice of the base flow, around which the system will be linearised, is crucial. If one wants to use the steady-state solution of the Navier-Stokes equations as base flow, a steady-state solver is implemented in *Nektar++*. The method used is the encapsulated formulation of the Selective Frequency Damping method [19]. Unstable steady base flows can be obtained with this method. The SFD method is based on the filtering and control of unstable temporal frequencies within the flow. The time continuous formulation of the SFD method is

$$\begin{cases} \dot{q} = NS(q) - \chi(q - \bar{q}), \\ \dot{\bar{q}} = \frac{q - \bar{q}}{\Delta}. \end{cases} \quad (11.26)$$

where q represents the problem unknown(s), the dot represents the time derivative, NS represents the Navier-Stokes equations, $\chi \in \mathbb{R}_+$ is the control coefficient, \bar{q} is a filtered version of q , and $\Delta \in \mathbb{R}_+^*$ is the filter width of a first-order low-pass time filter. The steady-state solution is reached when $q = \bar{q}$.

The convergence of the method towards a steady-state solution depends on the choice of the parameters χ and Δ . They have to be carefully chosen: if they are too small, the instabilities within the flow can not be damped; but if they are too large, the method may converge extremely slowly. If the dominant eigenvalue of the flow studied is known

(and given as input), the algorithm implemented can automatically select parameters that ensure a fast convergence of the SFD method. Most of the time, the dominant eigenvalue is not known, that is why an adaptive algorithm that adapts χ and Δ all along the solver execution is also implemented.

Note that this method can not be applied for flows with a pure exponential growth of the instabilities (*e.g.* jet flow within a pipe). In other words, if the frequency of the dominant eigenvalue is zero, then the SFD method is not a suitable tool to obtain a steady-state solution.

11.2 Usage

```
IncNavierStokesSolver session.xml
```

11.3 Session file configuration

In the following the possible options are shown for the incompressible Navier-Stokes. The **Expansion** section for an incompressible flow simulation can be set as for other solvers regardless of the projection type. Here an example for a 3D simulation (for 2D simulations the specified fields would be just `u,v,p`).

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="6" FIELDS="u,v,w,p" TYPE="MODIFIED" />
3 </EXPANSIONS>
```

In case of a simulation using the Direct Solver we need to set `FIELDS=u,v` as the pressure expansion order will be automatically set to fulfil the inf-sup condition. Possible choices for the expansion `TYPE` are:

Basis	TYPE
Modal	MODIFIED
Nodal	GLL_LAGRANGE
Nodal SEM	GLL_LAGRANGE_SEM

11.3.1 Solver Info

The following parameters can be specified in the `SOLVERINFO` section of the session file:

- `EqType`: sets the kind of equations we want to solve on the domain as:

```
1 <I PROPERTY="EQTYPE" VALUE="UnsteadyNavierStokes" />
```

Possible values are:

Equations	EQTYPE	Dim.	Projections	Alg.
Steady Stokes (SS)	SteadyStokes	All	CG	VCS
Steady Oseen (SO)	SteadyOseen	All	CG	DS
Unsteady Stokes (US)	UnsteadyStokes	All	CG	VCS
Steady Linearised NS (SLNS)	SteadyLinearisedNS	All	CG	DS
Unsteady Linearised NS (ULNS)	UnsteadyLinearisedNS	All	CG	VCS,DS
Unsteady NS (UNS)	UnsteadyNavierStokes	All	CG,CG-DG	VCS

- **SolverType**: sets the scheme we want to use to solve the set of equations as

```
1 <I PROPERTY="SolverType" VALUE="VelocityCorrectionScheme"/>
```

Possible values are:

Algorithm	SolverType	Dimensions	Projections
Velocity Correction Scheme (VCS)	VelocityCorrectionScheme	2D, Quasi-3D, 3D	CG, CG-DG
VCS with weak pressure	VCSWeakPressure	2D, Quasi-3D, 3D	CG, CG-DG
Direct solver	CoupledLinearisedNS	2D, Quasi-3D	CG

- **Driver**: this specifies the type of problem to be solved:

Driver	Description	Dimensions	Projections
Standard	Time integration of the equations	All	CG, DG
SteadyState	Steady-state solver (see Sec. 11.1.4)	All	CG, DG

- **Projection**: sets the Galerkin projection type as

```
1 <I PROPERTY="Projection" VALUE="Continuous"/>
```

Possible values are:

Galerkin Projection	Projection	Dimensions	Equations	Algorithms
Continuous (CG)	Continuous	All	All	All
Discontinuous (DG)	DisContinuous	All
Mixed CG and DG (CG-DG)	Mixed_CG_Discontinuous	2D,3D	just UNS	VCS-substepping

- **TimeIntegrationMethod**: sets the time integration method as

```
1 <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2"/>
```

Possible values are

Time-Integration Method	TimeIntegrationMethod	Dimensions	Equations	Projections
IMEX Order 1	IMEXOrder1	all	US, UNS	CG
IMEX Order 2	IMEXOrder2	all	US, UNS	CG
IMEX Order 3	IMEXOrder3	all	US, UNS	CG
Backward Euler	BackwardEuler	all	US, UNS	CG-DG
BDF Order 1	BDFImplicitOrder1	all	US, UNS	CG-DG
BDF Order 2	BDFImplicitOrder2	all	US, UNS	CG-DG

- **Extrapolation**: Specify the extrapolation method (standard or substepping) to be used in velocity correction scheme. Essentially this activates the sub-stepping routine which requires the mixed CG-DG projection

```
1 <I PROPERTY="Extrapolation" VALUE="SubStepping"/>
```

Possible values are **SubStepping** or **Standard** with “Standard” being the default value if nothing is specified.

- **SubStepIntScheme**: choose the substep DG time integration scheme so that a different order schemes can be used as compared to the overall time integration scheme.

```
1 <I PROPERTY="SubStepIntScheme" VALUE="RungeKutta2_ImprovedEuler"/>
```

Possible values are

Time-Integration Method	SubStepIntScheme
Forward Euler	ForwardEuler
RK 2	RungeKutta2_ImprovedEuler

This option is useful if you wish to use an overall scheme that is first order accurate for example with TimeIntegrationMethod as BDFImplicitOrder1 but using a second order RungeKutta2_ImprovedEuler for greater stability in the substep.

- **GlobalSysSoln**: sets the approach we use to solve the the linear systems of the type $Ax = b$ appearing in the solution steps, such as the Poisson equation for the pressure in the splitting-scheme. It can be set as

```
1 <I PROPERTY="GlobalSysSoln" VALUE="IterativeStaticCond"/>
```

Possible values are

System solution	GlobalSysSoln	Parallel
Direct Solver (DS)	DirectFull	just quasi-3D
DS with Static Condensation	DirectStaticCond	just Quasi-3D
DS with Multilevel Static Condensation	DirectMultiLevelStaticCond	just Quasi-3D
Iterative Solver (IS)	IterativeFull	just Quasi-3D
IS with Static Condensation	IterativeStaticCond	quasi-3D
IS with Multilevel Static Condensation	IterativeMultiLevelStaticCond	quasi-3D

Default values are `DirectMultiLevelStaticCond` in serial and `IterativeStaticCond` in parallel.

- `SmoothAdvection`: activates a stabilization technique which smooths the advection term using the pressure inverse mass matrix. It can be used just in combination with nodal expansion basis for efficiency reasons.

```
1 <I PROPERTY="SmoothAdvection" VALUE="True"/>
```

- `SpectralVanishingViscosity`: activates a stabilization technique which increases the viscosity on the modes with the highest frequencies.

```
1 <I PROPERTY="SpectralVanishingViscosity" VALUE="True"/>
```

In a Quasi-3D simulation, this will affect both the Fourier and the spectral/hp expansions. To activate them independently, use `SpectralVanishingViscositySpectralHP` and `SpectralVanishingViscosityHomo1D`.

There are three spectral vanishing viscosity kernels available:

SVV Kernel	SpectralVanishingViscosity
Exponential Kernel	True
Power Kernel	PowerKernel
DG Kernel	DGKernel

The Exponential kernel is based on the work of Maday et al. [28], its extension to 2D can be found in [22]. A diffusion coefficient can be specified which defines the base magnitude of the viscosity; this parameter is scaled by h/p . SVV viscosity is activated for expansion modes greater than the product of the cut-off ratio and the expansion order. The Power kernel is a smooth function with no cut-off frequency; it focusses on a narrower band of higher expansion modes as the polynomial order increases. The cut-off ratio parameter for the Power kernel corresponds to the power ratio, see Moura et al. [30]. The DG-Kernel is an attempt to match the dissipation of CG-SVV to DG schemes of lower expansion orders. This kernel does not require any parameters although the diffusion coefficient can still be modified.

- `DEALIASING`: activates the 3/2 padding rule on the advection term of a Quasi-3D simulation.

```
1 <I PROPERTY="DEALIASING" VALUE="True"/>
```

- `SPECTRALHPDEALIASING`: activates the spectral/hp dealiasing to stabilize the simulation. This method is based on the work of Kirby and Sherwin [7].

```
1 <I PROPERTY="SPECTRALHPDEALIASING" VALUE="True" />
```

11.3.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file:

- `TimeStep`: sets the timestep for the integration in time formula.
- `NumSteps`: sets the number of time-steps.
- `IO_CheckSteps`: sets the number of steps between successive checkpoint files.
- `IO_InfoSteps`: sets the number of steps between successive info stats are printed to screen.
- `Kinvis`: sets the cinematic viscosity coefficient formula.
- `SubStepCFL`: sets the CFL safety limit for the sub-stepping algorithm (default value = 0.5).
- `MinSubSteps`: perform a minimum number of substeps in sub-stepping algorithm (default is 1).
- `MaxSubSteps`: perform a maximum number of substeps in sub-stepping algorithm otherwise exit (default is 100).
- `SVVCutoffRatio`: sets the ratio of Fourier frequency not affected by the SVV technique (default value = 0.75, i.e. the first 75% of frequency are not damped).
- `SVVDiffCoeff`: sets the SVV diffusion coefficient (default value = 0.1 (Exponential and Power kernel), 1 (DG-Kernel)).
- `IO_CFLWriteFld`: sets a threshold value for the CFL number. If CFL exceeds this value, the flow field is written to file (only once). This is useful for debugging purposes, allowing to visually inspect a flow field that is becoming unstable.
- `IO_CFLWriteFldNumSteps`: sets the number of timesteps after which `IO_CFLWriteFld` becomes operational. This avoids writing the flow field at the beginning of a simulation when initialising a new geometry.

11.3.3 Womersley Boundary Condition

It is possible to define the time-dependent Womersley velocity profile for pulsatile flow in a pipe. The modulation of the velocity profile is based on the desired peak or centerline velocity which can be represented by a Fourier expansion $U_{max} = A(\omega_n)e^{i\omega_n t}$ where A are the Fourier modes and ω the frequency. The womersley solution is then defined as:

$$w(r, t) = A_0(1 - (r/R)^2) + \sum_{n=1}^N \tilde{A}_n \left[1 - \frac{J_0(i^{3/2}\alpha_n r/R)}{J_0(i^{3/2}\alpha)} \right] e^{i\omega_n t}$$

where the womersley number α is defined:

$$\alpha_n = R\sqrt{\frac{2\pi n}{T\nu}}$$

and \tilde{A}_n ($n = 1 : N$) are the Fourier coefficients scaled in the following way:

$$\tilde{A}_n = 2A_n/[1 - \frac{1}{J_0(i^{3/2}\alpha)}]$$

The Womersley velocity profile is implemented in the following way:

```

1 <REGION REF="0">
2   <D VAR="u" USERDEFINEDTYPE="Womersley:WomParams.xml" VALUE="0" />
3   <D VAR="v" USERDEFINEDTYPE="Womersley:WomParams.xml" VALUE="0" />
4   <D VAR="w" USERDEFINEDTYPE="Womersley:WomParams.xml" VALUE="0" />
5   <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
6 </REGION>

```

A file containing the Fourier coefficients, \tilde{A} , must be in the directory where the solver is called from. The name of the file is defined by the string given in the attribute `USERDEFINEDTYPE` after the “:” and contains the real and imaginary coefficients. This file has the format

```

1 <NEKTAR>
2   <WOMERSLEYBC>
3     <WOMPARAMS>
4       <W PROPERTY="Radius" VALUE="0.5" />
5       <W PROPERTY="Period" VALUE="1.0" />
6       <W PROPERTY="axisnormal" VALUE="0.0,0.0,1.0" />
7       <W PROPERTY="axispoint" VALUE="0.0,0.0,0.0" />
8     </WOMPARAMS>
9
10    <FOURIERCOEFFS>
11      <F ID="0"> 0.600393641193, 0.0 </F>
12      <F ID="1"> -0.277707172935, 0.0767582715413 </F>
13      <F ID="2"> -0.0229953131146, 0.0760936232478 </F>
14      <F ID="3"> 0.00858135174058, 0.017089888642 </F>
15      <F ID="4"> 0.0140332527651, 0.0171575122496 </F>
16      <F ID="5"> 0.0156970122129, -0.00547357750345 </F>
17      <F ID="6"> 0.00473626554238, -0.00498786519876 </F>
18      <F ID="7"> 0.00204434981523, -0.00614566561937 </F>
19      <F ID="8"> -0.000274697215201, 0.000153571881197 </F>
20      <F ID="9"> -0.000148037910774, 2.68919619581e-05 </F>
21    </FOURIERCOEFFS>
22  </WOMERSLEYBC>
23 </NEKTAR>

```

Each value of \tilde{A} is provided in the `FOURIERCOEFFS` section and provided as separate entries containing the real and imaginary components, i.e. the mean component provided above is 0.600393641193, 0.0.

Similarly in the `WOMPARAMS` section the key parameters of the boundary condition are also provided as:

- `RADIUS` is the radius of the boundary.
- `PERIOD` is the cycle time period,
- `AXISNORMAL` defines the normal direction to the boundary,
- `AXISPOINT` defines a coordinate in the boundary centre,

11.3.4 Forcing

11.3.4.1 MovingBody



Note

This force type is only supported for the Quasi-3D incompressible Navier-Stokes solver.

This force type allows the user to solve the interaction system of an incompressible fluid flowing past a flexible moving bodies [33]. By this forcing function, one can eliminate the difficulty of moving mesh by using body-fitted coordinates, so that an additional acceleration term (i.e., forcing term) is introduced to the momentum equations by the non-inertial transform from the deformed and moving coordinate system to non-deformed and stationary one.

```
1 <FORCE TYPE="MovingBody">
2 </FORCE>
```

Available options of the motion type for the moving body include free, constrained and forced vibrations, which can be specified in the `SOLVERINFO` section. The free type of motion allows the body to move in both streamwise and crossflow directions, while the constrained type limits the motion only in the crossflow direction. For the forced type, the vibration profiles of the body should be specified as a given function or read from input file in `MovingBody` section. For example:

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE"           VALUE="UnsteadyNavierStokes"  />
3   <I PROPERTY="SolverType"       VALUE="VelocityCorrectionScheme" />
4   <I PROPERTY="EvolutionOperator" VALUE="SkewSymmetric"        />
5   <I PROPERTY="Projection"       VALUE="Galerkin"           />
6   <I PROPERTY="GlobalSysSoln"    VALUE="DirectStaticCond"     />
7   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2"      />
8   <I PROPERTY="HOMOGENEOUS"      VALUE="1D"                 />
9   <I PROPERTY="USEFFT"           VALUE="FFT"                 />
10  <I PROPERTY="VibrationType"    VALUE="FREE"              />
11 </SOLVERINFO>
```

A moving body type boundary condition should be specified in `BOUNDARYCONDITIONS` for the velocities on the moving body,

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" USERDEFINEDTYPE="MovingBody" VALUE="0" />
4     <D VAR="v" USERDEFINEDTYPE="MovingBody" VALUE="0" />
5     <D VAR="w" VALUE="0" />
6     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

For the simulation of low mass ratio, there is an option to activate fictitious mass method for stabilizing explicit coupling between the fluid solver and structural dynamic solver. Here we need to specify the values of fictitious mass and damping in `PARAMETERS`, for example,

```

1 <SOLVERINFO>
2   <I PROPERTY="FictitiousMassMethod" VALUE="True" />
3 </SOLVERINFO>
4 <PARAMETERS>
5   <P> FictDamp      = 1000   </P>
6   <P> FictMass      = 1.5    </P>
7 </PARAMETERS>

```

A filter called `MovingBody` is encapsulated in this module to evaluate the aerodynamic forces along the moving body surface. The forces for each computational plane are projected along the Cartesian axes and the pressure and viscous contributions are computed in each direction.

The following parameters are supported:

Option name	Required	Default	Description
<code>OutputFile</code>	✗	<code>session</code>	Prefix of the output filename to which the forces are written.
<code>Frequency</code>	✗	1	Number of timesteps after which output is written.
<code>Boundary</code>	✓	-	Boundary surfaces on which the forces are to be evaluated.

To enable the filter, add the following to the `FORCE` tag::

```

1 <FORCE TYPE="MovingBody">
2   <PARAM NAME="OutputFile">DragLift</PARAM>
3   <PARAM NAME="OutputFrequency">10</PARAM>
4   <PARAM NAME="Boundary"> B[0] </PARAM>
5 </FORCE>

```

During the execution a file named `DragLift.fce` will be created and the value of the aerodynamic forces on boundary 0, defined in the `GEOMETRY` section, will be output every 10 time steps. evaluates the aerodynamic forces along the moving body surface. The forces for each computational plane are projected along the Cartesian axes and the pressure and viscous contributions are computed in each direction.

Also, to use this module a `MAPPING` needs to be specified, as described in section 11.6. In the case of free and constrained motion presented here, the functions defined by the mapping act as initial conditions. Also, when using the `MovingBody` forcing, it is not necessary to set the `TIMEDEPENDENT` property of the mapping.

11.3.5 Filters

The following filters are supported exclusively for the incompressible Navier-Stokes solver. Further filters from section 3.4 are also available for this solver.

- `Aerodynamic forces` (section 3.4.2)
- `Kinetic energy and enstrophy` (section 3.4.10)
- `Modal energy` (section 3.4.11)
- `Moving body` (section 3.4.12)
- `Reynolds stresses` (section 3.4.15)

11.4 Session file configuration: Linear stability analysis

Stability analyses of incompressible flow involves solving the linearised Navier-Stokes equations

$$\frac{\partial \mathbf{u}'}{\partial t} + \mathcal{L}(\mathbf{U}, \mathbf{u}') = -\nabla p + \nu \nabla^2 \mathbf{u}',$$

where \mathcal{L} is a linear operator, its adjoint form, or both. The evolution of the linearised Navier-Stokes operator, which evolves a solution from an initial state to a future time t , can be written as

$$u(t) = \mathcal{A}(t)u(0).$$

The adjoint evolution operator is denoted as \mathcal{A}^* . This section details the additional configuration options, in addition to the standard configuration options described earlier, relating to performing this task.

11.4.1 Solver Info

- **Ectype**: sets the type of equation to solve. For linear stability analysis this must be set to

Equation Type	Dimensions	Projections	Algorithms
UnsteadyNavierStokes	2D, Quasi-3D	Continuous	VCS,DS

- **EvolutionOperator**: sets the choice of the evolution operator:
 - **Nonlinear** (standard non-linear Navier-Stokes equations).
 - **Direct** (\mathcal{A} – linearised Navier-Stokes equations).
 - **Adjoint** (\mathcal{A}^* – adjoint Navier-Stokes equations).
 - **TransientGrowth** ($\mathcal{A}^* \mathcal{A}$ – transient growth evolution operator).
- **Driver**: specifies the type of problem to be solved:
 - **Standard** (normal time integration of the equations)
 - **ModifiedArnoldi** (computations of the leading eigenvalues and eigenmodes using modified Arnoldi method)
 - **Arpack** (computations of eigenvalues/eigenmodes using Implicitly Restarted Arnoldi Method (ARPACK)).
- **ArpackProblemType**: types of eigenvalues to be computed (for Driver Arpack only)
 - **LargestMag** (eigenvalues with largest magnitude).
 - **SmallestMag** (eigenvalues with smallest magnitude).
 - **LargestReal** (eigenvalues with largest real part).
 - **SmallestReal** (eigenvalues with smallest real part).
 - **LargestImag** (eigenvalues with largest imaginary part).
 - **SmallestIma** (eigenvalues with smallest imaginary part).
- **Homogeneous**: specifies the Fourier expansion in a third direction (optional)
 - **1D** (Fourier spectral method in z-direction).
- **ModeType**: this specifies the type of the quasi-3D problem to be solved.
 - **MultipleMode** (stability analysis with multiple modes, **HomModesZ** sets number of modes).
 - **SingleMode** (BiGlobal Stability Analysis: full-complex mode. Overrides **HomModesZ** to 1.).

