

XMDS2 Documentation

Release 2.2.2

Graham Dennis, Joe Hope and Mattias Johnsson

October 14, 2014

1	Introduction	1
2	Installation	3
2.1	Installers	3
2.2	Linux installer instructions	3
2.3	Mac OS X Installation	4
2.4	Manual installation from source	4
3	Quickstart Tutorial	9
4	Worked Examples	15
4.1	The nonlinear Schrödinger equation	15
4.2	Kubo Oscillator	19
4.3	Fibre Noise	21
4.4	Integer Dimensions	25
4.5	Wigner Function	28
4.6	Finding the Ground State of a BEC (continuous renormalisation)	31
4.7	Finding the Ground State of a BEC again	35
4.8	Multi-component Schrödinger equation	39
5	Reference section	43
5.1	Configuration, installation and runtime options	43
5.2	Useful XML Syntax	44
5.3	XMDS2 XML Schema	44
5.4	XMDS2 script elements	46
5.5	Modified Midpoint Method	74
5.6	Bulirsch-Stoer Algorithm	74
5.7	Error Scaling Behaviour	76
6	Advanced Topics	77
6.1	Importing data	77
6.2	Convolutions and Fourier transforms	79
6.3	‘Loose’ geometry_matching_mode	82
6.4	Dimension aliases	82
7	Frequently Asked Questions	83
7.1	XMDS scripts look complicated! How do I start?	83
7.2	Where can I get help?	83
7.3	How should I cite XMDS2?	83
7.4	I think I found a bug! Where should I report it?	83
7.5	How do I put time dependence into my vectors?	83
7.6	Can I specify the range of my domain and number of grid points at run-time?	84
7.7	When can I use IP operators (and why should I) and when must I use EX operators?	84
7.8	Visual Editors	84
8	Upgrading From XMDS 1.X	85

9	Optimisation Hints	87
9.1	Geometry and transform-based tricks	87
9.2	Reduce code complexity	87
9.3	Compiler and library tricks	89
9.4	Atom-optics-specific hints	91
10	xsil2graphics2	93
11	Developer Documentation	95
11.1	Test scripts	95
11.2	XMDS Documentation	96
11.3	How to update XMDS2 script validator (XML schema)	98
11.4	How to introduce a new integrator Stepper into the XMDS2 environment	101
11.5	Logical breakdown of XMDS2 Parsing Process	102
11.6	Directory layout	113
12	Licensing	115
12.1	GNU General Public License	115
12.2	GNU Free Documentation License	118
13	News	125
13.1	XMDS 2.2.1 “XMDS2 is a game of two halves” (September 30, 2014)	125
13.2	XMDS 2.2.0 “Out of cheese error” (January 13, 2014)	125
13.3	XMDS 2.1.4 “Well if this isn’t nice, I don’t know what is” (September 27, 2013)	126
13.4	XMDS 2.1.3 “Happy Mollusc” (June 7, 2013)	126
13.5	XMDS 2.1.2 “Happy Mollusc” (October 15, 2012)	126
13.6	XMDS 2.1 “Happy Mollusc” (June 14, 2012)	127
13.7	XMDS 2.0 “Shiny!” (September 13, 2010)	127
	Index	129

INTRODUCTION

Welcome to **XMDS2** (codenamed *xpdeint*), which is an all-new version of *XMDS*. Prepare for fast, easily-extended simulations with minimal code error.

Description: The purpose of XMDS2 is to simplify the process of creating simulations that solve systems of initial-value first-order partial and ordinary differential equations. Instead of going through the error-prone process of writing by hand thousands of lines of code, XMDS2 enables many problems to be described in a simple XML format. From this XML description XMDS2 writes a C++ simulation that solves the problem using fast algorithms. Anecdotally, the code generated by XMDS2 is as fast as, or faster than, code hand-written by an expert, but by using XMDS2 the time taken to produce the simulation is significantly reduced.

XMDS2 can be used to simulate almost any set of (coupled) (partial) (stochastic) differential equations in any number of dimensions. It can input and output data in a range of data formats, produce programs that can take command-line arguments, and produce parallelised code suitable for either modern computer architectures or distributed clusters.

If this is your first time with XMDS, then an ideal place to start is the *Quickstart Tutorial*, where we will show you how to write a basic simulation. *Installation* instructions should get you up and running and able to start playing with the large library of examples provided. The impatient will probably have good luck browsing the examples library included with the source, and the *Worked Examples* in this documentation for something that looks like their intended simulation.

If you are upgrading from **XMDS version 1.x**, then after following the installation instructions (*Installation*), you might want to have a quick read of the note for upgraders (*Upgrading From XMDS 1.X*). The syntax of the XML scripts has changed, but hopefully you will find the new scripts very intuitive.

Detailed advice on input/output issues, and ways to code more complicated simulations can be found in *Advanced Topics*.

XMDS2 should be cited as *Comput. Phys. Commun.* **184**, 201-208 (2013). **History:** XMDS was created in 1997 by Peter Drummond and Greg Collocutt, who conceived of the idea of using an XML-based code generator to simplify the process of integrating systems of equations with arbitrary dimension ¹. The first version was written in C, and featured a very flexible, strongly convergent stochastic algorithm: the *semi-implicit algorithm* ². Released under a public licence, it began to receive attention across several research groups. Over the next few years several people helped add new algorithms and features.

In 2003, the increased scope of the package prompted a complete rewrite by Greg Collocutt (using C++), which lead to **XMDS 1.0**. It was placed on sourceforge, and over a dozen developers contributed from 2003-2007 to help XMDS address a wider range of problems with a range of modern algorithms and support for parallel supercomputing. The documentation and installation method was improved enabling the software to be used in a wider context, and XMDS gained many users from across the world - in Australasia, Europe, America, India and China - over a variety of disciplines.

In 2008 a second complete rewrite was undertaken, largely by Graham Dennis (using Cheetah templates in python), leading to the current version **XMDS2**. This restructuring of the internal treatment of XML elements and the generated code allowed a new range of extensions to be explored. These included possibilities such as integrating multiple fields with different dimensionality, a more general set of differential equations that can be

¹ G.R.Collocutt and P.D.Drummond, *XmDs: eXtensible multi-dimensional simulator*, *Comput. Phys. Commun.* **142**, 219 (2001).

² M.J.Werner and P.D.Drummond, *Robust algorithms for solving stochastic partial differential equations*, *J. Comput. Phys.* **132**, 312 (1997).

solved efficiently, and multiple choices of transforms for transverse dimensions. This restructuring was paired with a equivalent effort to make the user experience more convenient through the use of single step installers for OS X and linux, as well as an extensive documentation and example library. Today, **XMDS2** has been downloaded approximately 40 000 times across 71 countries.

INSTALLATION

XMDS2 can be installed on any unix-like system including Linux, Tru64, and Mac OS X. It requires a C++ compiler, python, and several installed packages. Many of these packages are optional, but a good idea to obtain full functionality.

2.1 Installers

If you're using a Mac with OSX \geq 10.6, there is a simple drag-and-drop installer available.

If you are using a recent version of Ubuntu or Debian, XMDS2 is available as a package in your package / software manager. Note that if using the Ubuntu Software Centre rather than Synaptic, you must explicitly search for "xmids2" rather than "xmids", or you will only get XMDS1 as a result.

If you're using a version of Linux that's not Ubuntu / Debian, the easiest way to get started is to use the install shell script linked to below.

If we don't have an installer for your system, follow the *manual installation* instructions.

Linux (Ubuntu/Debian/Fedora/RedHat)	Download Linux shell script installer	Learn more
OS X 10.6 / 10.7 / 10.8 / 10.9	Download OS X Installer	Learn more
Other systems	Install from source	

If you have one of the supported operating systems listed above, but you find the installer doesn't work for you, please let us know by emailing `xmids-devel` <at> `lists.sourceforge.net`. If you'd like to tweak the linux installer to work on a distribution we haven't tested, we'd love you to do that and let us know!

2.2 Linux installer instructions

Note: If you're using Debian / Ubuntu, unless you want a developer install, you're probably better off installing XMDS2 via your package manager rather than this install script.

The linux installer has currently only been tested with Ubuntu, Debian, Fedora, and Red Hat. Download the installer here: http://svn.code.sf.net/p/xmids/code/trunk/xpdeint/admin/linux_installer.sh

Once you have downloaded it, make the installer executable and run it by typing the following into a terminal:

```
chmod u+x linux_installer.sh
./linux_installer.sh
```

Alternatively, if you wish to download and run the installer in a single step, you can use the following command:

```
/bin/bash -c "$(wget -qO - http://svn.code.sf.net/p/xmids/code/trunk/xpdeint/admin/linux_installer
```

The linux installer installs all XMDS2 dependencies from your native package manager where possible (`apt-get` for Ubuntu/Debian, `yum` for Fedora/Red Hat) but will download and compile the source code for libraries not available through the package manager. This means you'll need to be connected to the internet when running the installer. The installer should not be run with administrative privileges; it will ask you to enter your admin password at the appropriate point.

For instructions on how to install XMDS2 on systems where you lack administrative rights, see [Manual installation from source](#).

By default, this installer will install a known stable version of XMDS, which can be updated at any time by navigating to the XMDS directory and typing ‘make update’. To install the latest developer version at the beginning, simply run the installer with the `--develop` option.

Once XMDS2 has been installed, you can run it from the terminal by typing `xmds2`. See the [Quickstart Tutorial](#) for next steps.

2.3 Mac OS X Installation

2.3.1 Download

Mac OS X 10.6 (Snow Leopard) or later XMDS 2 installer: <http://sourceforge.net/projects/xmds/files/>

2.3.2 Using the Mac OS X Installer

A self-contained installer for Mac OS X 10.6 (Snow Leopard) and later is available from the link above. This installer is only compatible with Intel Macs. This means that the older PowerPC architecture is *not supported*. Xcode (Apple’s developer tools) is required to use this installer. Xcode is available for free from the Mac App Store for 10.7 or later, and is available on the install disk of earlier Macs as an optional install. For users of earlier operating systems (10.6.8 or earlier), it is possible to find a free copy of earlier versions of XCode on the Apple developer website (3.2.6 was the Snow Leopard compatible version). You will be prompted to install it if you haven’t already.

Once you have downloaded the XMDS installer, installation is as simple as dragging it to your Applications folder or any other location. Click the XMDS application to launch it, and press the “Launch XMDS Terminal” button to open a Terminal window customised to work with XMDS. The first time you do this, the application will complete the installation process. This process can take a few minutes, but is only performed once.

The terminal window launched by the XMDS application has environment variables set for using this installation of XMDS. You can run XMDS in this terminal by typing `xmds2`. See the [Quickstart Tutorial](#) for next steps.

To uninstall XMDS, drag the XMDS application to the trash. XMDS places some files in the directory `~/Library/XMDS`. Remove this directory to completely remove XMDS from your system.

This package includes binaries for [OpenMPI](#), [FFTW](#), [HDF5](#) and [GSL](#). These binaries are self-contained and do not overwrite any existing installations.

2.4 Manual installation from source

This installation guide will take you through a typical full install step by step. A large part of this procedure is obtaining and installing other libraries that XMDS2 requires, before installing XMDS2 itself.

While the instructions below detail these packages individually, if you have administrative privileges (or can request packages from your administrator) and if you are using an Ubuntu, Debian, Fedora or Red Hat linux distribution, you can install all required and optional dependencies (but not XMDS2 itself) via

Ubuntu / Debian:

```
sudo apt-get install build-essential subversion libopenmpi-dev openmpi-bin python-dev python-setuptools
```

Fedora / Red Hat:

```
sudo yum install gcc gcc-c++ make automake subversion openmpi-devel python-devel python-setuptools
```

You will still have to download and build FFTW 3.3 from source (see below) since prebuilt packages with MPI and AVX support are not currently available in the repositories.

Also note that this guide adds extra notes for users wishing to install XMDs2 using the SVN repository. This requires a few extra steps, but allows you to edit your copy, and/or update your copy very efficiently (with all the usual advantages and disadvantages of using unreleased material).

0. **You will need a copy of XMDs2.** The current release can be found at [Sourceforge](#), and downloaded as a single file. Download this file, and expand it in a directory where you want to keep the program files.

- Developer-only instructions: You can instead check out a working copy of the source using SVN. In a directory where you want to check out the repository, run: `svn checkout https://svn.code.sf.net/p/xmids/code/trunk/xpdeint .`

(Only do this once. To update your copy, type `svn up` or `make update` in the same directory, and then repeat any developer-only instructions below).

1. **You will need a working C++ compiler.** For Mac OS X, this means that the developer tools (XCode) should be installed. One common free compiler is `gcc`. It can be downloaded using your favourite package manager. XMDs2 can also use Intel's C++ compiler if you have it. Intel's compiler typically generates faster code than `gcc`, but it isn't free.

2. You will need a [python distribution](#).

- Mac OS X: It is pre-installed on Mac OS X 10.5 or later.
- Linux: It should be pre-installed. If not, install using your favourite package manager.

We require python 2.4 or greater. XMDs2 does not support Python 3.

3. **Install setuptools.** If you have root (sudo) access, the easy way to install this is by executing `ez_setup.py` from the repository. Simply type `sudo python ez_setup.py`

If you want to install into your home directory without root access, this is more complex:

- (a) First create the path `~/lib/python2.5/site-packages` (assuming you installed python version 2.5) and `~/bin`. Add “`export PYTHONPATH=~/lib/python2.5/site-packages:$PYTHONPATH`” and “`export PATH=~/bin:$PATH`” (if necessary) to your `.bashrc` file (and run “`~/bin`”).

- (b) If necessary install setuptools, by executing `ez_setup.py` from the repository. `python ez_setup.py --prefix=~`

If you use Mac OS X 10.5 or later, or installed the Enthought Python Distribution on Windows, then setuptools is already installed. Though if the next step fails, you may need to upgrade setuptools. To do that, type `sudo easy_install -U setuptools`

4. **Install HDF5 and FFTW3 (and optionally MPI).**

- (a) **HDF5 is a library for reading and writing the Hierarchical Data Format.** This is a standardised data format which it is suggested that people use in preference to the older ‘binary’ output (which is compatible with `xmids-1`). The advantage of HDF5 is that this data format is understood by a variety of other tools. `xsil2graphics2` provides support for loading data created in this format into Mathematica and Matlab.

XMDs2 only requires the single process version of HDF5, so there is no need to install the MPI version.

* Sidebar: Installing HDF5 from source follows a common pattern, which you may find yourself repeating later:

- i. After extracting the source directory, type `configure` and then add possible options.

(For HDF5, install with the `--prefix=/usr/local/` option if you want XMDs2 to find the library automatically. This is rarely needed for other packages.)

- ii. Once that is finished, type `make`. Then wait for that to finish, which will often be longer than you think.

iii. Finally, type `sudo make install` to install it into the appropriate directory.

- (b) **FFTW is the library XMDs2 uses for Fourier transforms.** This is the transform most people will use in their simulations. If you need support for MPI distributed simulations, you must configure FFTW to use MPI.

FFTW is available for free at the [FFTW website](#). To configure and compile it, follow the steps described in the HDF5 sidebar above. You may wish to add the `--enable-mpi` `--disable-fortran` options to the `configure` command.

- (c) **MPI is an API for doing parallel processing.** XMDs2 can use MPI to parallelise simulations on multi-processor/multi-core computers, or clusters of computers. Many supercomputing systems come with MPI libraries pre-installed. The [Open MPI](#) project has free distributions of this library available.

If you intend to take advantage of XMDs2's multi-processing features, you must install MPI, and configure FFTW3 to use it.

5. There are a range of optional installs. We recommend that you install them all if possible:

- (a) **A Matrix library like [ATLAS](#), Intel's [MKL](#) or the [GNU Scientific library \(GSL\)](#)** These libraries allow efficient implementation of transform spaces other than Fourier space. Mac OS X comes with its own (fast) matrix library.

- (b) **numpy is a tool that XMDs2 uses for automated testing.** It can be installed with `sudo easy_install numpy`.

Mac OS X 10.5 and later come with numpy.

- (c) **lxml is used to validate the syntax of scripts passed to XMDs2.** If you have root access, this can be installed with the command `sudo easy_install lxml`

You will need to have 'libxml2' and 'libxslt' installed (via your choice of package manager) to install lxml. Sufficient versions are preinstalled on Mac OS X 10.6.

If you don't have root access or want to install into your home directory, use:

```
easy_install --prefix=~ lxml
```

- (d) **h5py is needed for checking the results of XMDs2 tests that generate HDF5 output.** h5py requires numpy version 1.0.3 or later.

Upgrading [h5py](#) on Mac OS X is best done with the source of the package, as the `easy_install` option can get confused with multiple numpy versions. (Mac OS X Snow Leopard comes with version 1.2.1). After downloading the source, execute `python ./setup.py build` in the source directory, and then `python ./setup.py install` to install it.

6. **Install XMDs2 into your python path by running (in the `xmds-2.2.1/` directory):** `sudo ./setup.py develop`

If you want to install it into your home directory, type `./setup.py develop --prefix=~`

This step requires access to the net, as it downloads any dependent packages. If you are behind a firewall, you may need to set your `HTTP_PROXY` environment variable in order to do this.

- **Developer only instructions:** The Cheetah templates (*.tmpl) must be compiled into python. To do this, run `make` in the `xmds-2.2.1/` directory.
- **Developer-only instructions:** If you have 'numpy' installed, test XMDs2 by typing `./run_tests.py` in the `xmds-2.2.1/` directory. The package 'numpy' is one of the optional packages, with installation instructions below.
- **Developer-only instructions:** To build the user documentation, you first need to install sphinx, either via your package manager or: `sudo easy_install Sphinx`

Then, to build the documentation, in the `xmds-2.2.1/admin/userdoc-source/` directory run: `make html`

If this results in an error, you may need to run `sudo ./setup.py develop`

The generated html documentation will then be found at `xmds-2.2.1/documentation/index.html`

7. Configure XMDs2 by typing `xmds2 --reconfigure`. If XMDs2 is unable to find a library, you can tell XMDs2 where these libraries are located by adding `include` and `lib` search paths using the `--include-path` and `--lib-path` options. For example, if FFTW3 is installed in `/apps/fftw3` with headers in `/apps/fftw3/include/` and the libraries in `/apps/fftw3/lib`, (re)configure XMDs2 by typing:

- `xmds2 --reconfigure --include-path /apps/fftw3/include --lib-path /apps/fftw3/lib`.

If you need to use additional compiler or link flags for XMDs2 to use certain libraries, set the `CXXFLAGS` or `LINKFLAGS` environment variables before calling `xmds2 --reconfigure`. For example, to pass the compiler flag `-pedantic` and the link flag `-lm`, use:

- `CXXFLAGS="-pedantic" LINKFLAGS="-lm" xmds2 --reconfigure`.

Congratulations! You should now have a fully operational copy of `xmds2` and `xsil2graphics2`. You can test your copy using examples from the “`xmds-2.2.1/examples`” directory, and follow the worked examples in the [Quickstart Tutorial](#) and [Worked Examples](#).

QUICKSTART TUTORIAL

In this tutorial, we will create an XMDS2 script to solve the Lorenz Attractor, an example of a dynamical system that exhibits chaos. The equations describing this problem are

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x) \\ \frac{dy}{dt} &= x(\rho - z) - y \\ \frac{dz}{dt} &= xy - \beta z\end{aligned}$$

where we will solve with the parameters $\sigma = 10$, $\rho = 28$, $\beta = \frac{8}{3}$ and the initial condition $x(0) = y(0) = z(0) = 1$.

Below is a script that solves this problem (it's also saved as `examples/lorenz.xmgs` in your XMDS2 directory). Don't worry if it doesn't make sense yet, soon we'll break it down into easily digestible parts.

```
<?xml version="1.0" encoding="UTF-8"?>
<simulation xmgs-version="2">
  <name>lorenz</name>

  <!-- While not strictly necessary, the following two tags are handy. -->
  <author>Graham Dennis</author>
  <description>
    The Lorenz Attractor, an example of chaos.
  </description>

  <!--
  This element defines some constants. It can be used for other
  features as well, but we will go into that in more detail later.
  -->
  <features>
    <globals>
      <![CDATA[
        real sigma = 10.0;
        real b = 8.0/3.0;
        real r = 28.0;
      ]]>
    </globals>
  </features>

  <!--
  This part defines all of the dimensions used in the problem,
  in this case, only the dimension of 'time' is needed.
  -->
  <geometry>
    <propagation_dimension> t </propagation_dimension>
  </geometry>

  <!-- A 'vector' describes the variables that we will be evolving. -->
```

```

<vector name="position" type="real">
  <components>
    x y z
  </components>
  <initialisation>
    <![CDATA[
      x = y = z = 1.0;
    ]]>
  </initialisation>
</vector>

<sequence>
  <!--
  Here we define what differential equations need to be solved
  and what algorithm we want to use.
  -->
  <integrate algorithm="ARK89" interval="20.0" tolerance="1e-7">
    <samples>5000</samples>
    <operators>
      <integration_vectors>position</integration_vectors>
      <![CDATA[
        dx_dt = sigma*(y-x);
        dy_dt = r*x - y - x*z;
        dz_dt = x*y - b*z;
      ]]>
    </operators>
  </integrate>
</sequence>

<!-- This part defines what data will be saved in the output file -->
<output format="hdf5" filename="lorenz.xsil">
  <sampling_group initial_sample="yes">
    <moments>xR yR zR</moments>
    <dependencies>position</dependencies>
    <![CDATA[
      xR = x;
      yR = y;
      zR = z;
    ]]>
  </sampling_group>
</output>
</simulation>

```

You can compile and run this script with **XMDs2**. To compile the script, just pass the name of the script as an argument to **XMDs2**.

```

$ xmds2 lorenz.xmds
xmds2 version 2.1 "Happy Mollusc" (r2680)
Copyright 2000-2012 Graham Dennis, Joseph Hope, Mattias Johnsson
and the xmds team
Generating source code...
... done
Compiling simulation...
... done. Type './lorenz' to run.

```

Now we can execute the generated program 'lorenz'.

```

$ ./lorenz
Sampled field (for moment group #1) at t = 0.000000e+00
Sampled field (for moment group #1) at t = 4.000000e-03
Current timestep: 4.000000e-03
Sampled field (for moment group #1) at t = 8.000000e-03
Current timestep: 4.000000e-03

```

```
... many lines omitted ...

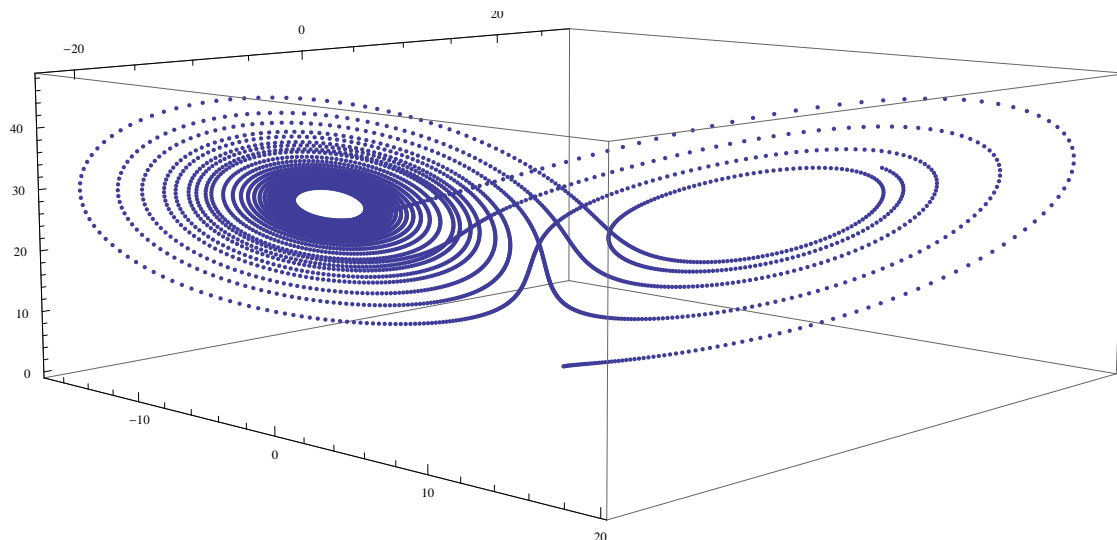
Current timestep: 4.000000e-03
Sampled field (for moment group #1) at t = 1.999600e+01
Current timestep: 4.000000e-03
Sampled field (for moment group #1) at t = 2.000000e+01
Current timestep: 4.000000e-03
Segment 1: minimum timestep: 9.997900e-06 maximum timestep: 4.000000e-03
  Attempted 7386 steps, 0.00% steps failed.
Generating output for lorenz
```

The program generated by **XMDs2** has now integrated your equations and produced two files. The first is the XML file “lorenz.xml”, which contains the all the information used to generate the simulation (including the XMDs2 code) and the metadata description of the output. The second file is named “lorenz.h5”, which is a [HDF5](#) file containing all of the output data. You can analysing these files yourself, or import them into your favourite visualisation/postprocessing tool. Here we will use the example of importing it into Mathematica. We run the included utility ‘xsil2graphics2’.

```
$ xsil2graphics2 -e lorenz.xml
xsil2graphics2 from xmds2 version 2.1 "Happy Mollusc" (r2680)
Generating output for Mathematica 6+.
Writing import script for 'lorenz.xml' to 'lorenz.nb'.
```

This has now generated the file ‘lorenz.nb’, which is a Mathematica notebook that loads the output data of the simulation. Loading it into Mathematica allows us to plot the points {xR1, yR1, zR1}:

```
ll = Transpose[{xR1, yR1, zR1}];
ListPointPlot3D[ll]
```



...and we see the lobes of the strange attractor. Now let us examine the code that produced this simulation.

First, we have the top level description of the code.

```
<?xml version="1.0" encoding="UTF-8"?>
<simulation xmds-version="2">
  <name>lorenz</name>

  <!-- While not strictly necessary, the following two tags are handy. -->
  <author>Graham Dennis</author>
  <description>
    The Lorenz Attractor, an example of chaos.
  </description>
```

One of the advantages of an XML format is that these tags are almost entirely self-explanatory. XMDs2 files

follow full XML syntax, so elements can be commented out using the `<!--` and `-->` brackets, and we have an example of that here.

The first line, `<?xml ...>`, just specifies the encoding and XML version. It is optional, but its presence helps some text editors perform the correct syntax highlighting.

The `<simulation>` element is mandatory, and encloses the entire simulation script.