- `HalfMode` (BiGlobal Stability Analysis: half-complex mode u.Re v.Re w.Im p.Re).

Note



For visualization of `Homogeneous` results with `FieldConvert` you can use `--output-points-hom-z` to set output number of modes to a desired value. To process results obtained with `HalfMode` you can convert to `SingleMode` using `FieldConvert` module `halfmodetofourier`.

11.4.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file:

- `kdim`: sets the dimension of the Krylov subspace κ . Can be used with: `ModifiedArnoldi` and `Arpack`. Default value: 16.
- `evtol`: sets the tolerance of the iterative eigenvalue algorithm. Can be used with: `ModifiedArnoldi` and `Arpack`. Default value: 1×10^{-6} .
- `nvec`: sets the number of converged eigenvalues sought. Can be used with: `ModifiedArnoldi` and `Arpack`. Default value: 2.
- `nits`: sets the maximum number of Arnoldi iterations to attempt. Can be used with: `ModifiedArnoldi` and `Arpack`. Default value: 500.
- `realShift`: provide a real shift to the direct solver eigenvalue problem by the specified value to improve convergence. Can be used with: `Arpack` only.
- `imagShift`: provide an imaginary shift to the direct solver eigenvalue problem by the specified value to improve convergence. Can be used with: `Arpack` only.
- `LZ`: sets the length in the spanswise direction L_z . Can be used with `Homogeneous` set to `1D`. Default value: 1.
- `HomModesZ`: sets the number of planes in the homogeneous directions. Can be used with `Homogeneous` set to `1D` and `ModeType` set to `MultipleModes`.
- `N_slices`: sets the number of temporal slices for Floquet stability analysis.
- `period`: sets the periodicity of the base flow.

11.4.3 Functions

When using the direct solver for stability analysis it is necessary to specify a Forcing function “StabilityCoupledLNS” in the form:

```
1 <FORCING>
2   <FORCE TYPE="StabilityCoupledLNS">
3   </FORCE>
4 </FORCING>
```

This is required since we need to tell the solver to use the existing field as a forcing function to the direct matrix inverse as part of the Arnoldi iteration.

Note



Examples of the set up of the direct solver stability analysis (and other incompressible Navier-Stokes solvers) can be found in the regression test directory `NEKTAR/solvers/IncNavierStokesSolver/Tests`. See for example the files `PPF_R15000_ModifiedArnoldi_Shift.tst` and `PPF_R15000_3D.xml` noting that some parameters are specified in the `.tst` files.

11.5 Session file configuration: Steady-state solver

In this section, we detail how to use the steady-state solver (that implements the selective frequency damping method, see Sec. 11.1.4). Two cases are detailed here: the execution of the classical SFD method and the *adaptive* SFD method, where the control coefficient χ and the filter width Δ of the SFD method are updated all along the solver execution. For the second case, the parameters of the SFD method do not need to be defined by the user (they will be automatically calculated all along the solver execution) but several session files must be defined in a very specific way.

11.5.1 Execution of the classical steady-state solver

11.5.1.1 Solver Info

The definition of `Eqtype`, `TimeIntegrationMethod` and `Projection` is similar as what is explained in 11.4.1. The use of the steady-state solver is enforced through the definition of the `Driver` which has to be `SteadyState`. `EvolutionOperator` does not need to be defined to run the unadapted SFD method (by default, it is set to `Nonlinear`).

11.5.1.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file:

- `Kinvis`: sets the kinematic viscosity ν . It is typically $1/Re$ if both the characteristic velocity and characteristic length are chosen to be 1.

- `ControlCoeff`: sets the control coefficient χ of the SFD method. Default value: 1.
- `FilterWidth`: sets the filter width Δ of the SFD method. Default value: 2.
- `GrowthRateEV` and `FrequencyEV`: if the growth rate and the frequency of the dominant eigenvalue are known, they can be given as input and the code will automatically select the optimum parameters χ and Δ (and overwrite the values of `ControlCoeff` and `GrowthRateEV` that may be given in the session file)
- `TOL`: sets the tolerance of the SFD method. The code will run until $\|q - \bar{q}\|_{inf} < TOL$. Default value: 10^{-8} .

Note that for the steady-state solver, the parameter `NumSteps` is not taken into account. The solver will run until a steady-state solution is found and not for a pre-defined number of time steps.

11.5.2 Execution of the adaptive steady-state solver

Running the adaptive selective frequency damping method requires to set up the session files in a very specific manner. First, the `Geometry` section must be in a separated archive file. If the test case studied is called "Session", the mesh file must be called `Session.xml.gz` (the linux command "gzip" can be used to obtain this file).

The requirements for the file `Session.xml` are similar as for the ones for the classical SFD method. The `Geometry` section being removed and placed in `Session.xml.gz`. This file defines the properties of the nonlinear problem solved (*i.e.* the flow for which we want a steady-state). Also, the `SOLVERINFO` section must contain the line:

```
1 <I PROPERTY="EvolutionOperator" VALUE="AdaptiveSFD" />
```

The adaptive SFD method used is coupled with a stability analysis method. Then `kdim`, `nvec`, `evtol` and `nits` should be defined into the `PARAMETERS` section of `Session.xml`. If not, these parameters will take the default values presented in Sec. 11.4.

The goal of running the stability analysis is to evaluate the dominant eigenvalue of a "partially converged" steady base flow. This approximation is then used by the steady-state solver to select a control coefficient χ and a filter width Δ then ensure a fast convergence towards a steady-state solution.

To define the linear stability problem, another file, that must be called `Session_LinNS.xml`, has to be defined. This file **must be an exact copy/paste of `Session.xml`**, only three things have to be modified:

1. The boundary conditions must be modified to be homogeneous (*i.e.* equal to zero) at all inflow boundaries.
2. A non-zero function `InitialConditions` has to be defined.

3. A random function `BaseFlow` has to be defined (it will be overwritten all along the solver execution). We recommend it to be a copy of `InitialConditions`.

Once these three files (the `Geometry` in `Session.xml.gz`, the nonlinear problem definition in `Session.xml` and the homogeneous linear problem in `Session_LinNS.xml`) are correctly defined, the adaptive SFD method must be executed using:

```
IncNavierStokesSolver Session.xml.gz Session.xml
```

11.6 Session file configuration: Coordinate transformations

This section describes how to include a coordinate transformation to the solution of the incompressible Navier-Stokes equations. In some cases, this approach allows a slightly deformed geometry to be mapped into a geometry with a homogeneous direction, which can be treated using a quasi-3D method. It is also useful for problems with a moving body, where otherwise a moving mesh would have to be employed.

11.6.1 Solver Info

To activate the mapping technique, `SolverType` needs to be set as:

```
1 <I PROPERTY="SolverType" VALUE="VCSMapping" />
```

Also, there are other optional properties in the `SolverInfo` section:

```
1 <I PROPERTY="MappingImplicitPressure" VALUE="TRUE"/> <!-- Default = FALSE -->
2 <I PROPERTY="MappingImplicitViscous" VALUE="TRUE"/> <!-- Default = FALSE -->
3
4 <I PROPERTY="MappingNeglectViscous" VALUE="FALSE"/> <!-- Default = FALSE -->
```

the first two options determine if the pressure and viscous terms resulting from the coordinate transformation are treated implicitly using an iterative procedure. If the last option is set to true, the viscous terms in the mapping are not computed. This leads to a faster solution, but the effect on the results need to be determined for the specific case. By default, all mapping terms are computed and treated explicitly.

11.6.2 Parameters

When treating the mapping terms implicitly, the following parameters can be set:

```
1 <P> MappingPressureTolerance = 1e-8 </P> <!-- Default = 1e-12 -->
2 <P> MappingViscousTolerance = 1e-8 </P> <!-- Default = 1e-12 -->
3 <P> MappingPressureRelaxation = 0.9 </P> <!-- Default = 1.0 -->
4 <P> MappingViscousRelaxation = 0.9 </P> <!-- Default = 1.0 -->
```

They determine the tolerance of the iterative solution of the equations, and a relaxation parameter which can improve the numerical stability of the method.

11.6.3 Mapping

The particular transformation employed is specified by:

```

1 <MAPPING TYPE="XYofZ">
2   <COORDS>Mapping</COORDS>
3   <VEL>MappingVel</VEL>
4   <TIMEDEPENDENT>True</TIMEDEPENDENT> <!-- Default is False -->
5 </MAPPING>

```

where `TIMEDEPENDENT` indicates if the transformation varies with time.

The available values for `TYPE`, and the transformations they represent, are:

Mapping type	\bar{x}	\bar{y}	\bar{z}
<code>Translation</code>	$x + f(t)$	$y + g(t)$	$z + h(t)$
<code>XofZ</code>	$x + f(z, t)$	y	z
<code>XofXZ</code>	$f(x, z, t)$	y	z
<code>XYofZ</code>	$x + f(z, t)$	$y + g(z, t)$	z
<code>XYofXY</code>	$f(x, y, t)$	$g(x, y, t)$	z
<code>General</code>	$f(x, y, z, t)$	$g(x, y, z, t)$	$h(x, y, z, t)$

where $(\bar{x}, \bar{y}, \bar{z})$ are the Cartesian (physical) coordinates and (x, y, z) are the transformed coordinates. Note that for quasi-3D problems, the z coordinate cannot be transformed.

11.6.4 Functions

The function `COORDS` (and `VEL` for time dependent mappings) indicated in the `MAPPING` section need to be defined, for example as:

```

1 <FUNCTION NAME="Mapping">
2   <E VAR="x" VALUE="x + cos(PI*z)" />
3   <E VAR="y" VALUE="y + cos(2*PI*t)" />
4 </FUNCTION>
5
6 <FUNCTION NAME="MappingVel">
7   <E VAR="vx" VALUE="0.0" />
8   <E VAR="vy" VALUE="-1.0*2*PI*sin(2*PI*t)" />
9 </FUNCTION>

```

the transformation defined by these functions need to be valid (non-zero Jacobian). By default, any component of `COORDS` that is not specified is taken as a trivial transformation, e.g. $\bar{x} = x$, and any velocity not specified is considered to be zero.

11.6.5 Boundary conditions

In case of a time-dependent mapping, a moving body boundary condition is available:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" USERDEFINEDTYPE="MovingBody" VALUE="0" />
4     <D VAR="v" USERDEFINEDTYPE="MovingBody" VALUE="0" />
5     <D VAR="w" VALUE="0" />
6     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
7   </REGION>
8 </BOUNDARYCONDITIONS>

```

When using the `MovingBody` boundary condition, the Dirichlet condition is relative to the boundary, while the regular Dirichlet boundary condition is taken in an absolute sense.

All Dirichlet boundary conditions are specified in the Cartesian (physical) space, and are automatically transformed to the computational frame of reference.

Note



Examples of the use of mappings can be found in the test directory `NEKTAR/solvers/IncNavierStokesSolver/Tests`. See for example the files `KovaFlow_3DH1D_P8_16modes_Mapping-implicit.xml` and `CylFlow_Mov_mapping.xml`.

11.7 Session file configuration: Adaptive polynomial order

An adaptive polynomial order procedure is available for 2D and Quasi-3D simulations. This procedure consists of the following steps:

- Advance the equations for a determined number of time steps
- Use the sensor defined in equation 9.9 to estimate an error measure (the variable used in the sensor can be specified). The error is defined here as the square of the sensor.
- Use the error to determine if the order in each element should be increased by one, decreased by one, or left unaltered.
- Project the solution in each element to the new polynomial order and use it as an initial condition to restart the equation, repeating all steps a given number of times.

It is important to note that restarting the simulation after the refinement can be an expensive operation (in a typical case 200 times the cost of a single time step). Therefore, the number of steps between successive refinements needs to be carefully chosen, since if this value is too low the procedure becomes inefficient, while if it is too high the refinement might not capture accurately structures that are convected by the flow.

11.7.1 Solver Info

The use of the adaptive polynomial order procedure is enforced through the definition of the `Driver` which has to be `Adaptive`.

11.7.2 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file:

- `NumSteps`: when using the adaptive order procedure, this parameter determines the number of time steps between successive refinements.
- `NumRuns`: this parameter defines the number of times the sequence of solving the equation and refining is performed. Therefore, the total number of time steps in the simulation is $NumSteps \times NumRuns$.
- `AdaptiveMaxModes`: sets the maximum number of modes (in each direction) that can be used in an element during the adaptive procedure. The solution will not be refined beyond this point, even if the error is higher than the tolerance. Default value: 12.
- `AdaptiveMinModes`: sets the minimal number of modes (in each direction) that can be used in an element during the adaptive procedure. Default value: 4.
- `AdaptiveUpperTolerance`: defines a constant tolerance. The polynomial order in an element is increased whenever the error is higher than this value. This can be replaced by a spatially-varying function, as described below. Default value: 10^{-6} .
- `AdaptiveLowerinModes`: defines a constant tolerance. The polynomial order in an element is decreased whenever the error is lower than this value. This can also be replaced by a spatially-varying function. Default value: 10^{-8} .
- `AdaptiveSensorVariable`: integer defining which variable will be considered when calculating the error. For example, if this parameter is set to 1 in the Incompressible Navier-Stokes Solver, the error will be estimated using the v velocity. Default value: 0.

11.7.3 Functions

Spatially varying tolerances can be specified by defining the functions `AdaptiveLowerinModes` and/or `AdaptiveUpperTolerance`. In this case, the tolerance in an element is taken as the average of the function in the quadrature points of the element. If these functions are defined, the values set for the tolerances in the `PARAMETERS` section are ignored.

11.7.4 Restarting the simulation

The simulation can be restarted using the final distribution of polynomial orders obtained from the adaptive procedure by setting the expansions as

```
1 <EXPANSIONS>
2   <F FILE="restartfile.fld" />
3 </EXPANSIONS>
```

note that this will only affect the polynomial order. The initial condition still needs to be set correctly, and does not need to come from the same file used for the expansions.

11.8 Advecting extra passive scalar fields

In some cases, it might be useful to advect passive scalar fields with the velocity obtained from the solution of the Navier-Stokes equation. For example, in study of mass transfer or heat transfer problems where getting analytical expression for advection velocity is not possible, the transport (advection-diffusion) equation needs to be solved along with the Navier-Stokes equation to get the scalar concentration or temperature distribution in the flow field.

In the input file, the extra field variables that are being advected need to be defined after the variables representing the velocity components. The pressure needs to be at the end of the list. For example, for a 2D simulation the expansions and variables would be defined as

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="5" FIELDS="u,v,c1,c2,p" TYPE="MODIFIED" />
3 </EXPANSIONS>
4
5 <VARIABLES>
6   <V ID="0"> u </V>
7   <V ID="1"> v </V>
8   <V ID="2"> c1 </V>
9   <V ID="3"> c2 </V>
10  <V ID="4"> p </V>
11 </VARIABLES>
```

where u , v are the velocity components, $c1$ and $c2$ are extra fields that are being advected and p is the pressure.

In addition, diffusion coefficients for each extra variable can be specified by adding a function `DiffusionCoefficient`

```
1 <FUNCTION NAME="DiffusionCoefficient">
2   <E VAR="c1" VALUE="0.1" />
3   <E VAR="c2" VALUE="0.01" />
4 </FUNCTION>
```

Boundary conditions for the extra fields are set up in the same way as the velocity and pressure

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="0" />
4     <D VAR="v" VALUE="0" />
5     <D VAR="c1" VALUE="1" />
6     <D VAR="c2" VALUE="0" />
7     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
8   </REGION>
9 </BOUNDARYCONDITIONS>

```

It should be noted that if the diffusion coefficient is too small, the transport equation becomes advection dominated. This leads to small grid spacing required to resolve all physical scales associated with the transport equation (the ratio of resolution required for transport to Navier Stokes equation scales with $Sc^{3/4}$, where Sc is the Schmidt number = kinematic viscosity/diffusion coefficient). Hence, small diffusion coefficient might lead to spurious oscillations if the mesh spacing is not small enough.

11.9 Imposing a constant flowrate

In some simulations, it may be desirable to drive the flow by fixing a value of the volumetric flux through a boundary surface. A common use case for this is a channel flow, where the inflow and outflow are treated using periodic boundary conditions, requiring a use of either a body force or a flowrate condition to drive the flow. Often, the use of a body force is sufficient, but in some cases (e.g. transitional or turbulent simulations), it may be difficult to determine the correct body force to use in order to attain a specific Reynolds number. The incompressible solver supports the use of an alternative forcing whereby the volumetric flux,

$$Q(\mathbf{u}) = \frac{1}{\mu(R)} \int_R \mathbf{u} \cdot d\mathbf{s},$$

through a user-defined surface R of area $\mu(R)$ is kept constant. This is supported for standard two- and three-dimensional simulations, where R is a boundary region, as well as three-dimensional homogeneous simulations. In the latter case, the forcing can be imposed either in the homogeneous direction (in the $x - y$ plane) or perpendicular to it (in the z direction).

The flowrate correction works by taking each timestep's velocity field \mathbf{u}^n , and computing a scalar α so that the corrected flow

$$\tilde{\mathbf{u}}^n = \mathbf{u}^n + \alpha \mathbf{u}_s$$

has the desired flowrate. Here, \mathbf{u}_s is a linear Stokes solution that is calculated once at the start of the simulation, so that the condition is not expensive to implement.

To enable flowrate corrections, three things must be defined in the session file:

- The `Flowrate` parameter in the parameters section, which defines the desired value of the volumetric flux $Q(\mathbf{u})$ through the reference region. To set a flux per unit surface of $Q = 1$ we would therefore define:

```
1 <PARAMETERS>
2   <P> Flowrate = 1.0 </P>
3 </PARAMETERS>
```

- A boundary condition must be tagged with the `Flowrate` user-defined type to define the reference region R . For example, a 2D channel flow with periodic boundary conditions might use the following arrangement:

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <P VAR="u" VALUE="[1]" USERDEFINEDTYPE="Flowrate" />
4     <P VAR="v" VALUE="[1]" />
5     <P VAR="p" VALUE="[1]" />
6   </REGION>
7   <REGION REF="1">
8     <P VAR="u" VALUE="[0]" />
9     <P VAR="v" VALUE="[0]" />
10    <P VAR="p" VALUE="[0]" />
11  </REGION>
12 </BOUNDARYCONDITIONS>
```

- a `FlowrateForce` function with components `ForceX`, `ForceY` and `ForceZ` that defines the direction in which the forcing will be applied. This should be a unit vector (i.e. of magnitude 1) and constant (i.e. not dependent on x , y , z or t). As an example, to impose a force in the x -direction we specify:

```
1 <FUNCTION NAME="FlowrateForce">
2   <E VAR="ForceX" VALUE="1.0" />
3   <E VAR="ForceY" VALUE="0.0" />
4   <E VAR="ForceZ" VALUE="0.0" />
5 </FUNCTION>
```

Importantly, note that in homogeneous simulations where the forcing is in the z -direction only the `Flowrate` parameter should be specified, and the reference area R is taken to be the homogeneous plane.

Optionally, the `IO_FlowSteps` parameter can be defined. If set to a non-zero integer, this produces a file `session.prs` which records the value of α used in the flowrate correction, every `IO_FlowSteps` steps.

11.10 Examples

11.10.1 Kovasznay Flow 2D

This example demonstrates the use of the velocity correction `o` solve the 2D Kovasznay flow at Reynolds number $Re = 40$. In the following we will numerically solve for the two dimensional velocity and pressure fields with steady boundary conditions.

11.10.1.1 Input file

The input for this example is given in the example file `KovaFlow_m8.xml`. The mesh consists of 12 quadrilateral elements.

We will use a 7th-order polynomial expansions ($N = 8$ modes) using the modified Legendre basis and therefore require the following expansion definition.

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="6" FIELDS="u,v,p" TYPE="MODIFIED" />
3 </EXPANSIONS>
```

We next specify the solver information for our problem. In particular, we select the velocity correction scheme formulation, using a continuous Galerkin projection. For this scheme, an implicit-explicit time-integration scheme must be used and we choose one of second order.

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE" VALUE="UnsteadyNavierStokes" />
3   <I PROPERTY="SolverType" VALUE="VelocityCorrectionScheme" />
4   <I PROPERTY="AdvectionForm" VALUE="Convective" />
5   <I PROPERTY="Projection" VALUE="Galerkin" />
6   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2" />
7 </SOLVERINFO>
```

The key parameters are listed below. Since the problem is unsteady we prescribe the time step and the total number of time steps. We also know the required Reynolds number, but we must prescribe the kinematic viscosity to the solver. We first define a *dummy* parameter for the Reynolds number, and then define the kinematic viscosity as the inverse of this. The value of λ is used when defining the boundary conditions and exact solution. Note that PI is a pre-defined constant.

```
1 <PARAMETERS>
2   <P> TimeStep = 0.001 </P>
3   <P> NumSteps = 100 </P>
4   <P> Re = 40 </P>
5   <P> Kinvis = 1/Re </P>
6   <P> LAMBDA = 0.5*Re-sqrt(0.25*Re*Re+4*PI*PI) </P>
7 </PARAMETERS>
```

We choose to impose a mixture of boundary condition types as defined below.

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="1-exp(LAMBDA*x)*cos(2*PI*y)" />
4     <D VAR="v" VALUE="(LAMBDA/2/PI)*exp(LAMBDA*x)*sin(2*PI*y)" />
5     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
6   </REGION>
7   <REGION REF="1">
8     <D VAR="u" VALUE="1-exp(LAMBDA*x)*cos(2*PI*y)" />
9     <D VAR="v" VALUE="(LAMBDA/2/PI)*exp(LAMBDA*x)*sin(2*PI*y)" />
10    <D VAR="p" VALUE="0.5*(1-exp(2*LAMBDA*x))" />
```

```

11 </REGION>
12 <REGION REF="2">
13   <N VAR="u" VALUE="0" />
14   <D VAR="v" VALUE="0" />
15   <N VAR="p" VALUE="0" />
16 </REGION>
17 </BOUNDARYCONDITIONS>

```

Initial conditions are obtained from the file `KovaFlow_m8.rst`, which is a *Nektar++* field file. This is the output of an earlier simulation, renamed with the extension `rst` to avoid being overwritten, and is used in this case to reduce the integration time necessary to obtain the steady flow.

```

1 <FUNCTION NAME="InitialConditions">
2   <F FILE="KovaFlow_m8.rst" />
3 </FUNCTION>

```

Note the use of the `F` tag to indicate the use of a file. In contrast, the exact solution is prescribed using analytic expressions which requires the use of the `E` tag.

```

1 <FUNCTION NAME="ExactSolution">
2   <E VAR="u" VALUE="1-exp(LAMBDA*x)*cos(2*PI*y)" />
3   <E VAR="v" VALUE="(LAMBDA/2/PI)*exp(LAMBDA*x)*sin(2*PI*y)" />
4   <E VAR="p" VALUE="0.5*(1-exp(2*LAMBDA*x))" />
5 </FUNCTION>

```

11.10.1.2 Running the simulation

Launch the simulation using the following command

```
IncNavierStokesSolver KovaFlow_m8.xml
```

After completing the prescribed 100 time-steps, the difference between the computed solution and the exact solution will be displayed. The actual mantissas may vary slightly, but the overall magnitude should be as shown.

```

L 2 error (variable u) : 3.75296e-07
L inf error (variable u) : 5.13518e-07
L 2 error (variable v) : 1.68897e-06
L inf error (variable v) : 2.23918e-06
L 2 error (variable p) : 1.46078e-05
L inf error (variable p) : 5.18682e-05

```

The output of the simulation is written to `KovaFlow_m8.fld`. This can be visualised by converting it to a visualisation format. For example, to use ParaView, convert the output into VTK format using the `tility`.

```
FieldConvert KovaFlow.xml KovaFlow.fld KovaFlow.vtu
```

The result should look similar to that shown in Figure 11.3.




Figure 11.3 Velocity profiles for the Kovasznay Flow (2D).

11.10.2 Kovasznay Flow 2D using high-order outflow boundary conditions

In this example, we solve the same case of 2D Kovasznay flow on severely-truncated computational domain but using a high-order outflow boundary condition, which is much more accurate and robust for unbounded flows [11]. The solver information and parameters used here are similar to the previous one. What only we need to modify in the input file is just the boundary condition type upon the outlet region shown as below

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="1-exp(KovLam*x)*cos(2*PI*y)" />
4     <D VAR="v" VALUE="(KovLam/2/PI)*exp(KovLam*x)*sin(2*PI*y)" />
5     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
6   </REGION>
7   <REGION REF="1">
8     <N VAR="u" USERDEFINEDTYPE="HOutflow"
9       VALUE="-Kinvis*KovLam*exp(KovLam*x)*cos(2*PI*y)
10        - 0.5*(1-exp(2*KovLam*x))-0.5*(((1-exp(KovLam*x)*cos(2*PI*y))
11        *(1-exp(KovLam*x)*cos(2*PI*y))+(KovLam/(2*PI)*exp(KovLam*x)
12        *sin(2*PI*y))*(KovLam/(2*PI)*exp(KovLam*x)*sin(2*PI*y))))
13        *(0.5*(1.0-tanh((1-exp(KovLam*x)*cos(2*PI*y))*20))))" />
14     <N VAR="v" USERDEFINEDTYPE="HOutflow"
15       VALUE="Kinvis*KovLam*KovLam/(2*PI)*exp(KovLam*x)*sin(2*PI*y)" />
16     <D VAR="p" USERDEFINEDTYPE="HOutflow"
17       VALUE="-Kinvis*KovLam*exp(KovLam*x)*cos(2*PI*y)
18        - 0.5*(1-exp(2*KovLam*x)) -0.5*(((1-exp(KovLam*x)*cos(2*PI*y))
19        *(1-exp(KovLam*x)*cos(2*PI*y))+(KovLam/(2*PI)*exp(KovLam*x)
20        *sin(2*PI*y))*(KovLam/(2*PI)*exp(KovLam*x)*sin(2*PI*y))))
21        *(0.5*(1.0-tanh((1-exp(KovLam*x)*cos(2*PI*y))*20))))" />
22   </REGION>
23   <REGION REF="2">

```

```

24     <N VAR="u" VALUE="0" />
25     <D VAR="v" VALUE="0" />
26     <N VAR="p" VALUE="0" />
27 </REGION>
28 </BOUNDARYCONDITIONS>

```

We note that in this example the “VALUE” property is set based on the analytic solution but this is not typically known and so often a VALUE of zero will be specified.

Instead of loading an initial condition from a specified file, we initialized the flow fields in this example by using following expressions

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="u" VALUE="(1-exp(KovLam*x)*cos(2*PI*y))" />
3   <E VAR="v" VALUE="(KovLam/(2*PI))*exp(KovLam*x)*sin(2*PI*y)" />
4   <E VAR="p" VALUE="0.5*(1-exp(2*KovLam*x))" />
5 </FUNCTION>

```

11.10.2.1 Running the simulation

We then launch the simulation by the same solver as that in the previous example

```
IncNavierStokesSolver KovaFlow_m8_short_HOBC.xml
```

The solution with errors displayed as below

```

L 2 error (variable u) : 2.51953e-08
L inf error (variable u) : 9.56014e-09
L 2 error (variable v) : 1.10694e-08
L inf error (variable v) : 9.47464e-08
L 2 error (variable p) : 5.59175e-08
L inf error (variable p) : 2.93085e-07

```

The physical solution visualized in velocity profiles is also illustrated in Figure 11.4.




Figure 11.4 Velocity profiles for the Kovasznay Flow in truncated domain (2D).

11.10.3 Steady Kovasznay Oseen Flow using the direct solver

In this example, we instead compute the steady Kovasznay Oseen flow using the direct solver. In contrast to the velocity correction scheme in which we time-step the solution to the final time, the direct solver computes the solution with a single solve.

11.10.3.1 Input file

We can begin with the same input file as for the previous example, but with the following modifications. For reference, the modified version is provided in the example `Oseen_m8.xml`.

In the solver information, we must instead select the Steady-Oseen equation type and choose to use the coupled linearised Navier-Stokes

```
1 <I PROPERTY="EQTYPE"      VALUE="SteadyOseen" />
2 <I PROPERTY="SolverType"  VALUE="CoupledLinearisedNS" />
```

Note



Since we are using a coupled system, we are not solving for the pressure. We should therefore remove all references to the variable `p` in the session. In particular, it should be removed from the `EXPANSIONS`, `VARIABLES`, `BOUNDARYCONDITIONS` and `FUNCTIONS` sections of the file.

Instead of loading an initial condition from file, we can simply prescribe a zero field.

```
1 <FUNCTION NAME="InitialConditions">
2   <E VAR="u" VALUE="0" />
3   <E VAR="v" VALUE="0" />
4 </FUNCTION>
```

We must also provide an advection velocity.

```
1 <FUNCTION NAME="AdvectionVelocity">
2   <E VAR="u" VALUE="(1-exp(-LAMBDA*x))*cos(2*PI*y)" />
3   <E VAR="v" VALUE="(-LAMBDA/(2*PI))*exp(-LAMBDA*x)*sin(2*PI*y)" />
4 </FUNCTION>
```

11.10.3.2 Running the simulation

Run the simulation using

```
IncNavierStokesSolver Oseen_m8.xml
```

The resulting flow field should match the solution from the previous example.

11.10.4 Laminar Channel Flow 2D

In this example, we will simulate the flow through a channel at Reynolds number 1 with fixed boundary conditions.

11.10.4.1 Input file

The input file for this example is given in `ChanFlow_m3_SKS.xml`. The geometry is a square channel with height and length $D = 1$, discretised using four quadrilateral elements. We use a quadratic expansion order, which is sufficient to capture the quadratic flow profile. In this example, we choose to use the skew-symmetric form of the advection term. This is chosen in the solver information section:

```
1 <I PROPERTY="EvolutionOperator" VALUE="SkewSymmetric" />
```

A first-order time integration scheme is used and we set the time-step and number of time integration steps in the parameters section. We also prescribe the kinematic viscosity $\nu = 1/Re = 1$.

Boundary conditions are defined on the walls (region 0) and at the inflow (regions 1) as Dirichlet for the velocity field and as high-order for the pressure. At the outflow the velocity is left free using Neumann boundary conditions and the pressure is pinned to zero.

```
1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="u" VALUE="0" />
4     <D VAR="v" VALUE="0" />
5     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
6   </REGION>
7   <REGION REF="1">
8     <D VAR="u" VALUE="y*(1-y)" />
9     <D VAR="v" VALUE="0" />
10    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
11  </REGION>
12  <REGION REF="2">
13    <N VAR="u" VALUE="0" />
14    <N VAR="v" VALUE="0" />
15    <D VAR="p" VALUE="0" />
16  </REGION>
17 </BOUNDARYCONDITIONS>
```

Initial conditions are set to zero. The exact solution is a parabolic profile with a pressure gradient dependent on the Reynolds number. This is defined to allow verification of the calculation.

```
1 <FUNCTION NAME="ExactSolution">
2   <E VAR="u" VALUE="y*(1-y)" />
3   <E VAR="v" VALUE="0" />
4   <E VAR="p" VALUE="-2*Kinvis*(x-1)" />
5 </FUNCTION>
```

11.10.4.2 Running the solver

```
IncNaverStokesSolver ChanFlow_m3_SKS.xml
```

The error in the solution should be displayed and be close to machine precision

```
L 2 error (variable u) : 4.75179e-16
L inf error (variable u) : 3.30291e-15
L 2 error (variable v) : 1.12523e-16
L inf error (variable v) : 3.32197e-16
L 2 error (variable p) : 1.12766e-14
L inf error (variable p) : 7.77156e-14
```

The solution should look similar to that shown in Figure 11.5.




Figure 11.5 Pressure and velocity profiles for the laminar channel flow (2D).

11.10.5 Laminar Channel Flow 3D

We now solve the incompressible Navier-Stokes equations on a three-dimensional domain. In particular, we solve the three-dimensional equivalent of the previous example. We will also solve the problem in parallel.

**Note**

In order to run the example, you must have a version of *Nektar++* compiled with MPI. This is the case for the packaged binary distributions.

11.10.5.1 Input file

The input file for this example is given in `Tet_channel_m8_par.xml`. In this example we use tetrahedral elements, indicated by the `(A)` element tags in the geometry section. All dimensions have length $D = 1$. We will use a 7th-order polynomial expansion. Since we now have three dimensions, and therefore three velocity components, the expansions section is now

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="8" FIELDS="u,v,w,p" TYPE="MODIFIED" />
3 </EXPANSIONS>
```

The solver information and parameters are similar to the previous example. Boundary conditions must now be defined on the six faces of the domain. Flow is prescribed in the z -direction through imposing a Poiseuille profile on the inlet and side walls. The outlet is zero-Neumann and top and bottom faces impose zero-Dirichlet conditions.

```
1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>      <!-- Inlet -->
3   <B ID="1"> C[6] </B>      <!-- Outlet -->
4   <B ID="2"> C[2] </B>      <!-- Wall -->
5   <B ID="3"> C[3] </B>      <!-- Wall left -->
6   <B ID="4"> C[4] </B>      <!-- Wall -->
7   <B ID="5"> C[5] </B>      <!-- Wall right -->
8 </BOUNDARYREGIONS>
9
10 <BOUNDARYCONDITIONS>
11   <REGION REF="0">
12     <D VAR="u" VALUE="0" />
13     <D VAR="v" VALUE="0" />
14     <D VAR="w" VALUE="y*(1-y)" />
15     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
16   </REGION>
17   <REGION REF="1">
18     <N VAR="u" VALUE="0" />
19     <N VAR="v" VALUE="0" />
20     <N VAR="w" VALUE="0" />
21     <D VAR="p" VALUE="0" />
22   </REGION>
23   <REGION REF="2">
24     <D VAR="u" VALUE="0" />
25     <D VAR="v" VALUE="0" />
26     <D VAR="w" VALUE="0" />
27     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
28   </REGION>
29   <REGION REF="3">
30     <D VAR="u" VALUE="0" />
31     <D VAR="v" VALUE="0" />
32     <D VAR="w" VALUE="y*(1-y)" />
33     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
34   </REGION>
35   <REGION REF="4">
36     <D VAR="u" VALUE="0" />
37     <D VAR="v" VALUE="0" />
```

```

38     <D VAR="w" VALUE="0" />
39     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
40 </REGION>
41 <REGION REF="5">
42     <D VAR="u" VALUE="0" />
43     <D VAR="v" VALUE="0" />
44     <D VAR="w" VALUE="y*(1-y)" />
45     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
46 </REGION>
47 </BOUNDARYCONDITIONS>

```

Initial conditions and exact solutions are also prescribed.

11.10.5.2 Running the solver

To run the solver in parallel, we use the `mpirun` command.

```
mpirun -np 2 IncNaverStokesSolver Tet_channel_m8_par.xml
```

The expected results are shown in Figure 11.6.




Figure 11.6 Pressure and velocity profiles for the laminar channel flow (full 3D).

11.10.6 Laminar Channel Flow Quasi-3D

For domains where at least one direction is geometrically homogeneous, a more efficient discretisation is to use a pure spectral discretisation, such as a Fourier expansion, in these directions. We use this approach to solve the same problem as in the previous example. We reuse the two-dimensional spectral/hp element mesh from the `nd` couple this with a Fourier expansion in the third component.

11.10.6.1 Input file

The input file for this example is `ChanFlow_3DH1D_MVM.xml`. We indicate that we are coupling the spectral/hp element domain with a pure spectral expansion using the following solver information

```
1 <I PROPERTY="HOMOGENEOUS" VALUE="1D"/>
```

We must also specify parameters to describe the particular spectral expansion

```
1 <P> HomModesZ      = 20      </P>
2 <P> LZ              = 1.0     </P>
```

The parameter `HomModesZ` specifies the number of Fourier modes to use in the homogeneous direction. The `LZ` parameter specifies the physical length of the domain in that direction.

Note



This example uses an in-built Fourier transform routine. Alternatively, one can use the FFTW library to perform Fourier transforms which typically offers improved performance. This is enabled using the following solver information

```
1 <I PROPERTY="USEFFT" VALUE="FFTW"/>
```

As with the spectral/hp element mesh consists of four quadrilateral elements with a second-order polynomial expansion. Since our domain is three-dimensional we have to now include the third velocity component

```
1 <E COMPOSITE="C[0]" NUMMODES="3" FIELDS="u,v,w,p" TYPE="MODIFIED" />
```

The remaining parameters and solver information is similar to previous examples.

Boundary conditions are specified as for the two-dimensional case (except with the addition of the third velocity component) since the side walls are now implicitly periodic. The initial conditions and exact solution are prescribed as for the fully three-dimensional case.

11.10.6.2 Running the solver

```
IncNaverStokesSolver ChanFlow_3DH1D_MVM.xml
```

The results can be post-processed and should match those of the fully three-dimensional case as shown in Figure 11.6.

11.10.7 Turbulent Channel Flow

In this example we model turbulence in a three-dimensional square channel at a Reynolds number of 2000.

**Note**

This example requires the FFTW Fast-Fourier transform library to be selected when compiling Nektar++.

11.10.7.1 Input file

The input file for this example is `TurbChF1_3DH1D.xml`. The geometry makes use of the homogeneous extension discussed in the previous example. The channel has height $D = 2$ and length $L = 4\pi$ and is discretised using quadratic quadrilateral elements in the spectral/hp element plane and a Fourier basis in the third coordinate direction. The elements are non-uniformly distributed so as to best capture the flow features with fewest degrees of freedom and is shown in Figure 11.7.




Figure 11.7 Mesh used for the turbulent channel flow.

The spanwise length of the channel is set using the `LZ` parameter and discretised with 32 Fourier modes by setting the value of `HomModesZ`.

```
1 <P> HomModesZ      = 32          </P>
2 <P> LZ              = 4*PI/3     </P>
```

A second-order IMEX scheme is used for time-integration scheme is used with a time-step of 0.0001. The length of the simulation is 1 time-unit (10,000 steps).

Periodicity is naturally enforced in the spanwise direction, so boundary conditions need only be provided for the upper and lower walls, inlet and outlet as denoted by the following `BOUNDARYREGIONS`.

```
1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B> //walls
3   <B ID="1"> C[2] </B> //inflow
4   <B ID="2"> C[3] </B> //outflow
5 </BOUNDARYREGIONS>
```

In this example, we will use a body force to drive the flow and so, in addition to the spanwise periodicity, enforce periodicity in the streamwise direction of the spectral/hp element mesh. This is achieved by imposing the following boundary conditions

```
1   <REGION REF="1">
2     <P VAR="u" VALUE="[2]" />
```

```

3  <P VAR="v" VALUE="[2]" />
4  <P VAR="w" VALUE="[2]" />
5  <P VAR="p" VALUE="[2]" />
6  </REGION>
7  <REGION REF="2">
8  <P VAR="u" VALUE="[1]" />
9  <P VAR="v" VALUE="[1]" />
10 <P VAR="w" VALUE="[1]" />
11 <P VAR="p" VALUE="[1]" />
12 </REGION>

```

Here, we use `P` to denote the boundary type is periodic, and the value in square brackets denotes the boundary region to which the given boundary is periodic with. In this case regions 1 and 2 are denoted periodic with each other.