The `<name>` element is optional, but recommended. It defines the name of the executable program that will be generated, as well as the default name of the output data files (although this can be over-ridden in the `<output>` element if desired). If `<name>` is not present, it will default to the filename of the script.

The next element we have used can be skipped entirely if you wish to use the default set of features and you don't want to define any global constants for your simulation.

```
<features>
  <globals>
    <![CDATA[
      real sigma = 10.0;
      real b = 8.0/3.0;
      real r = 28.0;
    ]]>
  </globals>
</features>
```

The `<features>` element can be used to choose a large number of features that will be discussed later, but here we have only used it to define a `<globals>` element. This element contains a block of text with `<![CDATA[` at the start and `]]>` at the end. These 'CDATA' blocks are used in several places in an XMD S script, and define a block of text that will be pasted directly into the generated C-code. They must therefore be formatted in legal C-syntax, and any legal C-syntax can be used. The `<globals>` element is placed at the top of the generated code, and can therefore be used to define any variables used in any other part of the simulation. Here we have defined our three real parameters. It is also possible to define variables that can be passed into the program at run-time, an example of which is given in the [Wigner Function](#) worked example.

The next element is the essential `<geometry>` element.

```
<geometry>
  <propagation_dimension> t </propagation_dimension>
</geometry>
```

This element is used to define all the dimensions in the problem. We only require the time dimension, which we are labelling 't', so this is a trivial example. We will discuss transverse dimensions in more detail in the next worked example ([The nonlinear Schrödinger equation](#)), where we deal with the integration of a partial differential equation rather than ordinary differential equations.

Next, we have the `<vector>` element.

```
<vector name="position" type="real">
  <components>
    x y z
  </components>
  <initialisation>
    <![CDATA[
      x = y = z = 1.0;
    ]]>
  </initialisation>
</vector>
```

We can define multiple vectors, but here we only need the variables that we wish to integrate. We named this vector "position", as it defines the position in phase space. These variables are real-valued (as opposed to, say, complex numbers), so we define `type="real"`. The `<components>` element defines the names of the elements of this vector, which we have called 'x', 'y' and 'z'. Finally, we provide the initial values of the variables in a CDATA block within the `<initialisation>` element.

Now we come to the heart of the simulation, where we define the evolution of our vector. This evolution is held

in the `<sequence>` element, which contains an ordered sequence of actions upon any defined vectors. Vectors can be altered with a `<filter>` element, or integrated in the propagation dimension with an `<integrate>` element.

```
<sequence>
  <integrate algorithm="ARK89" interval="20.0" tolerance="1e-7">
    <samples>5000</samples>
    <operators>
      <integration_vectors>position</integration_vectors>
      <![CDATA[
        dx_dt = sigma*(y-x);
        dy_dt = r*x - y - x*z;
        dz_dt = x*y - b*z;
      ]]>
    </operators>
  </integrate>
</sequence>
```

Here our sequence consists of a single `<integrate>` element. It contains several important pieces of information. At the heart, the `<operators>` element contains the equations of motion as described above, written in a very human-readable fashion. It also contains an `<integration_vectors>` element, which defines which vectors are used in this integrate block. We have only one vector defined in this simulation, so it is a trivial choice here.

All integrate blocks must define which algorithm is to be used - in this case the 8th (embedded 9th) order adaptive Runge-Kutta method, called “ARK89”. The details of different algorithms will be described later (FIXME: Link!), but for now all we need to know is that this algorithm requires a tolerance, and that smaller means more accurate, so we’ll make it 10^{-7} by setting `tolerance="1.0e-7"`. Finally, any integration will proceed a certain length in the propagation dimension, which is defined by the “interval” variable. This integrate block will therefore integrate the equations it contains with respect to the propagation dimension (‘t’) for 20.

The `<samples>` element says that the values of the output groups will be sampled 5000 times during this interval. The nature of the output is defined in the last element in the simulation: the `<output>` element.

```
<output format="hdf5" filename="lorenz.xsil">
  <sampling_group initial_sample="yes">
    <moments>xR yR zR</moments>
    <dependencies>position</dependencies>
    <![CDATA[
      xR = x;
      yR = y;
      zR = z;
    ]]>
  </sampling_group>
</output>
```

The two top-level arguments in the `<output>` element are “format” and “filename”. Here we define the output filename, although it would have defaulted to this value. We also choose the format to be HDF5, which is why the simulation resulted in the binary file “lorenz.h5” as well as “lorenz.xsil”. If we had instead said `format="ascii"`, then all of the output data would have been written in text form in “lorenz.xsil”.

The `<output>` element can contain any non-zero number of `<sampling_group>` elements, which specify the entire output of the program. They allow for subsampling, integration of some or all of the transverse dimensions, and/or conversion of some dimensions into Fourier space, but these will be described in more detail in the following examples. We have a `<dependencies>` element that specifies which vectors are needed for this output. We specify the list of output variables with a `<moments>` element, and then define them in CDATA block. In this case, we are simply defining the three variables that define our phase space.

And that’s it. This is quite a large framework to integrate three coupled ordinary differential equations, but the advantage of using XMD S2 is that vastly more complicated simulations can be performed without increasing the length or complexity of the XMD S2 script significantly. The *Worked Examples* section will provide more complicated examples with stochastic equations and partial differential equations. If you are moved to solve your own problem using XMD S2, then perhaps the most efficient method will be to take one of the worked examples

and adapt it to your needs. All of the examples in the documentation can be found in the “/examples” folder included with the installation.

WORKED EXAMPLES

One of the best ways to learn XMDS2 is to see several illustrative examples. Here are a set of example scripts and explanations of the code, which will be a good way to get started. As an instructional aid, they are meant to be read sequentially, but the adventurous could try starting with one that looked like a simulation they wanted to run, and adapt for their own purposes.

The nonlinear Schrödinger equation (partial differential equation)

Kubo Oscillator (stochastic differential equations)

Fibre Noise (stochastic partial differential equation using parallel processing)

Integer Dimensions (integer dimensions)

Wigner Function (two dimensional PDE using parallel processing, passing arguments in at run time)

Finding the Ground State of a BEC (continuous renormalisation) (PDE with continual renormalisation - computed vectors, filters, breakpoints)

Finding the Ground State of a BEC again (Hermite-Gaussian basis)

Multi-component Schrödinger equation (combined integer and continuous dimensions with matrix multiplication, aliases)

All of these scripts are available in the included “examples” folder, along with more examples that demonstrate other tricks. Together, they provide starting points for a huge range of different simulations.

4.1 The nonlinear Schrödinger equation

This worked example will show a range of new features that can be used in an **XMDS2** script, and we will also examine our first partial differential equation. We will take the one dimensional nonlinear Schrödinger equation, which is a common nonlinear wave equation. The equation describing this problem is:

$$\frac{\partial \phi}{\partial \xi} = \frac{i}{2} \frac{\partial^2 \phi}{\partial \tau^2} - \Gamma(\tau) \phi + i|\phi|^2 \phi$$

where ϕ is a complex-valued field, and $\Gamma(\tau)$ is a τ -dependent damping term. Let us look at an XMDS2 script that integrates this equation, and then examine it in detail.

```
<simulation xmds-version="2">
  <name>nlse</name>

  <author>Joe Hope</author>
  <description>
    The nonlinear Schrodinger equation in one dimension,
    which is a simple partial differential equation.
    We introduce several new features in this script.
  </description>
```

```
<features>
  <benchmark />
  <bing />
  <fftw plan="patient" />
  <openmp />
  <auto_vectorise />
  <globals>
    <![CDATA[
      const double energy = 4;
      const double vel = 0.3;
      const double hwhm = 1.0;
    ]]>
  </globals>
</features>

<geometry>
  <propagation_dimension> xi </propagation_dimension>
  <transverse_dimensions>
    <dimension name="tau" lattice="128" domain="(-6, 6)" />
  </transverse_dimensions>
</geometry>

<vector name="wavefunction" type="complex" dimensions="tau">
  <components> phi </components>
  <initialisation>
    <![CDATA[
      const double w0 = hwhm*sqrt(2/log(2));
      const double amp = sqrt(energy/w0/sqrt(M_PI/2));
      phi = amp*exp(-tau*tau/w0/w0)*exp(i*vel*tau);
    ]]>
  </initialisation>
</vector>

<vector name="dampingVector" type="real">
  <components> Gamma </components>
  <initialisation>
    <![CDATA[
      Gamma=1.0*(1-exp(-pow(tau*tau/4.0/4.0,10)));
    ]]>
  </initialisation>
</vector>

<sequence>
  <integrate algorithm="ARK45" interval="20.0" tolerance="1e-7">
    <samples>10 100 10</samples>
    <operators>
      <integration_vectors>wavefunction</integration_vectors>
      <operator kind="ex">
        <operator_names>Ltt</operator_names>
        <![CDATA[
          Ltt = -i*ktau*ktau*0.5;
        ]]>
      </operator>
      <![CDATA[
        dphi_dxi = Ltt[phi] - phi*Gamma + i*mod2(phi)*phi;
      ]]>
      <dependencies>dampingVector</dependencies>
    </operators>
  </integrate>
</sequence>

<output>
  <sampling_group basis="tau" initial_sample="yes">
```

```

    <moments>density</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
        density = mod2(phi);
    ]]>
</sampling_group>

<sampling_group basis="tau(0)" initial_sample="yes">
    <moments>normalisation</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
        normalisation = mod2(phi);
    ]]>
</sampling_group>

<sampling_group basis="ktau(32)" initial_sample="yes">
    <moments>densityK</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
        densityK = mod2(phi);
    ]]>
</sampling_group>

</output>
</simulation>

```

Let us examine the new items in the `<features>` element that we have demonstrated here. The existence of the `<benchmark>` element causes the simulation to be timed. The `<bing>` element causes the computer to make a sound upon the conclusion of the simulation. The `<fftw>` element is used to pass options to the [FFTW libraries for fast Fourier transforms](#), which are needed to do spectral derivatives for the partial differential equation. Here we used the option `plan="patient"`, which makes the simulation test carefully to find the fastest method for doing the FFTs. More information on possible choices can be found in the [FFTW documentation](#).

Finally, we use two tags to make the simulation run faster. The `<auto_vectorise>` element switches on several loop optimisations that exist in later versions of the GCC compiler. The `<openmp>` element turns on threaded parallel processing using the OpenMP standard where possible. These options are not activated by default as they only exist on certain compilers. If your code compiles with them on, then they are recommended.

Let us examine the `<geometry>` element.

```

<geometry>
    <propagation_dimension> xi </propagation_dimension>
    <transverse_dimensions>
        <dimension name="tau" lattice="128" domain="(-6, 6)" />
    </transverse_dimensions>
</geometry>

```

This is the first example that includes a transverse dimension. We have only one dimension, and we have labelled it “tau”. It is a continuous dimension, but only defined on a grid containing 128 points (defined with the lattice variable), and on a domain from -6 to 6. The default is that transforms in continuous dimensions are fast Fourier transforms, which means that this dimension is effectively defined on a loop, and the “tau=-6” and “tau=6” positions are in fact the same. Other transforms are possible, as are discrete dimensions such as an integer-valued index, but we will leave these advanced possibilities to later examples.

Two vector elements have been defined in this simulation. One defines the complex-valued wavefunction “phi” that we wish to evolve. We define the transverse dimensions over which this vector is defined by the `dimensions` tag in the description. By default, it is defined over all of the transverse dimensions in the `<geometry>` element, so even though we have omitted this tag for the second vector, it also assumes that the vector is defined over all of tau.

The second vector element contains the component “Gamma” which is a function of the transverse variable tau, as specified in the equation of motion for the field. This second vector could have been avoided in two ways. First, the function could have been written explicitly in the `integrate` block where it is required, but calculating it once

and then recalling it from memory is far more efficient. Second, it could have been included in the “wavefunction” vector as another component, but then it would have been unnecessarily complex-valued, it would have needed an explicit derivative in the equations of motion (presumably $d\Gamma_{dxi} = 0$), and it would have been Fourier transformed whenever the ϕ component was transformed. So separating it as its own vector is far more efficient.

The `<integrate>` element for a partial differential equation has some new features:

```
<integrate algorithm="ARK45" interval="20.0" tolerance="1e-7">
  <samples>10 100 10</samples>
  <operators>
    <integration_vectors>wavefunction</integration_vectors>
    <operator kind="ex">
      <operator_names>Ltt</operator_names>
      <![CDATA[
        Ltt = -i*ktau*ktau*0.5;
      ]]>
    </operator>
    <![CDATA[
      dphi_dxi = Ltt[phi] - phi*Gamma + i*mod2(phi)*phi;
    ]]>
    <dependencies>dampingVector</dependencies>
  </operators>
</integrate>
```

There are some trivial changes from the tutorial script, such as the fact that we are using the ARK45 algorithm rather than ARK89. Higher order algorithms are often better, but not always. Also, since this script has multiple output groups, we have to specify how many times each of these output groups are sampled in the `<samples>` element, so there are three numbers there. Besides the vectors that are to be integrated, we also specify that we want to use the vector “dampingVector” during this integration. This is achieved by including the `<dependencies>` element inside the `<operators>` element.

The equation of motion as written in the CDATA block looks almost identical to our desired equation of motion, except for the term based on the second derivative, which introduces an important new concept. Inside the `<operators>` element, we can define any number of operators. Operators are used to define functions in the transformed space of each dimension, which in this case is Fourier space. The derivative of a function is equivalent to multiplying by $i * k$ in Fourier space, so the $\frac{i}{2} \frac{\partial^2 \phi}{\partial \tau^2}$ term in our equation of motion is equivalent to multiplying by $-\frac{i}{2} k_\tau^2$ in Fourier space. In this example we define “Ltt” as an operator of exactly that form, and in the equation of motion it is applied to the field “phi”.

Operators can be explicit (`kind="ex"`) or in the interaction picture (`kind="ip"`). The interaction picture can be more efficient, but it restricts the possible syntax of the equation of motion. Safe utilisation of interaction picture operators will be described later, but for now let us emphasise that **explicit operators should be used** unless the user is clear what they are doing. That said, XMD52 will generate an error if the user tries to use interaction picture operators incorrectly. The `constant="yes"` option in the operator block means that the operator is not a function of the propagation dimension “xi”, and therefore only needs to be calculated once at the start of the simulation.

The output of a partial differential equation offers more possibilities than an ordinary differential equation, and we examine some in this example.

For vectors with transverse dimensions, we can sample functions of the vectors on the full lattice or a subset of the points. In the `<sampling_group>` element, we must add a string called “basis” that determines the space in which each transverse dimension is to be sampled, optionally followed by the number of points to be sampled in parentheses. If the number of points is not specified, it will default to a complete sampling of all points in that dimension. If a non-zero number of points is specified, it must be a factor of the lattice size for that dimension.

```
<sampling_group basis="tau" initial_sample="yes">
  <moments>density</moments>
  <dependencies>wavefunction</dependencies>
  <![CDATA[
    density = mod2(phi);
  ]]>
</sampling_group>
```

The first output group samples the mod square of the vector “phi” over the full lattice of 128 points.

If the lattice parameter is set to zero points, then the corresponding dimension is integrated.

```
<sampling_group basis="tau(0)" initial_sample="yes">
  <moments>normalisation</moments>
  <dependencies>wavefunction</dependencies>
  <![CDATA[
    normalisation = mod2(phi);
  ]]>
</sampling_group>
```

This second output group samples the normalisation of the wavefunction $\int d\tau |\phi(\tau)|^2$ over the domain of τ . This output requires only a single real number per sample, so in the integrate element we have chosen to sample it many more times than the vectors themselves.

Finally, functions of the vectors can be sampled with their dimensions in Fourier space.

```
<sampling_group basis="ktau(32)" initial_sample="yes">
  <moments>densityK</moments>
  <dependencies>wavefunction</dependencies>
  <![CDATA[
    densityK = mod2(phi);
  ]]>
</sampling_group>
```

The final output group above samples the mod square of the Fourier-space wavefunction phi on a sample of 32 points.

4.2 Kubo Oscillator

This example demonstrates the integration of a stochastic differential equation. We examine the Kubo oscillator, which is a complex variable whose phase is evolving according to a Wiener noise. In a suitable rotating frame, the equation of motion for the variable is

$$\frac{dz}{dt} = iz \eta$$

where $\eta(t)$ is the Wiener differential, and we interpret this as a Stratonovich equation. In other common notation, this is sometimes written:

$$dz = iz \circ dW$$

Most algorithms employed by XMDs require the equations to be input in the Stratonovich form. Ito differential equations can always be transformed into Stratonovich equations, and in this case the difference is equivalent to the choice of rotating frame. This equation is solved by the following XMDs2 script:

```
<simulation xmds-version="2">
  <name>kubo</name>
  <author>Graham Dennis and Joe Hope</author>
  <description>
    Example Kubo oscillator simulation
  </description>

  <geometry>
    <propagation_dimension> t </propagation_dimension>
  </geometry>

  <driver name="multi-path" paths="10000" />
```

```

<features>
  <error_check />
  <benchmark />
</features>

<noise_vector name="drivingNoise" dimensions="" kind="wiener" type="real" method="dsfmt" seed="
  <components>eta</components>
</noise_vector>

<vector name="main" type="complex">
  <components> z </components>
  <initialisation>
    <![CDATA[
      z = 1.0;
    ]]>
  </initialisation>
</vector>

<sequence>
  <integrate algorithm="SI" interval="10" steps="1000">
    <samples>100</samples>
    <operators>
      <integration_vectors>main</integration_vectors>
      <dependencies>drivingNoise</dependencies>
      <![CDATA[
        dz_dt = i*z*eta;
      ]]>
    </operators>
  </integrate>
</sequence>

<output>
  <sampling_group initial_sample="yes">
    <moments>zR zI</moments>
    <dependencies>main</dependencies>
    <![CDATA[
      zR = z.Re();
      zI = z.Im();
    ]]>
  </sampling_group>
</output>
</simulation>

```

The first new item in this script is the `<driver>` element. This element enables us to change top level management of the simulation. Without this element, XMD S2 will integrate the stochastic equation as described. With this element and the option `name="multi-path"`, it will integrate it multiple times, using different random numbers each time. The output will then contain the mean values and standard errors of your output variables. The number of integrations included in the averages is set with the `paths` variable.

In the `<features>` element we have included the `<error_check>` element. This performs the integration first with the specified number of steps (or with the specified tolerance), and then with twice the number of steps (or equivalently reduced tolerance). The output then includes the difference between the output variables on the coarse and the fine grids as the ‘error’ in the output variables. This error is particularly useful for stochastic integrations, where algorithms with adaptive step-sizes are less safe, so the number of integration steps must be user-specified.

We define the stochastic elements in a simulation with the `<noise_vector>` element.

```

<noise_vector name="drivingNoise" dimensions="" kind="wiener" type="real" method="dsfmt" seed="31
  <components>eta</components>
</noise_vector>

```

This defines a vector that is used like any other, but it will be randomly generated with particular statistics and

characteristics rather than initialised. The name, dimensions and type tags are defined just as for normal vectors. The names of the components are also defined in the same way. The noise is defined as a Wiener noise here (`kind = "wiener"`), which is a zero-mean Gaussian random noise with an average variance equal to the discretisation volume (here it is just the step size in the propagation dimension, as it is not defined over transverse dimensions). Other noise types are possible, including uniform and Poissonian noises, but we will not describe them in detail here.

We may also define a noise method to choose a non-default pseudo random number generator, and a seed for the random number generator. Using a seed can be very useful when debugging the behaviour of a simulation, and many compilers have pseudo-random number generators that are superior to the default option (`posix`).

The `integrate` block is using the semi-implicit algorithm (`algorithm="SI"`), which is a good default choice for stochastic problems, even though it is only second order convergent for deterministic equations. More will be said about algorithm choice later, but for now we should note that adaptive algorithms based on Runge-Kutta methods are not guaranteed to converge safely for stochastic equations. This can be particularly deceptive as they often succeed, particularly for almost any problem for which there is a known analytic solution.

We include elements from the noise vector in the equation of motion just as we do for any other vector. The default SI and Runge-Kutta algorithms converge to the *Stratonovich* integral. Ito stochastic equations can be converted to Stratonovich form and vice versa.

Executing the generated program ‘`kubo`’ gives slightly different output due to the “multi-path” driver.

```
$ ./kubo
Beginning full step integration ...
Starting path 1
Starting path 2

... many lines omitted ...

Starting path 9999
Starting path 10000
Beginning half step integration ...
Starting path 1
Starting path 2

... many lines omitted ...

Starting path 9999
Starting path 10000
Generating output for kubo
Maximum step error in moment group 1 was 4.942549e-04
Time elapsed for simulation is: 2.71 seconds
```

The maximum step error in each moment group is given in absolute terms. This is the largest difference between the full step integration and the half step integration. While a single path might be very stochastic:

The average over multiple paths can be increasingly smooth.

4.3 Fibre Noise

This simulation is a stochastic partial differential equation, in which a one-dimensional damped field is subject to a complex noise. This script can be found in `examples/fibre.xmds`.

$$\frac{\partial \psi}{\partial t} = -i \frac{\partial^2 \psi}{\partial x^2} - \gamma \psi + \beta \frac{1}{\sqrt{2}} (\eta_1(x) + i \eta_2(x))$$

where the noise terms $\eta_j(x, t)$ are Wiener differentials and the equation is interpreted as a Stratonovich differential equation. On a finite grid, these increments have variance $\frac{1}{\Delta x \Delta t}$.

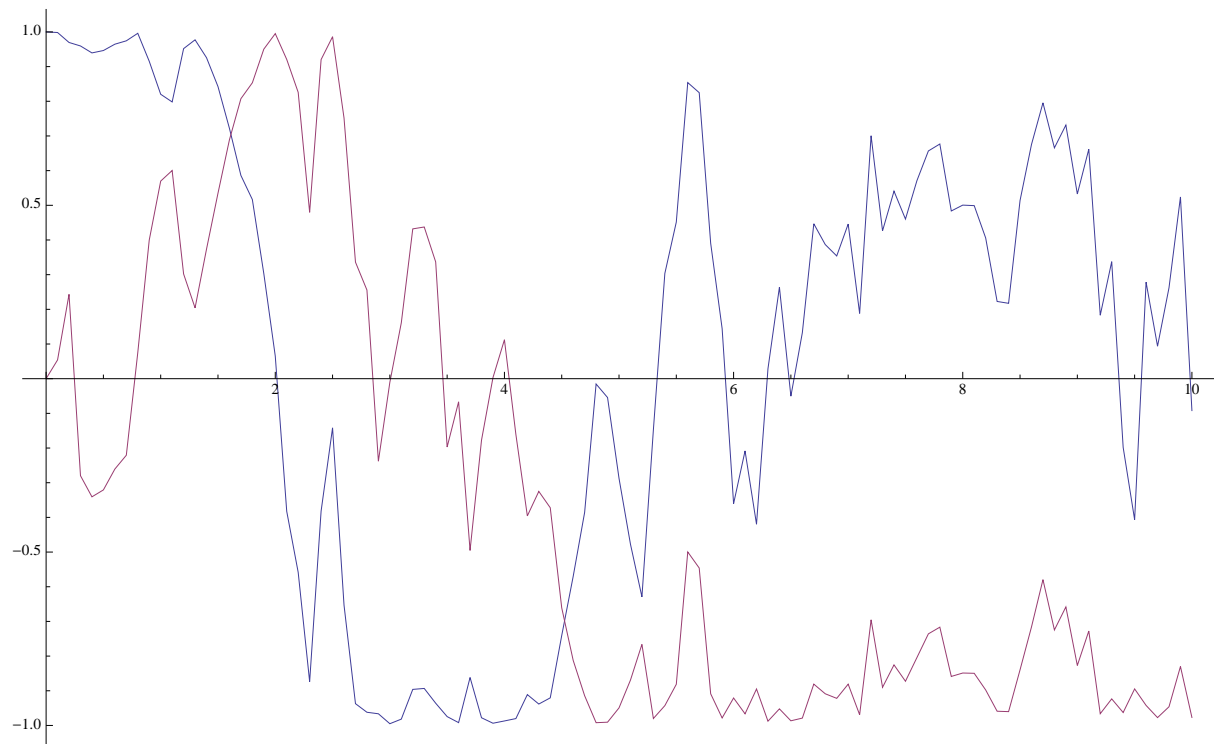


Figure 4.1: The mean value of the real and imaginary components of the z variable for a single path of the simulation.