A streamwise plug-profile initial condition is prescribed such that $u = 1$ everywhere, except the wall boundaries. The body force requires two components in the XML file. The first specifies the type of forcing to apply and appears directly within the `NEKTAR` tag.

```

1 <FORCING>
2   <FORCE TYPE="Body">
3     <BODYFORCE> BodyForce </BODYFORCE>
4   </FORCE>
5 </FORCING>

```

The second defines the `BodyForce` function which will be used and is located within the `CONDITIONS` section,

```

1 <FUNCTION NAME="BodyForce">
2   <E VAR="u" VALUE="2*Kinvis" />
3   <E VAR="v" VALUE="0" />
4   <E VAR="w" VALUE="0" />
5 </FUNCTION>

```

To improve numerical stability, we also enable dealiasing of the advection term. This uses additional points to perform the quadrature and then truncates the higher-order terms when projecting back onto the polynomial space, thereby removing spurious oscillations. It is enabled by setting the solver information tag

```

1 <I PROPERTY="DEALIASING" VALUE="True" />

```

This feature is only available when using the FFTW library is used, so we enable this using

```

1 <I PROPERTY="USEFFT" VALUE="FFTW" />

```

11.10.7.2 Running the solver

To run the solver, we use the following command

```
IncNaverStokesSolver TurbChFl_3DH1D.xml
```

The result after transition has occurred is illustrated in Figure 11.8.




Figure 11.8 Velocity profile of the turbulent channel flow (quasi-3D).

11.10.8 Turbulent Pipe Flow

In this example we simulate flow in a pipe at Reynolds number 3000 using a mixed spectral/hp element and Fourier discretisation. The Fourier expansion is used in the streamwise direction in this case and the spectral/hp elements are used to capture the circular cross-section.

11.10.8.1 Input File

The circular pipe has diameter $D = 1$, the 2D mesh is composed of 64 quadrilateral elements and the streamwise direction is discretised with 128 Fourier modes. An illustrative diagram of the discretisation is given in Figure 11.9.

The input file for this example is `Pipe_turb.xml`. We use 7th-order lagrange polynomials through the Gauss-Lobatto-Legendre points for the quadrilateral expansions.

```
1 <E COMPOSITE="C[0]" NUMMODES="8" FIELDS="u,v,w,p" TYPE="GLL_LAGRANGE_SEM" />
```

We set the Fourier options, as in the previous example, except using 128 modes and a length of 5 non-dimensional units. A small amplitude noise is also added to the initial condition, which is a plug profile, to help stimulate transition. Since the streamwise direction is the Fourier direction, we must necessarily use a body force to drive the flow.




Figure 11.9 Domain for the turbulent pipe flow problem.

11.10.8.2 Running the solver

In this example we will run the solver in parallel. Due to the large number of Fourier modes and relatively few elements, it is more efficient to parallelise in the streamwise direction. We can specify this by providing an additional flag to the solver, `-npz`. This indicates the number of partitions in the z -coordinate. In this example, we will only run two processes. We therefore would specify `-npz 2` to ensure parallelisation only occurs in the Fourier direction.

To improve the efficiency of the solver further, we would prefer to solve the Helmholtz problems within the spectral/hp element planes using a direct solver (since no parallelisation is necessary). The default when running in parallel is to use an iterative solver, so we explicitly specify the type of algorithm to use in the session file solver information:

```
1 <I PROPERTY="GlobalSysSoln" VALUE="DirectStaticCond" />
```

The solver can now be run as follows

```
mpirun -np 2 IncNavierStokesSolver --npz 2 Pipe_turb.xml
```

When the pipe transitions, the result should look similar to Figure 11.10.

11.10.9 Aortic Blood Flow

The following example demonstrates the application of the [incompressible Navier-Stokes solver](#) using the [Velocity Correction Scheme](#) algorithm for modelling viscous Newtonian blood flow in a region of a rabbit descending thoracic aorta with intercostal branch pairs. Such studies are necessary to understand the effect local blood flow changes have on cardiovascular diseases such as atherosclerosis.




Figure 11.10 Velocity profile of the turbulent pipe flow (quasi-3D).

In the following we will numerically solve for the three dimensional velocity and pressure field for steady boundary conditions. The Reynolds number under consideration is 300, which is physiologically relevant.

Geometry

The geometry under consideration is a segment of a rabbit descending aorta with two pairs of intercostal arteries branching off. The inlet has a diameter $D = 3.32mm$.




Figure 11.11 Reduced region of rabbit descending thoracic aorta.

In order to capture the physics of the flow in the boundary layer, a thin layer consisting of prismatic elements is created adjacent to the surface, and curved using spherigons. The interior consist of tetrahedral elements.

Input parameters




Figure 11.12 Surface mesh indicating curved surface elements at a branch location.

11.10.9.1 Expansion:

In this example we will use a fourth order polynomial expansion. There are two composites defined here since we have both prismatic and tetrahedral elements.

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="5" TYPE="MODIFIED" FIELDS="u,v,w,p" />
3   <E COMPOSITE="C[1]" NUMMODES="5" TYPE="MODIFIED" FIELDS="u,v,w,p" />
4 </EXPANSIONS>

```

11.10.9.2 Solver information:

```

1 <SOLVERINFO>
2   <I PROPERTY="SolverType"           VALUE="VelocityCorrectionScheme" />
3   <I PROPERTY="EQTYPE"               VALUE="UnsteadyNavierStokes" />
4   <I PROPERTY="AdvectionForm"        VALUE="Convective" />
5   <I PROPERTY="Projection"           VALUE="Galerkin" />
6   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder1" />
7 </SOLVERINFO>

```

11.10.9.3 Parameters:

Since we are prescribing a Reynolds number of 300, and to simplify the problem definition, we set the mean inlet velocity to 1, this allows us to define the kinematic viscosity as $\nu = \frac{UD}{Re} = \frac{3.32}{300} = 1/90.36$.

```

1 <PARAMETERS>
2   <P> TimeStep      = 0.0005    </P>
3   <P> NumSteps      = 1600      </P>
4   <P> IO_CheckSteps = 200       </P>
5   <P> IO_InfoSteps  = 50        </P>
6   <P> Kinvis        = 1.0/90.36 </P>
7 </PARAMETERS>

```

11.10.9.4 Boundary conditions:

For the purpose of this example a blunted inlet velocity profile has been prescribed. Ideally to obtain more significant results, the velocity profile at the inlet would be obtained from previous simulations on the complete rabbit aorta (including aortic root, aortic arch, and descending aorta with all 5 pairs of intercostal arteries), where a blunted profile at the aortic root is a better representation of reality.

Dirichlet boundary conditions are imposed for the velocity at the inlet, as well as on the wall to account for the no-slip condition. Neumann boundary conditions are imposed for the velocity field at the outlets where fully developed flow is imposed.

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[2] </B>           <!-- Inlet -->
3   <B ID="1"> C[3,4,5,6] </B>      <!-- intercostal outlets -->
4   <B ID="2"> C[7] </B>           <!-- outlet -->
5   <B ID="3"> C[8] </B>           <!-- wall -->
6 </BOUNDARYREGIONS>
7
8 <BOUNDARYCONDITIONS>
9   <REGION REF="0">
10    <D VAR="u" VALUE="0.024" />
11    <D VAR="v" VALUE="-0.064" />
12    <D VAR="w" VALUE="-0.998" />
13    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
14  </REGION>
15  <REGION REF="1">
16    <N VAR="u" VALUE="0" />
17    <N VAR="v" VALUE="0" />
18    <N VAR="w" VALUE="0" />
19    <D VAR="p" VALUE="0" />
20  </REGION>
21  <REGION REF="2">
22    <N VAR="u" VALUE="0" />
23    <N VAR="v" VALUE="0" />
24    <N VAR="w" VALUE="0" />
25    <D VAR="p" VALUE="0" />
26  </REGION>
27  <REGION REF="3">
28    <D VAR="u" VALUE="0" />
29    <D VAR="v" VALUE="0" />
30    <D VAR="w" VALUE="0" />
31    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
32  </REGION>
33 </BOUNDARYCONDITIONS>

```

11.10.9.5 Functions:

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="u" VALUE="0" />
3   <E VAR="v" VALUE="0" />
4   <E VAR="w" VALUE="0" />
5   <E VAR="p" VALUE="0" />
6 </FUNCTION>

```

11.10.9.6 Results

We can visualise the internal velocity field by applying a volume render filter in ParaView.




Figure 11.13 The solved-for velocity field.

It is possible to visualise the wall shear stress distribution by running the FldAddWSS utility.




Figure 11.14 Non-dimensional wall shear stress distribution.

11.10.10 finite-strip modeling of flow past flexible cables

As a computationally efficient model, strip theory-based modeling technique has been proposed previously to predict vortex-induced vibration (VIV) for higher Reynolds number flows. In the strip theory-based model, the fluid flow solution is obtained on a series of 2D computational planes (also called as “strips”) along the riser’s axis direction. These strips then are coupled with each other through structural dynamic model of the riser, and then VIV response prediction is achieved by the strip-structure interactions. In the 2D strip theory, it is assumed that the flow is purely two-dimensional without spanwise correlation, which allows the problems to be split into various 2D planes. A consequence

of 2D strip solution under this assumption is that it is unable to reflect the influence of spanwise wake turbulence on the structural dynamics. In order to overcome this shortcoming, we proposed a new module in the framework of Nektar++, in which a spanwise scale is locally allocated to each one of the strips, so that the spanwise velocity correlation is reconstructed in the flow field within each strips. In particular, this model lets the fluid domain to be divided in N strips with thickness ratio of L_z/D and evenly distributed along the spanwise (z) direction. The gap between the neighboring strips, represented by L_g , satisfies relation $L_c = N(L_z + L_g)$. Since the strip in this model has finite scale in the z -direction, we named it as finite strip to distinguish from traditional 2D strip plane. Next, the flow dynamics within each individual strips are modeled by viscous incompressible Navier-Stokes equations, while a tensioned beam model is employed to govern the dynamics of the flexible structures. In this example, we will show how to perform a finite-strip model to predict the vortex-induced vibration responses of flexible cables. Let us consider a vortex-induced vibration of a slender cable with an aspect ratio of $L_z/D=4\pi$, which is immersed in uniform flows at $\text{Re}=100$.

11.10.10.1 Input File

The cable with a mass ratio (defined as the ratio of the total oscillating mass to the mass of displaced fluid) of 1 has diameter $D = 1$, the 2D mesh is composed of 284 quadrilateral elements. The spanwise direction is split in 16 strips with thickness ratio of $L_c/D=\pi/8$ and one pair of complex Fourier modes for each one of the strips. We will use a sixth order polynomial expansion for the spectral element and the input file for this example is `CylFlow_HomoStrip.xml`.

```
1 <E COMPOSITE="C[73]" NUMMODES="6" TYPE="MODIFIED" FIELDS="u,v,w,p" />
```

To use the finite strip routines we need just to insert a flag of `"HomoStrip"` in the solver information as below, in addition, we need to specify the types of vibration and support ends for the cables. In this case, the vibration type is specified as `VALUE="CONSTRAINED"`, which means that the cable's vibration is constrained only in the crossflow direction. Other options include `VALUE="FREE"` and `"FORCED"`, respectively corresponding to the free vibrations in both streamwise and crossflow directions and forced vibration by specified functions given in input file. For the support ends of the cable, another option of `VALUE="PINNED-PINNED"` is available for the simulations, which satisfies the condition of zero values of displacements on the support ends.

11.10.10.2 Solver information:

```
1 <SOLVERINFO>
2   <I PROPERTY="HomoStrip" VALUE="True" />
3   <I PROPERTY="VibrationType" VALUE="CONSTRAINED" />
4   <I PROPERTY="SupportType" VALUE="FREE-FREE" />
5 </SOLVERINFO>
```

11.10.10.3 Parameters

All the simulation parameters are specified in the section as follows.

```

1  <PARAMETERS>
2    <P> LZ      = PI/8      </P> <!--thickness ratio-->
3    <P> LC      = 4*PI      </P> <!--aspect ratio-->
4    <P> A       = 0.025     </P>
5    <P> omega   = 1.0       </P>
6    <P> PROC_Z  = 16        </P>
7    <P> Strip_Z = 16        </P> <!--number of the strips-->
8    <P> DistStrip = PI/4    </P> <!--distance of the strips-->
9    <P> StructStiff = 0.02  </P>
10   <P> StructRho = 2.0     </P>
11   <P> CableTension = 8.82 </P>
12   <P> BendingStiff = 0.0  </P>
13   <P> FictDamp   = 0.0    </P>
14   <P> FictMass   = 3.0    </P>
15 </PARAMETERS>
16

```

11.10.10.4 Running the solver

In this example we will run the solver in parallel. We can specify the number of the strips by providing an additional flag to the solver, `-nsz`. In this example, we will run 16 strips, therefore it would be specified as `-nsz 16`. The solver can now be run as follows

```

mpirun -np 16 IncNavierStokesSolver CylFlow_HomoStrip.xml --npz 16 --nsz
16

```

The simulation results are illustrated in spanwise vorticity contours in Figure 11.15. The wake response of the cable appears as standing wave pattern in the earlier stage and then it transitions into travelling wave response, as shown in this figure.




Figure 11.15 Spanwise vorticity contours in standing wave and travelling wave patterns predicted in finite strip modeling.

11.10.11 2D direct stability analysis of the channel flow

In this example, it will be illustrated how to perform a direct stability analysis using Nektar++. Let us consider a canonical stability problem, the flow in a channel which is confined in the wall-normal direction between two infinite plates (Poiseuille flow) at Reynolds number 7500. This problem is a particular case of the stability solver for the IncNavierStokesSolver.

11.10.11.1 Background

We consider the linearised Incompressible Navier-Stokes equations:

$$\frac{\partial \mathbf{u}'}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{u}' + \mathbf{u}' \cdot \nabla \mathbf{U} = -\nabla p + \nu \nabla^2 \mathbf{u}' + \mathbf{f} \quad (11.27a)$$

$$\nabla \cdot \mathbf{u}' = 0 \quad (11.27b)$$

We are interested to compute the leading eigenvalue of the system using the Arnoldi method.

11.10.11.2 Geometry

The geometry under consideration is a 2D channel.

11.10.11.3 Mesh Definition

In the **GEOMETRY** section, the dimensions of the problem are defined. Then, the coordinates (**XSCALE**, **YSCALE**, **ZSCALE**) of each vertices of each element are specified. As this input file defines a two-dimensional problem: **ZSCALE** = 0.

```
1 <GEOMETRY DIM="2" SPACE="2">
2   <VERTEX>
3     <V ID="0">3.142e+00 1.000e+00 0.000e+00</V>
4     ...
5     <V ID="62">-3.142e+00 -1.000e+00 0.000e+00</V>
6   </VERTEX>
```

Edges can now be defined by two vertices.

```
1 <EDGE>
2   <E ID="0"> 0 1 </E>
3   ...
4   <E ID="109"> 62 55 </E>
5 </EDGE>
```

In the **ELEMENT** section, the tag **T** and **Q** define respectively triangular and quadrilateral element. Triangular elements are defined by a sequence of three edges and quadrilateral elements by a sequence of four edges.


```

1      <ELEMENT>
2          <Q ID="0">    0    1    2    3 </Q>
3          ...
4          <Q ID="47"> 107   108   109   95 </Q>
5      </ELEMENT>
6

```

Finally, collections of elements are listed in the **COMPOSITE** section and the **DOMAIN** section specifies that the mesh is composed by all the triangular and quadrilateral elements. The other composites will be used to enforce boundary conditions.

```

1      <COMPOSITE>
2          <C ID="0"> Q[0-47] </C>
3          <C ID="1"> E[17,31,44,57,70,83,96,109,0,19,32,45,58,71,84,97] </C> //
wall
4          <C ID="2"> E[3,6,9,12,15,18] </C> //inflow
5          <C ID="3"> E[98,100,102,104,106,108] </C> //outflow
6      </COMPOSITE>
7      <DOMAIN> C[0] </DOMAIN>
8 </GEOMETRY>
9

```

11.10.11.4 Expansion

This section defines the polynomial expansions used on each composites. For this example we will use a 10th order polynomial, i.e. $P = 11$.

```

1 <EXPANSIONS>
2     <E COMPOSITE="C[0]" NUMMODES="11" FIELDS="u,v,p" TYPE="MODIFIED" />
3 </EXPANSIONS>
4

```

11.10.11.5 Solver Info

In this example the **EvolutionOperator** must be **Direct** to consider the linearised Navier-Stokes equations and the **Driver** was set up to **ModifiedArnoldi** for the solution of the eigenproblem.

```

1 <SOLVERINFO>
2     <I PROPERTY="SolverType"          VALUE="VelocityCorrectionScheme" />
3     <I PROPERTY="EQTYPE"              VALUE="UnsteadyNavierStokes" />
4     <I PROPERTY="EvolutionOperator"    VALUE="Direct" />
5     <I PROPERTY="Projection"          VALUE="Galerkin" />
6     <I PROPERTY="Driver"               VALUE="ModifiedArnoldi" />
7     <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder1" />
8 </SOLVERINFO>
9

```

11.10.11.6 Parameters

All the stability parameters are specified in this section.

```

1 <PARAMETERS>
2   <P> TimeStep      = 0.002 </P>
3   <P> NumSteps      = 500 </P>
4   <P> IO_CheckSteps = 1000 </P>
5   <P> IO_InfoSteps  = 10 </P>
6   <P> Re            = 7500 </P>
7   <P> Kinvis        = 1./Re </P>
8   <P> kdim          = 16 </P>
9   <P> nvec          = 2 </P>
10  <P> evtol         = 1e-5 </P>
11  <P> nits          = 5000 </P>
12 </PARAMETERS>
13

```

11.10.11.7 Boundary Conditions

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[2] </B>
4   <B ID="2"> C[3] </B>
5 </BOUNDARYREGIONS>
6
7 <BOUNDARYCONDITIONS>
8   <REGION REF="0">
9     <D VAR="u" VALUE="0" />
10    <D VAR="v" VALUE="0" />
11    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
12  </REGION>
13  <REGION REF="1">
14    <P VAR="u" VALUE="[2]" />
15    <P VAR="v" VALUE="[2]" />
16    <P VAR="p" VALUE="[2]" />
17  </REGION>
18  <REGION REF="2">
19    <P VAR="u" VALUE="[1]" />
20    <P VAR="v" VALUE="[1]" />
21    <P VAR="p" VALUE="[1]" />
22  </REGION>
23 </BOUNDARYCONDITIONS>
24

```

11.10.11.8 Function

We need to set up the base flow that can be specified as a function **BaseFlow**. In case the base flow is not analytical, it can be generated by means of the **Nonlinear** evolution operator using the same mesh and polynomial expansion. The initial guess is specified in the **InitialConditions** functions and can be both analytical or a file. In this example it is read from a file.

```

1 <FUNCTION NAME="BaseFlow">
2   <F VAR="u,v,p" FILE="ChanStability.bse" />
3 </FUNCTION>
4

```

```

5      <FUNCTION NAME="InitialConditions">
6          <F VAR="u,v,p" FILE="ChanStability.rst" />
7      </FUNCTION>
8

```

11.10.11.9 Usage

IncNavierStokesSolver ChanStability.xml

11.10.11.10 Results

The stability simulation takes about 250 iterations to converge and the dominant eigenvalues (together with the respective eigenvectors) will be printed. In this case it was found $\lambda_{1,2} = 1.000224 \times e^{\pm 0.24984i}$. Therefore, since the magnitude of the eigenvalue is larger than 1, the flow is absolutely unstable. It is possible to visualise the eigenvectors using the post-processing utilities. The figure shows the profile of the two eigenmode component, which shows the typical Tollmien - Schlichting waves that arise in viscous boundary layers.




Figure 11.16

11.10.12 2D adjoint stability analysis of the channel flow

In this example, it will be illustrated how to perform an adjoint stability analysis using Nektar++. Let us consider a canonical stability problem, the flow in a channel which is confined in the wall-normal direction between two infinite plates (Poiseuille flow) at Reynolds number 7500




Figure 11.17

11.10.12.1 Background

We consider the equations:

$$-\frac{\partial \mathbf{u}^*}{\partial t} + (\mathbf{U} \cdot \nabla) \mathbf{u}^* + (\nabla \mathbf{U})^T \cdot \mathbf{u}^* = -\nabla p^* + \frac{1}{Re} \nabla^2 \mathbf{u} \quad (11.28a)$$

$$\nabla \cdot \mathbf{u}^* = 0 \quad (11.28b)$$

We are interested in computing the leading eigenvalue of the system using the Arnoldi method.

11.10.12.2 Geometry & Mesh

The geometry and mesh are the same ones used for the direct stability analysis in the previous example.