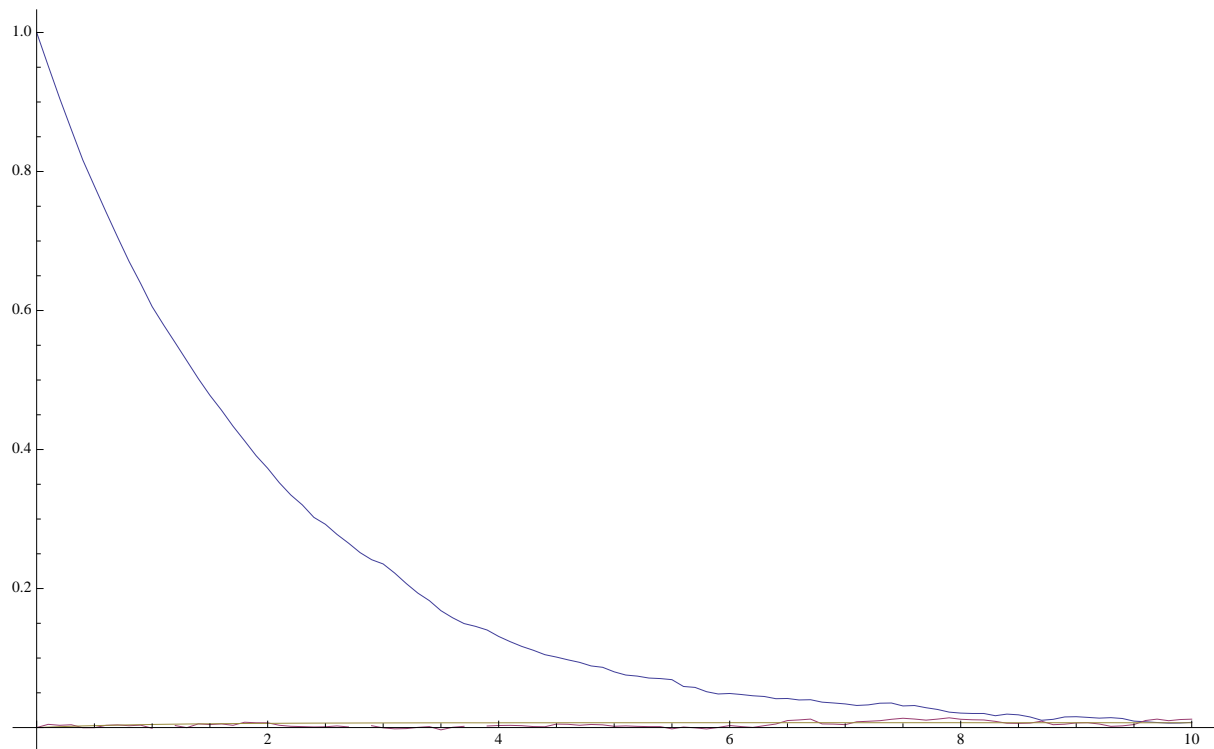


Figure 4.2: The mean and standard error of the z variable averaged over 10000 paths, as given by this simulation. It agrees within the standard error with the expected result of $\exp(-t/2)$.

```

<simulation xmds-version="2">
  <name>fibre</name>
  <author>Joe Hope and Graham Dennis</author>
  <description>
    Example fibre noise simulation
  </description>

  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="64" domain="(-5, 5)" />
    </transverse_dimensions>
  </geometry>

  <driver name="mpi-multi-path" paths="8" />

  <features>
    <auto_vectorise />
    <benchmark />
    <error_check />
    <globals>
      <![CDATA[
        const real ggamma = 1.0;
        const real beta = sqrt(M_PI*ggamma/10.0);
      ]]>
    </globals>
  </features>

  <noise_vector name="drivingNoise" dimensions="x" kind="wiener" type="complex" method="dsfmt" seed="123456789">
    <components>Eta</components>
  </noise_vector>

  <vector name="main" initial_basis="x" type="complex">
    <components>phi</components>
    <initialisation>
      <![CDATA[
        phi = 0.0;
      ]]>
    </initialisation>
  </vector>

  <sequence>
    <integrate algorithm="SI" iterations="3" interval="2.5" steps="200000">
      <samples>50</samples>
      <operators>
        <operator kind="ex">
          <operator_names>L</operator_names>
          <![CDATA[
            L = -i*kx*kx;
          ]]>
        </operator>
        <dependencies>drivingNoise</dependencies>
        <integration_vectors>main</integration_vectors>
        <![CDATA[
          dphi_dt = L[phi] - ggamma*phi + beta*Eta;
        ]]>
      </operators>
    </integrate>
  </sequence>

  <output>
    <sampling_group basis="kx" initial_sample="yes">
      <moments>pow_dens</moments>
    </sampling_group>
  </output>

```

```
<dependencies>main</dependencies>
<![CDATA[
    pow_dens = mod2(phi);
]]>
</sampling_group>
</output>
</simulation>
```

Note that the noise vector used in this example is complex-valued, and has the argument `dimensions="x"` to define it as a field of delta-correlated noises along the x-dimension.

This simulation demonstrates the ease with which XMD S2 can be used in a parallel processing environment. Instead of using the stochastic driver “multi-path”, we simply replace it with “mpi-multi-path”. This instructs XMD S2 to write a parallel version of the program based on the widespread [MPI standard](#). This protocol allows multiple processors or clusters of computers to work simultaneously on the same problem. Free open source libraries implementing this standard can be installed on a linux machine, and come standard on Mac OS X. They are also common on many supercomputer architectures. Parallel processing can also be used with deterministic problems to great effect, as discussed in the later example *Wigner Function*.

Executing this program is slightly different with the MPI option. The details can change between MPI implementations, but as an example:

```
$xmds2 fibre.xmds
xmds2 version 2.1 "Happy Mollusc" (r2543)
Copyright 2000-2012 Graham Dennis, Joseph Hope, Mattias Johnsson
                    and the xmds team
Generating source code...
... done
Compiling simulation...
... done. Type './fibre' to run.
```

Note that different compile options (and potentially a different compiler) are used by XMD S2, but this is transparent to the user. MPI simulations will have to be run using syntax that will depend on the MPI implementation. Here we show the version based on the popular open source [Open-MPI](#) implementation.

```
$ mpirun -np 4 ./fibre
Found enlightenment... (Importing wisdom)
Planning for x <---> kx transform... done.
Beginning full step integration ...
Rank[0]: Starting path 1
Rank[1]: Starting path 2
Rank[2]: Starting path 3
Rank[3]: Starting path 4
Rank[3]: Starting path 8
Rank[0]: Starting path 5
Rank[1]: Starting path 6
Rank[2]: Starting path 7
Rank[3]: Starting path 4
Beginning half step integration ...
Rank[0]: Starting path 1
Rank[2]: Starting path 3
Rank[1]: Starting path 2
Rank[3]: Starting path 8
Rank[0]: Starting path 5
Rank[2]: Starting path 7
Rank[1]: Starting path 6
Generating output for fibre
Maximum step error in moment group 1 was 4.893437e-04
Time elapsed for simulation is: 20.99 seconds
```

In this example we used four processors. The different processors are labelled by their “Rank”, starting at zero. Because the processors are working independently, the output from the different processors can come in a randomised order. In the end, however, the .xsil and data files are constructed identically to the single processor

outputs.

The analytic solution to the stochastic averages of this equation is given by

$$\langle |\psi(k, t)|^2 \rangle = \exp(-2\gamma t) |\psi(k, 0)|^2 + \frac{\beta^2 L_x}{4\pi\gamma} (1 - \exp(-2\gamma t))$$

where L_x is the length of the x domain. We see that a single integration of these equations is quite chaotic:

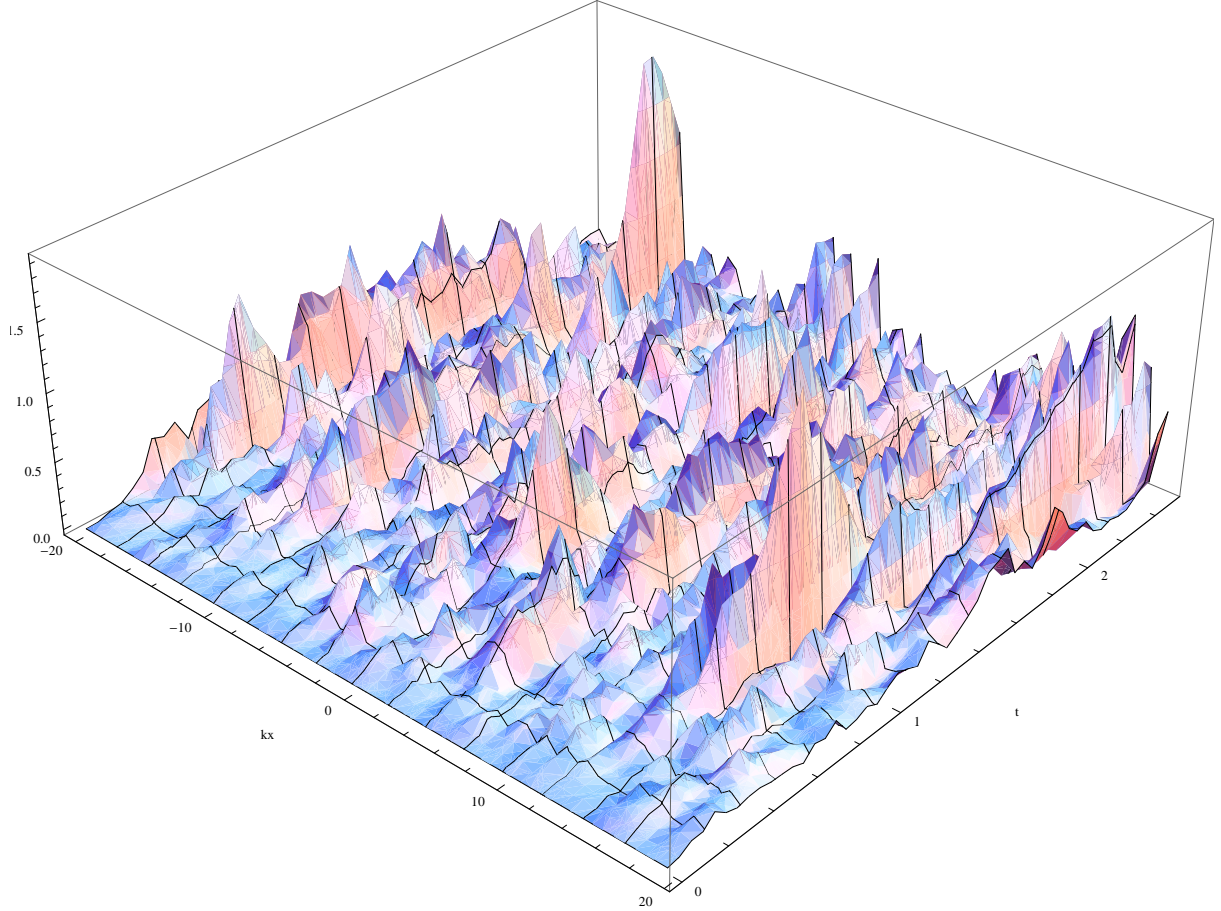


Figure 4.3: The momentum space density of the field as a function of time for a single path realisation.

while an average of 1024 paths (change `paths="8"` to `paths="1024"` in the `<driver>` element) converges nicely to the analytic solution:

4.4 Integer Dimensions

This example shows how to handle systems with integer-valued transverse dimensions. We will integrate the following set of equations

$$\frac{dx_j}{dt} = x_j (x_{j-1} - x_{j+1})$$

where x_j are complex-valued variables defined on a ring, such that $j \in \{0, j_{max}\}$ and the $x_{j_{max}+1}$ variable is identified with the variable x_0 , and the variable x_{-1} is identified with the variable $x_{j_{max}}$.

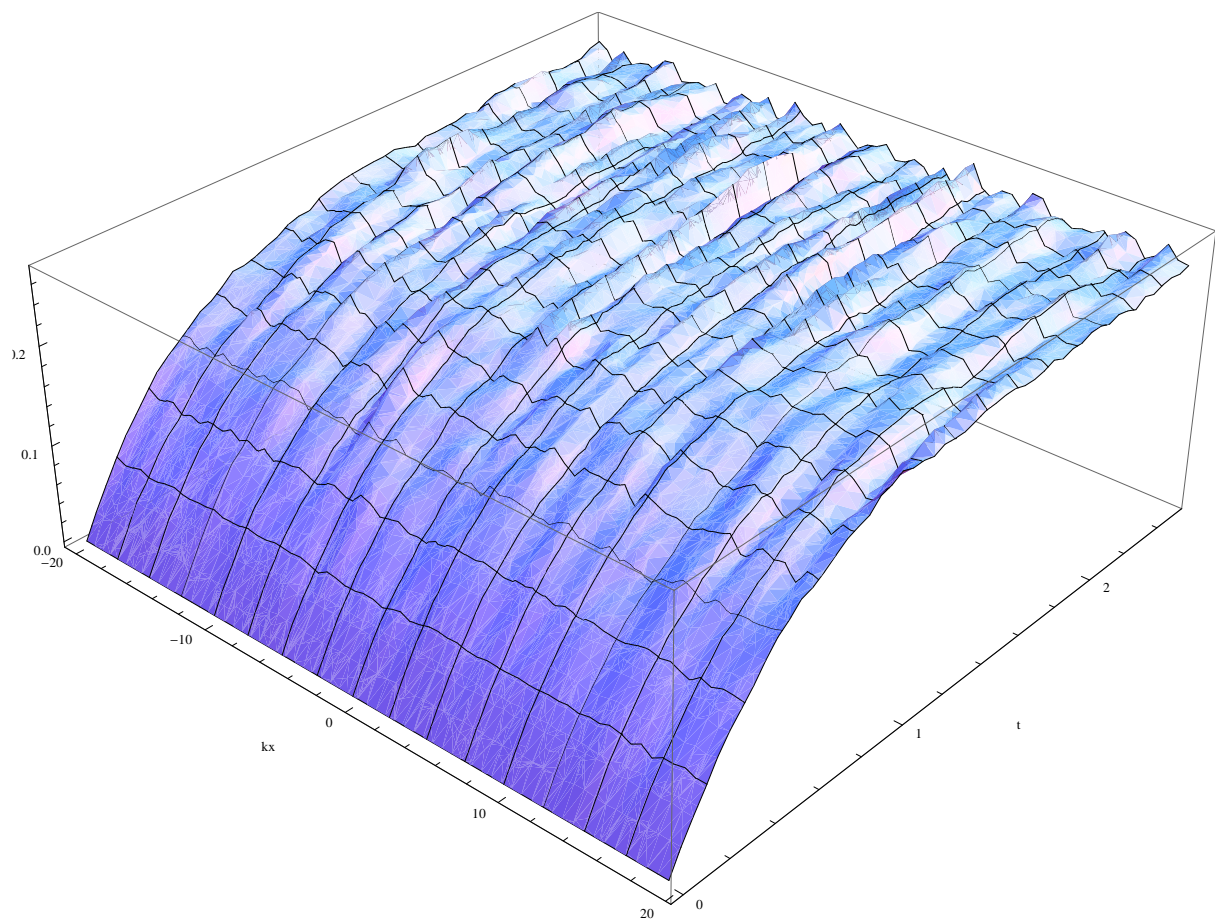


Figure 4.4: The momentum space density of the field as a function of time for an average of 1024 paths.

```

<simulation xmds-version="2">
  <name>integer_dimensions</name>
  <author>Graham Dennis</author>
  <description>
    XMDs2 script to test integer dimensions.
  </description>

  <features>
    <benchmark />
    <error_check />
    <bing />
    <diagnostics /> <!-- This will make sure that all nonlocal accesses of dimensions are safe -->
  </features>

  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="j" type="integer" lattice="5" domain="(0,4)" />
    </transverse_dimensions>
  </geometry>

  <vector name="main" type="complex">
    <components> x </components>
    <initialisation>
      <![CDATA[
        x = 1.0e-3;
        x(j => 0) = 1.0;
      ]]>
    </initialisation>
  </vector>

  <sequence>
    <integrate algorithm="ARK45" interval="60" steps="25000" tolerance="1.0e-9">
      <samples>1000</samples>
      <operators>
        <integration_vectors>main</integration_vectors>
        <![CDATA[
          long j_minus_one = (j-1) % _lattice_j;
          if (j_minus_one < 0)
            j_minus_one += _lattice_j;
          long j_plus_one = (j+1) % _lattice_j;
          dx_dt(j => j) = x(j => j)*(x(j => j_minus_one) - x(j => j_plus_one));
        ]]>
      </operators>
    </integrate>
  </sequence>

  <output>
    <sampling_group basis="j" initial_sample="yes">
      <moments>xR</moments>
      <dependencies>main</dependencies>
      <![CDATA[
        xR = x.Re();
      ]]>
    </sampling_group>
  </output>
</simulation>

```

The first extra feature we have used in this script is the `<diagnostics>` element. It performs run-time checking that our generated code does not accidentally attempt to access a part of our vector that does not exist. Removing this tag will increase the speed of the simulation, but its presence helps catch coding errors.

The simulation defines a vector with a single transverse dimension labelled “j”, of type “integer” (“int” and “long”

can also be used as synonyms for “integer”). In the absence of an explicit type, the dimension is assumed to be real-valued. The dimension has a “domain” argument as normal, defining the minimum and maximum values of the dimension’s range. The lattice element, if specified, is used as a check on the size of the domain, and will create an error if the two do not match.

Integer-valued dimensions can be called non-locally. Real-valued dimensions are typically coupled non-locally only through local operations in the transformed space of the dimension, but can be called non-locally in certain other situations as described in [the reference](#). The syntax for calling integer dimensions non-locally can be seen in the initialisation CDATA block:

```
x = 1.0e-3;
x(j => 0) = 1.0;
```

where the syntax `x(j => 0)` is used to reference the variable x_0 directly. We see a more elaborate example in the integrate CDATA block:

```
dx_dt(j => j) = x(j => j)*(x(j => j_minus_one) - x(j => j_plus_one));
```

where the vector “x” is called using locally defined variables. This syntax is chosen so that multiple dimensions can be addressed non-locally with minimal possibility for confusion.

4.5 Wigner Function

This example integrates the two-dimensional partial differential equation

$$\frac{\partial W}{\partial t} = \left[\left(\omega + \frac{U_{int}}{\hbar} (x^2 + y^2 - 1) \right) \left(x \frac{\partial}{\partial y} - y \frac{\partial}{\partial x} \right) - \frac{U_{int}}{16\hbar} \left(x \left(\frac{\partial^3}{\partial x^2 \partial y} + \frac{\partial^3}{\partial y^3} \right) - y \left(\frac{\partial^3}{\partial y^2 \partial x} + \frac{\partial^3}{\partial x^3} \right) \right) \right] W(x, y, t)$$

with the added restriction that the derivative is forced to zero outside a certain radius. This extra condition helps maintain the long-term stability of the integration. The script can be found in `examples/wigner_arguments_mpi.xmgs` under your XMDS2 installation directory.

```
<simulation xmgs-version="2">
  <name>wigner</name>
  <author>Graham Dennis and Joe Hope</author>
  <description>
    Simulation of the Wigner function for an anharmonic oscillator with the initial state
    being a coherent state.
  </description>
  <features>
    <benchmark />
    <globals>
      <![CDATA[
        real Uint_hbar_on16;
      ]]>
    </globals>
    <arguments>
      <argument name="omega" type="real" default_value="0.0" />
      <argument name="alpha_0" type="real" default_value="3.0" />
      <argument name="absorb" type="real" default_value="8.0" />
      <argument name="width" type="real" default_value="0.3" />
      <argument name="Uint_hbar" type="real" default_value="1.0" />
      <![CDATA[
        /* derived constants */
        Uint_hbar_on16 = Uint_hbar/16.0;
      ]]>
  </arguments>
</simulation>
```

```

</arguments>
<bing />
<fftw plan="patient" />
<openmp />
</features>

<driver name="distributed-mpi" />

<geometry>
  <propagation_dimension> t </propagation_dimension>
  <transverse_dimensions>
    <dimension name="x" lattice="128" domain="(-6, 6)" />
    <dimension name="y" lattice="128" domain="(-6, 6)" />
  </transverse_dimensions>
</geometry>

<vector name="main" initial_basis="x y" type="complex">
  <components> W </components>
  <initialisation>
    <![CDATA[
      W = 2.0/M_PI * exp(-2.0*(y*y + (x-alpha_0)*(x-alpha_0)));
    ]]>
  </initialisation>
</vector>

<vector name="dampConstants" initial_basis="x y" type="real">
  <components>damping</components>
  <initialisation>
    <![CDATA[
      if (sqrt(x*x + y*y) > _max_x-width)
        damping = 0.0;
      else
        damping = 1.0;
    ]]>
  </initialisation>
</vector>

<sequence>
  <integrate algorithm="ARK89" tolerance="1e-7" interval="7.0e-4" steps="100000">
    <samples>50</samples>
    <operators>
      <operator kind="ex">
        <operator_names>Lx Ly Lxxx Lxxy Lxyy Lyyy</operator_names>
        <![CDATA[
          Lx = i*kx;
          Ly = i*ky;
          Lxxx = -i*kx*kx*kx;
          Lxxy = -i*kx*kx*ky;
          Lxyy = -i*kx*ky*ky;
          Lyyy = -i*ky*ky*ky;
        ]]>
      </operator>
      <integration_vectors>main</integration_vectors>
      <dependencies>dampConstants</dependencies>
      <![CDATA[
        real rotation = omega + Uint_hbar*(-1.0 + x*x + y*y);

        dW_dt = damping * ( rotation * (x*Ly[W] - y*Lx[W])
          - Uint_hbar_on16*( x*(Lxxy[W] + Lyyy[W]) - y*(Lxyy[W] + Lxxx[W]) )
        );
      ]]>
    </operators>
  </integrate>

```

```

</sequence>

<output>
  <sampling_group basis="x y" initial_sample="yes">
    <moments>WR WI</moments>
    <dependencies>main</dependencies>
    <![CDATA[
      _SAMPLE_COMPLEX(W);
    ]]>
  </sampling_group>
</output>
</simulation>

```

This example demonstrates two new features of XMDs2. The first is the use of parallel processing for a deterministic problem. The FFTW library only allows MPI processing of multidimensional vectors. For multidimensional simulations, the generated program can be parallelised simply by adding the `name="distributed-mpi"` argument to the `<driver>` element.

```

$ xmds2 wigner_argument_mpi.xmds
xmds2 version 2.1 "Happy Mollusc" (r2680)
Copyright 2000-2012 Graham Dennis, Joseph Hope, Mattias Johnsson
                        and the xmds team
Generating source code...
... done
Compiling simulation...
... done. Type './wigner' to run.

```

To use multiple processors, the final program is then called using the (implementation specific) MPI wrapper:

```

$ mpirun -np 2 ./wigner
Planning for (distributed x, y) <----> (distributed ky, kx) transform... done.
Planning for (distributed x, y) <----> (distributed ky, kx) transform... done.
Sampled field (for moment group #1) at t = 0.000000e+00
Current timestep: 5.908361e-06
Sampled field (for moment group #1) at t = 1.400000e-05
Current timestep: 4.543131e-06

...

```

The possible acceleration achievable when parallelising a given simulation depends on a great many things including available memory and cache. As a general rule, it will improve as the simulation size gets larger, but the easiest way to find out is to test. The optimum speed up is obviously proportional to the number of available processing cores.

The second new feature in this simulation is the `<arguments>` element in the `<features>` block. This is a way of specifying global variables with a given type that can then be input at run time. The variables are specified in a self explanatory way

```

<arguments>
  <argument name="omega" type="real" default_value="0.0" />
  ...
  <argument name="Uint_hbar" type="real" default_value="1.0" />
</arguments>

```

where the “default_value” is used as the valuable of the variable if no arguments are given. In the absence of the generating script, the program can document its options with the `--help` argument:

```

$ ./wigner --help
Usage: wigner --omega <real> --alpha_0 <real> --absorb <real> --width <real> --Uint_hbar <real>

Details:
Option          Type          Default value
-o, --omega      real          0.0
-a, --alpha_0    real          3.0

```

```
-b, --absorb      real      8.0
-w, --width      real      0.3
-U, --Uint_hbar  real      1.0
```

We can change one or more of these variables' values in the simulation by passing it at run time.

```
$ mpirun -np 2 ./wigner --omega 0.1 --alpha_0 2.5 --Uint_hbar 0
Found enlightenment... (Importing wisdom)
Planning for (distributed x, y) <---> (distributed ky, kx) transform... done.
Planning for (distributed x, y) <---> (distributed ky, kx) transform... done.
Sampled field (for moment group #1) at t = 0.000000e+00
Current timestep: 1.916945e-04
```

...

The values that were used for the variables, whether default or passed in, are stored in the output file (wigner.xsil).

<info>

```
Script compiled with XMDs2 version 2.1 "Happy Mollusc" (r2680)
See http://www.xmds.org for more information.
```

Variables that can be specified on the command line:

```
Command line argument omega = 1.000000e-01
Command line argument alpha_0 = 2.500000e+00
Command line argument absorb = 8.000000e+00
Command line argument width = 3.000000e-01
Command line argument Uint_hbar = 0.000000e+00
```

</info>

Finally, note the shorthand used in the output group

```
<![CDATA[
  _SAMPLE_COMPLEX(W);
]]>
```

which is short for

```
<![CDATA[
  WR = W.Re();
  WI = W.Im();
]]>
```

4.6 Finding the Ground State of a BEC (continuous renormalisation)

This simulation solves another partial differential equation, but introduces several powerful new features in XMDs2. The nominal problem is the calculation of the lowest energy eigenstate of a non-linear Schrödinger equation:

$$\frac{\partial \phi}{\partial t} = i \left[\frac{1}{2} \frac{\partial^2}{\partial y^2} - V(y) - U_{int} |\phi|^2 \right] \phi$$

which can be found by evolving the above equation in imaginary time while keeping the normalisation constant. This causes eigenstates to exponentially decay at the rate of their eigenvalue, so after a short time only the state with the lowest eigenvalue remains. The evolution equation is straightforward:

$$\frac{\partial \phi}{\partial t} = \left[\frac{1}{2} \frac{\partial^2}{\partial y^2} - V(y) - U_{int} |\phi|^2 \right] \phi$$

but we will need to use new XMDs2 features to manage the normalisation of the function $\phi(y, t)$. The normalisation for a non-linear Schrödinger equation is given by $\int dy |\phi(y, t)|^2 = N_{particles}$, where $N_{particles}$ is the number of particles described by the wavefunction.

The code for this simulation can be found in `examples/groundstate_workedexamples.xml`:

```
<simulation xmds-version="2">
  <name>groundstate</name>
  <author>Joe Hope</author>
  <description>
    Calculate the ground state of the non-linear Schrodinger equation in a harmonic magnetic trap
    This is done by evolving it in imaginary time while re-normalising each timestep.
  </description>

  <features>
    <auto_vectorise />
    <benchmark />
    <bing />
    <fftw plan="exhaustive" />
    <globals>
      <![CDATA[
        const real Uint = 2.0;
        const real Nparticles = 5.0;
      ]]>
    </globals>
  </features>

  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="y" lattice="256" domain="(-15.0, 15.0)" />
    </transverse_dimensions>
  </geometry>

  <vector name="potential" initial_basis="y" type="real">
    <components> V1 </components>
    <initialisation>
      <![CDATA[
        V1 = 0.5*y*y;
      ]]>
    </initialisation>
  </vector>

  <vector name="wavefunction" initial_basis="y" type="complex">
    <components> phi </components>
    <initialisation>
      <![CDATA[
        if (fabs(y) < 3.0) {
          phi = 1.0;
          // This will be automatically normalised later
        } else {
          phi = 0.0;
        }
      ]]>
    </initialisation>
  </vector>

  <computed_vector name="normalisation" dimensions="" type="real">
    <components> Ncalc </components>
    <evaluation>
      <dependencies basis="y">wavefunction</dependencies>
      <![CDATA[
        // Calculate the current normalisation of the wave function.
        Ncalc = mod2(phi);
      ]]>
    </evaluation>
  </computed_vector>
</simulation>
```

```

]]>
</evaluation>
</computed_vector>

<sequence>
  <filter>
    <![CDATA[
      printf("Hello world from a filter segment!\n");
    ]]>
  </filter>

  <filter>
    <dependencies>normalisation wavefunction</dependencies>
    <![CDATA[
      phi *= sqrt(Nparticles/Ncalc);
    ]]>
  </filter>

  <integrate algorithm="ARK45" interval="1.0" steps="4000" tolerance="1e-10">
    <samples>25 4000</samples>
    <filters where="step end">
      <filter>
        <dependencies>wavefunction normalisation</dependencies>
        <![CDATA[
          // Correct normalisation of the wavefunction
          phi *= sqrt(Nparticles/Ncalc);
        ]]>
      </filter>
    </filters>
    <operators>
      <operator kind="ip">
        <operator_names>T</operator_names>
        <![CDATA[
          T = -0.5*ky*ky;
        ]]>
      </operator>
      <integration_vectors>wavefunction</integration_vectors>
      <dependencies>potential</dependencies>
      <![CDATA[
        dphi_dt = T[phi] - (V1 + Uint*mod2(phi))*phi;
      ]]>
    </operators>
  </integrate>

  <breakpoint filename="groundstate_break.xsil">
    <dependencies basis="ky">wavefunction </dependencies>
  </breakpoint>
</sequence>

<output>
  <sampling_group basis="y" initial_sample="yes">
    <moments>norm_dens</moments>
    <dependencies>wavefunction normalisation</dependencies>
    <![CDATA[
      norm_dens = mod2(phi);
    ]]>
  </sampling_group>

  <sampling_group initial_sample="yes">
    <moments>norm</moments>
    <dependencies>normalisation</dependencies>
    <![CDATA[

```

```

        norm = Ncalc;
    ]]>
</sampling_group>
</output>
</simulation>

```

We have used the `plan="exhaustive"` option in the `<fftw>` element to ensure that the absolute fastest transform method is found. Because the FFTW package stores the results of its tests (by default in the `~/.xmds/wisdom` directory), this option does not cause significant computational overhead, except perhaps on the very first run of a new program.

This simulation introduces the first example of a very powerful feature in XMD S2: the `<computed_vector>` element. This has syntax like any other vector, including possible dependencies on other vectors, and an ability to be used in any element that can use vectors. The difference is that, much like noise vectors, computed vectors are recalculated each time they are required. This means that a computed vector can never be used as an integration vector, as its values are not stored. However, computed vectors allow a simple and efficient method of describing complicated functions of other vectors. Computed vectors may depend on other computed vectors, allowing for spectral filtering and other advanced options. See for example, the *Advanced Topics* section on *Convolutions and Fourier transforms*.

The difference between a computed vector and a stored vector is emphasised by the replacement of the `<initialisation>` element with an `<evaluation>` element. Apart from the name, they have virtually identical purpose and syntax.

```

<computed_vector name="normalisation" dimensions="" type="real">
  <components> Ncalc </components>
  <evaluation>
    <dependencies basis="y">wavefunction</dependencies>
    <![CDATA[
      // Calculate the current normalisation of the wave function.
      Ncalc = mod2(phi);
    ]]>
  </evaluation>
</computed_vector>

```

Here, our computed vector has no transverse dimensions and depends on the components of “wavefunction”, so the extra transverse dimensions are integrated out. This code therefore integrates the square modulus of the field, and returns it in the variable “Ncalc”. This will be used below to renormalise the “phi” field. Before we examine that process, we have to introduce the `<filter>` element.

The `<filter>` element can be placed in the `<sequence>` element, or inside `<integrate>` elements as we will see next. Elements placed in the `<sequence>` element are executed in the order they are found in the .xmds file. Filter elements place the included CDATA block directly into the generated program at the designated position. If the element does not contain any dependencies, like in our first example, then the code is placed alone:

```

<filter>
  <![CDATA[
    printf("Hello world from a filter segment!\n");
  ]]>
</filter>

```

This filter block merely prints a string into the output when the generated program is run. If the `<filter>` element contains dependencies, then the variables defined in those vectors (or computed vectors, or noise vectors) will be available, and the CDATA block will be placed inside loops that run over all the transverse dimensions used by the included vectors. The second filter block in this example depends on both the “wavefunction” and “normalisation” vectors:

```

<filter>
  <dependencies>normalisation wavefunction</dependencies>
  <![CDATA[
    phi *= sqrt(Nparticles/Ncalc);
  ]]>
</filter>

```

Since this filter depends on a vector with the transverse dimension “y”, this filter will execute for each point in “y”. This code multiplies the value of the field “phi” by the factor required to produce a normalised function in the sense that $\int dy |\phi(y, t)|^2 = N_{\text{particles}}$.

The next usage of a `<filter>` element in this program is inside the `<integrate>` element, where all filters are placed inside a `<filters>` element.

```
<filters where="step end">
  <filter>
    <dependencies>wavefunction normalisation</dependencies>
    <![CDATA[
      // Correct normalisation of the wavefunction
      phi *= sqrt(Nparticles/Ncalc);
    ]]>
  </filter>
</filters>
```

Filters placed in an integration block are applied each integration step. The “where” flag is used to determine whether the filter should be applied directly before or directly after each integration step. The default value for the where flag is `where="step start"`, but in this case we chose “step end” to make sure that the final output was normalised after the last integration step.

At the end of the sequence element we introduce the `<breakpoint>` element. This serves two purposes. The first is a simple matter of convenience. Often when we manage our input and output from a simulation, we are interested solely in storing the exact state of our integration vectors. A breakpoint element does exactly that, storing the components of any vectors contained within, taking all the normal options of the `<output>` element but not requiring any `<sampling_group>` elements as that information is assumed.

```
<breakpoint filename="groundstate_break.xsil">
  <dependencies basis="ky">wavefunction</dependencies>
</breakpoint>
```

If the filename argument is omitted, the output filenames are numbered sequentially. Any given `<breakpoint>` element must only depend on vectors with identical dimensions.

This program begins with a very crude guess to the ground state, but it rapidly converges to the lowest eigenstate.

4.7 Finding the Ground State of a BEC again

Here we repeat the same simulation as in the *Finding the Ground State of a BEC (continuous renormalisation)* example, using a different transform basis. While spectral methods are very effective, and Fourier transforms are typically very efficient due to the Fast Fourier transform algorithm, it is often desirable to describe nonlocal evolution in bases other than the Fourier basis. The previous calculation was the Schrödinger equation with a harmonic potential and a nonlinear term. The eigenstates of such a system are known analytically to be Gaussians multiplied by the Hermite polynomials.

$$\left[-\frac{\hbar}{2m} \frac{\partial^2}{\partial x^2} + \frac{1}{2} \omega^2 x^2 \right] \phi_n(x) = E_n \phi_n(x)$$

where

$$\phi_n(x, t) = \sqrt{\frac{1}{2^n n!}} \left(\frac{m\omega}{\hbar\pi} \right)^{\frac{1}{4}} e^{-\frac{m\omega x^2}{2\hbar}} H_n \left(\sqrt{\frac{m\omega}{\hbar}} x \right), \quad E_n = \left(n + \frac{1}{2} \right) \omega$$

where $H_n(u)$ are the physicist’s version of the Hermite polynomials. Rather than describing the derivatives as diagonal terms in Fourier space, we therefore have the option of describing the entire $-\frac{\hbar}{2m} \frac{\partial^2}{\partial x^2} + \frac{1}{2} \omega^2 x^2$ term as a diagonal term in the Hermite-Gaussian basis. Here is an XMD S2 simulation that performs the integration in this basis. The following is a simplified version of the `examples/hermitegauss_groundstate.xmds` script.

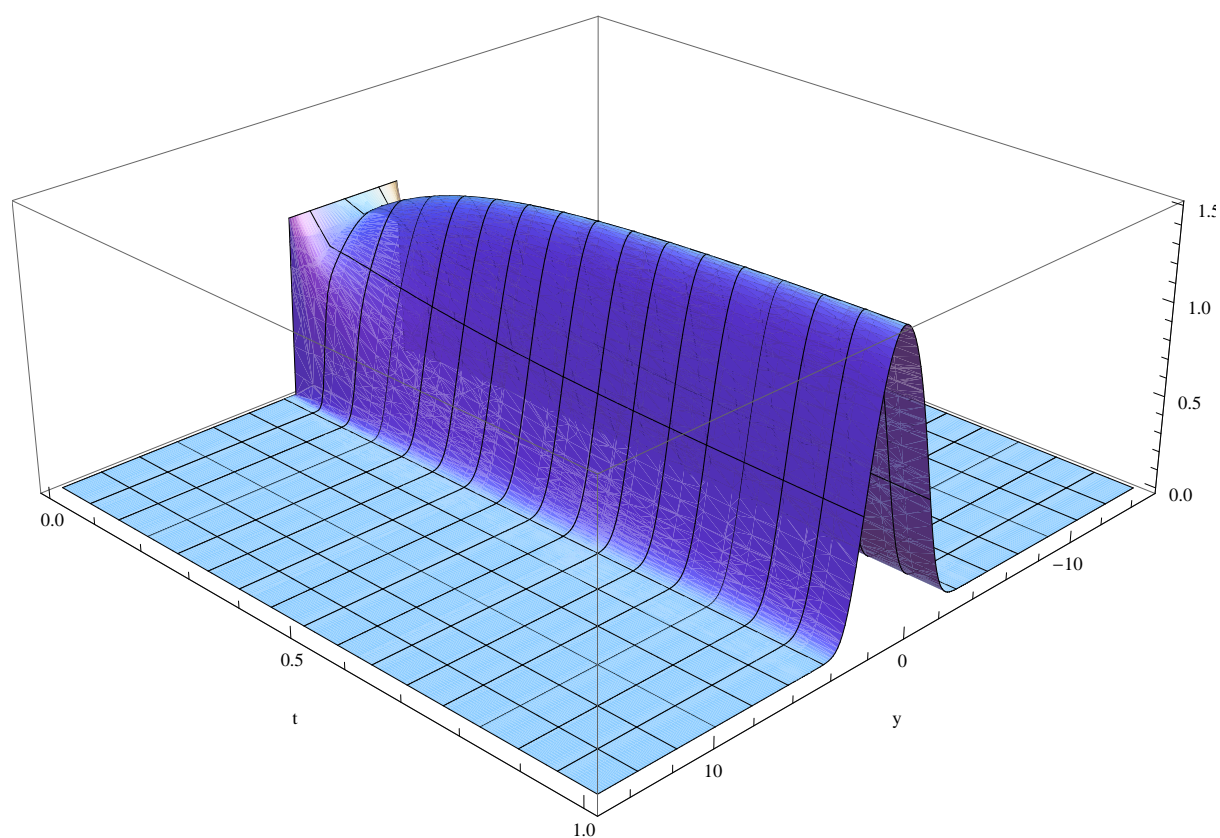


Figure 4.5: The shape of the ground state rapidly approaches the lowest eigenstate. For weak nonlinearities, it is nearly Gaussian.

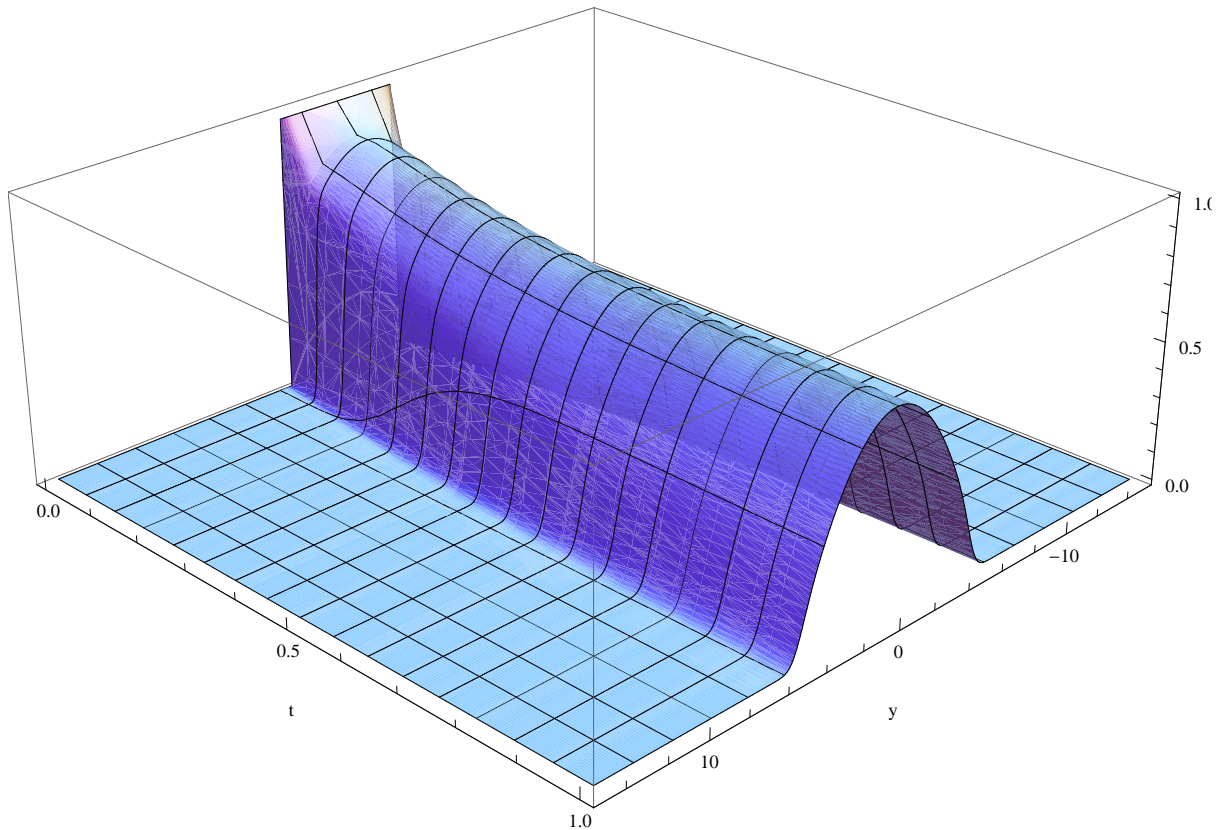


Figure 4.6: When the nonlinear term is larger ($U = 20$), the ground state is wider and more parabolic.

```
<simulation xmds-version="2">
  <name>hermitegauss_groundstate</name>
  <author>Graham Dennis</author>
  <description>
    Solve for the groundstate of the Gross-Pitaevskii equation using the Hermite-Gauss basis.
  </description>

  <features>
    <benchmark />
    <bing />
    <validation kind="run-time" />
    <globals>
      <![CDATA[
        const real omegaz = 2*M_PI*20;
        const real omegarho = 2*M_PI*200;
        const real hbar = 1.05457148e-34;
        const real M = 1.4095392000000000e-25;
        const real g = 9.8;
        const real scatteringLength = 5.57e-9;
        const real transverseLength = 1e-5;
        const real Uint = 4.0*M_PI*hbar*hbar*scatteringLength/M/transverseLength/transverseLength;
        const real Nparticles = 5.0e5;

        /* offset constants */
        const real EnergyOffset = 0.3*pow(pow(3.0*Nparticles/4*omegarho*Uint,2.0)*M/2.0,1/3.0);
      ]]>
    </globals>
  </features>
```

```
<geometry>
  <propagation_dimension> t </propagation_dimension>
  <transverse_dimensions>
    <dimension name="x" lattice="100" length_scale="sqrt(hbar/(M*omegarho))" transform="hermite" />
  </transverse_dimensions>
</geometry>

<vector name="wavefunction" initial_basis="x" type="complex">
  <components> phi </components>
  <initialisation>
    <![CDATA[
      phi = sqrt(Nparticles) * pow(M*omegarho/(hbar*M_PI), 0.25) * exp(-0.5*(M*omegarho/hbar)*x*x)
    ]]>
  </initialisation>
</vector>

<computed_vector name="normalisation" dimensions="" type="real">
  <components> Ncalc </components>
  <evaluation>
    <dependencies basis="x">wavefunction</dependencies>
    <![CDATA[
      // Calculate the current normalisation of the wave function.
      Ncalc = mod2(phi);
    ]]>
  </evaluation>
</computed_vector>

<sequence>
  <integrate algorithm="ARK45" interval="1.0e-2" steps="4000" tolerance="1e-10">
    <samples>100 100</samples>
    <filters>
      <filter>
        <dependencies>wavefunction normalisation</dependencies>
        <![CDATA[
          // Correct normalisation of the wavefunction
          phi *= sqrt(Nparticles/Ncalc);
        ]]>
      </filter>
    </filters>
    <operators>
      <operator kind="ip" type="real">
        <operator_names>L</operator_names>
        <![CDATA[
          L = EnergyOffset/hbar - (nx + 0.5)*omegarho;
        ]]>
      </operator>
      <integration_vectors>wavefunction</integration_vectors>
      <![CDATA[
        dphi_dt = L[phi] - Uint/hbar*mod2(phi)*phi;
      ]]>
    </operators>
  </integrate>

  <filter>
    <dependencies>normalisation wavefunction</dependencies>
    <![CDATA[
      phi *= sqrt(Nparticles/Ncalc);
    ]]>
  </filter>

  <breakpoint filename="hermitegauss_groundstate_break.xsil" format="ascii">
    <dependencies basis="nx">wavefunction</dependencies>
  </breakpoint>
```

```

</sequence>

<output>
  <sampling_group basis="x" initial_sample="yes">
    <moments>dens</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
      dens = mod2(phi);
    ]]>
  </sampling_group>
  <sampling_group basis="kx" initial_sample="yes">
    <moments>dens</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
      dens = mod2(phi);
    ]]>
  </sampling_group>
</output>
</simulation>

```

The major difference in this simulation code, aside from the switch back from dimensionless units, is the new transverse dimension type in the `<geometry>` element.

```

<dimension name="x" lattice="100" length_scale="sqrt(hbar/(M*omegarho))" transform="hermite-gauss

```

We have explicitly defined the “transform” option, which by default expects the Fourier transform. The `transform="hermite-gauss"` option requires the ‘mpmath’ package installed, just as Fourier transforms require the FFTW package to be installed. The “lattice” option details the number of hermite-Gaussian eigenstates to include, and automatically starts from the zeroth order polynomial and increases. The number of hermite-Gaussian modes fully determines the irregular spatial grid up to an overall scale given by the `length_scale` parameter.

The `length_scale="sqrt(hbar/(M*omegarho))"` option requires a real number, but since this script defines it in terms of variables, XMD S2 is unable to verify that the resulting function is real-valued at the time of generating the code. XMD S2 will therefore fail to compile this program without the feature:

```

<validation kind="run-time" />

```

which disables many of these checks at the time of writing the C-code.

4.8 Multi-component Schrödinger equation

This example demonstrates a simple method for doing matrix calculations in XMD S2. We are solving the multi-component PDE

$$\frac{\partial \phi_j(x, y)}{\partial t} = \frac{i}{2} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) \phi_j(x, y) - iU(x, y) \sum_k V_{jk} \phi_k(x, y)$$

where the last term is more commonly written as a matrix multiplication. Writing this term out explicitly is feasible for a small number of components, but when the number of components becomes large, or perhaps V_{jk} should be precomputed for efficiency reasons, it is useful to be able to perform this sum over the integer dimensions automatically. This example show how this can be done naturally using a computed vector. The XMD S2 script is as follows:

```

<simulation xmds-version="2">
  <name>2DMSse</name>

  <author>Joe Hope</author>
  <description>

```

```
Schroedinger equation for multiple internal states in two spatial dimensions.
</description>

<features>
  <benchmark />
  <bing />
  <fftw plan="patient" />
  <openmp />
  <auto_vectorise />
</features>

<geometry>
  <propagation_dimension> t </propagation_dimension>
  <transverse_dimensions>
    <dimension name="x" lattice="32" domain="(-6, 6)" />
    <dimension name="y" lattice="32" domain="(-6, 6)" />
    <dimension name="j" type="integer" lattice="2" domain="(0,1)" aliases="k"/>
  </transverse_dimensions>
</geometry>

<vector name="wavefunction" type="complex" dimensions="x y j">
  <components> phi </components>
  <initialisation>
    <![CDATA[
      phi = j*sqrt(2/sqrt(M_PI/2))*exp(-(x*x+y*y)/4)*exp(i*0.1*x);
    ]]>
  </initialisation>
</vector>

<vector name="spatialInteraction" type="real" dimensions="x y">
  <components> U </components>
  <initialisation>
    <![CDATA[
      U=exp(-(x*x+y*y)/4);
    ]]>
  </initialisation>
</vector>

<vector name="internalInteraction" type="real" dimensions="j k">
  <components> V </components>
  <initialisation>
    <![CDATA[
      V=3*(j*(1-k)+(1-j)*k);
    ]]>
  </initialisation>
</vector>

<computed_vector name="coupling" dimensions="x y j" type="complex">
  <components>
    VPhi
  </components>
  <evaluation>
    <dependencies basis="x y j k">internalInteraction wavefunction</dependencies>
    <![CDATA[
      // Calculate the current normalisation of the wave function.
      VPhi = V*phi(j => k);
    ]]>
  </evaluation>
</computed_vector>

<sequence>
  <integrate algorithm="ARK45" interval="2.0" tolerance="1e-7">
    <samples>20 100</samples>
  </integrate>
</sequence>
```

```

<operators>
  <integration_vectors>wavefunction</integration_vectors>
  <operator kind="ip" dimensions="x">
    <operator_names>Lx</operator_names>
    <![CDATA[
      Lx = -i*kx*kx*0.5;
    ]]>
  </operator>
  <operator kind="ip" dimensions="y">
    <operator_names>Ly</operator_names>
    <![CDATA[
      Ly = -i*ky*ky*0.5;
    ]]>
  </operator>
  <![CDATA[
    dphi_dt = Lx[phi] + Ly[phi] -i*U*VPhi;
  ]]>
  <dependencies>spatialInteraction coupling</dependencies>
</operators>
</integrate>
</sequence>

<output>
  <sampling_group basis="x y j" initial_sample="yes">
    <moments>density</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
      density = mod2(phi);
    ]]>
  </sampling_group>
  <sampling_group basis="x(0) y(0) j" initial_sample="yes">
    <moments>normalisation</moments>
    <dependencies>wavefunction</dependencies>
    <![CDATA[
      normalisation = mod2(phi);
    ]]>
  </sampling_group>
</output>
</simulation>

```

The only truly new feature in this script is the “aliases” option on a dimension. The integer-valued dimension in this script indexes the components of the PDE (in this case only two). The V_{jk} term is required to be a square array of dimension of this number of components. If we wrote the k-index of V_{jk} using a separate `<dimension>` element, then we would not be enforcing the requirement that the matrix be square. Instead, we note that we will be using multiple ‘copies’ of the j-dimension by using the “aliases” tag.

```
<dimension name="j" type="integer" lattice="2" domain="(0,1)" aliases="k"/>
```

This means that we can use the index “k”, which will have exactly the same properties as the “j” index. This is used to define the “V” function in the “internalInteraction” vector. Now, just as we use a computed vector to perform an integration over our fields, we use a computed vector to calculate the sum.

```

<computed_vector name="coupling" dimensions="x y j" type="complex">
  <components>
    VPhi
  </components>
  <evaluation>
    <dependencies basis="x y j k">internalInteraction wavefunction</dependencies>
    <![CDATA[
      // Calculate the current normalisation of the wave function.
      VPhi = V*phi(j => k);
    ]]>
  </evaluation>

```

```
</computed_vector>
```

Since the output dimensions of the computed vector do not include a “k” index, this index is integrated. The volume element for this summation is the spacing between neighbouring values of “j”, and since this spacing is one, this integration is just a sum over k, as required.

This example also demonstrates an optimisation for the IP operators by separating the x and y parts of the operator (see *Optimising with the Interaction Picture (IP) operator*). This gives an approximately 30% speed improvement over the more straightforward implementation:

```
<operator kind="ip">
  <operator_names>L</operator_names>
  <![CDATA[
    L = -i*(kx*kx + ky*ky)*0.5;
  ]]>
</operator>
<![CDATA[
  dphi_dt = L[phi] - i*U*VPhi;
]]>
```

By this point, we have introduced most of the important features in XMD52. More details on other transform options and rarely used features can be found in the *Advanced Topics* section.

REFERENCE SECTION

Contents:

:: index:: Configure, Reconfigure, XMDS2 runtime options

5.1 Configuration, installation and runtime options

Running the ‘xmds2’ program with the option ‘–help’, gives several options that can change its behaviour at runtime. These

- ‘-o’ or ‘–output’, which overrides the name of the output file to be generated
- ‘-n’ or ‘–no-compile’, which generates the C code for the simulation, but does not try to compile it
- ‘-v’ or ‘–verbose’, which gives verbose output about compilation flags.
- ‘-g’ or ‘–debug’, which compiles the simulation in debug mode (compilation errors refer to lines in the source, not the .xmds file). This option implies ‘-v’. This option is mostly useful when debugging XMDS code generation.
- ‘–waf-verbose’, which makes waf be very verbose when configuring XMDS or compiling simulations. This option is intended for developer use only to aid in diagnosing problems with waf.

It also has commands to configure XMDS2 and recheck the installation. If your program requires extra paths to compile, you can configure XMDS2 to include those paths by default. Simply use the command

```
$ xmds2 --configure --include-path /path/to/include --lib-path /path/to/lib
```

Alternatively, you can set the CXXFLAGS or LINKFLAGS environment variables before calling xmds2 --reconfigure. For example, to pass the compiler flag -pedantic and the link flag -lm using the bash shell, use:

```
$ export CXXFLAGS="-pedantic"
$ export LINKFLAGS="-lm"
$ xmds2 --reconfigure ''
```

This method can also be used to change the default compilers for standard and parallel processing, using the CXX and MPICXX flags respectively.

Running XMDS2 with the ‘–configure’ option also searches for packages that have been installed since you last installed or configured XMDS2. If you wish to run ‘xmds2 –configure’ with the same extra options as last time, simply use the command:

```
$ xmds2 --reconfigure
```

A detailed log of the checks is saved in the file ‘~/xmds/waf_configure/config.log’. This can be used to identify issues with packages that XMDS2 is not recognised, but you think that you have successfully installed on your system.

5.2 Useful XML Syntax

Standard XML placeholders can be used to simplify some scripts. For example, the following (abbreviated) code ensures that the limits of a domain are symmetric.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE simulation [
<!ENTITY Npts      "64">
<!ENTITY L         "3.0e-5">
]>
  <simulation xmds-version="2">

    . . .