11.10.12.3 Solver Info

This section defines the problem solved. In this example the `EvolutionOperator` must be `Adjoint` to consider the adjoint Navier-Stokes equations and the `Driver` was set up to `ModifiedArnoldi` for the solution of the eigenproblem.

```

1 <SOLVERINFO>
2   <I PROPERTY="SolverType"           VALUE="VelocityCorrectionScheme" />
3   <I PROPERTY="EQTYPE"               VALUE="UnsteadyNavierStokes" />
4   <I PROPERTY="EvolutionOperator"    VALUE="Adjoint" />
5   <I PROPERTY="Projection"           VALUE="Galerkin" />
6   <I PROPERTY="Driver"               VALUE="ModifiedArnoldi" />
7   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder1" />
8 </SOLVERINFO>

```

```

9  \end{subequations}
10
11 \textbf{Parameters}
12
13 \begin{lstlisting}[style=XMLStyle]
14 <PARAMETERS>
15   <P> TimeStep      = 0.002 </P>
16   <P> NumSteps      = 500 </P>
17   <P> IO_CheckSteps = 1000 </P>
18   <P> IO_InfoSteps  = 10 </P>
19   <P> Re            = 7500 </P>
20   <P> Kinvis        = 1./Re </P>
21   <P> kdim          = 16 </P>
22   <P> nvec          = 2 </P>
23   <P> evtol         = 1e-5 </P>
24   <P> nits          = 5000 </P>
25 </PARAMETERS>
26

```

11.10.12.4 Boundary Conditions

```

1  <BOUNDARYREGIONS>
2    <B ID="0"> C[1] </B>
3    <B ID="1"> C[2] </B>
4    <B ID="2"> C[3] </B>
5  </BOUNDARYREGIONS>
6
7  <BOUNDARYCONDITIONS>
8    <REGION REF="0">
9      <D VAR="u" VALUE="0" />
10     <D VAR="v" VALUE="0" />
11     <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
12   </REGION>
13   <REGION REF="1">
14     <P VAR="u" VALUE="[2]" />
15     <P VAR="v" VALUE="[2]" />
16     <P VAR="p" VALUE="[2]" />
17   </REGION>
18   <REGION REF="2">
19     <P VAR="u" VALUE="[1]" />
20     <P VAR="v" VALUE="[1]" />
21     <P VAR="p" VALUE="[1]" />
22   </REGION>
23 </BOUNDARYCONDITIONS>
24

```

11.10.12.5 Functions

We need to set up the base flow that can be specified as a function `BaseFlow`. In case the base flow is not analytical, it can be generated by means of the `Nonlinear` evolution operator using the same mesh and polynomial expansion.

```

1  <FUNCTION NAME="BaseFlow">
2    <F VAR="u,v,p" FILE="ChanStability.bse" />

```

```

3      </FUNCTION>
4

```

The initial guess is specified in the `InitialConditions` functions and can be both analytical or a file. In this example it is read from a file.

```

1      <FUNCTION NAME="InitialConditions">
2          <F VAR="u,v,p" FILE="ChanStability.rst" />
3      </FUNCTION>
4

```

11.10.12.6 Usage

`IncNavierStokesSolver ChanStability_adj.xml`

11.10.12.7 Results

The equations will then be evolved backwards in time (consistently with the negative sign in front of the time derivative) and the leading eigenvalues will be computed after about 300 iterations. It is interesting to note that their value is the same one computed for the direct problem, but the eigenmodes present a different shape.




Figure 11.18

11.10.13 2D Transient Growth analysis of a flow past a backward-facing step

In this section it will be described how to perform a transient growth stability analysis using Nektar++. Let us consider a flow past a backward-facing step (figure 11.20). This is an important case because it allows us to understand the effects of separation caused by abrupt changes in the geometry and it is a common geometry in several studies of flow control and turbulence of separated flow.




Figure 11.19




Figure 11.20

11.10.13.1 Background

Transient growth analysis allows us to study the presence of convective instabilities that can arise in stable flows. Despite the fact that these instabilities will decay for a long time (due to the stability of the flow), they can produce significant increases in the energy of perturbations. The phenomenon of transient growth is associated with the non-normality of the linearised Navier-Stokes equations and it consists in computing the perturbation that leads to the highest energy growth for a fixed time horizon.

11.10.13.2 Input Parameters

In the **GEOMETRY** section, the dimensions of the problem are defined. Then, the coordinates (**XSCALE**, **YSCALE**, **ZSCALE**) of each vertices are specified. As this input file defines a two-dimensional problem: **ZSCALE** = 0.

```

1 <GEOMETRY DIM="2" SPACE="2">
2   <VERTEX>
3     <V ID="0">3.000e+00 -1.000e+00 0.000e+00</V>
4     ...
5     <V ID="399">-1.000e+01 0.000e+00 0.000e+00</V>
6   </VERTEX>

```

7

Edges can now be defined by two vertices.

```

1 <EDGE>
2   <E ID="0">    0  1  </E>
3   ...
4   <E ID="828"> 399 394 </E>
5 </EDGE>
6

```

In the **ELEMENT** section, the tag **T** and **Q** define respectively triangular and quadrilateral element. Triangular elements are defined by a sequence of three edges and quadrilateral elements by a sequence of four edges.

```

1 <ELEMENT>
2   <T ID="0">    0    1    2 </T>
3   ...
4   <T ID="209"> 333   314   332 </T>
5   <Q ID="210"> 334   335   336    0 </Q>
6   ...
7   <Q ID="429"> 826   827   828   818 </Q>
8 </ELEMENT>
9

```

Finally, collections of elements are listed in the **COMPOSITE** section and the **DOMAIN** section specifies that the mesh is composed by all the triangular and quadrilateral elements. The other composites will be used to enforce boundary conditions.

```

1 <COMPOSITE>
2   <C ID="0"> T[0-209] </C>
3   <C ID="1"> Q[210-429] </C>
4   <C ID="2"> E[2-3,7,10,16,21,2,...,828] </C>
5   <C ID="3"> E[821,823,825,827] </C>
6   <C ID="4"> E[722,724,726,728] </C>
7 </COMPOSITE>
8
9   <DOMAIN> C[0,1] </DOMAIN>
10 </GEOMETRY>
11

```

11.10.13.3 Expansion

For this example we will use a 6th order polynomial, i.e. $P = 7$:

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="7" FIELDS="u,v,p" TYPE="MODIFIED" />
3   <E COMPOSITE="C[1]" NUMMODES="7" FIELDS="u,v,p" TYPE="MODIFIED" />
4 </EXPANSIONS>
5

```


11.10.13.4 Solver Information

This section defines the problem solved. In this example the `EvolutionOperator` must be `TransientGrowth` and the `Driver` was set up to `Arpack` for the solution of the eigenproblem.

```

1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE"           VALUE="UnsteadyNavierStokes"  />
3   <I PROPERTY="EvolutionOperator" VALUE="TransientGrowth"      />
4   <I PROPERTY="Projection"        VALUE="Galerkin"              />
5   <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2"         />
6   <I PROPERTY="SOLVERTYPE"        VALUE="VelocityCorrectionScheme" />
7   <I PROPERTY="Driver"            VALUE="Arpack"                />
8   <I PROPERTY="ArpackProblemType" VALUE="LargestMag"           />
9 </SOLVERINFO>
10

```

11.10.13.5 Parameters

```

1 <PARAMETERS>
2   <P> FinalTime      = 0.1                </P>
3   <P> TimeStep       = 0.005              </P>
4   <P> NumSteps       = FinalTime/TimeStep </P>
5   <P> IO_CheckSteps  = 1/TimeStep         </P>
6   <P> IO_InfoSteps   = 1                  </P>
7   <P> Re             = 500                 </P>
8   <P> Kinvis         = 1.0/Re             </P>
9   <P> kdim           = 4                   </P>
10  <P> nvec           = 1                   </P>
11  <P> evtol          = 1e-4                </P>
12 </PARAMETERS>
13

```

11.10.13.6 Boundary Conditions

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[2] </B>      <!-- Wall -->
3   <B ID="1"> C[3] </B>      <!-- Inlet -->
4   <B ID="2"> C[4] </B>      <!-- Outlet -->
5 </BOUNDARYREGIONS>
6
7 <BOUNDARYCONDITIONS>
8   <REGION REF="0">
9     <D VAR="u" VALUE="0" />
10    <D VAR="v" VALUE="0" />
11    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
12  </REGION>
13  <REGION REF="1">
14    <D VAR="u" VALUE="0" />
15    <D VAR="v" VALUE="0" />
16    <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
17  </REGION>
18  <REGION REF="2">
19    <D VAR="u" VALUE="0" />
20    <D VAR="v" VALUE="0" />

```

```

21      <N VAR="p" USERDEFINEDTYPE="H" VALUE="0" />
22      </REGION>
23      </BOUNDARYCONDITIONS>
24

```

11.10.13.7 Functions

We need to set up the base flow that can be specified as a function `BaseFlow`. In case the base flow is not analytical, it can be generated by means of the `Nonlinear` evolution operator using the same mesh and polynomial expansion.

```

1 <FUNCTION NAME="BaseFlow">
2     <F VAR="u,v,p" FILE="bfs_tg-AR.bse" />
3 </FUNCTION>
4

```

The initial guess is specified in the `InitialConditions` functions and in this case is read from a file.

```

1 <FUNCTION NAME="InitialConditions">
2     <F VAR="u,v,p" FILE="bfs_tg-AR.rst" />
3 </FUNCTION>
4

```

11.10.13.8 Usage

`IncNavierStokesSolver bfs_tg-AR.xml`

11.10.13.9 Results

The solution will be evolved forward in time using the operator \mathcal{A} , then backward in time through the adjoint operator \mathcal{A}^* . The leading eigenvalue is $\lambda = 3.236204$). This corresponds to the largest possible transient growth at the time horizon $\tau = 1$. The leading eigenmode is shown below. This is the optimal initial condition which will lead to the greatest growth when evolved under the linearised Navier-Stokes equations.

It is possible to visualise the transient growth plotting the energy evolution over time if the system is initially perturbed with the leading eigenvector. This analysis was performed for a time horizon $\tau = 60$. It can be seen that the energy grows in time reaching its maximum value at $x = 24$ and then decays, almost disappearing after 100 temporal units.

11.10.14 BiGlobal Floquet analysis of a of flow past a cylinder

In this example it will be described how to run a BiGlobal stability analysis for a time-periodic base flow using Nektar++. Let us consider a flow past a circular cylinder at $Re = 220$ has a 2D time-periodic wake that is unstable to a 3D synchronous "mode A" instability.




Figure 11.21




Figure 11.22




Figure 11.23




Figure 11.24

11.10.14.1 Background

The numerical solution of the fully three-dimensional linear eigenvalue problem is often computationally demanding and may not have significant advantages over performing a direct numerical simulation. Therefore, some simplifications are required; the most radical consist in considering that the base flow depends only on one spatial coordinate, assuming that the other two spatial coordinates are homogenous. While this method offers a good prediction for the instability of boundary layers, it is not able to predict the instability of Hagen-Poiseuille flow in a pipe at all Reynolds numbers. Between a flow that depends upon one and three-spatial directions, it is possible to consider a steady or time-periodic base flow depending upon two spatial directions and impose three-dimensional disturbances that are periodic in the third homogeneous spatial direction. This approach is known as BiGlobal stability analysis and it represents the extension of the classic linear stability theory; let us consider a base flow \mathbf{U} that is function of only two spatial coordinates: $\mathbf{U}(x, y, t)$. The perturbation velocity can \mathbf{u}' can be expressed in a similar form, but with the dependence on the third homogeneous direction incorporated through the Fourier mode: $\mathbf{u}' = \hat{\mathbf{u}}'(x, y, t)e^{i\beta z}$, where $\beta = 2\pi/L$ and L is the length in the homogeneous direction.

11.10.14.2 Input parameters

In this example we use a mesh of 500 quadrilateral elements with a 6th order polynomial expansion. The base flow has been computed using the `Nonlinear` evolution operator with appropriate boundary conditions. From its profile, it was possible to determine the periodicity of the flow sampling the velocity profile over time. In order to reconstruct the temporal behaviour of the flow, 32 time slices were considered over one period. Using these data it is possible to set up the stability simulation for a specified β , for example $\beta = 1.7$. Let us note that while the base flow is 2D, the stability simulation that we are performing is 3D.

11.10.14.3 Expansion

In this example we will use a 6th order polynomial expansion, i.e. $P = 7$.

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="7" TYPE="GLL_LAGRANGE_SEM" FIELDS="u,v,w,p" />
3 </EXPANSIONS>
4
```

11.10.14.4 Parameters

```

1 <SOLVERINFO>
2   <I PROPERTY="SolverType"           VALUE="VelocityCorrectionScheme" />
3   <I PROPERTY="EQTYPE"               VALUE="UnsteadyNavierStokes" />
4   <I PROPERTY="EvolutionOperator"    VALUE="Direct" />
5   <I PROPERTY="Projection"           VALUE="Galerkin" />
6   <I PROPERTY="ModeType"              VALUE="HalfMode" />
7   <I PROPERTY="Driver"                VALUE="ModifiedArnoldi" />
8   <I PROPERTY="HOMOGENEOUS"          VALUE="1D" />
```

```

9      <I PROPERTY="TimeIntegrationMethod" VALUE="IMEXOrder2" />
10 </SOLVERINFO>
11

```

11.10.14.5 Functions

```

1 <FUNCTION NAME="BaseFlow">
2     <F VAR="u,v,p" FILE="cyinder_floq" />
3 </FUNCTION>
4

```

11.10.14.6 Usage

IncNavierStokesSolver session.xml

11.10.14.7 Results

The stability simulation takes about 20 cycles to converge and the leading eigenvalue is $\lambda = 1.2670$ with a growth rate $\sigma = 4.7694e - 02$. The figure below shows the profile of the magnitude of the eigenmode at $z = 2$.




Figure 11.25

Linear elasticity solver

12.1 Synopsis

The `LinearElasticSolver` is a solver for solving the linear elasticity equations in two and three dimensions. Whilst this may be suitable for simple solid mechanics problems, its main purpose is for use for mesh deformation and high-order mesh generation, whereby the finite element mesh is treated as a solid body, and the deformation is applied at the boundary in order to curve the interior of the mesh.

Currently the following equation systems are supported:

Value	Description
<code>LinearElasticSystem</code>	Solves the linear elastic equations.
<code>IterativeElasticSystem</code>	A multi-step variant of the elasticity solver, which breaks a given deformation into multiple steps, and applies the deformation to a mesh.

12.1.1 The linear elasticity equations

The linear elasticity equations model how a solid body deforms under the application of a ‘small’ deformation or strain. The formulation starts with the equilibrium of forces represented by the equation

$$\nabla \cdot \mathbf{S} + \mathbf{f} = \mathbf{0} \quad \text{in } \Omega \quad (12.1)$$

where \mathbf{S} is the stress tensor and \mathbf{f} denotes a spatially-varying force. We further assume that the stress tensor \mathbf{S} incorporates elastic and, optionally, thermal stresses that can be switched on to assist in mesh deformation applications. We assume these can be decomposed so that \mathbf{S} is written as

$$\mathbf{S} = \mathbf{S}_e + \mathbf{S}_t,$$

where the subscripts e and t denote the elastic and thermal terms respectively. We adopt the usual linear form of the elastic stress tensor as

$$\mathbf{S}_e = \lambda \text{Tr}(\mathbf{E}) \mathbf{I} + \mu \mathbf{E},$$

where λ and μ are the Lamé constants, \mathbf{E} represents the strain tensor, and \mathbf{I} is the identity tensor. For small deformations, the strain tensor \mathbf{E} is given as

$$\mathbf{E} = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^t) \quad (12.2)$$

where \mathbf{u} is the two- or three-dimensional vector of displacements. The boundary conditions required to close the problem consist of prescribed displacements at the boundary $\partial\Omega$, i.e.

$$\mathbf{u} = \hat{\mathbf{u}} \quad \text{in } \partial\Omega. \quad (12.3)$$

We further express the Lamé constants in terms of the Young's modulus E and Poisson ratio ν as

$$\lambda = \frac{\nu E}{(1 + \nu)(1 - 2\nu)}, \quad \mu = \frac{E}{2(1 + \nu)}.$$

The Poisson ratio, valid in the range $\nu < \frac{1}{2}$, is a measure of the compressibility of the body, and the Young's modulus $E > 0$ is a measure of its stiffness.

12.2 Usage

```
LinearElasticSolver [arguments] session.xml [another.xml] ...
```

12.3 Session file configuration

12.3.1 Solver Info

- **EqType** Specifies the PDE system to solve, based on the choices in the table above.
- **Temperature** Specifies the form of the thermal stresses to use. The choices are:
 - **None**: No stresses (default).
 - **Jacobian**: Sets $\mathbf{S}_t = \beta J \mathbf{I}$, where β is a parameter defined in the parameters section, J is the elemental Jacobian determinant and \mathbf{I} is the identity matrix.
 - **Metric**: A more complex term, based on the eigenvalues of the metric tensor. This can only be used for simplex elements (triangles and tetrahedra). Controlled again by the parameter β .
- **BCType** Specifies the type of boundary condition to apply when the **IterativeElasticSystem** is being used.

- **Normal**: The boundary conditions are split into **NumSteps** steps, as defined by a parameter in the session file (default).
- **Repeat**: As the geometry is updated, re-evaluate the boundary conditions. This enables, for example, a circle to be rotated continuously.

12.3.2 Parameters

The following parameters can be specified in the **PARAMETERS** section of the session file:

- **nu**: sets the Poisson ratio ν .
Default value: 0.25.
- **E**: sets the Young's modulus E .
Default value: 1.
- **beta**: sets the thermal stress coefficient β .
Default value: 1.
- **NumSteps**: sets the number of steps to use in the case that the iterative elastic system is enabled. Should be greater than 0.
Default value: 0.

12.4 Examples

12.4.1 L-shaped domain

The first example is the classic L-shaped domain, in which an exact solution is known, which makes it an ideal test case [23]. The domain is the polygon formed from the vertices

$$(-1, -1), (0, -2), (2, 0), (0, 2), (-1, -1), (0, 0).$$

The exact solution for the displacements is known in polar co-ordinates (r, θ) as

$$\begin{aligned} u_r(r, \theta) &= \frac{r^\alpha}{2\mu} [C_1(C_2 - \alpha - 1) \cos((\alpha - 1)\theta) - (\alpha + 1) \cos((\alpha + 1)\theta)] \\ u_\theta(r, \theta) &= \frac{r^\alpha}{2\mu} [(\alpha + 1) \sin((\alpha + 1)\theta) + C_1(C_2 + \alpha - 1) \sin((\alpha - 1)\theta)] \end{aligned}$$

where $\alpha \approx 0.544483737\dots$ is the solution of the equation $\alpha \sin(2\omega) + \sin(2\omega\alpha) = 0$,

$$C_1 = -\frac{\cos((\alpha + 1)\omega)}{\cos((\alpha - 1)\omega)}, \quad C_2 = 2\frac{\lambda + 2\mu}{\lambda + \mu}$$

with λ and μ being the Lamé constants and $\omega = 3\pi/4$. Boundary conditions are set to be the exact solution and $\mathbf{f} = \mathbf{0}$. The solution has a singularity at the origin, and so in order to test convergence h -refinement is required.




Figure 12.1 Solution of the u displacement field for the L-shaped domain.

A simple example of how the linear elastic solver can be set up can be found in the `Tests/L-shaped.xml` session file in the linear elastic solver directory. A more refined domain with the obtained u solution is shown in figure 12.1. The solver can be run using the command:

```
LinearElasticSolver L-domain.xml
```

The obtained solution `L-domain.fld` can be applied to the mesh to obtain a deformed XML file using the `deform` module in `FieldConvert`:

```
FieldConvert -m deform L-domain.xml L-domain.fld L-domain-deformed.xml
```

12.4.2 Boundary layer deformation

In this example we use the iterative elastic system to apply a large deformation to a triangular boundary layer mesh of a square mesh $\Omega = [0, 1]^2$. At the bottom edge, we apply a Dirichlet condition $g = \frac{1}{2} \sin(\pi x)$ that is enforced by splitting it into N substeps, so that at each step we solve the system with the boundary condition $g^n(x) = g(x)/N$. The process is depicted in figure 12.2.

The setup is very straightforward. The geometry can be found inside the file `Examples/bl-mesh.xml`




Figure 12.2 Figures that show the initial domain (left), after 50 steps (middle) and final deformation of the domain (right).

and the conditions inside `Examples/bl-conditions.xml`. The solver can be set up using the following parameters, with `NumSteps` denoting N :

```

1  <SOLVERINFO>
2    <I PROPERTY="EQTYPE" VALUE="IterativeElasticSystem" />
3  </SOLVERINFO>
4
5  <PARAMETERS>
6    <P> nu      = 0.3 </P>
7    <P> E       = 1.0 </P>
8    <P> NumSteps = 100 </P>
9  </PARAMETERS>

```

To identify the boundary that we intend to split up into substeps, we must assign the `WALL` tag to our boundary regions:

```

1  <BOUNDARYCONDITIONS>
2    <REGION REF="0">
3      <D VAR="u" VALUE="0" USERDEFINEDTYPE="Wall" />
4      <D VAR="v" VALUE="0.5*sin(PI*x)" USERDEFINEDTYPE="Wall" />
5    </REGION>
6    <REGION REF="1">
7      <D VAR="u" VALUE="0" />
8      <D VAR="v" VALUE="0" />
9    </REGION>
10 </BOUNDARYCONDITIONS>

```

The solver can then be run using the command:

```
LinearElasticSolver bl-mesh.xml bl-conditions.xml
```

This will produce a series of meshes `bl-mesh-%d.xml`, where `%d` is an integer running between 0 and 100. If at any point the mesh becomes invalid, that is, a negative Jacobian is detected, execution will cease.

Pulse Wave Solver

13.1 Synopsis

1D modelling of the vasculature (arterial network) represents an insightful and efficient tool for tackling problems encountered in arterial biomechanics as well as other engineering problems. In particular, 3D modelling of the vasculature is relatively expensive. 1D modelling provides an alternative in which the modelling assumptions provide a good balance between physiological accuracy and computational efficiency. To describe the flow and pressure in this network we consider the conservation of mass and momentum applied to an impermeable, deformable tube filled with an incompressible fluid, the nonlinear system of partial differential equations presented in non-conservative form is given by

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{H} \frac{\partial \mathbf{U}}{\partial x} = \mathbf{S} \quad (13.1)$$

$$\mathbf{U} = \begin{bmatrix} U \\ A \end{bmatrix}, \quad \mathbf{H} = \begin{bmatrix} U & A \\ \rho \frac{\partial p}{\partial A} & U \end{bmatrix}, \quad \mathbf{S} = \begin{bmatrix} 0 \\ \frac{1}{\rho} \left(\frac{f}{A} - s \right) \end{bmatrix}$$

in which A is the Area (related to pressure), x is the axial coordinate along the vessel, $U(x, t)$ the axial velocity, $P(x, t)$ is the pressure in the tube, ρ is the density and finally f the frictional force per unit length. The unknowns in equation 13.1 are u , A and p ; hence we must provide an explicit algebraic relationship to close this system. Typically, closure is provided by an algebraic relationship between A and p . For a thin elastic tube this is given by

$$p = p_0 + \beta \left(\sqrt{A} - \sqrt{A_0} \right), \quad \beta = \frac{\sqrt{\pi h E}}{(1 - \nu^2) A_0} \quad (13.2)$$

where p_0 is the external pressure, A_0 is the initial cross-sectional area, E is the Young's modulus, h is the vessel wall thickness and ν is the Poisson's ratio. Other more elaborate pressure - area relationships are currently being implemented into the framework. Application of Riemann's method of characteristics to equations 13.1 and 13.2 indicates that velocity and area are propagated through the system by forward and backward travelling waves. These waves are reflected and within the network by appropriate treatment of interfaces and boundaries. In the following, we will explain the usage of the blood flow solver on the basis of a single-artery problem and also on an arterial network consisting of 55 arteries.

13.2 Usage

```
PulseWaveSolver session.xml
```

13.3 Session file configuration

13.3.1 Pulse Wave Solver mesh connectivity

Typically 1D arterial networks are made up of a connection of different base units: segments, bifurcations and merging junctions. The input format in the PulseWaveSolver means these connections are handled naturally from the mesh topology; hence care must be taken when designing the 1D domain. The figure below outlines the structure of a bifurcation, which is a common reoccurring structure in the vasculature.




Figure 13.1 Model of bifurcating artery. The bifurcation is made of three domains and 15 vertices. Vertex $V[0]$ is the inlet and vertices $V[10]$ and $V[15]$ the outlets.

To represent this topology in the xml file we specify the following vertices under the section `VERTEX` (the extents are: $-100 \geq x \leq 100$ and $-100 \geq y \leq 100$)

```
1 <VERTEX>
2 <V ID="0">-1.000e+02 0.000e+00 0.000e+00</V>
3 <V ID="1">-8.000e+01 0.000e+00 0.000e+00</V>
4 <V ID="2">-6.000e+01 0.000e+00 0.000e+00</V>
5 <V ID="3">-4.000e+01 0.000e+00 0.000e+00</V>
```

```

6 <V ID="4">-2.000e+01 0.000e+00 0.000e+00</V>
7 <V ID="5"> 0.000e+00 0.000e+00 0.000e+00</V>
8
9 <V ID="6"> 2.000e+01 2.000e+01 0.000e+00</V>
10 <V ID="7"> 4.000e+01 4.000e+01 0.000e+00</V>
11 <V ID="8"> 6.000e+01 6.000e+01 0.000e+00</V>
12 <V ID="9"> 8.000e+01 8.000e+01 0.000e+00</V>
13 <V ID="10"> 1.000e+02 1.000e+02 0.000e+00</V>
14
15 <V ID="11"> 2.000e+01 -2.000e+01 0.000e+00</V>
16 <V ID="12"> 4.000e+01 -4.000e+01 0.000e+00</V>
17 <V ID="13"> 6.000e+01 -6.000e+01 0.000e+00</V>
18 <V ID="14"> 8.000e+01 -8.000e+01 0.000e+00</V>
19 <V ID="15"> 1.000e+02 -1.000e+02 0.000e+00</V>
20 </VERTEX>

```

The elements from these vertices are then constructed under the section `ELEMENT` by defining

```

1 <ELEMENT>
2 <!-- Parent artery -->
3 <S ID="0"> 0 1 </S>
4 <S ID="1"> 1 2 </S>
5 <S ID="2"> 2 3 </S>
6 <S ID="3"> 3 4 </S>
7 <S ID="4"> 4 5 </S>
8 <!-- Daughter artery 1 -->
9 <S ID="5"> 5 6 </S>
10 <S ID="6"> 6 7 </S>
11 <S ID="7"> 7 8 </S>
12 <S ID="8"> 8 9 </S>
13 <S ID="9"> 9 10 </S>
14 <!-- Daughter artery 2 -->
15 <S ID="11"> 5 11 </S>
16 <S ID="12"> 11 12 </S>
17 <S ID="13"> 12 13 </S>
18 <S ID="14"> 13 14 </S>
19 <S ID="15"> 14 15 </S>
20 </ELEMENT>

```

The composites, which represent groups of elements and boundary regions are defined under the section `COMPOSITE` by

```

1 <COMPOSITE>
2 <C ID="0"> S[0-4] </C> <!-- Parent artery -->
3 <C ID="1"> V[0] </C> <!-- Inlet to domain -->
4
5 <C ID="3"> S[5-9] </C> <!-- Daughter artery 1 -->
6 <C ID="4"> V[10] </C> <!-- Outlet of daughter artery 1 -->
7
8 <C ID="6"> S[11-15] </C> <!-- Daughter artery 2 -->
9 <C ID="8"> V[15] </C> <!-- Outlet of daughter artery 2 -->
10 </COMPOSITE>

```

Each of the segments can be then represented under the section `DOMAIN` by

```
1 <DOMAIN>
2   <D ID="0"> C[0] </D>   <!-- Parent artery -->
3   <D ID="1"> C[3] </D>   <!-- Daughter artery 1 -->
4   <D ID="2"> C[6] </D>   <!-- Daughter artery 2 -->
5 </DOMAIN>
```

We will use the different domains later to define variable material properties and cross-sectional areas.

13.3.2 Session Info

The PulseWaveSolver is specified through the `EquationType` option in the session file. This can be set as follows:

- `Projection`: Only a discontinuous projection can be specified using the following option:
 - `Discontinuous` for a discontinuous Galerkin (DG) projection.
- `TimeIntegrationMethod`
- `UpwindTypePulse`:
 - `UpwindPulse`

13.3.3 Parameters

The following parameters can be specified in the `PARAMETERS` section of the session file.

- `TimeStep` is the time-step size;
- `FinTime` is the final physical time at which the simulation will stop;
- `NumSteps` is the equivalent of `FinTime` but instead of specifying the physical final time the number of time-steps is defined;
- `IO_CheckSteps` sets the number of steps between successive checkpoint files;
- `IO_InfoSteps` sets the number of steps between successive info stats are printed to screen;
- `rho` density of the fluid. Default value = 1.0;
- `nue` Poisson's ratio. Default value = 0.5 ;
- `pest` external pressure. Default value = 0;
- `h0` wall thickness Default value = 1.0;

13.3.4 Boundary conditions

In this section we can specify the boundary conditions for our problem. First we need to define the variables under the section `VARIABLES`.

```
1 <VARIABLES>
2   <V ID="0"> A </V>
3   <V ID="1"> u </V>
4 </VARIABLES>
```

The composites that we want to apply our boundary conditions then need to be defined in the `BOUNDARYREGIONS`, for example if we had three composites (C[1], C[4] and C[8]) that correspond to three vertices of the computational mesh we would define:

```
1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[4] </B>
4   <B ID="2"> C[8] </B>
5 </BOUNDARYREGIONS>
```

Finally we can specify the boundary conditions on the regions specified under `BOUNDARYREGIONS`.

The Pulse Wave Solver comes with a number of boundary conditions that are unique to this solver. Boundary conditions must be provided for both the area and velocity at the inlets and outlets of the domain. Examples of the different boundary conditions will be provided in the following.

13.3.4.0.1 Inlet boundary condition: The inlet condition may be specified algebraically in four different ways: as an area variation (`A-inflow`); a velocity profile (`U-inflow`); a volume flux (`Q-inflow`); or by prescribing the forward characteristic (`TimeDependent`). When prescribing a volume flux, it must be specified in the input file via the area, as illustrated below. Note that $u = 1.0$.

```
1 <REGION REF="0">
2   <D VAR="A" USERDEFINEDTYPE="Q-inflow" VALUE="(7.112e-4)*(sin(7.854*t)
3 -0.562)*(1/(1+exp(-400*(sin(7.854*t)-0.562))))" />
4   <D VAR="u" USERDEFINEDTYPE="Q-inflow" VALUE="1.0" />
5 </REGION>
```

13.3.4.0.2 Terminal boundary conditions: At the outlets of the domain there are four possible boundary conditions: reflection (`Terminal`), terminal resistance (`R-terminal`), Two element windkessel (CR) (`CR-terminal`), and three element windkessel (RCR) (`RCR-terminal`). An example of the outflow boundary condition of the RCR terminal is given below

```
1 <REGION REF="1">
2   <D VAR="A" USERDEFINEDTYPE="RCR-terminal" VALUE="RT" />
3   <D VAR="u" USERDEFINEDTYPE="RCR-terminal" VALUE="C" />
4 </REGION>
```


Where `RT` is the total peripheral resistance used in the the `R-terminal`, `CR-terminal` and `RCR-terminal` models

13.3.5 Functions

The following functions can be specified inside the `CONDITIONS` section of the session file:

- `MaterialProperties`: specifies the material properties for each domain.
- `A_0`: Initial area of each domain.
- `AdvectionVelocity`: specifies the advection velocity v .
- `InitialConditions`: specifies the initial condition for unsteady problems.
- `Forcing`: specifies the forcing function f

As an example to specify the material properties for each domain in the previous bifurcation example we would enter:

```
1 <FUNCTION NAME="MaterialProperties">
2   <E VAR="beta" DOMAIN="0" VALUE="97" />
3   <E VAR="beta" DOMAIN="1" VALUE="87" />
4   <E VAR="beta" DOMAIN="2" VALUE="233" />
5 </FUNCTION>
```

The values of `beta` are used in the pressure-area relationship (equation 13.2).


13.4 Examples

13.4.1 Human Vascular Network

The Pulse Wave Solver is also capable of handling more complex networks, such as a complete human arterial tree proposed by Westerhof et al. [45]. In this example, we will use the refined data from [40] and set up the network shown in the figure in the right. We will explain how bifurcations are set correctly and how each arterial segment gets its correct physiological data.

First, we will set up the mesh where each arterial segment is represented by one element and two vertices respectively. Then, we will subdivide the mesh into different subdomains by using the `<COMPOSITE>` section. Here, each arterial segment is described by the contained elements and its first and last vertex.

The mesh connectivity is specified during the creation of elements by indicating the starting vertex and ending vertex of each individual artery segment. Shared vertices are used to describe bifurcations, junctions and mergers between different artery segments in the network.



#	Artery	Length (cm)	Area (cm ²)	β (kg s ⁻² cm ⁻²)	R_t
1	Ascending Aorta	4.0	5.983	97	-
2	Aortic Arch I	2.0	5.147	87	-
3	Brachiocephalic	3.4	1.219	233	-
4	R. Subclavian I	3.4	0.562	423	-
5	R. Carotid	17.7	0.432	516	-
6	R. Vertebral	14.8	0.123	2590	0.906
7	R. Subclavian II	42.2	0.510	466	-
8	R. Radial	23.5	0.106	2866	0.82
9	R. Ulnar I	6.7	0.145	2246	-
10	R. Interosseous	7.9	0.031	12894	0.956
11	R. Ulnar II	17.1	0.133	2446	0.893
12	R. Internal Carotid	17.6	0.121	2644	0.784
13	R. External Carotid	17.7	0.121	2467	0.79
14	Aortic Arch II	3.9	3.142	130	-
15	L. Carotid	20.8	0.430	519	-
16	L. Internal Carotid	17.6	0.121	2644	0.784
17	L. External Carotid	17.7	0.121	2467	0.791
18	Thoracic Aorta I	5.2	3.142	124	-
19	L. Subclavian I	3.4	0.562	416	-
20	Vertebral	14.8	0.123	2590	0.906
21	L. Subclavian II	42.2	0.510	466	-
22	L. Radial	23.5	0.106	2866	0.821
23	L. Ulnar I	6.7	0.145	2246	-
24	L. Interosseous	7.9	0.031	12894	0.956
25	L. Ulnar II	17.1	0.133	2446	0.893
26	Intercostals	8.0	0.196	885	0.627
27	Thoracic Aorta II	10.4	3.017	117	-
28	Abdominal I	5.3	1.911	167	-
29	Celiac I	2.0	0.478	475	-
30	Celiac II	1.0	0.126	1805	-
31	Hepatic	6.6	0.152	1142	0.925
32	Gastric	7.1	0.102	1567	0.921
33	Splenic	6.3	0.238	806	0.93
34	Superior Mesenteric	5.9	0.430	569	0.934
35	Abdominal II	1.0	1.247	227	-
36	L. Renal	3.2	0.332	566	0.861
37	Abdominal III	1.0	1.021	278	-
38	R. Renal	3.2	0.159	1181	0.861
39	Abdominal IV	10.6	0.697	381	-
40	Inferior Mesenteric	5.0	0.080	1895	0.918
41	Abdominal V	1.0	0.578	399	-
42	R. Common Iliac	5.9	0.328	649	-
43	L. Common Iliac	5.8	0.328	649	-
44	L. External Iliac	14.4	0.252	1493	-
45	L. Internal Iliac	5.0	0.181	3134	0.925
46	L. Femoral	44.3	0.139	2559	-
47	L. Deep Femoral	12.6	0.126	2652	0.885
48	L. Posterior Tibial	32.1	0.110	5808	0.724
49	L. Anterior Tibial	34.3	0.060	9243	0.716
50	R. External Iliac	14.5	0.252	1493	-
51	R. Internal Iliac	5.1	0.181	3134	0.925
52	R. Femoral	44.4	0.139	2559	-
53	R. Deep Femoral	12.7	0.126	2652	0.888
54	L. Posterior Tibial	32.2	0.110	5808	0.724
55	R. Anterior Tibial	34.4	0.060	9243	0.716

The composites are then used to specify the two adjoining segments of an artery, where the first segment merely allows for description of the connectivity.

```

1 <GEOMETRY DIM="1" SPACE="1">
2   <VERTEX>
3     <V ID="0"> 0.000e+00 0.000e+00 0.000e+00</V> <!-- 1 -->
4     <V ID="1"> 4.000e+00 0.000e+00 0.000e+00</V>
5
6     <V ID="2"> 4.000e+00 0.000e+00 0.000e+00</V> <!-- 2 -->
7     <V ID="3"> 6.000e+00 0.000e+00 0.000e+00</V>
8
9     <V ID="4"> 4.000e+00 0.000e+00 0.000e+00</V> <!-- 3 -->
10    <V ID="5"> 7.400e+00 0.000e+00 0.000e+00</V>
11    .
12    .
13    .
14    <V ID="108"> 109.100e+00 -45.000e+00 0.000e+00</V> <!-- 55 -->
15    <V ID="109"> 143.500e+00 -45.000e+00 0.000e+00</V>
16  </VERTEX>
17  <ELEMENT>
18    <S ID="0"> 0 1 </S>
19    <S ID="1"> 1 2 </S>
20    <S ID="2"> 1 4 </S>
21    <S ID="3"> 2 3 </S>
22    <S ID="4"> 4 5 </S>

```

```

23      <S ID="5">      5      6 </S>
24      <S ID="6">      5      8 </S>
25      <S ID="7">      6      7 </S>
26      <S ID="8">      8      9 </S>
27      .
28      .
29      .
30      <S ID="106">    103    108 </S>
31      <S ID="107">    108    109 </S>
32      <S ID="108">    85     98 </S>
33      <ELEMENT>
34      <COMPOSITE>
35          <C ID="0"> S[0] </C> <!-- 1 -->
36          <C ID="1"> V[0] </C>
37          <C ID="2"> V[1] </C>
38
39          <C ID="3"> S[1,3] </C> <!-- 2 -->
40          <C ID="4"> V[2] </C>
41          <C ID="5"> V[3] </C>
42
43          <C ID="6"> S[2,4] </C> <!-- 3 -->
44          <C ID="7"> V[4] </C>
45          <C ID="8"> V[5] </C>
46      .
47      .
48      .
49      <C ID="162"> S[106,107] </C> <!-- 55 -->
50      <C ID="163"> V[108] </C>
51      <C ID="164"> V[109] </C>
52      </COMPOSITE>
53 </GEOMETRY>

```

Then the choice of polynomial order, solver information, area of the arteries and other parameters are specified.

```

1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="5" FIELDS="A,u" TYPE="MODIFIED" />
3   <E COMPOSITE="C[3]" NUMMODES="5" FIELDS="A,u" TYPE="MODIFIED" />
4   ...
5
6   <E COMPOSITE="C[162]" NUMMODES="5" FIELDS="A,u" TYPE="MODIFIED" />
7 </EXPANSIONS>
8
9 <CONDITIONS>
10
11   <PARAMETERS>
12
13     <P> TimeStep      = 1e-4          </P>
14     <P> FinTime       = 1.0           </P>
15     <P> NumSteps      = FinTime/TimeStep </P>
16     <P> IQ_CheckSteps = NumSteps/50    </P>
17     ...
18     <P> A53           = 0.126          </P>
19     <P> A54           = 0.110          </P>
20     <P> A55           = 0.060          </P>


```

```

21 </PARAMETERS>
22
23 <SOLVERINFO>
24   <I PROPERTY="EQTYPE"           VALUE="PulseWavePropagation" />
25   <I PROPERTY="Projection"        VALUE="DisContinuous"         />
26   <I PROPERTY="TimeIntegrationMethod" VALUE="RungeKutta2_ImprovedEuler" />
27   <I PROPERTY="UpwindTypePulse"    VALUE="UpwindPulse"           />
28 </SOLVERINFO>
29
30 <VARIABLES>
31   <V ID="0"> A </V>
32   <V ID="1"> u </V>
33 </VARIABLES>

```

The vertices where the network terminates are specified as boundary regions based on their subsequent composite ids.