    <geometry>
      <propagation_dimension> t </propagation_dimension>
      <transverse_dimensions>
        <dimension name="x" lattice="&Npts;" domain="(-&L;, &L;)" />
      </transverse_dimensions>
    </geometry>
```

5.3 XMDS2 XML Schema

There are many, many XML tags that can make up an XMDS2 script. Most of them are optional, or have default values if not specified. It is, however, useful to know which elements are possible, and their position and relationship to other elements in the script. Shown below is the full XML tree for XMDS2, which shows all possible elements and their position in the tree. An ellipsis (...) outside an element indicates the element above can be repeated indefinitely, and an ellipsis inside an element (<element> ... </element>) indicates that the structure of that element has already been shown previously.

The syntax <element /> can be used for lowest-level elements that have attributes but no content, and are shorthand for <element> </element>. This shorthand notation can also be used for elements which can only contain the content “yes” or “no”; in this case the presence of <element /> is equivalent to <element> yes </element>, and the absence of such an element is equivalent to <element> no </element>

The possible attributes and attribute values for each element are not shown; see the individual entries in the Reference section for details.

```
<?xml version="1.0" encoding="UTF-8"?>
<simulation xmds-version="2">
  <name> </name>
  <author> <author>
  <description> </description>

  <features>
    <arguments>
      <argument />
      <argument />
      ...
    </arguments>
    <auto_vectorise />
    <benchmark />
    <bing />
    <cflags> </cflags>
    <chunked_output />
    <diagnostics />
    <error_check />
    <halt_non_finite />
    <fftw />
    <globals> </globals>
```

```

    <openmp />
    <precision> </precision>
    <validation />
</features>

<driver />

<geometry>
  <propagation_dimension> </propagation_dimension>
  <transverse_dimensions>
    <dimension />
    <dimension />
    ...
  </transverse_dimensions>
</geometry>

<vector>
  <components> </components>
  <initialisation>
    <dependencies> </dependencies>
    <filename>
      <![CDATA [
      ]]>
    </initialisation>
</vector>

<vector> ... </vector>
<vector> ... </vector>
...

<computed_vector>
  <components> </components>
  <evaluation>
    <dependencies> </dependencies>
    <![CDATA [
    ]]>
  </evaluation>
</computed_vector>

<computed_vector> ... </computed_vector>
<computed_vector> ... </computed_vector>
...

<noise_vector>
  <components> </components>
</noise_vector>

<noise_vector> ... </noise_vector>
<noise_vector> ... </noise_vector>
...

<sequence>

  <filter>
    <dependencies> </dependencies>
    <![CDATA [
    ]]>
  </filter>

  <integrate>
    <samples> </samples>

    <computed_vector> ... </computed_vector>

```

```

<filters>
  <filter> ... </filter>
  <filter> ... </filter>
  ...
</filters>

<operators>

  <operator>
    <boundary_condition>
      <dependencies> </dependencies>
      <![CDATA [
]]>
    </boundary_condition>
    <operator_names> </operator_names>
    <dependencies> </dependencies>
    <![CDATA [
]]>
  </operator>

  <operator> ... </operator>
  <operator> ... </operator>
  ...

  <integration_vectors> </integration_vectors>
  <dependencies> </dependencies>
  <![CDATA [
]]>

</operators>

</integrate>

<breakpoint>
  <dependencies> </dependencies>
</breakpoint>

</sequence>

<output>
  <sampling_group>
    <dependencies> </dependencies>
    <moments> </moments>
    <operator> ... </operator>
    <![CDATA [
]]>
  </sampling_group>

  <sampling_group> ... </sampling_group>
  <sampling_group> ... </sampling_group>
  ...

</output>

</simulation>

```

5.4 XMD S2 script elements

This section outlines all the elements and options available in an XMD S2 script. This is very much a **work in progress**, beginning with placeholders in most cases, as we have prioritised the tutorials for new users. One

of the most productive ways that non-developer veterans can contribute to the project is to help develop this documentation.

5.4.1 Simulation element

The `<simulation>` element is the single top level element in an XMDs2 simulation, and contains all the other elements. All XMDs scripts must contain exactly one simulation element, and it must have the `xmds-version="2"` attribute defined.

Example syntax:

```
<simulation xmds-version="2">
  <!-- Rest of simulation goes here -->
</simulation>
```

5.4.2 Name element

The name of your simulation. This element is optional, but recommended. If it is set, it will be the name of the executable file generated from this script. It will also be the name of the output file (with an appropriate extension) if the `filename` attribute is not given a value in the `<output>` element.

Example syntax:

```
<name> funky_solver </name>
```

5.4.3 Author element

The author(s) of this script. This element is optional, but can be useful if you need to find the person who has written an incomprehensible script and thinks comments are for the weak.

Example syntax:

```
<author> Ima Mollusc </author>
```

5.4.4 Description element

A description of what the simulation does. Optional, but recommended, in case you (or someone else) has to revisit the script at some distant point in the future.

Example syntax:

```
<description>
  Calculate the 3D ground state of a Rubidium BEC in a harmonic magnetic trap assuming
  cylindrical symmetry about the z axis and reflection symmetry about z=0.
  This permits us to use the cylindrical Bessel functions to expand the solution transverse
  to z and a cosine series to expand the solution along z.
</description>
```

5.4.5 Features Elements

Features elements are where simulation-wide options are specified. The `<features>` element wraps one or more elements describing features. There are many possible feature elements. Currently, a full list of the features supported is:

- *arguments*
- *auto_vectorise*

- *benchmark*
- *bing*
- *cflags*
- *chunked_output*
- *diagnostics*
- *error_check*
- *halt_non_finite*
- *fftw*
- *globals*
- *OpenMP*
- *precision*
- *validation*

Example syntax:

```
<simulation xmds-version="2">
  <features>
    <bing />
    <precision> double </precision>
    ...
  </features>
</simulation>
```

Arguments Element

The `<arguments>` element is optional, and allows defining variables that can be passed to the simulation at run time. These variables are then globally accessible throughout the simulation script. Each of the variables must be defined in an `<argument>` element (see below). The variables can then be passed to the simulation executable as options on the command line. For example, one could define the variables `size`, `number`, and `pulse_shape`

```
<name> arguments_test </name>
<features>
  <arguments>
    <argument name="size" type="real" default_value="20.0"/>
    <argument name="number" type="integer" default_value="7"/>
    <argument name="pulse_shape" type="string" default_value="gaussian"/>
  </arguments>
</features>
```

When XMD S2 is run on this script the executable `arguments_test` is created. The values of `size`, `number`, and `pulse_shape` can then be set to whatever is desired at runtime via

```
./arguments_test --size=1.3 --number=2 --pulse_shape=lorentzian
```

It is also possible to include an optional CDATA block inside the `<arguments>` block. This code will run after the arguments have been initialised with the values passed from the command line. This code block could be used, for example, to sanity check the parameters passed in, or for assigning values to global variables based on those parameters. Any references to variables defined in an `<argument>` element should be made here rather than in the *Globals* element, or else the variables will only have their default values. For example, one could have the following

```
<features>
  <globals>
    <![CDATA[
      real atom_kick;
```

```

]]>
<globals>
<arguments>
  <argument name="bragg_order" type="integer" default_value="2"/>
  <![CDATA[
    atom_kick = bragg_order * 2*M_PI / 780e-9;
  ]]>
</arguments>
</features>

```

The arguments and their values can be added to the filename of the output files by using the `append_args_to_output_filename` attribute. For example:

```

<name> arguments_test </name>
<features>
  <arguments append_args_to_output_filename="yes">
    <argument name="size" type="real" default_value="20.0"/>
    <argument name="number" type="integer" default_value="7"/>
    <argument name="pulse_shape" type="string" default_value="gaussian"/>
  </arguments>
</features>

```

When the `arguments_test` executable is run, it will create output files named `arguments_test.number_7.pulse_shape_gaussian.size_20.0.xsil` and `arguments_test.number_7.pulse_shape_gaussian.size_20.0.h5`.

Argument element

Each `<argument>` element describes one variable that can be passed to the simulation at runtime via the command line. There are three mandatory attributes: `name`, `type`, and `default_value`. `name` is the name by which you can refer to that variable later in the script, as well as the name of the command line parameter. `type` defines the data type of the variable, and `default_value` is the value to which the variable is set if it is not given a value on the command line.

Auto_vectorise element

The `<auto_vectorise />` feature attempts to activate automatic vectorisation for large loops, if it is available in the compiler. This should make some simulations go faster.

Benchmark

The `<benchmark />` feature includes a timing routine in the generated code, so that it is possible to see how long the simulations take to run.

Bing

The `<bing />` feature causes the simulation to make an invigorating sound when the simulation finishes executing.

C Flags

The `<cflags>` feature allows extra flags to be passed to the compiler. This can be useful for optimisation, and also using specific external libraries. The extra options to be passed are defined with a `'CDATA'` block. The compile options can be made visible by running XMD52 either with the `"-v"` (verbose) option, or the `"-g"` (debug) option.

Example syntax:

```
<cflags>
  <![CDATA[
    -O4
  ]]>
</cflags>
```

Chunked Output

By default, XMDs2 keeps the contents of all output moment groups in memory until the end of the simulation when they are written to the output file. This can be a problem if your simulation creates a very large amount of output. `<chunked_output />` causes the simulation to save the output data in chunks as the simulation progresses. For some simulations this can significantly reduce the amount of memory required. The amount of data in a chunk can be specified with the `size` attribute where the suffixes “KB” (kilobytes), “MB” (megabytes), “GB” (gigabytes) and “TB” (terabytes) are understood. Note that `size` specifies the chunk size per output sampling group, per MPI process. So a chunk size of 4MB for a distributed-MPI simulation using 20 processes will cause each process to save up 4MB of data, and data to be written to the output file 80MB at a time.

Limitations (XMDs will give you an error if you violate any of these):

- This feature cannot be used with the ASCII output file format due to limitations in the file format.
- This feature cannot be used with the `multi-path` drivers because all sampling data is required to compute the mean and standard error statistics.
- Neither is this feature compatible with the `error_check` feature as that relies on all sampling data being available to compute the error.

Example syntax:

```
<simulation xmds-version="2">
  <features>
    <chunked_output size="5MB" />
  </features>
</simulation>
```

Diagnostics

The `<diagnostics />` feature causes a simulation to output more information as it executes. This should be useful when a simulation is dying / giving bad results to help diagnose the cause. Currently, it largely outputs step error information.

Error Check

It’s often important to know whether you’ve got errors. This feature runs each integration twice: once with the specified error tolerance or defined lattice spacing in the propagation dimension, and then again with half the lattice spacing, or an equivalently lower error tolerance. Each component of the output then shows the difference between these two integrations as an estimate of the error. This feature is particularly useful when integrating stochastic equations, as it treats the noise generation correctly between the two runs, and thus makes a reasonable estimate of the strong convergence of the equations.

Example syntax:

```
<simulation xmds-version="2">
  <features>
    <error_check />
  </features>
</simulation>
```

Halt_Non_Finite

The `<halt_non_finite />` feature is used to stop computations from continuing to run after the vectors stop having numerical values. This can occur when a number is too large to represent numerically, or when an illegal operation occurs. Processing variables with non-numerical values is usually much slower than normal processing, and the results are meaningless. Of course, there is a small cost to introducing a run-time check, so this feature is optional.

fftw element

The `<fftw \>` feature can be used to pass options to the [Fast Fourier Transform library](#) used by XMDs. This library tests algorithms on each architecture to determine the fastest method of solving each problem. Typically this costs very little overhead, as the results of all previous tests are stored in the directory “`~/.xmds/wisdom`”. The level of detail for the search can be specified using the `plan` attribute, which can take values of “`estimate`”, “`measure`”, “`patient`”, or “`exhaustive`”, in order of the depth of the search. The number of threads for threaded FFTs can be specified with the `threads` attribute, which must be a positive integer.

Example syntax:

```
<fftw plan="patient" threads="3" />
```

Globals

The `globals` feature places the contents of a ‘`CDATA`’ block near the top of the generated program. Amongst other things, this is useful for defining variables that are then accessible throughout the entire program.

Example syntax:

```
<globals>
  <![CDATA[
    const real omegaz = 2*M_PI*20;
    long Nparticles = 50000;

    /* offset constants */
    real frequency = omegaz/2/M_PI;
  ]]>
</globals>
```

OpenMP

The `<openmp />` feature instructs compatible compilers to parallelise key loops using the [OpenMP API](#) standard. By default the simulation will use all available CPUs. The number of threads used can be restricted by specifying the number of threads in the script with `<openmp threads="2"/>`, or by setting the `OMP_NUM_THREADS` environment variable at run-time like so:

```
OMP_NUM_THREADS=2 ./simulation_name
```

Precision

This specifies the precision of the XMDs2 `real` and `complex` datatypes, as well as the precision used when computing transforms. Currently two values are accepted: `single` and `double`. If this feature isn’t specified, XMDs2 defaults to using double precision for its variables and internal calculations.

Single precision has approximately 7.2 decimal digits of accuracy, with a minimum value of 1.4×10^{-45} and a maximum of 3.8×10^{34} . Double precision has approximately 16 decimal digits of accuracy, a minimum value of 4.9×10^{-324} and a maximum value of 1.8×10^{308} .

Using single precision can be attractive, as it can be more than twice as fast, depending on whether a simulation is CPU bound, memory bandwidth bound, MPI bound or bottlenecked elsewhere, although in some situations you may see no speed-up at all. Caution should be exercised, however. Keep in mind how many timesteps your simulation requires, and take note of the tolerance you have set per step, to see if the result will lie within your acceptable total error - seven digit precision isn't a lot. Quite apart from the precision, the range of single precision can often be inadequate for many physical problems. In atomic physics, for example, intermediate values below 1.4×10^{-45} are easily obtained, and will be taken as zero. Similarly, values above 3.8×10^{34} will result in NaNs and make the simulation results invalid.

Also note that when using an adaptive step integrator, setting a tolerance close to limits of the precision can lead to very slow performance.

A further limitation is that not all the combinations of random number generators and probability distributions that are supported in double precision are supported in single precision. For example, the `solirte` generator does not support single precision gaussian distributions. `dsfmt`, however, is one of the fastest generators, and does support single precision.

WARNING: Single precision mode has not been tested anywhere near as thoroughly as the default double precision mode, and there is a higher chance you will run into bugs.

Example syntax:

```
<simulation xmds-version="2">
  <features>
    <precision> single </precision>
  </features>
</simulation>
```

Validation

XMD52 makes a large number of checks in the code generation process to verify that the values for all parameters are safe choices. Sometimes we wish to allow these parameters to be specified by variables. This opens up many possibilities, but requires that any safety checks for parameters be performed during the execution of the program itself. The `<validation>` feature activates that option, with allowable attributes being "run-time", "compile-time" and "none".

As an example, one may wish to define the number of grid points and the range of the grid at run-time rather than explicitly define them in the XMD52 script. To accomplish this, one could do the following:

```
<name> validation_test </name>
<features>
  <validation kind="run-time" />
  <arguments>
    <argument name="xmin" type="real" default_value="-1.0"/>
    <argument name="xmax" type="real" default_value="1.0"/>
    <argument name="numGridPoints" type="integer" default_value="128"/>
  </arguments>
</features>

<geometry>
  <propagation_dimension> t </propagation_dimension>
  <transverse_dimensions>
    <dimension name="x" lattice="numGridPoints" domain="(xmin, xmax)" />
  </transverse_dimensions>
</geometry>
```

and then run the resulting executable with:

```
./validation_test --xmin=-2.0 --xmax=2.0 --numGridPoints=64
```

This approach means that when XMD52 is parsing the script it is unable to tell, for example, if the number of sampling points requested is less than or equal to the lattice size. Consequently it will create an executable with

“numGridPoints” as an internal variable, and make the check at run-time, when it knows the value of “numGridPoints” rather than at compile time, when it doesn’t.

..index:: MPI

5.4.6 Driver Element

The driver element controls the overall management of the simulation, including how many paths of a stochastic simulation are to be averaged, and whether or not it is to be run using distributed memory parallelisation. If it is not included, then the simulation is performed once without using MPI parallelisation. If it is included, it must have a `name` attribute.

The `name` attribute can have values of “none” (which is equivalent to the default option of not specifying a driver), “distributed-mpi”, “multi-path”, “mpi-multi-path” or “adaptive-mpi-multi-path”.

Choosing the `name="distributed-mpi"` option allows a single integration over multiple processors. The resulting executable can then be run according to your particular implementation of MPI. The FFTW library only allows MPI processing of multidimensional vectors, as otherwise shared memory parallel processing requires too much inter-process communication to be efficient. Maximally efficient parallelisation occurs where evolution is entirely local in one transverse dimension (see *transverse dimensions* below). In that case, that dimension should be listed first in the `<geometry>` element. As noted in the worked example *Wigner Function*, it is wise to test the speed of the simulation using different numbers of processors.

The `name="multi-path"` option is used for stochastic simulations, which are typically run multiple times and averaged. It requires a `paths` attribute with the number of iterations of the integration to be averaged. The output will report the averages of the desired samples, and the standard error in those averages. The `name="mpi-multi-path"` option integrates separate paths on different processors, which is typically a highly efficient process. The `name="adaptive-mpi-multi-path"` option integrates separate paths on different processors with load balancing.

Example syntax:

```
<simulation xmds-version="2">
  <driver name="distributed-mpi" />
  <!-- or -->
  <driver name="multi-path" paths="10" />
  <!-- or -->
  <driver name="mpi-multi-path" paths="1000" />
  <!-- or -->
  <driver name="adaptive-mpi-multi-path" paths="1000" />
</simulation>
```

5.4.7 Geometry Element

The `<geometry>` element describes the dimensions used in your simulation, and is required. The only required element inside is the `<propagation_dimension>` element, which defines the name of the dimension along which your simulation will integrate. Nothing else about this dimension is specified, as requirements for the lattice along the integration dimension is specified by the `<integrate>` blocks themselves, as described in section *Integrate element*.

If there are other dimensions in your problem, they are called “transverse dimensions”, and are described in the `<transverse_dimensions>` element. Each dimension is then described in its own `<dimension>` element. A transverse dimension must have a unique name defined by a `name` attribute. If it is not specified, the type of dimension will default to “real”, otherwise it can be specified with the `type` attribute. Allowable types (other than “real”) are “long”, “int”, and “integer”, which are actually all synonyms for an integer-valued dimension.

Each transverse dimension must specify how many points or modes it requires, and the range over which it is defined. This is done by the `lattice` and `domain` attributes respectively. The `lattice` attribute is an integer, and is optional for integer dimensions, where it can be defined implicitly by the `domain`. The `domain` attribute is specified as a pair of numbers (e.g. `domain="(-17, 3)"`) defining the minimum and maximum of the grid.

Any dimension can have a number of aliases. These act exactly like copies of that dimension, but must be included explicitly in the definition of subsequent vectors (i.e. they are not included in the default list of dimensions for a new vector). The list of aliases for a dimension are included in an `aliases` attribute. They are useful for non-local reference of variables. See `groundstate_gaussian.xmds` and `2DMultistateSE.xmds` as examples.

Integrals over a dimension can be multiplied by a common prefactor, which is specified using the `volume_prefactor` attribute. For example, this allows the automatic inclusion of a factor of two due to a reflection symmetry by adding the attribute `volume_prefactor="2"`. In very specific cases, you may wish to refer to volume elements explicitly. This will lead to grid-dependent behaviour, which is sometimes required in certain stochastic field simulations, for example. In this case, the volume element for each variable is described by a `d` prefix (e.g. `lambda` would be referred to as `dlambda`). These volume elements contain any implicit prefactors (for example, the radial coordinate for dimensions defined using *Bessel transforms*), including the `volume_prefactor` element.

If you are using the `distributed-mpi` driver to parallelise the simulation, place the dimension you wish to split over multiple processors first. The most efficient parallelisation would involve distributing a dimension with only local evolution, as the different memory blocks would not need to communicate. Nonlocal evolution that is local in Fourier space is the second preference, as the Fourier transform can also be successfully parallelised with minimum communication.

Each transverse dimension can be associated with a transform. This allows the simulation to manipulate vectors defined on that dimension in the transform space. The default is Fourier space (with the associated transform being the discrete Fourier transform, or “dft”), but others can be specified with the `transform` attribute. The other options are “none”, “dst”, “dct”, “bessel”, “spherical-bessel”, “bessel-neumann” and “hermite-gauss”. Using the right transform can dramatically improve the speed of a calculation.

An advanced feature discussed further in *Dimension aliases* are dimension aliases, which are specified by the `aliases` attribute. This feature is useful for example, when calculating correlation functions.

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <!-- A real-valued dimension from -1.5 to 1.5 -->
      <dimension name="x" lattice="128" domain="(-1.5, 1.5)" />

      <!-- An integer-valued dimension with the 6 values -2, -1, 0, 1, 2, 3 -->
      <dimension name="j" domain="(-2,3)" type="integer" />

      <!-- A real-valued dimension using the bessel transform for a radial coordinate -->
      <dimension name="r" lattice="64" domain="(0, 5)" transform="bessel" volume_prefactor="2" />
    </transverse_dimensions>
  </geometry>
</simulation>
```

The “dft” transform

The “dft” transform is performed using the the normal discrete Fourier transform, which means that it enforces periodic boundary conditions on vectors defined on that dimension. Another implication is that it can only be used with complex-valued vectors. The discrete Fourier transform is almost exactly the same as a standard Fourier transform. The standard Fourier transform is

$$\mathcal{F}[f(x)](k) = \frac{1}{\sqrt{2\pi}} \int_{x_{\min}}^{x_{\max}} f(x) e^{-ikx} dx$$

The discrete Fourier transform has no information about the domain of the lattice, so the XMDS2 transform is equivalent to

$$\begin{aligned}\tilde{\mathcal{F}}[f(x)](k) &= \frac{1}{\sqrt{2\pi}} \int_{x_{\min}}^{x_{\max}} f(x) e^{-ik(x+x_{\min})} dx \\ &= e^{-ix_{\min}k} \mathcal{F}[f(x)](k)\end{aligned}$$

The standard usage in an XMDS simulation involves moving to Fourier space, applying a transformation, and then moving back. For this purpose, the two transformations are entirely equivalent as the extra phase factor cancels. However, when fields are explicitly defined in Fourier space, care must be taken to include this phase factor explicitly. See section [Convolutions and Fourier transforms](#) in the Advanced Topics section.

When a dimension uses the “dft” transform, then the Fourier space variable is defined as the name of the dimension prefixed with a “k”. For example, the dimensions “x”, “y”, “z” and “tau” will be referenced in Fourier space as “kx”, “ky”, “kz” and “ktau”.

Fourier transforms allow easy calculation of derivatives, as the n^{th} derivative of a field is proportional to the n^{th} moment of the field in Fourier space:

$$\mathcal{F}\left[\frac{\partial^n f(x)}{\partial x^n}\right](k_x) = (i k_x)^n \mathcal{F}[f(x)](k_x)$$

This identity can be used to write the differential operator $\mathcal{L} = \frac{\partial}{\partial x}$ as an IP or EX operator as `L = i*kx`; (see [Operators and operator elements](#) for more details).

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <!-- transform="dft" is the default, omitting it wouldn't change anything -->
      <dimension name="x" lattice="128" domain="(-1.5, 1.5)" transform="dft" />
    </transverse_dimensions>
  </geometry>
</simulation>
```

The “dct” transform

The “dct” (discrete cosine transform) is a Fourier-based transform that implies different boundary conditions for associated vectors. XMDS uses the type-II DCT, often called “the DCT”, and its inverse, which is also called the type-III DCT. This transform assumes that any vector using this dimension is both periodic, and also even around a specific point within each period. The grid is therefore only defined across a half period in order to sample each unique point once, and can therefore be of any shape where all the odd derivatives are zero at each boundary. This is a very different boundary condition compared to the DFT, which demands periodic boundary conditions, and is therefore suitable for different simulations. For example, the DCT is a natural choice when implementing zero Neumann boundary conditions.

As the DCT transform can be defined on real data rather than only complex data, it can also be superior to DFT-based spectral methods for simulations of real-valued fields where boundary conditions are artificial.

XMDS labels the cosine transform space variables the same as for [Fourier transforms](#) and all the even derivatives can be calculated the same way. Odd moments of the cosine-space variables are in fact *not* related to the corresponding odd derivatives by an inverse cosine transform.

Discrete cosine transforms allow easy calculation of even-order derivatives, as the $2n^{\text{th}}$ derivative of a field is

proportional to the $2n^{\text{th}}$ moment of the field in DCT-space:

$$\mathcal{F}_{\text{DCT}} \left[\frac{\partial^{2n} f(x)}{\partial x^{2n}} \right] (k_x) = (-k_x^2)^n \mathcal{F}_{\text{DCT}} [f(x)] (k_x)$$

This identity can be used to write the differential operator $\mathcal{L} = \frac{\partial^2}{\partial x^2}$ as an IP or EX operator as $L = -k_x * k_x$; (see *Operators and operator elements* for more details).

For problems where you are defining the simulation domain over only half of the physical domain to take advantage of reflection symmetry, consider using `volume_prefactor="2.0"` so that all volume integrals are over the entire physical domain, not just the simulation domain. i.e. integrals would be over -1 to 1 instead of 0 to 1 if the domain was specified as `domain="(0, 1)"`.

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="128" domain="(-1.5, 1.5)" transform="dct" />
      <!-- Or to cause volume integrals to be multiplied by 2 -->
      <dimension name="y" lattice="128" domain="(0, 1)" transform="dct" volume_prefactor="2" />
    </transverse_dimensions>
  </geometry>
</simulation>
```

The “dst” transform

The “dst” (discrete sine transform) is a counterpart to the DCT transform. XMDs uses the type-II DST and its inverse, which is also called the type-III DST. This transform assumes that fields are periodic in this dimension, but also that they are also odd around a specific point within each period. The grid is therefore only defined across a half period in order to sample each unique point once, and can therefore be of any shape where all the even derivatives are zero at each boundary.

The DST transform can be defined on real-valued vectors. As odd-valued functions are zero at the boundaries, this is a natural transform to use when implementing zero Dirichlet boundary conditions.

XMDs labels the sine transform space variables the same as for *Fourier transforms* and all the even derivatives can be calculated the same way. Odd moments of the sine-space variables are in fact *not* related to the corresponding odd derivatives by an inverse sine transform.