```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B> <B ID="1"> C[17] </B> <B ID="2"> C[23] </B>
3   ...
4   <B ID="28"> C[164] </B>
5 </BOUNDARYREGIONS>

```

In the boundary conditions section the inflow and outflow conditions are set up. Here we use an inflow boundary condition for the area at the beginning of the ascending aorta taken from [40] and plotted on the right. Potential choices for inflow boundary conditions include Q-Inflow and Time-Dependent inflow. The outflow conditions for the terminal regions of the network could be specified by different models including eTerminal, R, CR, RCR and Time-Dependant outflow.

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0"> <!-- Inflow -->
3     <D VAR="A" USERDEFINEDTYPE="TimeDependent"
4       VALUE="5.983*(1+0.597*(sin(6.28*t + 0.628) - 0.588)*
5         (1./(1+exp(-2*200*(sin(6.28*t + 0.628) - 0.588)))))" />
6     <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="0.0" />
7   </REGION>
8   <REGION REF="1">
9     <D VAR="A" USERDEFINEDTYPE="TimeDependent" VALUE="A6" />

```

```

10     <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="0.0" />
11 </REGION>
12 <REGION REF="2">
13     <D VAR="A" USERDEFINEDTYPE="TimeDependent" VALUE="A8" />
14     <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="0.0" />
15 </REGION>
16 <REGION REF="3">
17     <D VAR="A" USERDEFINEDTYPE="TimeDependent" VALUE="A10" />
18     <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="0.0" />
19 </REGION>
20 ....
21 <REGION REF="28">
22     <D VAR="A" USERDEFINEDTYPE="TimeDependent" VALUE="A55" />
23     <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="0.0" />
24 </REGION>
25 </BOUNDARYCONDITIONS>

```

Again, for the initial conditions we start our simulation from static equilibrium conditions $A = A_0$ and for u being initially at rest. The following lines show how we specify A_0 and β for different arterial segments.


```

1 <FUNCTION NAME="InitialConditions">
2     <E VAR="A" DOMAIN="0" VALUE="5.983" />
3     <E VAR="u" DOMAIN="0" VALUE="0.0" />
4 </FUNCTION>
5 ...
6 <FUNCTION NAME="InitialConditions">
7     <E VAR="A" DOMAIN="54" VALUE="A55" />
8     <E VAR="u" DOMAIN="54" VALUE="0.0" />
9 </FUNCTION>
10
11 <FUNCTION NAME="A_0">
12     <E VAR="A_0" DOMAIN="0" VALUE="A1" />
13     ...
14     <E VAR="A_0" DOMAIN="54" VALUE="A55" />
15 </FUNCTION>
16
17 <FUNCTION NAME="MaterialProperties">
18     <E VAR="beta" DOMAIN="0" VALUE="97" />
19     ...
20     <E VAR="beta" DOMAIN="54" VALUE="9243" />
21 </FUNCTION>

```

Our simulation is started as described before and the results show the time history for the conservative variables A and u , as well as for the characteristic variables $W1$ and $W2$ at the beginning of the ascending aorta (Artery 1). We can see that physically correct the shape of the inflow boundary condition appears in the forward travelling characteristic $W1$. As we do not have a terminal resistance at the outflow, one would normally expect $W2$ to be constant. However this is not the case, as bifurcations cause reflections if the radii of parent and daughter vessels are not well matching, leading to changes in $W2$. The shapes of A and u result from this facts and show the values for the physiological variables during one cardiac cycle. We may annotate that this values slightly differ from

in vivo measurements due to the missing terminal resistance, which will be added in future.



These short examples should give an insight to the functionality of our PulseWaveSolver and show that results such as luminal area and pressure within the artery can be simulated. These results can contribute to understanding the physiology of the human vascular system and they can be used for patient-specific planning of medical interventions.

13.4.2 Stented Artery

13.4.3 Stented Artery

In the following we will explain the usage of the Pulse Wave solver to model the flow and pressure variation through a stented artery - a cardiovascular procedure in which a small mesh tube is inserted into an artery to restore blood flow through a constricted region. Due to the implantation of the stent this region will have different material properties compared to the surrounding unstented tissue; hence will influence the propagation of waves through this system. The stent scenario to be modelled is a straight arterial segment with a stent situated between $x = a_1$ and $x = a_2$ as shown below.




Figure 13.2 Model of straight artery with a stent in the middle.

13.4.3.0.1 Geometry: In the following we describe the geometry setup for modelling 1D flow in a stent. This is done by defining vertices, elements and composites. The vertices of the domain are shown below, consisting of 30 elements (Ω) and 31 vertices ($V[n]$).




Figure 13.3 1D arterial domain consisting of 30 elements and 31 vertices.

To represent the above in the xml file, we define 31 vertices as follows:

```

1 <VERTEX>
2   <V ID="0"> 0.000e+00 0.000e+00 0.000e+00</V>
3   .
4   .
5   .
6   <V ID="30">30.000e+00 0.000e+00 0.000e+00</V>
7 </VERTEX>

```

and the connectivity of these vertices to make up the 30 elements:

```

1 <ELEMENT>
2   <S ID="0"> 0 1 </S>
3   .
4   .
5   .
6   <S ID="29"> 29 30 </S>
7 </ELEMENT>

```

These elements are combined to three different composites (shown below): composite 0 represents all the elements; composite 1 the inflow boundary and composite 2 the outflow boundary.




Figure 13.4 Three composites ($C[0]$, $C[1]$ and $C[2]$) for the stented artery.

The above composites are specified as follows:

```

1 <COMPOSITE>
2   <C ID="0"> S[0-29] </C>
3   <C ID="1"> V[0] </C>
4   <C ID="2"> V[30] </C>
5 </COMPOSITE>

```

Finally the domain is specified by the first composite by

```
1 <DOMAIN>
2   <D ID="0"> C[0] </D>
3 </DOMAIN>
```

13.4.3.0.2 Expansion: For the expansions we use 4th-order polynomials which define our two variables A and u on the domain.

```
1 <EXPANSIONS>
2   <E COMPOSITE="C[0]" NUMMODES="5" FIELDS="A,u" TYPE="MODIFIED" />
3 </EXPANSIONS>
```

13.4.3.0.3 Solver Information: The Discontinuous Galerkin Method is used as projection scheme and the time-integration is performed by a simple Forward Euler scheme. A full list of possible time integration scheme is given in the parameter section of the Pulse Wave Solver

```
1 <SOLVERINFO>
2   <I PROPERTY="EQTYPE" VALUE="PulseWavePropagation" />
3   <I PROPERTY="Projection" VALUE="DisContinuous" />
4   <I PROPERTY="TimeIntegrationMethod" VALUE="ForwardEuler" />
5   <I PROPERTY="UpwindTypePulse" VALUE="UpwindPulse" />
6 </SOLVERINFO>
```

13.4.3.0.4 Parameters: Parameters used for the simulation are taken from [40]

```
1 <PARAMETERS>
2   <P> TimeStep      = 2e-6          </P>
3   <P> FinTime       = 0.25          </P>
4   <P> NumSteps      = FinTime/TimeStep </P>
5   <P> IO_CheckSteps = NumSteps/50    </P>
6   <P> IO_InfoSteps  = 100           </P>
7   <P> T             = 0.33          </P>
8   <P> h0            = 1.0           </P>
9   <P> rho           = 1.0           </P>
10  <P> nue           = 0.5           </P>
11  <P> pext          = 0.0           </P>
12  <P> a1            = 10.0          </P>
13  <P> a2            = 20.0          </P>
14  <P> kappa         = 100.0         </P>
15  <P> Y0            = 1.9099e+5     </P>
16  <P> k             = 2             </P>
17  <P> k1            = 200           </P>
18 </PARAMETERS>
```

13.4.3.0.5 Boundary conditions: At the inflow we apply a pressure boundary condition as shown in the figure below. This condition models the pressure variation during one heartbeat. A simple absorbing outflow boundary condition is applied the right end of the tube.

These are defined in the xml file as follows,




Figure 13.5 Pressure profile applied at the inlet of the artery

```

1 <BOUNDARYREGIONS>
2   <B ID="0"> C[1] </B>
3   <B ID="1"> C[2] </B>
4 </BOUNDARYREGIONS>
5
6 <BOUNDARYCONDITIONS>
7   <REGION REF="0">
8     <D VAR="A" USERDEFINEDTYPE="TimeDependent" VALUE=
9       "(2000*sin(2*PI*t/T)*1./(1+exp(-2*k1*(T/2-t))-pext)/451352+1)^2" />
10    <D VAR="u" USERDEFINEDTYPE="TimeDependent" VALUE="1.0" />
11  </REGION>
12  <REGION REF="1">
13    <D VAR="A" VALUE="1.0" />
14    <D VAR="u" VALUE="1.0" />
15  </REGION>
16 </BOUNDARYCONDITIONS>

```

The simulation starts from the static equilibrium of the vessel with normalised area and velocity.

```

1 <FUNCTION NAME="InitialConditions">
2   <E VAR="A" DOMAIN="0" VALUE="1.0" />
3   <E VAR="u" DOMAIN="0" VALUE="1.0" />
4 </FUNCTION>
5
6 <FUNCTION NAME="A_0">
7   <E VAR="A" DOMAIN="0" VALUE="1.0" />
8 </FUNCTION>

```

13.4.3.0.6 Functions: The stent is introduced by applying a variable material properties function (β - see equation 13.2) along the vessel in the x direction, shown graphically below and defined in the xml file by

```

1 <FUNCTION NAME="MaterialProperties">
2   <E VAR="EO" DOMAIN="0" VALUE=

```




Figure 13.6 material property variation along the artery. The stiff region in the middle represents the stent.

```
3      "Y0*(1.0-kappa/(1+exp(-2*k*(a1-x)))+kappa/(1+exp(-2*k*(a2-x))))" />
4 </FUNCTION>
```

13.4.3.1 Simulation

The simulation is started by running

```
PulseWaveSolver Test_1.xml
```

It will take about 60 seconds on a 2.4GHz Intel Core 2 Duo processor and therefore is computationally realisable at every clinical site.

13.4.3.2 Results

As a result we get a 3-dimensional interpretation of the aortic cross-sectional area varying in axial direction both for the stented and non-stented vessel. In case of the stent, the rigid metal mesh will restrict the deformation of the area in that specific part of the artery compared to the normal vessel (Fig. 13.7).

Also, if we look at the pressure at three points within the artery (P, M, D) we will recognize that there are major differences between the stented and normal vessel. While in the normal vessel (left) the pressure wave applied at the inflow is propagated without any losses, this does not hold for the stented artery (right). Here, the stiffening at the stent causes reflections and thus there are losses for total pressure at the medial (M) and distal (D) point.





Figure 13.7



13.5 Further Information

The PulseWaveSolver has been developed with contributions by various students and researchers at the Department of Aeronautics, Imperial College London. Further information on the solver and its underlying mathematical framework can be found in [37, 36].

13.6 Future Development

The PulseWaveSolver is a useful tool for computational modelling of one-dimensional blood flow in the human body. However, there are several ideas for future development which include:

1. Inclusion of a pre-processor and post-processor.
2. Profiling the code to improve performance.
3. Cleaning up the input file to make the input format more user-friendly.
4. Modelling of valves and alternative pressure-area laws for models of venous flow.
5. Incorporating a model of the heart.

Shallow Water Solver

14.1 Synopsis

The ShallowWaterSolver is a solver for depth-integrated wave equations of shallow water type. Presently the following equations are supported:

Value	Description
<code>LinearSWE</code>	Linearized SWE solver in primitive variables (constant still water depth)
<code>NonlinearSWE</code>	Nonlinear SWE solver in conservative variables (constant still water depth)

14.1.1 The Shallow Water Equations

The shallow water equations (SWE) is a two-dimensional system of nonlinear partial differential equations of hyperbolic type that are fundamental in hydraulic, coastal and environmental engineering. In deriving the SWE the vertical velocity is considered negligible and the horizontal velocities are assumed uniform with depth. The SWE are hence valid when the water depth can be considered small compared to the characteristic length scale of the problem, as typical for flows in rivers and shallow coastal areas. Despite the limiting restrictions the SWE can be used to describe many important phenomena, for example storm surges, tsunamis and river flooding.

The two-dimensional SWE is stated in conservation form as

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U})$$

where $\mathbf{F}(\mathbf{U}) = [\mathbf{E}(\mathbf{U}), \mathbf{G}(\mathbf{U})]$ is the flux vector and the vector of conserved variables read $\mathbf{U} = [H, Hu, Hv]^T$. Here $H(\mathbf{x}, t) = \zeta(\mathbf{x}, t) + d(\mathbf{x})$ is the total water depth, $\zeta(\mathbf{x}, t)$ is the free surface elevation and $d(\mathbf{x})$ is the still water depth. The depth-averaged velocity is

denoted by $\mathbf{u}(\mathbf{x}, t) = [u, v]^T$, where u and v are the velocities in the x - and y -directions, respectively. The content of the flux vector is

$$\mathbf{E}(U) = \begin{bmatrix} Hu \\ Hu^2 + gH^2/2 \\ Huv \end{bmatrix}, \quad \mathbf{G}(U) = \begin{bmatrix} Hv \\ Hvu \\ Hv^2 + gH^2/2 \end{bmatrix},$$

in which g is the acceleration due to gravity. The source term $\mathbf{S}(U)$ accounts for, e.g., forcing due to bed friction, bed slope, Coriolis force and higher-order dispersive effects (Boussinesq terms). In the distributed version of the ShallowWaterSolver only the Coriolis force is included.

14.2 Usage

```
ShallowWaterSolver session.xml
```

14.3 Session file configuration

14.3.1 Solver Info

- `Eqtype`: Specifies the equation to solve. This should be set to `NonlinearSWE`.
- `UpwindType`
- `Projection`
- `TimeIntegrationScheme`

14.3.2 Parameters

- `Gravity`

14.3.3 Functions

- `Coriolis`: Specifies the Coriolis force (variable name: 'f')
- `WaterDepth`: Specifies the water depth (variable name: 'd')

14.4 Examples

14.4.1 Rossby modon case

This example, provided in `RossbyModon_Nonlinear_DG.xml` is of a discontinuous Galerkin simulation of the westward propagation of an equatorial Rossby modon.

14.4.1.1 Input Options

For what concern the ShallowWaterSolver the `<SOLVERINFO>` section allows us to specify the solver, the type of projection (continuous or discontinuous), the explicit time integration scheme to use and (in the case the discontinuous Galerkin method is used) the choice of numerical flux. A typical example would be:

```

1 <SOLVERINFO>
2   <I PROPERTY="EqType"           VALUE="NonlinearSWE"           />
3   <I PROPERTY="Projection"       VALUE="DisContinuous"       />
4   <I PROPERTY="TimeIntegrationMethod" VALUE="ClassicalRungeKutta4" />
5   <I PROPERTY="UpwindType"       VALUE="HLLC"                 />
6 </SOLVERINFO>

```

In the `<PARAMETERS>` section we, in addition to the normal setting of time step etc., also define the acceleration of gravity by setting the parameter "Gravity":

```

1 <PARAMETERS>
2   <P> TimeStep      = 0.04          </P>
3   <P> NumSteps      = 1000          </P>
4   <P> IO_CheckSteps = 100           </P>
5   <P> IO_InfoSteps  = 100           </P>
6   <P> Gravity       = 1.0           </P>
7 </PARAMETERS>

```

We specify f which is the Coriolis parameter and d denoting the still water depth as analytic functions:

```

1 <FUNCTION NAME="Coriolis">
2   <E VAR="f" VALUE="0+1*y" />
3 </FUNCTION>
4
5 <FUNCTION NAME="WaterDepth">
6   <E VAR="d" VALUE="1" />
7 </FUNCTION>

```

Initial values and boundary conditions are given in terms of primitive variables (please note that also the output files are given in terms of primitive variables). For the discontinuous Galerkin we typically enforce any slip wall boundaries weakly using symmetry technique. This is given by the `USERDEFINEDTYPE="Wall"` choice in the `<BOUNDARYCONDITIONS>` section:

```

1 <BOUNDARYCONDITIONS>
2   <REGION REF="0">
3     <D VAR="eta" USERDEFINEDTYPE="Wall" VALUE="0" />
4     <D VAR="u"   USERDEFINEDTYPE="Wall" VALUE="0" />
5     <D VAR="v"   USERDEFINEDTYPE="Wall" VALUE="0" />
6   </REGION>
7 </BOUNDARYCONDITIONS>

```

14.4.1.2 Running the code

After the input file has been copied to the build directory of the `ShallowWaterSolver` the code can be executed by:


```
./ShallowWaterSolver Rossby_Nonlinear_DG.xml
```

14.4.1.3 Post-proceesing

After the final time step the solver will write an output file `RossbyModon_Nonlinear_DG.fld`. We can convert it to tecplot format by using the `FieldConvert` utility. Thus we execute the following command:

```
FieldConvert RossbyModon_Nonlinear_DG.xml RossbyModon_Nonlinear_DG.fld \
RossbyModon_Nonlinear_DG.dat
```

This will generate a file called `RossbyModon_Nonlinear_DG.dat` that can be loaded directly into tecplot:



Part IV

Reference

Optimisation

One of the most frequently asked questions when performing almost any scientific computation is: how do I make my simulation faster? Or, equivalently, why is my simulation running so slowly?

The spectral element method is no exception to this rule. The purpose of this chapter is to highlight some of the easiest parameters that can be tuned to attain optimum performance for a given simulation.

Details are kept as untechnical as possible, but some background information on the underlying numerical methods is necessary in order to understand the various options available and the implications that they can have on your simulation.

15.1 Operator evaluation strategies

When discretising a PDE using most variants of the spectral element method, the resulting problem is usually expressed as a matrix equation. In traditional linear finite element codes, the matrix is usually represented as a large sparse global matrix, which represents the action of a particular operator such as the Laplacian matrix across the whole domain.

However, when we consider spectral element methods, in which the polynomial order representing the expansion can be far higher, this method becomes far less optimal. We can instead consider the action of an operator locally on each element, and then perform an assembly operation. This is mathematically equivalent to the global matrix approach and gives exactly the same answer, but at high polynomial orders it is far more efficient on modern CPU architectures.

Furthermore, this local approach can be represented in one of two ways: either as a dense matrix for each element, which is typically more efficient at intermediate polynomial orders, or in the *hp* element case as a tensor product of smaller dense matrices via an approach deemed *sum-factorisation*, which is used at very high polynomial orders. Figure ?? gives an overview of these three different operator strategies.

A goal of *Nektar++* is to support not only high order expansions, but *all* orders from low (where element size h is the dominant factor) to high (where p dominates); a procedure we have dubbed “from h to p efficiently”.

15.1.1 Selecting an operator strategy

An obvious question is: “which strategy should I select?” The most important factors in this decision are:

1. what the operator is;
2. polynomial order p ;
3. element type and dimension of the problem;
4. underlying hardware architecture;
5. the number of operator calls in the solver;
6. BLAS implementation speed.

Generally you can use results from three publications [44, 7, 6] which outline results for two- and three-dimensional elements.

In general, the best approach is to perform some preliminary timings by changing the appropriate variables in the session file, which is outlined below. As a very rough guide, for $1 \leq p \leq 2$ you should use the global approach; for $3 \leq p \leq 7$ use the local approach; and for $p \geq 8$ use sum-factorisation. However, these guidelines will vary due to the parameters noted above. In future releases of *Nektar++* we hope to tune these variables automatically to make this decision easier to make.

15.1.2 XML syntax

Operator evaluation strategies can be configured in the `GLOBALOPTIMISATIONPARAMETERS` tag, which lies inside the root `NEKTAR` tag:

```

1 <NEKTAR>
2   <GLOBALOPTIMISATIONPARAMETERS>
3     <BwdTrans>
4       <DO_GLOBAL_MAT_OP VALUE="0" />
5       <DO_BLOCK_MAT_OP TRI="1" QUAD="1" TET="1"
6         PYR="1" PRISM="1" HEX="1" />
7     </BwdTrans>
8     <IPProductWRTBase>
9       <DO_GLOBAL_MAT_OP VALUE="0" />
10      <DO_BLOCK_MAT_OP TRI="1" QUAD="1" TET="1"
11        PYR="1" PRISM="1" HEX="1" />
12    </IPProductWRTBase>
13    <HelmholtzMatrixOp>

```

```

14 <DO_GLOBAL_MAT_OP VALUE="0" />
15 <DO_BLOCK_MAT_OP TRI="1" QUAD="1" TET="1"
16 PYR="1" PRISM="1" HEX="1" />
17 </HelmholtzMatrixOp>
18 <MassMatrixOp>
19 <DO_GLOBAL_MAT_OP VALUE="0" />
20 <DO_BLOCK_MAT_OP TRI="1" QUAD="1" TET="1"
21 PYR="1" PRISM="1" HEX="1" />
22 </MassMatrixOp>
23 </GLOBALOPTIMIZATIONPARAMETERS>
24 </NEKTAR>

```

15.1.3 Selecting different operator strategies

Operator evaluation is supported for four operators: backward transform, inner product, Helmholtz and mass operators. It is possible to specify the following optimisation flags for different operators:

1. `DO_GLOBAL_MAT_OP`: If `VALUE` is `1`, the globally assembled system matrix will be used to evaluate the operator. If `VALUE` is `0`, the operator will be evaluated elementally.
2. `DO_BLOCK_MAT_OP`: If `VALUE` is `1`, the elemental evaluation will be done using the elemental/local matrices (which are all concatenated in a block matrix, hence the name). If `VALUE` is `0`, the elemental evaluation will be done using the sum-factorisation technique.

Each element type (triangle, quadrilateral, etc) has its own `VALUE`, since break-even points for sum-factorisation and the local matrix approach will differ depending on element type. Note that due to a small shortcoming in the code, all element types must be defined; so three-dimensional elements must still be defined even if the simulation is two-dimensional.

Note that global takes precedence over block, so if `VALUE` is set to `1` for both then the operator will be global.

For very complex operators – in particular `HelmholtzMatrixOp` – always set `DO_BLOCK_MAT_OP` to `1` as sum-factorisation for these operator types can be costly.

15.2 Collections

The Collections library adds optimisations to perform certain elemental operations collectively by applying an operator using a matrix-matrix operation, rather than a sequence of matrix-vector multiplications. Certain operators benefit more than others from this treatment, so the following implementations are available:

- StdMat: Perform operations using collated matrix-matrix type elemental operation.
- SumFac: Perform operation using collated matrix-matrix type sum factorisation operations.
- IterPerExp: Loop through elements, performing matrix-vector operation.
- NoCollections: Use the original LocalRegions implementation to perform the operation.

All configuration relating to Collections is given in the `COLLECTIONS` XML element within the `NEKTAR` XML element.

15.2.1 Default implementation

The default implementation for all operators may be chosen through setting the `DEFAULT` attribute of the `COLLECTIONS` XML element to one of `StdMat`, `SumFac`, `IterPerExp` or `NoCollection`. For example, the following uses the collated matrix-matrix type elemental operation for all operators and expansion orders:

```
1 <COLLECTIONS DEFAULT="StdMat" />
```

15.2.2 Auto-tuning

The choice of implementation for each operator, for the given mesh and expansion orders, can be selected automatically through auto-tuning. To enable this, add the following to the *Nektar++* session file:

```
1 <COLLECTIONS DEFAULT="auto" />
```

This will collate elements from the given mesh and given expansion orders, run and time each implementation strategy in turn, and select the fastest performing case. Note that the selections will be mesh- and order- specific. The selections made via auto-tuning are output if the `-verbose` command-line switch is given.

15.2.3 Manual selection

The choice of implementation for each operator may be set manually within the `COLLECTIONS` tag as shown in the following example. Different implementations may be chosen for different element shapes and expansion orders. Specifying `*` for `ORDER` sets the default implementation for any expansion orders not explicitly defined.

```
1 <COLLECTIONS>
2   <OPERATOR TYPE="BwdTrans">
3     <ELEMENT TYPE="T" ORDER="*" IMPTYPE="IterPerExp" />
4     <ELEMENT TYPE="T" ORDER="1-5" IMPTYPE="StdMat" />
5   </OPERATOR>
6   <OPERATOR TYPE="IProductWRTBase">
7     <ELEMENT TYPE="Q" ORDER="*" IMPTYPE="SumFac" />
```

```
8 </OPERATOR>  
9 </COLLECTIONS>
```

Manual selection is intended to document the optimal selections on a given hardware platform after extensive prior testing, to avoid the need to run the auto-tuning for each run.

15.2.4 Collection size

The maximum number of elements within a single collection can be enforced using the `MAXSIZE` attribute.

Command-line Options

- verbose**
Displays extra info.
- version**
Displays software version, and source control information if applicable.
- help**
Displays help information about the available command-line options for the executable.
- parameter [key]=[value]**
Override a parameter (or define a new one) specified in the XML file.
- solverinfo [key]=[value]**
Override a solverinfo (or define a new one) specified in the XML file.
- io-format [format]**
Determines the output format for writing *Nektar++* field files that are used to store, for example, checkpoint and solution field files. The default for **format** is **Xml**, which is an XML-based format, which is written as one file per process. If *Nektar++* is compiled with HDF5 support, then an alternative option is **Hdf5**, which will write one file for all processes and can be more efficient for very large-scale parallel jobs.
- npx [int]**
When using a fully-Fourier expansion, specifies the number of processes to use in the x-coordinate direction.
- npy [int]**
When using a fully-Fourier expansion or 3D expansion with two Fourier directions, specifies the number of processes to use in the y-coordinate direction.

`--npz [int]`

When using Fourier expansions, specifies the number of processes to use in the z-coordinate direction.

`--part-info`

Prints detailed information about the generated partitioning, such as number of elements, number of local degrees of freedom and the number of boundary degrees of freedom.

`--part-only [int]`

Partition the mesh only into the specified number of partitions, write to file and exit. This can be used to pre-partition a very large mesh on a single high-memory node, prior to being executed on a multi-node cluster.

`--use-metis`

Forces the use of METIS for mesh partitioning. Requires the `NEKTAR_USE_METIS` option to be set.

`--use-scotch`

Forces the use of Scotch for mesh partitioning. If *Nektar++* is compiled with METIS support, the default is to use METIS.

Frequently Asked Questions

17.1 Compilation and Testing

Q. I compile Nektar++ successfully but, when I run ctest, all the tests fail. What might be wrong?

On Linux or Mac, if you compile the ThirdParty version of Boost, rather than using version supplied with your operating system (or MacPorts on a Mac), the libraries will be installed in the `ThirdParty/dist/lib` subdirectory of your Nektar++ directory. When Nektar++ executables are run, the Boost libraries will not be found as this path is not searched by default. To allow the Boost libraries to be found set the following environmental variable, substituting `$NEKTAR_HOME` with the absolute path of your Nektar++ directory:

- On Linux (sh, bash, etc)

```
export LD_LIBRARY_PATH=${NEKTAR_HOME}/ThirdParty/dist/lib
```

or (csh, etc)

```
setenv LD_LIBRARY_PATH ${NEKTAR_HOME}/ThirdParty/dist/lib
```

- On Mac

```
export DYLD_LIBRARY_PATH=${NEKTAR_HOME}/ThirdParty/dist/lib
```

Q. How to I compile Nektar++ to run in parallel?

Parallel execution of all Nektar++ solvers is available using MPI. To compile using MPI, enable the `NEKTAR_USE_MPI` option in the CMake configuration. On recent versions of

MPI, the solvers can still be run in serial when compiled with MPI. More information on Nektar++ compilation options is available in Section 1.3.5.

Q. When compiling Nektar++, I receive the following error:

```
CMake Error: The following variables are used in this
object, but they are set to NOTFOUND.
Please set them or make sure they are set and tested
correctly in the CMake files:
NATIVE_BLAS (ADVANCED)
  linked by target "LibUtilities" in directory
  /path/to/nektar++/library/LibUtilities
NATIVE_LAPACK (ADVANCED)
  linked by target "LibUtilities" in directory
  /path/to/nektar++/library/LibUtilities
```

This is caused by one of two problems:

- The BLAS and LAPACK libraries and development files are not installed. On Linux systems, both the LAPACK library package (usually called liblapack3 or lapack) and the development package (usually called liblapack-dev or lapack-devel) must be installed. Often the latter is missing.
- An alternative BLAS/LAPACK library should be used. HPC systems frequently use the Intel compilers (icc, icpc) and the Intel Math Kernel Library (MKL). This software should be made available (if using the modules environment) and the option `NEKTAR_USE_MKL` should be enabled.

Q. When I compile Nektar++ I receive an error

```
error: #error "SEEK_SET is #defined but must not be for
the C++ binding of MPI. Include mpi.h before stdio.h"
```

This can be fixed by including the flags

```
-DMPICH_IGNORE_CXX_SEEK -DMPICH_SKIP_MPICXX
```

in the `CMAKE_CXX_FLAGS` option within the `ccmake` configuration.

Q. After installing Nektar++ on my local HPC cluster, when I run the 'ctest' command, all the parallel tests fail. Why is this?

The parallel tests are those which include the word `parallel` or `par`. On many HPC systems, the MPI binaries used to execute jobs are not available on the login nodes, to prevent inadvertent parallel runs outside of the queuing system. Consequently, these

tests will not execute. To fully test the code, you can submit a job to the queuing system using a minimum of two cores, to run the `ctest` command.

Q. When running any Nektar++ executable on Windows, I receive an error that `zlib.dll` cannot be found. How do I fix this?

Windows searches for DLL files in directories specified in the `PATH` environmental variable. You should add the location of the ThirdParty files to your path. To fix this (example for Windows XP):

- As an administrator, open "System Properties" in control panel, select the "Advanced" tab, and select "Environment Variables".
- Edit the system variable 'path' and append
`C:\path\to\nektar++\ThirdParty\dist\bin`
 to the end, replacing `path\to\nektar++` appropriately.

Q. When compiling Nektar++ Thirdparty libraries I get an error "CMake Error: Problem extracting tar"

Nektar++ tries to download the appropriate ThirdParty libraries. However if the download protocols are restricted on your computer this may fail leading to the error "CMake Error: Problem extracting tar". These libraries are available from

`http://www.nektar.info/thirdparty/`

and can be downloaded directly into the `$NEKTAR_HOME/ThirdParty` directory

17.2 Usage

Q. How do I run a solver in parallel?

In a desktop environment, simply prefix the solver executable with the `mpirun` helper. For example, to run the Incompressible Navier-Stokes solver on a 4-core desktop computer, you would run

```
mpirun -np 4 IncNavierStokesSolver Cyl.xml
```

In a cluster environment, using PBS for example, the `mpiexec` command should be used.

Q. How can I generate a mesh for use with Nektar++?

Nektar++ supports a number of mesh input formats. These are converted to the Nektar++ native XML format (see Section 3) using the NekMesh utility (see Section 4. Supported formats include:

- Gmsh (.msh)
- Polygon (.ply)
- Nektar (.rea)
- Sentex (.sem)

Bibliography

- [1] M Ainsworth and S Sherwin. Domain decomposition preconditioners for p and hp finite element approximation of stokes equations. *COMPUTER METHODS IN APPLIED MECHANICS AND ENGINEERING*, 175:243–266, 1999.
- [2] R. R. Aliev and A. V. Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7:293–301, 1996.
- [3] Ivo Babuška and Manil Suri. The p and h-p versions of the finite element method, basic principles and properties. *SIAM review*, 36(4):578–632, 1994.
- [4] Y. Bao, R. Palacios, M. Graham, and S.J. Sherwin. Generalized “thick” strip modelling for vortex-induced vibration of long flexible cylinders. *J. Comp. Phys*, 321:1079–1097, 2016.
- [5] P-E Bernard, J-F Remacle, Richard Comblen, Vincent Legat, and Koen Hillewaert. High-order discontinuous galerkin schemes on general 2d manifolds applied to the shallow water equations. *Journal of Computational Physics*, 228(17):6514–6535, 2009.
- [6] C. D. Cantwell, S. J. Sherwin, R. M. Kirby, and P. H. J. Kelly. From h to p efficiently: selecting the optimal spectral/hp discretisation in three dimensions. *Math. Mod. Nat. Phenom.*, 6:84–96, 2011.
- [7] C. D. Cantwell, S. J. Sherwin, R. M. Kirby, and P. H. J. Kelly. From h to p efficiently: strategy selection for operator evaluation on hexahedral and tetrahedral elements. *Computers & Fluids*, 43:23–28, 2011.
- [8] CD Cantwell, D Moxey, A Comerford, A Bolis, G Rocco, G Mengaldo, D De Grazia, S Yakovlev, J-E Lombard, D Ekelschot, et al. Nektar++: An open-source spectral/hp element framework. *Computer Physics Communications*, 192:205–219, 2015.
- [9] D. De Grazia, G. Mengaldo, D. Moxey, P. E. Vincent, and S. J. Sherwin. Connections between the discontinuous galerkin method and high-order flux reconstruction

- schemes. *International Journal for Numerical Methods in Fluids*, 75(12):860–877, 2014.
- [10] S. Dong. A convective-like energy-stable open boundary condition for simulation of incompressible flows. *Journal of Computational Physics*, 302:300–328, 2015.
 - [11] S. Dong, G. E. Karniadakis, and C. Chrysosostomidis. A robust and accurate outflow boundary condition for incompressible flow simulations on severely-truncated unbounded domains. *Journal of Computational Physics*, 261:95–136, 2014.
 - [12] Niederer "et al.". Verification of cardiac tissue electrophysiology simulators using an n-version benchmark. *Philos Transact A Math Phys Eng Sci*, 369:4331–51, 2011.
 - [13] Roland Ewert and Wolfgang Schröder. Acoustic perturbation equations based on flow decomposition via source filtering. *Journal of Computational Physics*, 188(2):365–398, 7 2003.
 - [14] Abel Gargallo-Peiró, Xevi Roca, Jaime Peraire, and Josep Sarrate. Distortion and quality measures for validating and generating high-order tetrahedral meshes. *Engineering with Computers*, 31(3):423–437, 2015.
 - [15] Georg Geiser, Holger Nawroth, Arash Hosseinzadeh, Feichi Zhang, Henning Bockhorn, Peter Habisreuther, Johannes Janicka, Christian O. Paschereit, and Wolfgang Schröder. Thermoacoustics of a turbulent premixed flame. In *20th AIAA/CEAS Aeroacoustics Conference*. American Institute of Aeronautics and Astronautics.
 - [16] David Gottlieb, Steven A Orszag, and CAMBRIDGE HYDRODYNAMICS INC MA. *Numerical analysis of spectral methods*. SIAM, 1977.
 - [17] J.L. Guermond and J. Shen. Velocity-correction projection methods for incompressible flows. *SIAM J. Numer. Anal.*, 41:112–134, 2003.
 - [18] Jan S Hesthaven and Tim Warburton. Nodal high-order methods on unstructured grids: I. time-domain solution of maxwell's equations. *Journal of Computational Physics*, 181(1):186–221, 2002.
 - [19] B. E. Jordi, C. J. Cotter, and S. J. Sherwin. Encapsulated formulation of the selective frequency damping method. *Phys. Fluids*, 2014.
 - [20] G. E. Karniadakis, M. Israeli, and S. A. Orszag. High-order splitting methods for the incompressible Navier–Stokes equations. 97(2):414–443, 1991.
 - [21] G. E. Karniadakis and S. J. Sherwin. *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford Science Publications, 2005.
 - [22] Robert M Kirby and Spencer J Sherwin. Stabilisation of spectral/hp element methods through spectral vanishing viscosity: Application to fluid mechanics modelling. *Computer methods in applied mechanics and engineering*, 195(23):3128–3144, 2006.

- [23] Jonas Koko. Vectorized matlab codes for linear two-dimensional elasticity. *Scientific Programming*, 15(3):157–172, 2007.
- [24] Kilian Lackhove. *Hybrid Noise Simulation for Enclosed Configurations*. Doctoral thesis, Technische Universität Darmstadt, 2018.
- [25] C. H. Luo and Y. Rudy. A model of the ventricular cardiac action potential. depolarization repolarization and their interaction. *Circulation research*, 68:1501–1526, 1991.
- [26] R. J. Ramirez M. Courtemanche and S. Nattel. Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology*, 275:H301–H321, 1998.
- [27] Y. Maday, A. T. Patera, and E.M. Ronquist. An operator-integration-factor splitting method for time-dependent problems: Application to incompressible fluid flow. *J. Sci. Comp.*, 4:263–292, 1990.
- [28] Yvon Maday, Sidi M Ould Kaber, and Eitan Tadmor. Legendre pseudospectral viscosity method for nonlinear conservation laws. *SIAM Journal on Numerical Analysis*, 30(2):321–342, 1993.
- [29] Gianmarco Mengaldo, Daniele De Grazia, Freddie Witherden, Antony Farrington, Peter Vincent, Spencer Sherwin, and Joaquim Peiro. *A Guide to the Implementation of Boundary Conditions in Compact High-Order Methods for Compressible Aerodynamics*. American Institute of Aeronautics and Astronautics, 2014/08/10 2014.
- [30] RC Moura, SJ Sherwin, and Joaquim Peiró. Eigensolution analysis of spectral/hp continuous galerkin approximations to advection–diffusion problems: Insights into spectral vanishing viscosity. *Journal of Computational Physics*, 307:401–422, 2016.
- [31] D. Moxey, M. Hazan, J. Peiró, and S. J. Sherwin. An isoparametric approach to high-order curvilinear boundary-layer meshing. *Comp. Meth. Appl. Mech. Eng.*, 2014.
- [32] D. Moxey, M. Hazan, J. Peiró, and S. J. Sherwin. On the generation of curvilinear meshes through subdivision of isoparametric elements. to appear in proceedings of Tetrahedron IV, 2014.
- [33] David J Newman and George Em Karniadakis. A direct numerical simulation study of flow past a freely vibrating cable. *Journal of Fluid Mechanics*, 344:95–136, 1997.
- [34] Anthony T Patera. A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of computational Physics*, 54(3):468–488, 1984.
- [35] P.-O. Persson and J. Peraire. Sub-cell shock capturing for Discontinuous Galerkin methods. In *44th AIAA Aerospace Sciences Meeting and Exhibit*, page 112, 2006.

- [36] N Pignier. One-dimensional modelling of blood flow in the cardiovascular system, 2012.
- [37] CJ Roth. Pulse wave propagation in the human vascular system, 2012.
- [38] S Sherwin. A substepping navier-stokes splitting scheme for spectral/hp element discretisations. pages 43–52. Elsevier Science, 2003.
- [39] SJ Sherwin and M Ainsworth. Unsteady navier-stokes solvers using hybrid spectral/hp element methods. *APPLIED NUMERICAL MATHEMATICS*, 33:357–363, 2000.
- [40] SJ Sherwin, L Formaggia, J Peiró, and V Franke. Computational modelling of 1d blood flow with variable mechanical properties and its application to the simulation of wave propagation in the human arterial system. *Int. J. Numer. Meth. Fluids*, 43:673–700, 2003.
- [41] SJ Sherwin and G Em Karniadakis. Tetrahedral $\langle i \rangle$ hp $\langle /i \rangle$ finite elements: Algorithms and flow simulations. *Journal of Computational Physics*, 124(1):14–45, 1996.
- [42] K. H. W. J. ten Tusscher and A. V. Panfilov. Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology-Heart and Circulatory Physiology*, 291:H1088–H1100, 2006.
- [43] M Turner, J Peiró, and D Moxey. A Variational Framework for High-Order Mesh Generation. In *25th International Meshing Roundtable*, volume 163, pages 340–352, 2016.
- [44] Peter EJ Vos, Spencer J Sherwin, and Robert M Kirby. From h to p efficiently: Implementing finite and spectral/hp element methods to achieve optimal performance for low-and high-order discretisations. *Journal of Computational Physics*, 229(13):5161–5181, 2010.
- [45] N Westerhof. Anatomic studies of the human systemic arterial tree. *J. Biomech.*, 2:121–143, 1969.
- [46] D Xiu, SJ Sherwin, S Dong, and GE Karniadakis. Strong and auxiliary forms of the semi-lagrangian method for incompressible flows. *J. Sci. Comp.*, 25:323–346, 2005.
- [47] Olgierd Cecil Zienkiewicz and Robert Leroy Taylor. *Basic formulation and linear problems*. McGraw-Hill, 1989.