Discrete sine transforms allow easy calculation of even-order derivatives, as the $2n^{\text{th}}$ derivative of a field is proportional to the $2n^{\text{th}}$ moment of the field in DST-space:

$$\mathcal{F}_{\text{DST}} \left[\frac{\partial^{2n} f(x)}{\partial x^{2n}} \right] (k_x) = (-k_x^2)^n \mathcal{F}_{\text{DST}} [f(x)] (k_x)$$

This identity can be used to write the differential operator $\mathcal{L} = \frac{\partial^2}{\partial x^2}$ as an IP or EX operator as $L = -k_x * k_x$; (see *Operators and operator elements* for more details).

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="128" domain="(0, 1.5)" transform="dst" />
    </transverse_dimensions>
  </geometry>
</simulation>
```

The “bessel” and “bessel-neumann” transforms

Just as the Fourier basis is useful for finding derivatives in Euclidean geometry, the basis of Bessel functions is useful for finding certain common operators in cylindrical co-ordinates. In particular, we use the Bessel functions of the first kind, $J_m(u)$. The relevant transform is the Hankel transform:

$$F_m(k) = \mathcal{H}_m[f](k) = \int_0^\infty r f(r) J_m(kr) dr$$

which has the inverse transform:

$$f(r) = \mathcal{H}_m^{-1}[F_m](r) = \int_0^\infty k F_m(k) J_m(kr) dk$$

This transform pair has the useful property that the Laplacian in cylindrical co-ordinates is diagonal in this basis:

$$\nabla^2 (f(r) e^{im\theta}) = \left(\frac{\partial^2 f}{\partial r^2} + \frac{1}{r} \frac{\partial f}{\partial r} - \frac{m^2}{r^2} f \right) e^{im\theta} = \{ \mathcal{H}_m^{-1} [(-k^2) F_m(k)](r) \} e^{im\theta}$$

XMDS labels the variables in the transformed space with a prefix of ‘k’, just as for *Fourier transforms*. The order m of the transform is defined by the `order` attribute in the `<dimension>` element, which must be assigned as a non-negative integer. If the order is not specified, it defaults to zero which corresponds to the solution being independent of the angular coordinate θ .

The difference between the “bessel” and “bessel-neumann” transforms is that the “bessel” transform enforces Dirichlet boundary conditions at the edge of the computational domain ($f(R) = 0$), while “bessel-neumann” enforces Neumann boundary conditions ($\frac{\partial}{\partial r} f(r)|_{r=R} = 0$).

It can often be useful to have a different sampling in normal space and Hankel space. Reducing the number of modes in either space dramatically speeds simulations. To set the number of lattice points in Hankel space to be different to the number of lattice points for the field in its original space, use the attribute `spectral_lattice`. The Bessel space lattice is chosen such that the boundary condition at the edge of the domain is zero. This ensures that all of the Bessel modes are orthogonal. The spatial lattice is also chosen in a non-uniform manner so that Gaussian quadrature methods can be used for spectrally accurate transforms.

Hankel transforms allow easy calculation of the Laplacian of fields with cylindrical symmetry. Applying the operator $\mathcal{L} = -\mathbf{k}r \star \mathbf{k}r$ in Hankel space is therefore equivalent to applying the operator

$$\mathcal{L} = \left(\frac{\partial^2}{\partial r^2} + \frac{1}{r} \frac{\partial}{\partial r} - \frac{m^2}{r^2} \right)$$

in coordinate space.

In non-Euclidean co-ordinates, integrals have non-unit volume elements. For example, in cylindrical co-ordinates with a radial co-ordinate ‘r’, integrals over this dimension have a volume element $r dr$. When performing integrals along a dimension specified by the “bessel” transform, the factor of the radius is included implicitly. If you are using a geometry with some symmetry, it is common to have prefactors in your integration. For example, for a two-dimensional volume in cylindrical symmetry, all integrals would have a volume element of $2\pi r dr$. This extra factor of 2π can be included for all integrals by specifying the attribute `volume_prefactor="2*M_PI"`. See the example `bessel_cosine_groundstate.xmgs` for a demonstration.

Example syntax:

```
<simulation xmgs-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="r" lattice="128" domain="(0, 3)" transform="bessel" volume_prefactor="2*M_PI">
```

```

        </transverse_dimensions>
    </geometry>
</simulation>

```

The “spherical-bessel” transform

When working in spherical coordinates, it is often useful to use the spherical Bessel functions $j_l(x) = \sqrt{\frac{\pi}{2x}} J_{l+\frac{1}{2}}(x)$ as a basis. These are eigenfunctions of the radial component of Laplace’s equation in spherical coordinates:

$$\nabla^2 [j_l(kr) Y_l^m(\theta, \phi)] = \left[\frac{\partial^2}{\partial r^2} + \frac{2}{r} \frac{\partial}{\partial r} - \frac{l(l+1)}{r^2} \right] j_l(kr) Y_l^m(\theta, \phi) = -k^2 j_l(kr) Y_l^m(\theta, \phi)$$

Just as the Bessel basis above, the transformed dimensions are prefixed with a ‘k’, and it is possible (and usually wise) to use the `spectral_lattice` attribute to specify a different lattice size in the transformed space. Also, the spacing of these lattices are again chosen in a non-uniform manner to Gaussian quadrature methods for spectrally accurate transforms. Finally, the `order` attribute can be used to specify the order l of the spherical Bessel functions used.

If we denote the transformation to and from this basis by \mathcal{SH} , then we can write the useful property:

$$\frac{\partial^2 f}{\partial r^2} + \frac{2}{r} \frac{\partial f}{\partial r} - \frac{l(l+1)}{r^2} = \mathcal{SH}_l^{-1} [(-k^2) F_l(k)](r)$$

Spherical Bessel transforms allow easy calculation of the Laplacian of fields with spherical symmetry. Applying the operator $\mathbb{L} = -kr*kr$ in Spherical Bessel space is therefore equivalent to applying the operator

$$\mathcal{L} = \left(\frac{\partial^2}{\partial r^2} + \frac{2}{r} \frac{\partial}{\partial r} - \frac{l(l+1)}{r^2} \right)$$

in coordinate space.

In non-Euclidean co-ordinates, integrals have non-unit volume elements. For example, in spherical co-ordinates with a radial co-ordinate ‘r’, integrals over this dimension have a volume element $r^2 dr$. When performing integrals along a dimension specified by the “spherical-bessel” transform, the factor of the square of the radius is included implicitly. If you are using a geometry with some symmetry, it is common to have prefactors in your integration. For example, for a three-dimensional volume in spherical symmetry, all integrals would have a volume element of $4\pi r^2 dr$. This extra factor of 4π can be included for all integrals by specifying the attribute `volume_prefactor="4*M_PI"`. This is demonstrated in the example `bessel_transform.xmds`.

Example syntax:

```

<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="r" lattice="128" domain="(0, 3)" transform="spherical-bessel" volume_prefactor="4*M_PI">
    </transverse_dimensions>
  </geometry>
</simulation>

```

The “hermite-gauss” transform

The “hermite-gauss” transform allows transformations to and from the basis of Hermite functions $\psi_n(x)$:

$$\psi_n(x) = (2^n n! \sigma \sqrt{\pi})^{-1/2} e^{-x^2/2\sigma^2} H_n(\sigma x)$$

where the functions $H_n(x)$ are the Hermite polynomials:

$$H_n(x) = (-1)^n e^{x^2} \frac{d^n}{dx^n} (e^{-x^2})$$

which are eigenfunctions of the Schrodinger equation for a harmonic oscillator:

$$-\frac{\hbar^2}{2m} \frac{\partial^2 \psi_n}{\partial x^2} + \frac{1}{2} m \omega^2 x^2 \psi_n(x) = \hbar \omega \left(n + \frac{1}{2} \right) \psi_n(x),$$

with $\sigma = \sqrt{\frac{\hbar}{m\omega}}$.

This transform is different to the others in that it requires a `length_scale` attribute rather than a `domain` attribute, as the range of the lattice will depend on the number of basis functions used. The `length_scale` attribute defines the scale of the domain as the standard deviation σ of the lowest order Hermite function $\psi_0(x)$:

$$\psi_0(x) = (\sigma^2 \pi)^{-1/4} e^{-x^2/2\sigma^2}$$

When a dimension uses the “hermite-gauss” transform, then the variable indexing the basis functions is defined as the name of the dimension prefixed with an “n”. For example, when referencing the basis function indices for the dimensions “x”, “y”, “z” and “tau”, use the variable “nx”, “ny”, “nz” and “ntau”.

Applying the operator $\mathcal{L} = nx + 0.5$ in Hermite space is therefore equivalent to applying the operator

$$\mathcal{L} = \left(-\frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + \frac{1}{2\sigma^2} x^2 \right)$$

in coordinate space.

The Hermite-Gauss transform permits one to work in energy-space for the harmonic oscillator. The normal Fourier transform of “hermite-gauss” dimensions can also be referenced using the dimension name prefixed with a “k”. See the examples `hermitegauss_transform.xmds` and `hermitegauss_groundstate.xmds` for examples.

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="128" length_scale="1.0" transform="hermite-gauss" />
    </transverse_dimensions>
  </geometry>
</simulation>
```

5.4.8 Vector Element

Vectors are arrays of data, defined over any subset of the transverse dimensions defined in your *Geometry Element*. These dimensions are listed in the attribute `dimensions`, which can be an empty string if you wish the vector to not be defined on any dimensions. If you do not include a `dimensions` attribute then the vector defaults to being a function of all transverse dimensions, not including any aliases. Vectors are used to store static or dynamic variables, but you do not have to specify their purpose when they are defined. They can then be referenced and/or changed by sequence elements, as described below.

Each `<vector>` element has a unique name, defined by a `name` attribute. It is either complex-valued (the default) or real-valued, which can be specified using the `type="real"` attribute.

A vector contains a list of variables, each defined by name in the `<components>` element. The name of each component is the name used to reference it later in the simulation.

Vectors are initialised at the beginning of a simulation, either from code or from an input file. The basis choice for this initialisation defaults to the normal space as defined in the `<geometry>` element, but any transverse dimension can be initialised in their transform basis by specifying them in an `initial_basis` attribute. The `initial_basis` attribute lists dimensions either by their name as defined by the `<geometry>` element, or by their transformed name. For example, to initialise a two-dimensional vector defined with `dimensions="x y"` in Fourier space for the y-dimension, we would include the attribute `initial_basis="x ky"`, or just `initial_basis="ky"`.

When initialising the vector within the XMDs script, the appropriate code is placed in a 'CDATA' block inside an `<initialisation>` element. This code is in standard C-syntax, and should reference the components of the vector by name. XMDs defines a few useful *shorthand macros* for this C-code. If you wish to initialise all the components of the vector as zeros, then it suffices simply to add the attribute `kind="zero"` or to omit the `<initialisation>` element entirely.

While the default XMDs behaviour is to reference all variables locally, any vector can be referenced non-locally. The notation for referencing the value of a vector 'phi' with a dimension 'j' at a value of 'j=jk' is `phi(j => jk)`. Multiple non-local dimensions are addressed by adding the references in a list, e.g. `phi(j => jk, x => y)`. See `2DMultistateSE.xmds` for an example.

Dimensions can only be accessed non-locally if one of the following conditions is true:

- The dimension is an `integer` dimension,
- The dimension is accessed with an *alias* of that dimension. For example, `phi(x => y)` if the dimension `x` has `y` as an alias, or vice-versa.
- The dimension is a Fourier transform dimension (`dft`), used in the spectral basis (i.e. `kx` for an `x` dimension) and it is accessed with the negative of that dimension. For example `phi(kx => -kx)`.
- The dimension is uniformly spaced (i.e. corresponds to the spatial basis of a dimension with a transform of `dft`, `dct`, `dst` or `none`), the dimension is symmetric about zero and it is accessed with the negative of the dimension name. For example `phi(x => -x)` for a dimension with domain of `(-1.2, 1.2)`.
- The dimension is uniformly spaced (i.e. corresponds to the spatial basis of a dimension with a transform of `dft`, `dct`, `dst` or `none`), and it is accessed with the lower limit of that dimension. For example, `phi(x => -1.2)` for a dimension with a domain of `(-1.2, 1.2)`. Note that the dimension must be accessed with the exact characters used in the definition of the domain. For the previous example `phi(x => -1.20)` does not satisfy this condition.
- **Advanced behaviour:** The value of a variable at an arbitrary point can be accessed via the integer index for that dimension. For example `phi(x_index => 3)` accesses the value of `phi` at the grid point with index 3. As `x_index` is zero-based, this will be the *fourth* grid point. It is highly recommended that the *diagnostics* feature be used when writing simulations using this feature. Once the simulation has been tested, `<diagnostics>` can be turned off for data-taking runs.

Note that a dimension cannot be accessed non-locally in `distributed-mpi` simulations if the simulation is distributed across that dimension.

If you wish to initialise from a file, then you can choose to initialise from an `hdf5` file using `kind="hdf5"` in the `<initialisation>` element, and then supply the name of the input file with the `filename` element. This is a standard data format which can be generated from XMDs, or from another program. An example for generating a file in another program for input into XMDs is detailed in the Advanced topic: *Importing data*.

When initialising from a file, the default is to require the lattice of the transverse dimensions to exactly match the lattice defined by XMDs. There is an option to import data defined on a subset or superset of the lattice points. Obviously, the dimensionality of the imported field still has to be correct. This option is activated by defining the attribute `geometry_matching_mode="loose"`. The default option is defined as `geometry_matching_mode="strict"`. A requirement of the initialisation geometry is that the lattice points of the input file are spaced identically to those of the simulation grid. This allows expanding or contracting a domain between simulations. If used in Fourier space, this feature can be used for coarsening or refining a simulation grid. See '*Loose*' *geometry_matching_mode* for details.

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="128" domain="(-1, 1)" />
    </transverse_dimensions>
  </geometry>

  <!-- A one-dimensional vector with dimension 'x' -->
  <vector name="wavefunction" initial_basis="x" type="complex">
    <components> phi </components>
    <initialisation>
      <![CDATA[
        // 'cis(x)' is cos(x) + i * sin(x)
        phi = exp(-0.5 * x * x) * cis(40 * x);
      ]]>
    </initialisation>
  </vector>

  <!-- A zero-dimensional real vector with components u and v -->
  <vector name="zero_dim" dimensions="" type="real">
    <components>
      u v
    </components>
    <initialisation kind="hdf5">
      <filename>data.h5</filename>
    </initialisation>
  </vector>
</simulation>
```

The dependencies element

Often a vector, computed vector, filter, integration operator or output group will reference the values in one or more other vectors, computed vectors or noise vectors. These dependencies are defined via a `<dependencies>` element, which lists the names of the vectors. The components of those vectors will then be available for use in the 'CDATA' block, and can be referenced by their name.

For a vector, the basis of the dependent vectors, and therefore the basis of the dimensions available in the 'CDATA' block, are defined by the `initial_basis` of the vector. For a `<computed_vector>`, `<filter>` `<integration_vector>`, or moment group vector, the basis of the dependencies can be specified by a `basis` attribute in the `<dependencies>` element. For example, `basis="x ny kz"`.

Any transverse dimensions that appear in the `<dependencies>` element that do not appear in the `dimensions` attribute of the vector are integrated out. For integer dimensions, this is simply an implicit sum over the dimension. For real-valued dimensions, this is an implicit integral over the range of that dimension.

Example syntax:

```
<simulation xmds-version="2">
  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="128" domain="(-1, 1)" />
      <dimension name="y" lattice="10" domain="(-3, 2)" transform="dct" />
    </transverse_dimensions>
  </geometry>

  <!-- A one-dimensional vector with dimension 'x' -->
  <vector name="wavefunction" dimensions="x" initial_basis="x" type="complex">
    <components> phi </components>
    <initialisation>
```

```

<!--
    The initialisation of the vector 'wavefunction' depends on information
    in the 'two_dim' vector. The vector two_dim is DCT-transformed into the
    (x, ky) basis, and the ky dimension is implicitly integrated over in the
    following initialisation code
-->
<dependencies basis="x ky">two_dim</dependencies>
<![CDATA[
    // 'cis(x)' is cos(x) + i * sin(x)
    phi = exp(-0.5 * x * x + v) * cis(u * x);
]]>
</initialisation>
</vector>

<!-- A two-dimensional real vector with components u and v -->
<vector name="two_dim" type="real">
    <components>
        u v
    </components>
    <initialisation kind="hdf5">
        <filename>data.h5</filename>
    </initialisation>
</vector>
</simulation>

```

5.4.9 Computed Vector Element

Computed vectors are arrays of data much like normal `<vector>` elements, but they are always calculated as they are referenced, so they cannot be initialised from file. It is defined with a `<computed_vector>` element, which has a name attribute, optional dimensions and type attributes, and a `<components>` element, just like a `<vector>` element. Instead of an `<initialisation>` element, it has an `<evaluation>` element that serves the same purpose. The `<evaluation>` element contains a `<dependencies>` element (see above `<Dependencies>`), and a 'CDATA' block containing the code that defines it.

As it is not being stored, a `<computed_vector>` does not have or require an `initial_basis` attribute, as it will be transformed into an appropriate basis for the element that references it. The basis for its evaluation will be determined entirely by the `basis` attribute of the `<dependencies>` element.

Example syntax:

```

<simulation xmds-version="2">
    <geometry>
        <propagation_dimension> t </propagation_dimension>
        <transverse_dimensions>
            <dimension name="x" lattice="128" domain="(-1, 1)" />
        </transverse_dimensions>
    </geometry>

    <!-- A one-dimensional vector with dimension 'x' -->
    <vector name="wavefunction" type="complex">
        <components> phi </components>
        <initialisation>
            <![CDATA[
                // 'cis(x)' is cos(x) + i * sin(x)
                phi = exp(-0.5 * x * x) * cis(40 * x);
            ]]>
        </initialisation>
    </vector>

    <!-- A zero-dimensional real computed vector with components Ncalc -->
    <computed_vector name="zero_dim" dimensions="" type="real">
        <components>

```

```

        Ncalc
    </components>
    <evaluation>
        <dependencies>wavefunction</dependencies>
        <![CDATA[
            // Implicitly integrating over the dimension 'x'
            Ncalc = mod2(phi);
        ]]>
    </evaluation>
</computed_vector>
</simulation>

```

5.4.10 Noise Vector Element

Noise vectors are used like computed vectors, but when they are evaluated they generate arrays of random numbers of various kinds. They do not depend on other vectors, and are not initialised by code. They are defined by a `<noise_vector>` element, which has a name attribute, and optional dimensions, `initial_basis` and `type` attributes, which work identically as for normal vectors.

The choice of pseudo-random number generator (RNG) can be specified with the `method` attribute, which has options “posix” (the default), “mkl”, “solirte” and “dsfmt”. It is only possible to use any particular method if that library is available. Although “posix” is the default, it is also the slowest, and produces the lowest quality random numbers (although this is typically not a problem). “mkl” refers to the Intel Math Kernel Library, and is only available if installed. “solirte” and “dsfmt” are fast, hardware-accelerated random number sources that should work on most systems. “mkl”, “solirte” and “dsfmt” have comparable performance.

The random number generators can be provided with a seed using the `seed` attribute, which should typically consist of a list of three integers. All RNGs require positive integers as seeds. It is possible to use the `<validation kind="run-time">` feature to use passed variables as seeds. It is advantageous to use fixed seeds rather than timer-based seeds, as the `<error_check>` element can test for strong convergence if the same seeds are used for both integrations. If the `seed` attribute is not specified, then seeds will be generated at the time the simulation is run. Different executions of the same simulation will therefore give different results. However, results can be reproduced by examining the `.xsil` file produced by the simulation which contains the generated seeds. If these seeds are used for the `seed` attribute, the same results can be reproduced. Unless you need to reproduce particular results, it is unnecessary to specify the `seed` attribute.

The different types of noise vectors are defined by a mandatory `kind` attribute, which must take the value of ‘gauss’, ‘gaussian’, ‘wiener’, ‘poissonian’, ‘jump’ or ‘uniform’.

Example syntax:

```

<simulation xmds-version="2">
    <geometry>
        <propagation_dimension>t </propagation_dimension>
        <transverse_dimensions>
            <dimension name="x" lattice="128" domain="(-1, 1)" />
        </transverse_dimensions>
    </geometry>

    <!--
        A one-dimensional complex wiener noise vector.
        This noise is appropriate for using in the complex
        random-walk equation of motion:
            dz_dt = eta;
    -->
    <noise_vector name="noise" kind="wiener">
        <components>
            eta
        </components>
    </vector>
</simulation>

```

Uniform noise

Uniform noises defined over any transverse dimensions are simply uniformly distributed random numbers between zero and one. This noise is an example of a “static” noise, i.e. one suitable for initial conditions of a field. If it were included in the equations of motion for a field, then the effect of the noise would depend on the lattice spacing of the propagation dimension. XMDs therefore does not allow this noise type to be used in integration elements.

Example syntax:

```
<simulation xmds-version="2">
  <noise_vector name="drivingNoise" dimensions="x" kind="uniform" type="complex" method="dsfmt"
    <components>Eta</components>
  </noise_vector>
</simulation>
```

Gaussian noise

Noise generated with the “gaussian” method is gaussian distributed with zero mean. For a real-valued noise vector, the variance at each point is the inverse of the volume element of the transverse dimensions in the vector. This volume element for a single transverse dimension is that used to perform integrals over that dimension. For example, it would include a factor of r^2 for a dimension “r” defined with a spherical-bessel transform. It can be non-uniform for dimensions based on non-Fourier transforms, and will include the product of the volume_prefactor attribute as specified in the *Geometry* element. The volume element for an integer-type dimension is unity (i.e. where the integral is just an unweighted sum). The volume element for a noise_vector with multiple dimensions is simply the product of the volume elements of the individual dimensions.

This lattice-dependent variance is typical in most applications of partial differential equations with stochastic initial conditions, as the physical quantity is the variance of the field over some finite volume, which does not change if the variance at each lattice site varies as described above.

For complex-valued noise vector, the real and imaginary parts of the noise are independent, and each have half the variance of a real-valued noise. This means that the modulus squared of a complex-valued noise vector has the same variance as a real-valued noise vector at each point.

Gaussian noise vectors are an example of a “static” noise, i.e. one suitable for initial conditions of a field. If they were included in the equations of motion for a field, then the effect of the noise would depend on the lattice spacing of the propagation dimension. XMDs therefore does not allow this noise type to be used in integration elements.

Example syntax:

```
<simulation xmds-version="2">
  <noise_vector name="initialNoise" dimensions="x" kind="gauss" type="real" method="posix" seed=
    <components>fuzz</components>
  </noise_vector>
</simulation>
```

Wiener noise

Noise generated with the “wiener” method is gaussian distributed with zero mean and the same variance as the static “gaussian” noise defined above, multiplied by a factor of the lattice step in the propagation dimension. This means that these noise vectors can be used to define Wiener noises for standard stochastic ordinary or partial differential equations. Most integrators in XMDs effectively interpret these noises as Stratonovich increments.

As a dynamic noise, a Wiener process is not well-defined except in an integrate element.

Example syntax:

```
<simulation xmds-version="2">
  <noise_vector name="diffusion" dimensions="x" kind="wiener" type="real" method="solirte" seed=
    <components>dW</components>
  </noise_vector>
</simulation>
```

```

</noise_vector>
</simulation>

```

Poissonian noise

A noise vector using the “poissonian” method generates a random variable from a Poissonian distribution. While the Poisson distribution is integer-valued, the variable will be cast as a real number. The rate of the Poissonian distribution is defined by the `mean` or `mean-density` attributes. These are synonyms, and must be defined as positive real numbers. For Poissonian noises defined over real-valued transverse dimensions, the rate is given by the product of this `mean-density` attribute and the volume element at that point, taking into account all transverse dimensions, including their `volume_prefactor` attributes. The result is that the integral over each volume in space is a sample from a Poissonian distribution of that rate.

Poissonian noise vectors are an example of a “static” noise, i.e. one suitable for initial conditions of a field. If they were included in the equations of motion for a field, then the effect of the noise would depend on the lattice spacing of the propagation dimension. XMD S therefore does not allow this noise type to be used in integration elements.

Example syntax:

```

<simulation xmds-version="2">
  <noise_vector name="initialDistribution" dimensions="x" kind="poissonian" type="real" mean-density="1.0">
    <components>Pdist</components>
  </noise_vector>
</simulation>

```

Jump noise

A noise vector using the “jump” method is the dynamic version of the poissonian noise method, and must have the `mean-rate` attribute specified as a positive real number. The variable at each point is chosen from a Poissonian distribution with a mean equal to the product of three variables: the `mean-rate` attribute; the volume of the element as defined by its transverse dimensions (including their `volume_prefactor` attributes); and the step size in the propagation dimension. Normally defined in the limit where the noise value is zero almost always, with a few occurrences where it is unity, and none of any higher value, this type of noise is commonly used in differential equations with a Poissonian jump process.

It is common to wish to vary the mean rate of a jump process, which means that the `mean-rate` attribute must be a variable or a piece of code. These cannot be verified to be a positive real number at compile time, so they must be used with the `<validation>` feature with either the `kind="none"` or `kind="run-time"` attributes.

As a dynamic noise, a jump process is not well-defined except in an `integrate` element. The only algorithm that currently integrates jump noises correctly is the Euler algorithm. This can be implemented by the “SI” method with the option `iterations="1"`.

Example syntax:

```

<simulation xmds-version="2">
  <noise_vector name="initialDistribution" dimensions="" kind="jump" type="real" mean-rate="2.7">
    <components>dN</components>
  </noise_vector>
</simulation>

```

5.4.11 Sequence Element

All processing of vectors happens in sequence elements. Each simulation must have exactly one main sequence element, but it can then contain any number of nested sequence elements. A sequence element can contain any number of `<sequence>`, `<filter>`, `<integrate>` and/or `<breakpoint>` elements, which are executed in the order they are written. A sequence can be repeated a number of times by using the `cycles` attribute. For example, `<sequence cycles="10">` will execute the elements in that sequence 10 times.

Example syntax:

```
<simulation xmds-version="2">
  <sequence cycles="2">
    <sequence> ... </sequence>
    <filter> ... </filter>
    <integrate> ...</integrate>
  </sequence>
</simulation>
```

5.4.12 Filter element

A `<filter>` element can be placed inside a `<sequence>` element or an `<integrate>` element. It contains a ‘CDATA’ block and an optional `<dependencies>` element, which may give access to variables in other `<vector>`, `<computed_vector>` or `<noise_vector>` elements. The code inside the ‘CDATA’ block is executed over the combined tensor product space of the dependencies, or simply once if there is no dependencies element. This element therefore allows arbitrary execution of C-code.

Sometimes it is desirable to apply a filter conditionally. The most efficient way of doing this is to call the function from the piece of code that contains the conditional statement (likely another `<filter>` element) rather than embed the conditional function in the filter itself, as the latter method can involve the conditional statement being evaluated multiple times over the transverse dimensions. For this reason, it is possible to give a filter a name attribute, and the filter can thenceforth be called in CDATA blocks by that name. For example: `<filter name="filterName">` allows the function to be called using the C-function `filterName()`.

One of the common uses of a filter element is to apply discontinuous changes to the vectors and variables of the simulation.

Example syntax:

```
<sequence>
  <filter>
    <![CDATA[
      printf("Hello world from the first filter segment! This filter rather wastefully calls t
      fname();
    ]]>
  </filter>

  <filter name="fname">
    <dependencies>normalisation wavefunction</dependencies>
    <![CDATA[
      phi *= sqrt(Nparticles/Ncalc);
    ]]>
  </filter>
</sequence>
```

5.4.13 Integrate element

The `<integrate>` element is at the heart of most XMDs simulations. It is used to integrate a set of (potentially stochastic) first-order differential equations for one or more of the vectors defined using the `<vector>` element along the propagation dimension. At the beginning of the simulation, the value of the propagation dimension is set to zero, and the vectors are initialised as defined in the `<vector>` element. As successive sequence elements change these variables, each integrate element simply integrates onward from the current values.

The length of the integration is defined by the `interval` attribute, which must be a positive real number. An `<integrate>` element must have an `algorithm` attribute defined, which defines the integration method. Current methods include *SI*, *SIC*, *RK4*, *RK9*, *ARK45*, and *ARK89*. Fixed step algorithms require a `steps` attribute, which must be a positive integer that defines the number of (evenly spaced) integration steps. Adaptive stepsize algorithms require a `tolerance` attribute that must be a positive real number much smaller than one, which defines the allowable relative error per integration step. If the `steps` attribute is specified for an adaptive stepsize algorithm, then it is used to generate the initial stepsize estimate.

The optional `<samples>` element is used to track the evolution of one or more vectors or variables during an integration. This element must contain a non-negative integer for each `<sampling_group>` element defined in the simulation's `<output>` element. The list of integers then defines the number of times that the moments defined in those groups will be sampled. For a fixed step algorithm, each non-zero number of samples must be a factor of the total number of steps.

The vectors to be integrated and the form of the differential equations are defined in the `<operators>` element (or elements). Filters to be applied each step can be defined with optional `<filters>` elements.

Computed vectors can be defined with the `<computed_vector>` element. These act exactly like a globally defined *Computed Vector Element*, but are only available within the single `<integrate>` element.

Example syntax:

```
<integrate algorithm="ARK89" interval="1e-4" steps="10000" tolerance="1e-8">
  <samples>20</samples>
  <filters>
    <filter>
      <dependencies>wavefunction normalisation</dependencies>
      <![CDATA[
        phi *= sqrt(Nparticles/Ncalc);    // Correct normalisation of the wavefunction
      ]]>
    </filter>
  </filters>
  <operators>
    <operator kind="ip">
      <operator_names>T</operator_names>
      <![CDATA[
        T = -0.5*hbar/M*ky*ky;
      ]]>
    </operator>
    <dependencies>potential</dependencies>
    <![CDATA[
      dphi_dt = T[phi] - (Vl + Uint/hbar*mod2(phi))*phi;
    ]]>
    <integration_vectors>wavefunction</integration_vectors>
  </operators>
</integrate>
```

Operators and operator elements

An `<integrate>` element must contain one or more `<operators>` elements, which define both which vectors are to be integrated, and their derivative in the propagation dimension. When all vectors to be integrated have the same dimensionality, they can all be defined within a single `<operators>` element, and when vectors with different dimension are to be integrated, each set of vectors with the same dimensionality should be placed in separate `<operators>` elements.

Within each `<operators>` element, the vectors that are to be integrated are listed by name in the `<integration_vectors>` element, and the differential equations are written in a 'CDATA' block. The derivative of each component of the integration vectors must be defined along the propagation dimension. For example, if the integration vectors have components 'phi' and 'beta', and the propagation dimension is labelled 'tau', then the 'CDATA' block must define the variables 'dphi_dtau' and 'dbeta_dtau'. These derivatives can be any function of the available variables, including any components from other vectors, computed vectors or noise vectors that are listed in the optional `<dependencies>` element. These dependent vectors must be defined on a subset of the dimensions of the integration vectors.

When noise vectors are referenced, equations with Wiener noises should be written as though the equations are in differential form, as described in the worked examples *Kubo Oscillator* and *Fibre Noise*. Jump-based Poisson noises will also be written in an equivalent form, as modelled by the example `photodetector.xmds`.

By default, the name of each component references the local value of the vector, but *nonlocal variables* can be accessed using the standard syntax. However, typically the most common (and most efficient) method of

referencing nonlocal variables is to reference variables that are local in the *transformed space* for a given transverse dimension. This is done using `<operator>` elements.

There are three kinds of `<operator>` elements. The first is denoted with a `kind="functions"` attribute, and contains a 'CDATA' block that will be executed in the order that it is defined. This is useful when you wish to calculate functions that do not depend on the transverse dimensions. Defining these along with the main equations of motion causes them to be recalculated separately for each point. The second kind of `<operator>` element is used to define an operation in a transformed space. This is often an efficient method of calculating common nonlocal terms such as derivatives. The third kind is used to define integration of one or more vectors along a transverse dimension.

Example syntax:

```
<operator kind="functions">
  <![CDATA[
    f = cos(t);
  ]]>
</operator>
```

The second kind of operator element defines a list of operators in an `<operator_names>` element. The basis of these operators defaults to the transform space unless a different basis is specified using the `basis` attribute. These operators must then be defined in a 'CDATA' block, using any *dependencies* as normal. The operators defined in these elements can then be used in the 'CDATA' block that defines the equations of motion. The application of operator 'L' to vector 'psi' is denoted `L[psi]`. Operators can be applied to functions of vectors using the same notation, such as `L[psi*psi]`. Aside from the example above, many examples can be found in the examples folder, and the *Worked Examples* section of the documentation.

Operators of this second kind have the `kind="IP"` or `kind="EX"` attribute, standing for 'interaction picture' and 'explicit' operators respectively. Explicit operators can be used in all situations, and simply construct and calculate a new vector of the form in the square brackets. IP operators use less memory and can improve speed by allowing larger timesteps, but have two important restrictions. **Use of IP operators without understanding these restrictions can lead to incorrect code.** The first restriction is that IP operators can only be applied to named components of one of the integration vectors, and not functions of those components. The second restriction is that the equations of motion must be written such that the term with the operator is not multiplied by any quantity or used inside a function. (For those interested, the reason for this is that the IP algorithm applies the operator separately to the rest of the evolution, and therefore the actual text of the `L[psi]` term is replaced by the numeral zero.) If you must break either of those rules, then you need to use the EX algorithm.

If the IP or EX operator is constant across the integration, then the attribute `constant="yes"` may be set to ensure that it is precalculated at the start of integration, otherwise the `constant="no"` attribute ensures that the operator is recalculated at each step. The `constant` attribute is optional and a sensible default is chosen if the attribute is omitted. Note that for EX operators the default is `constant="no"` because the EX operator is typically cheap to calculate and not precomputing it reduces memory bandwidth requirements, usually leading to faster simulations. If your simulation has a computationally expensive EX operator, it may benefit from adding the `constant="yes"` attribute.

Example syntax:

```
<operator kind="ex" constant="yes">
  <operator_names>T</operator_names>
  <![CDATA[
    T = -0.5*hbar/M*ky*ky;
  ]]>
</operator>
```

The third kind of operator element is used to define an integration along a transverse dimension. This kind of evolution is called "cross-propagation", and is described briefly in the examples 'tla.xmds', 'tla_sic.xmds' and 'sine_cross.xmds'. This class of equations have a subset of vectors that have an initial condition on one side of a transverse dimension, and a differential equation defined in that dimension, and as such, this kind of operator element has much of the structure of an entire `<integrate>` element.

An operator element with the `kind="cross_propagation"` attribute must specify the transverse dimension along which the integration would proceed with the `propagation_dimension` attribute. It must also specify

its own `<integration_vectors>` element, its own `<operators>` elements (of the second kind), and may define an optional `<dependencies>` element. The algorithm to be used for the transverse integration is specified by the `algorithm` attribute, with options being `algorithm="SI"` and `algorithm="RK4"`. The derivatives in the cross propagation direction are defined in a 'CDATA' block, just as for a normal `<integrate>` element.

The boundary conditions are specified by a `<boundary_conditions>` element, which requires the `kind="left"` or `kind="right"` attribute to specify on which side of the grid that the boundary conditions are specified. The boundary conditions for the `<integration_vectors>` are then specified in a 'CDATA' block, which may refer to vectors in an optional `<dependencies>` element that can be contained in the `<boundary_conditions>` element.

Example syntax:

```
<operator kind="cross_propagation" algorithm="RK4" propagation_dimension="t">
  <integration_vectors>cross</integration_vectors>
  <dependencies>constants</dependencies>
  <boundary_condition kind="left">
    <![CDATA[
      v = 1.0;
      w = 1.0;
    ]]>
  </boundary_condition>

  <operator kind="ip" constant="yes">
    <operator_names>L</operator_names>
    <![CDATA[
      L = i;
    ]]>
  </operator>

  <![CDATA[
    dv_dt = i*v;
    dw_dt = L[w];
  ]]>
</operator>
```

Algorithms

The stability, efficiency and even convergence of a numerical integration can depend on the method. Due to the varying properties of different sets of equations, it is impossible to define the best method for all equations, so XMD S provides an option to use different algorithms. These include fixed step algorithms, which divide the integration region into equal steps, and adaptive stepsize algorithms, which attempt to estimate the error in the simulation in order to choose an appropriate size for the next step. As a first guess, a good method for a deterministic integration would be *ARK89*, and a good guess for a stochastic method would be the *SI and SIC algorithms*.

For the purposes of the descriptions below, we will assume that we are considering the following set of coupled differential equations for the vector of variables $\mathbf{x}(t)$:

$$\frac{dx_j}{dt} = f_j(\mathbf{x}(t), t)$$

SI and SIC algorithms

The SI algorithm is a semi-implicit fixed-step algorithm that finds the increment of the vector by solving

$$x_j(t + \Delta t) = x_j(t) + f_j\left(\mathbf{x}\left(t + \frac{\Delta t}{2}\right), t + \frac{\Delta t}{2}\right) \Delta t$$

using a simple iteration to find the values of the vector at the midpoint of the step self-consistently. The number of iterations can be set using the `iterations` attribute, and it defaults to `iterations="3"`. The choice of `iterations="1"` is therefore fully equivalent to the Euler algorithm, where

$$x_j(t + \Delta t) = x_j(t) + f_j(\mathbf{x}(t), t) \Delta t.$$

The Euler algorithm is the only safe algorithm for direct integration of *jump-based Poisson processes*. Efficient numerical solution of those types of equations is best done via a process of triggered filters, which will be described in the *Advanced Topics* section. Integrating using the Euler algorithm computes the Ito integral, as opposed to the Stratonovich integral, which all the other algorithms compute.

When SI integration is used in conjunction with SI cross-propagation, a slight variant of the SI algorithm can be employed where the integration in both directions is contained within the iteration process. This is activated by using `algorithm="SIC"` rather than `algorithm="SI"`.

The SI algorithm is correct to second order in the step-size for deterministic equations, and first order in the step-size for Stratonovich stochastic equations with Wiener noises. This makes it the highest order stochastic algorithm in XMDs, although there are many sets of equations that integrate more efficiently with lower order algorithms. When called with the `iterations="1"` option (the Euler algorithm), it is correct to first order in the step-size for deterministic equations, and one-half order in the step-size for Ito stochastic equations with Wiener noises.

Runge-Kutta algorithms

Runge-Kutta algorithms are the workhorse of numerical integration, and XMDs employs two fixed step versions: `algorithm="RK4"`, which is correct to fourth-order in the step size, and `algorithm="RK9"`, which is correct to ninth order in the step size. It must be strongly noted that a higher order of convergence does not automatically mean a superior algorithm. RK9 requires several times the memory of the RK4 algorithm, and each step requires significantly more computation.

All Runge-Kutta algorithms are convergent for Stratonovich stochastic equations at the order of the square root of the step-size. This ‘half-order’ convergence may seem very weak, but for some classes of stochastic equation this improves up to one half of the deterministic order of convergence. Also, the convergence of some stochastic equations is limited by the ‘deterministic part’, which can be improved dramatically by using a higher order Runge-Kutta method.

Adaptive Runge-Kutta algorithms

Fixed step integrators can encounter two issues. First, as the equations or parameters of a simulation are changed, the minimum number of steps required to integrate it may change. This means that the convergence must be re-tested multiple times for each set of parameters, as overestimating the number of steps required to perform an integration to a specified error tolerance can be very inefficient. Second, even if the minimum acceptable number of steps required is known for a given simulation, it may be that there are regions of integration that are of wildly varying difficulty. For a fixed step integrator, this means that the step-size must be small enough to handle the most difficult region, and is therefore inefficiently small for the easier regions. Adaptive step-size algorithms get around this problem by testing the convergence during the integration, and adjusting the step-size until it reaches some target tolerance.

XMDs employs two adaptive step-size algorithms based on ‘embedded Runge-Kutta’ methods. These are Runge-Kutta methods that can output multiple variables that have different convergence. The difference between the higher-order and the lower-order solutions gives an estimate of the error in each step, which can then be used to estimate an appropriate size for the next step. We use `algorithm="ARK45"`, which contains fourth and fifth order solutions, and `algorithm="ARK89"`, which contains eighth and ninth order solutions. Each algorithm converges with the order of the lowest order solution (fourth and eighth order respectively). The overheads involved in estimating the error and step-size make the adaptive algorithms slower than fixed step integration using the same step-size, but overall there is typically a significant performance gain from being able to avoid doing this optimisation manually.

All adaptive stepsize algorithms require a `tolerance` attribute, which must be a positive real number that defines the allowable error per step. It is also possible to specify a `max_iterations` attribute, which is a positive

integer that stops the integrator from trying too many times to find an acceptable stepsize. The integrator will abort with an error if the number of attempts for a single step exceeds the maximum specified with this attribute.

As all Runge-Kutta solutions have equal order of convergence for stochastic equations, *if the step-size is limited by the stochastic term then the step-size estimation is entirely unreliable*. Adaptive Runge-Kutta algorithms are therefore not appropriate for stochastic equations.

Richardson Extrapolation Algorithms and the Bulirsch-Stoer Method

The Richardson Extrapolation technique begins with a large initial interval and uses another stepper algorithm to compute the solution for this interval. It does this by subdividing the interval into increasing subintervals (i.e. with smaller and smaller stepsizes) and uses rational extrapolation to produce a higher order result than would be obtained using the other stepper on its own. The number of extrapolations performed is controllable via the `extrapolations` attribute which is a positive integer (defaults to 4).

Richardson Extrapolation provides the best trade off between computational effort and accuracy when paired with the Modified Midpoint stepper. This stepper is notable as its error scaling function contains only even powers of two. This means that each extrapolation performed in the Richardson technique gains two orders rather than one order as is expected for most other steppers. This combined with the low computational overhead of the Modified Midpoint makes it a powerful tool. The combination of Richardson Extrapolation and the Modified Midpoint stepper is known as the Bulirsch-Stoer method.

A number of combinations of fixed-step, fixed-order Richardson Extrapolation are available in XMDs2. The most notable is the Bulirsch-Stoer method which can be selected using `algorithm="BS"` or alternatively `algorithm="REMM"`. Other combinations include 'RERK4', 'RERK9' and 'RESI' (for stochastic equations). Please note that these additional combinations have not been tested as strictly as the 'REMM' combination and so care should be taken to ensure the results are sane. Additionally the Modified Midpoint stepper is available as a standalone stepper under the mnemonic 'MM', although this is probably not useful outside of testing.

Richardson Extrapolation in general uses more memory than other integrators as multiple result vectors must be stored at the same time, which is something users should be aware of if the `extrapolations` attribute is set too high (generally < 10 should be sufficient).

See the section on the *Bulirsch-Stoer Algorithm* for more details.

Filters element

Filter elements are used inside *sequence elements* to execute arbitrary code, or make discontinuous changes in the vectors. Sometimes it is desirable to perform a filter element at the beginning or end of each step in an integration. This can be done by placing `<filter>` elements in a `<filters>` element within the `<integrate>` element. The `<filters>` element specifies whether the filters are to be executed at the end of each step or the beginning of each step with the `where="step end"` and `where="step start"` attributes respectively. Each filter is then executed in the order found in the `<filters>` element.

Example syntax:

```
<integrate algorithm="ARK45" interval="100000.0" steps="10000000" tolerance="1e-8">
  <samples>5000 100</samples>
  <filters where="step end">
    <filter>
      <dependencies>vector1 vector2</dependencies>
      <![CDATA[
        x = 1;
        y *= ynorm;
      ]]>
    </filter>
  </filters>

  <operators>
    <integration_vectors>vector1</integration_vectors>
    <![CDATA[
```

```
    dx_dt = alpha;  
    dy_dt = beta*y;  
  ]]  
</operators>  
</integrate>
```

5.4.14 Breakpoint element

The `<breakpoint>` element is used to output the full state of one or more vectors. Unlike sampled output, it executes immediately rather than at the end of a program, and can therefore be used to examine the current state of an ongoing simulation. The vectors to be output are defined via a `<dependencies>` element, and the basis is chosen by the `basis` attribute supplied to that `<dependencies>` element, as usual. A single `<breakpoint>` element must only contain vectors of equal dimension. The data format is specified by the `format` attribute, with current options being “ascii”, “binary” and the recommended: “hdf5”. The filename for the output can be specified by a `filename` attribute, in which case the same filename will be used each time the element is executed. If the `filename` attribute is not specified, then the first output will default to “1.xsil”, and subsequent executions of the same breakpoint will increment the number by one.

Example syntax:

```
<breakpoint filename="groundstate_break.xsil" format="hdf5">  
  <dependencies basis="ky">wavefunction</dependencies>  
</breakpoint>
```

5.4.15 Output element

The `<output>` element describes the output of the program. It is often inefficient to output the complete state of all vectors at all times during a large simulation, so the purpose of this function is to define subsets of the information required for output. Each different format of information is described in a different `<sampling_group>` element inside the output element. The `<output>` element may contain any number of `<sampling_group>` elements. The format of the output data can be specified by the optional `format` attribute, which may take values of “ascii”, “binary”, and “hdf5” (the default). The filename can be specified with the optional `filename` attribute, which otherwise defaults to the simulation name with the ‘.xsil’ suffix.

The `<samples>` inside `<integrate>` elements defines a string of integers, with exactly one for each `<sampling_group>` element. During that integration, the variables described in each `<sampling_group>` element will be sampled and stored that number of times.

Sampling Group Element

A `<sampling_group>` element defines a set of variables that we wish to output, and typically they are functions of some subset of vectors. The names of the desired variables are listed in a `<moments>` element, just like the `<components>` element of a vector. They are defined with a ‘*CDATA*’ block, accessing any components of vectors and computed vectors that are defined in a `<dependencies>` element, also just like a vector. *Computed vectors* and `<operator>` elements can be defined and used in the definition, just like in an `<integrate>` element.

The basis of the output is specified by the `basis` attribute. This overrides any basis specification in the `<dependencies>` element. Because we often wish to calculate these vectors on a finer grid than we wish to output, it is possible to specify that the output on a subset of the points defined for any transverse dimension. This is done by adding a number in parentheses after that dimension in the basis string, e.g. `basis="x y (32) kz (64) "`. If the number is zero, then that dimension is integrated out. If that number is one or more, then that dimension will be sampled on a subset of points in that space.

The `initial_sample` attribute, which must be “yes” or “no”, determines whether the moment group will be sampled before any integration occurs.

Example syntax:

```

<output format="hdf5" filename="SimOutput.xsil">
  <sampling_group basis="x y" initial_sample="yes">
    <computed_vector name="filter3" dimensions="" type="complex">
      <components>sparemomentagain</components>
      <evaluation>
        <dependencies basis="kx ky">integrated_u main</dependencies>
        <![CDATA[
          sparemomentagain = mod2(u);
        ]]>
      </evaluation>
    </computed_vector>
    <operator kind="ex">
      <operator_names>L</operator_names>
      <![CDATA[
        L = -T*kx*kx/mu;
      ]]>
    </operator>
    <moments>amp ke</moments>
    <dependencies>main filter1</dependencies>
    <![CDATA[
      amp = mod2(u + moment);
      ke = mod2(L[u]);
    ]]>
  </sampling_group>

  <sampling_group basis="kx(0) ky(64)" initial_sample="yes">
    <moments>Dens_P </moments>
    <dependencies>fields </dependencies>
    <![CDATA[
      Dens_P = mod2(psi);
    ]]>
  </sampling_group>
</output>

```

5.4.16 XMD S2-specific C syntax

Sampling complex numbers can be written more efficiently using:

```

<![CDATA[
  __SAMPLE_COMPLEX(W);
]]>

```

which is short for

```

<![CDATA[
  WR = W.Re();
  WI = W.Im();
]]>

```

Various properties of dimensions are available. For example, for a dimension called `x`:

- The number of points is accessible with the variable `__lattice_x`,
- The minimum range of that dimension is `__min_x`,
- The maximum range of that dimension is `__max_x`,
- The step size of a dimension is `dx`, and if it is constant, also available using `__dx`, but note that the latter does not include the effect of any `volumePrefix` you may have set!

5.5 Modified Midpoint Method

Although the modified midpoint can be used standalone as an ordinary differential equation integrator, it is regarded as much more powerful when used as a stepper to complement the Bulirsch-Stoer technique.

The modified midpoint method advances a vector of dependent variables $y(x)$ from a point x , to a point $x + H$ by a sequence of n substeps, each of size $h = H/n$.

The number of right-hand side evaluations required by the modified midpoint method is $n + 1$. The formulas for the method are

$$\begin{aligned} z_0 &= y(x) \\ z_1 &= z_0 + hf(x, z_0) \\ z_{m+1} &= z_{m-1} + 2hf(x + mh, z_m) \text{ for } m = 1, 2, \dots, n-1 \\ y(x + H) &\approx y_n = \frac{1}{2} [z_n + z_{n-1} + hf(x + H, z_n)] \end{aligned}$$

The error of this, expressed as a power series in h , the stepsize, contains only even powers of h :

$$y_n - y(x + H) = \sum_{i=1}^{\infty} \alpha_i h^{2i}$$

where H is held constant, but h changes y by varying the n in $h = H/n$.

The importance of this even power series is that using Richardson Extrapolation to combine steps and knock out higher-order error terms gains us two orders at a time.

The modified midpoint method is a second-order method, but holds an advantage over second order Runge-Kutta, as it only requires one derivative evaluation per step, instead of the two evaluations that Runge-Kutta necessitates.

5.6 Bulirsch-Stoer Algorithm

The Bulirsch-Stoer algorithm utilizes three core concepts in its design.

First, the usage of Richardson Extrapolation.

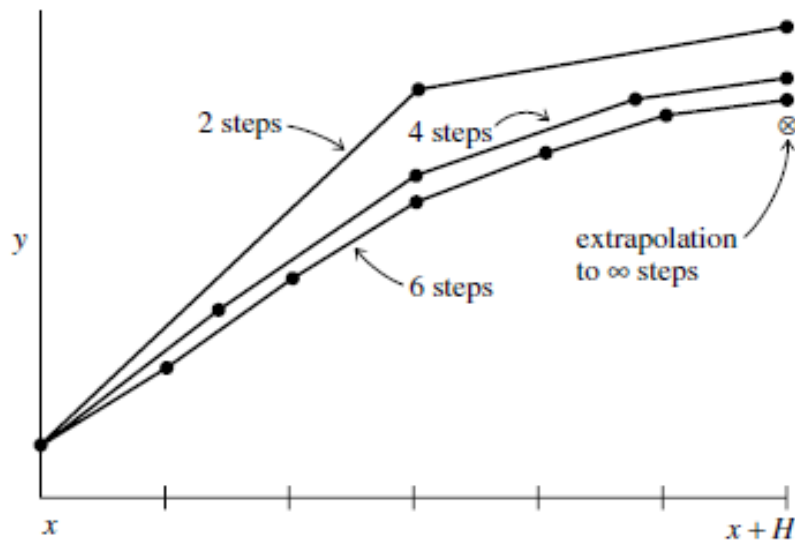


Figure 17.3.1. Richardson extrapolation as used in the Bulirsch-Stoer method. A large interval H is spanned by different sequences of finer and finer substeps. Their results are extrapolated to an answer that is supposed to correspond to infinitely fine substeps. In the Bulirsch-Stoer method, the integrations are done by the modified midpoint method, and the extrapolation technique is polynomial extrapolation.

Richardson Extrapolation considers the final answer of a numerical calculation as being an analytic function of an adjustable parameter such as the stepsize h . That analytic function can be probed by performing the calculation with various values of h , none of them being necessarily small enough to yield the accuracy that we desire. When we know enough about the function, we fit it to some analytic form and then evaluate it at the point where $h = 0$.

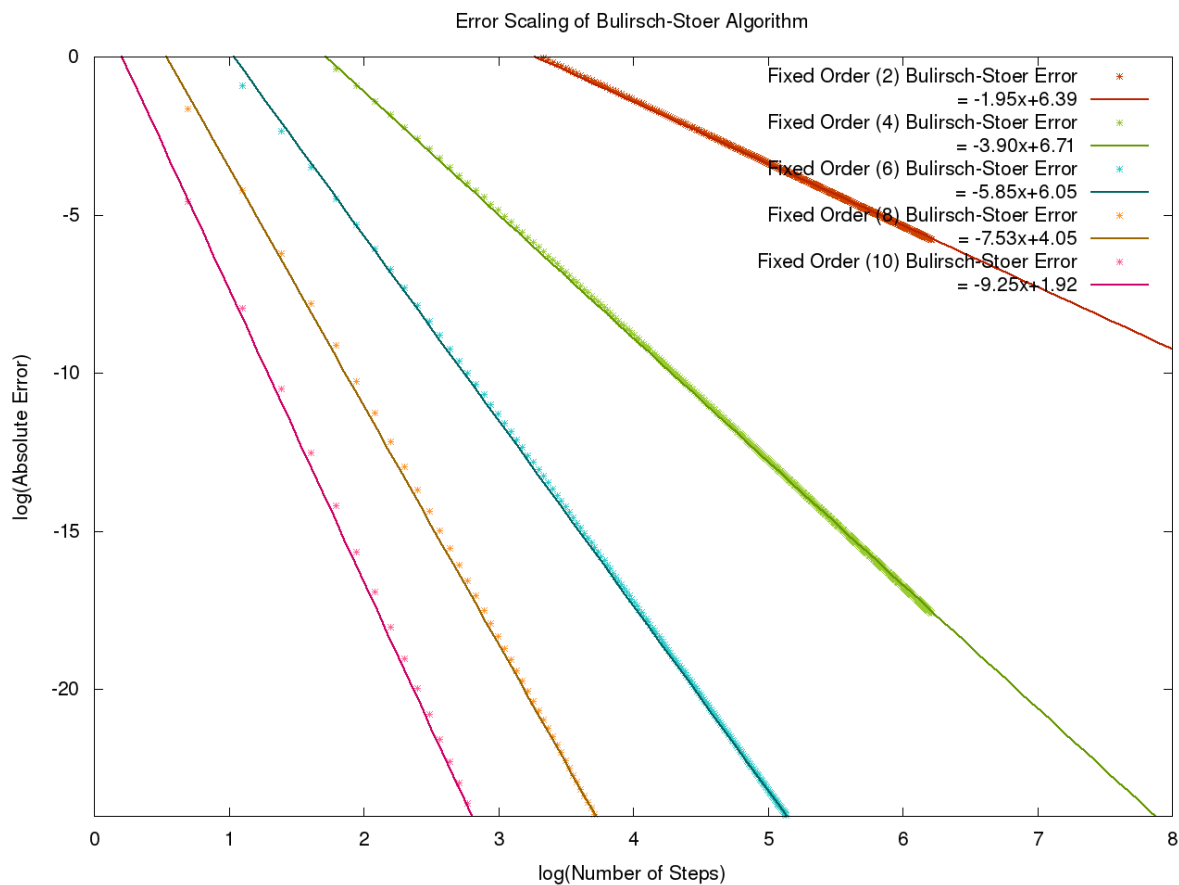
Secondly, the usage of rational function extrapolation in Richardson-type applications. Rational function fits can remain good approximations to analytic functions even after the various terms in powers of h , all have comparable magnitudes. In other words, h can be so large as to make the whole notion of the “order” of the method meaningless — and the method can still work superbly.

The third idea is to use an integration method whose error function is strictly even, allowing the rational function or polynomial approximation to be in terms of the variable h^2 instead of just h .

These three ideas give us the Bulirsch-Stoer method, where a single step takes us from x to $x + H$, where H is supposed to be a significantly large distance. That single step consists of many substeps of the modified midpoint method, which is then extrapolated to zero stepsize.

(Excerpts derived from **Numerical Recipes: The Art of Scientific Computing**, Third Edition (2007), p1256; Cambridge University Press; ISBN-10: 0521880688, <http://www.nr.com/>)

5.7 Error Scaling Behaviour



The graph above shows the error scaling behaviour for the Bulirsch-Stoer method. This was generated using data from XMDs2 for a simple problem whose analytical solution was known. For more information and to generate this plot yourself see the `testsuite/integrators/richardson_extrapolation/error_scaling` directory. There you will find the .xmds files for generating the data and a python script to generate the plot above (requires gnuplot).

ADVANCED TOPICS

This section has further details on some important topics.

Importing data (importing data into XMDS2, and data formats used in the export)

Convolutions and Fourier transforms (extra information on the Fourier transforms used in XMDS2, how this applies to convolutions)

Dimension aliases (dimensions which are declared to be identical, useful for correlation functions)

6.1 Importing data

There are many cases where it is advantageous to import previously acquired data into XMDS2. For example, the differential equation you wish to solve may depend on a complicated functional form, which is more easily obtained via an analytical package such as Mathematica or Maple. Furthermore, importing data from another source can be quicker than needlessly performing calculations in XMDS2. In this tutorial, we shall consider an example of importing into XMDS2 a function generated in Mathematica, version 6.0. Note, however, that in order to do this it is required that hdf5 is installed (see <http://www.hdfgroup.org/HDF5/>).

Suppose we want to import the following function into XMDS2:

$$f(x) = x^2.$$

The first step is to create an hdf5 file, from XMDS2, which specifies the dimensions of the grid for the x dimension. Create and save a new XMDS2 file. For the purposes of this tutorial we shall call it “grid_specifier.xmls” with name “grid_specifier”. Within this file, enter the following “dummy” vector - which we shall call “gen_dummy” - which depends on the x dimension:

```
<vector type="real" name="gen_dummy" dimensions="x">
  <components>dummy</components>
  <initialisation>
    <![CDATA[
      dummy = x;
    ]]>
  </initialisation>
</vector>
```

What “dummy” is not actually important. It is only necessary that it is a function of x . However, it is important that the domain and lattice for the x dimension are identical to those in the XMDS2 you plan to import the function into. We output the following xsil file (in hdf5 format) by placing a breakpoint in the sequence block as follows:

```
<sequence>
  <breakpoint filename="grid.xmls" format="hdf5">
    <dependencies>
      gen_dummy
    </dependencies>
  </breakpoint>
```

In terminal, compile the file “grid_specifier.xmds” in XMD S2 and run the C code as usual. This creates two files - “grid.xsil” and “grid.h5”. The file “grid.h5” contains the list of points which make up the grids for the x dimensions. This data can now be used to ensure that the function $f(x)$ which we will import into XMD S2 is compatible with the the specified grid in your primary XMD S2 file.

In order to read the “grid.h5” data into Mathematica version 6.0, type the following command into terminal:.. code-block:

```
xsil2graphics2 -e grid.xsil
```

This creates the Mathematica notebook “grid.nb”. Open this notebook in Mathematica and evaluate the first set of cells. This has loaded the grid information into Mathematica. For example, suppose you have specified that the x dimension has a lattice of 128 points and a domain of (-32, 32). Then calling “x1” in Mathematica should return the following list:

```
{-32., -31.5, -31., -30.5, -30., -29.5, -29., -28.5, -28., -27.5,
-27., -26.5, -26., -25.5, -25., -24.5, -24., -23.5, -23., -22.5,
-22., -21.5, -21., -20.5, -20., -19.5, -19., -18.5, -18., -17.5,
-17., -16.5, -16., -15.5, -15., -14.5, -14., -13.5, -13., -12.5,
-12., -11.5, -11., -10.5, -10., -9.5, -9., -8.5, -8., -7.5, -7.,
-6.5, -6., -5.5, -5., -4.5, -4., -3.5, -3., -2.5, -2., -1.5, -1.,
-0.5, 0., 0.5, 1., 1.5, 2., 2.5, 3., 3.5, 4., 4.5, 5., 5.5, 6., 6.5,
7., 7.5, 8., 8.5, 9., 9.5, 10., 10.5, 11., 11.5, 12., 12.5, 13.,
13.5, 14., 14.5, 15., 15.5, 16., 16.5, 17., 17.5, 18., 18.5, 19.,
19.5, 20., 20.5, 21., 21.5, 22., 22.5, 23., 23.5, 24., 24.5, 25.,
25.5, 26., 26.5, 27., 27.5, 28., 28.5, 29., 29.5, 30., 30.5, 31.,
31.5}
```

This is, of course, the list of points which define our grid.

We are now in a position to define the function $f(x)$ in Mathematica. Type the following command into a cell in the Mathematica notebook “grid.nb”:

```
f[x_] := x^2
```

At this stage this is an abstract mathematical function as defined in Mathematica. What we need is a list of values for $f(x)$ corresponding to the specified grid points. We will call this list “func”. This achieved by simply acting the function on the list of grid points “x1”:

```
func := f[x1]
```

For the example grid mentioned above, calling “func” gives the following list:

```
{1024., 992.25, 961., 930.25, 900., 870.25, 841., 812.25, 784.,
756.25, 729., 702.25, 676., 650.25, 625., 600.25, 576., 552.25, 529.,
506.25, 484., 462.25, 441., 420.25, 400., 380.25, 361., 342.25, 324.,
306.25, 289., 272.25, 256., 240.25, 225., 210.25, 196., 182.25, 169.,
156.25, 144., 132.25, 121., 110.25, 100., 90.25, 81., 72.25, 64.,
56.25, 49., 42.25, 36., 30.25, 25., 20.25, 16., 12.25, 9., 6.25, 4.,
2.25, 1., 0.25, 0., 0.25, 1., 2.25, 4., 6.25, 9., 12.25, 16., 20.25,
25., 30.25, 36., 42.25, 49., 56.25, 64., 72.25, 81., 90.25, 100.,
110.25, 121., 132.25, 144., 156.25, 169., 182.25, 196., 210.25, 225.,
240.25, 256., 272.25, 289., 306.25, 324., 342.25, 361., 380.25, 400.,
420.25, 441., 462.25, 484., 506.25, 529., 552.25, 576., 600.25, 625.,
650.25, 676., 702.25, 729., 756.25, 784., 812.25, 841., 870.25, 900.,
930.25, 961., 992.25}
```

The next step is to export the list “func” as an h5 file that XMD S2 can read. This is done by typing the following command into a Mathematica cell:

```
SetDirectory[NotebookDirectory[]];
Export["func.h5", {func, x1}, {"Datasets", { "function_x", "x"}}
```

In the directory containing the notebook “grid.nb” you should now see the file “func.h5”. This file essentially contains the list {func, x1}. However, the hdf5 format stores func and x1 as separate entities called “Datasets”.

For importation into XMDs2 it is necessary that these datasets are named. This is precisely what the segment `{"Datasets", { "function_x", "x"}}` in the above Mathematica command does. The dataset corresponding to the grid `x1` needs to be given the name of the dimension that will be used in XMDs2 - in our case this is `"x"`. It does not matter what the name of the dataset corresponding to the list `"func"` is; in our case it is `"function_x"`.

The final step is to import the file `"func.h5"` into your primary XMDs2 file. This data will be stored as a vector called `"gen_function_x"`, in component `"function_x"`.

```
<vector type="real" name="gen_function_x" dimensions="x">
  <components>function_x</components>
  <initialisation kind="hdf5">
    <filename> function_x.h5 </filename>
  </initialisation>
</vector>
```

You're now done. Anytime you want to use $f(x)$ you can simply refer to `"function_x"` in the vector `"gen_function_x"`.

The situation is slightly more complicated if the function you wish to import depends on more than one dimension. For example, consider

$$g(x, y) = x \sin(y)$$

As for the single dimensional case, we need to export an hdf5 file from XMDs2 which specifies the dimensions of the grid. As in the one dimensional case, this is done by creating a dummy vector which depends on all the relevant dimensions:

```
<vector type="real" name="gen_dummy" dimensions="x y">
  <components>dummy</components>
  <initialisation>
    <![CDATA[
      dummy = x;
    ]]>
  </initialisation>
</vector>
```

and exporting it as shown above.

After importing the grid data into Mathematica, define the multi-dimensional function which you wish to import into XMDs2:

```
g[x_, y_] := x*Sin[y]
```

We need to create a 2x2 array of data which depends upon the imported lists `x1` and `y1`. This can be done by using the `Table` function:

```
func := Table[g[x, p], {x, x1}, {p, p1}]
```

This function can be exported as an h5 file,

```
SetDirectory[NotebookDirectory[]];
Export["func.h5", {func, x1, y1}, {"Datasets", { "function_x", "x", "y"}}
```

and imported into XMDs2 as outlined above.

6.2 Convolutions and Fourier transforms

When evaluating a numerical Fourier transform, XMDs2 doesn't behave as expected. While many simulations have ranges in their spatial coordinate (here assumed to be `x`) that range from some negative value x_{\min} to some positive value x_{\max} , the Fourier transform used in XMDs2 treats all spatial coordinates as starting at zero. The

result of this is that a phase factor of the form $e^{-ix_{\min}k}$ is applied to the Fourier space functions after all forward (from real space to Fourier space) Fourier transforms, and its conjugate is applied to the Fourier space functions before all backward (from Fourier space to real space) Fourier transforms.

The standard Fourier transform is

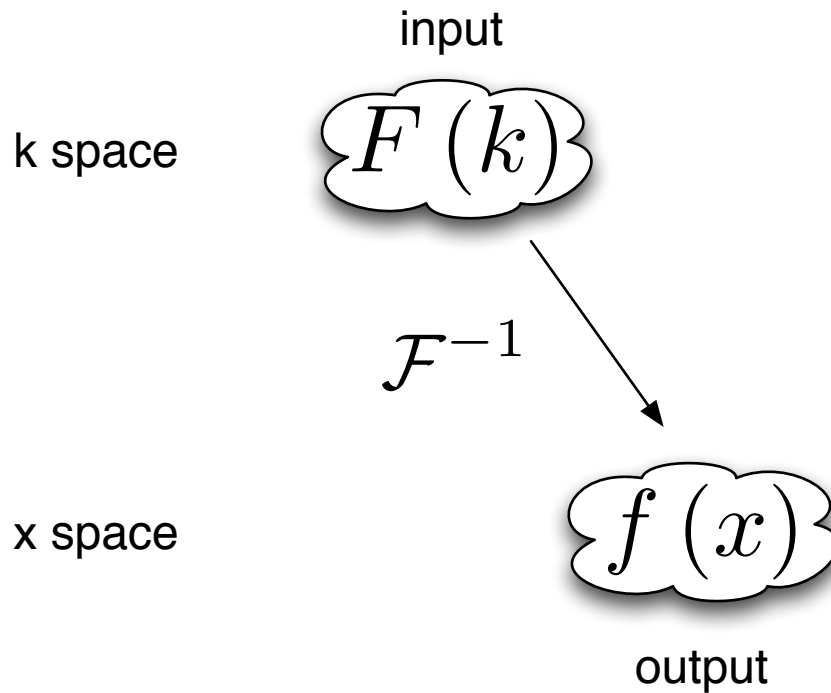
$$\mathcal{F}[f(x)](k) = \frac{1}{\sqrt{2\pi}} \int_{x_{\min}}^{x_{\max}} f(x) e^{-ikx} dx$$

The XMDS2 Fourier transform is

$$\begin{aligned} \tilde{\mathcal{F}}[f(x)](k) &= \frac{1}{\sqrt{2\pi}} \int_{x_{\min}}^{x_{\max}} f(x) e^{-i(x+x_{\min})k} dx \\ &= e^{-ix_{\min}k} \mathcal{F}[f(x)](k) \end{aligned}$$

When the number of forward Fourier transforms and backwards Fourier transforms are unequal a phase factor is required. Some examples of using Fourier transforms in XMDS2 are shown below.

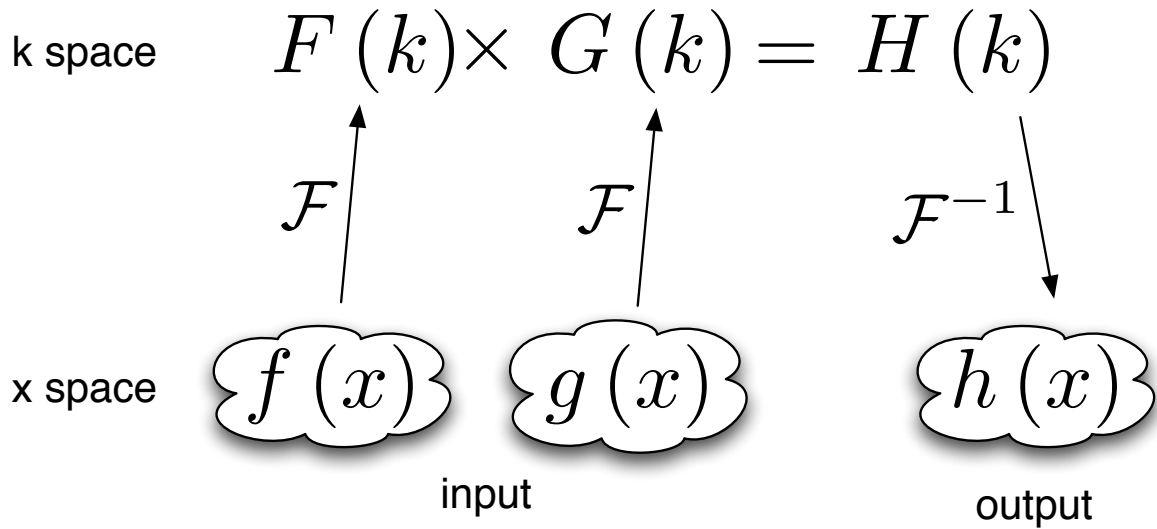
6.2.1 Example 1



When data is input in Fourier space and output in real space there is one backwards Fourier transform is required. Therefore the Fourier space data must be multiplied by a phase factor before the backwards Fourier transform is applied.

$$\mathcal{F}^{-1}[F(k)](x) = \tilde{\mathcal{F}}[e^{ix_{\min}k} F(k)](x)$$

6.2.2 Example 2

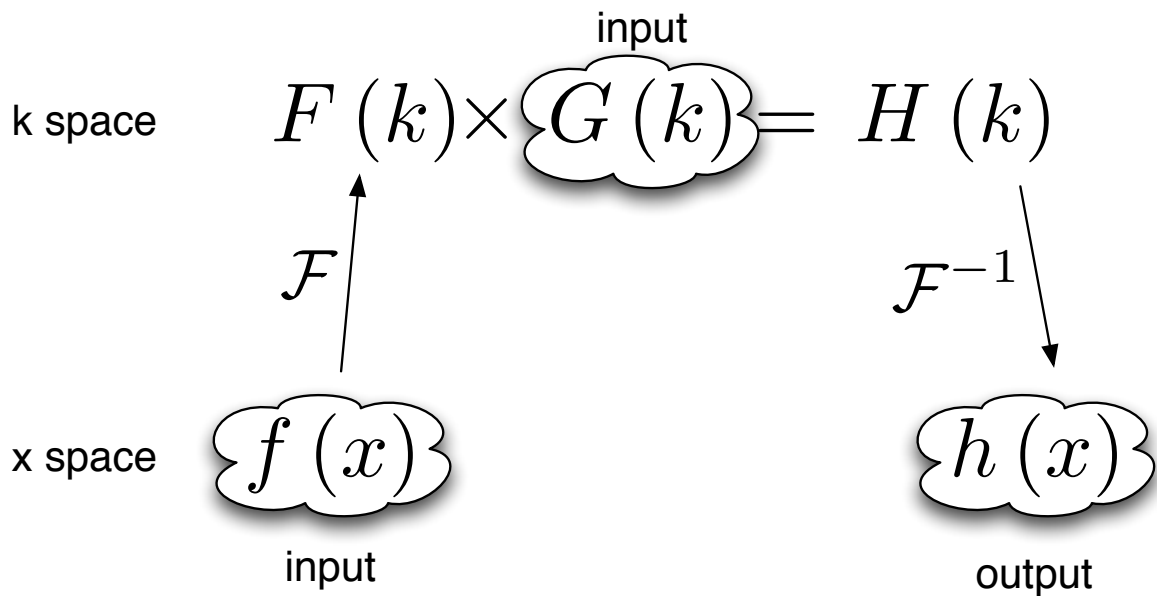


Functions of the form $h(x) = \int f(x')g(x-x')dx'$ can be evaluated using the convolution theorem:

$$\mathcal{F}[h(x)](k) = \mathcal{F}[f(x)](k) \times \mathcal{F}[g(x)](k)$$

This requires two forward Fourier transforms to get the two functions f and g into Fourier space, and one backwards Fourier transform to get the resulting product back into real space. Thus in Fourier space the product needs to be multiplied by a phase factor $e^{-ix_{\min}k}$

6.2.3 Example 3



Sometimes when the convolution theorem is used one of the forward Fourier transforms is calculated analytically and input in Fourier space. In this case only one forward numerical Fourier transform and one backward numerical Fourier transform is used. The number of forward and backward transforms are equal, so no phase factor is required.

6.3 ‘Loose’ geometry_matching_mode

6.4 Dimension aliases

Dimension aliases specify that two or more dimensions have exactly the same lattice, domain and transform. This can be useful in situations where the problem enforces this, for example when computing correlation functions or representing square matrices.

Dimension aliases are not just a short-hand for defining an additional dimension, they also permit dimensions to be accessed *non-locally*, which is essential when computing spatial correlation functions.

Here is how to compute a spatial correlation function $g^{(1)}(x, x') = \psi^*(x)\psi(x')$ of the quantity `psi`:

```
<simulation xmds-version="2">

  <!-- name, features block -->

  <geometry>
    <propagation_dimension> t </propagation_dimension>
    <transverse_dimensions>
      <dimension name="x" lattice="1024" domain="(-1.0, 1.0)" aliases="xp" />
    </transverse_dimensions>
  </geometry>

  <vector name="wavefunction" type="complex" >
    <components> psi </components>
    <initialisation>
      <!-- initialisation code -->
    </initialisation>
  </vector>

  <computed_vector name="correlation" dimensions="x xp" type="complex" >
    <components> g1 </components>
    <evaluation>
      <dependencies> wavefunction </dependencies>
      <![CDATA[
        g1 = conj(psi(x => x)) * psi(x => xp);
      ]]>
    </evaluation>
  </computed_vector>

  <!-- integration and sampling code -->

</simulation>
```

In this simulation note that the vector `wavefunction` defaults to only having the dimension “x” even though “xp” is also a dimension (implicitly declared through the `aliases` attribute). `vector`’s without an explicit `dimensions` attribute will only have the dimensions that are explicitly listed in the `transverse_dimensions` block, i.e. this will not include aliases.

See the example `groundstate_gaussian.xmds` for a complete example.

FREQUENTLY ASKED QUESTIONS

7.1 XMDS scripts look complicated! How do I start?

If you're unfamiliar with XMDS2, writing a script from scratch might seem difficult. In most cases, however, the best approach is to take an existing script and modify it for your needs. At the most basic level, you can simply take a script from the `/examples` directory that is similar to what you want to do, change the name of the integration variable(s) and replace the line describing the differential equation to use your DE instead. That's all you need to do, and will ensure all the syntax is correct and all the required XML blocks are present.

You can then incrementally change things such as the number of output points, what quantities get output, number of grid points, and so on. Many XMDS2 users have never written a script from scratch, and just use their previous scripts and example scripts as a scaffold when they create a script for a new problem.

7.2 Where can I get help?

The documentation on this website is currently incomplete, but it still covers a fair bit and is worth reading. Similarly, the example scripts in the `/examples` directory cover most of the functionality of XMDS2, so it's worth looking through them to see if any of them do something similar to what you're trying to do.

You should also feel free to email questions to the XMDS users' mailing list at xmds-users@lists.sourceforge.net, where the developers and other users can assist you. You can join the mailing list by going to <http://sourceforge.net/projects/xmds/> and clicking on "mailing lists." Also, if you look through the mailing list archives, your particular problem may already have been discussed.

7.3 How should I cite XMDS2?

If you publish work that has involved XMDS2, please cite it as: *Comput. Phys. Commun.* 184, 201-208 (2013).

7.4 I think I found a bug! Where should I report it?

Please report bugs to the developer mailing list at xmds-devel@lists.sourceforge.net. In your email, please include a description of the problem and attach the XMDS2 script that triggers the bug.

7.5 How do I put time dependence into my vectors?

Standard vectors can't have time dependence (or, more accurately, depend on the `propagation_dimension` variable), but computed vectors can. So, for example, if you have set your `propagation_dimension` as "t", you can simply use the variable "t" in your computed vector and it will work.

Alternatively, you can explicitly use the `propagation_dimension` variable in your differential equation inside the `<operators>` block.

7.6 Can I specify the range of my domain and number of grid points at run-time?

Yes, you can. In your script, specify the domain and number of grid points as arguments to be passed in at run-time, use those variables in your `<geometry>` block rather than explicitly specifying them, and use the `<validation kind="run-time" />` feature. See the [Validation](#) entry in the Reference section for an example.

While the domain can always be specified in this way, specifying the lattice size at run-time is currently only allowed with the following transforms: 'dct', 'dst', 'dft' and 'none' (see [Transforms](#) in the Reference section).

Also note that for some multi-dimensional spaces using different transforms, XMD52 will sometimes optimise the code it generates based on the relative sizes of the dimensions. If one or more of the lattices are specified at run-time it is unable to do this and will have to make guesses. In some situations this may result in slightly slower code.

7.7 When can I use IP operators (and why should I) and when must I use EX operators?

An `<operator>` that specifies named operators to be used in integration equations can have the `kind="IP"` or `kind="EX"` attribute, standing for 'interaction picture' and 'explicit' operators respectively. Explicit operators can be used in all situations, and simply construct and calculate a new vector of the form in the square brackets. IP operators use less memory and can improve speed by allowing larger timesteps, but have two important restrictions. **Use of IP operators without understanding these restrictions can lead to incorrect code.**

Some explanation is in order. The IP algorithm applies the operator separately to the rest of the evolution. The reason this can be so effective is that the separate evolution can be performed exactly. The solution of the equation $\frac{d\psi}{dt} = L\psi$ is $\psi(t + \Delta t) = \exp(L\Delta t)\psi(t)$ for arbitrarily large timestep Δt . For a diagonal linear L , the matrix exponential is straightforward. Also, when it is constant, then the exponential can be computed and stored prior to the integration, which makes the implementation of this operator very cheap. Thus, when IP operators are defined, XMD52 reads the equations as written by the user, and determines which operators to apply to which fields. It then implements these operators separately, and the text describing the operator inside the equations (in this example, the `L[psi]` term) is replaced by the numeral zero.

Therefore, the limitations of IP operators themselves means that they can only be applied to named components of one of the integration vectors, and not functions of those components. Furthermore, an IP operator acting on a component must only be used in the derivative for that particular component. Secondly, due to the implementation of IP operators in XMD52, it is not safe to use them in comments, or in conjunction with declared variables. It is also not safe to multiply or divide them by any factors, functions or vectors. They must turn up in a purely additive way when defining the derivative of a component of an integration vector. The XMD52 parser attempts to catch possible violations of these rules, and will produce warnings in some cases.

7.8 Visual Editors

In this section goes stuff about how to set up TextMate (or other editors to highlight xpdeint scripts).

UPGRADING FROM XMDS 1.X

While **XMDS2** is a complete rewrite of the **XMDS** project, much of the syntax has remained very similar. That said, your code will have to be rewritten as an XMDS2 program. We recommend that you work through the *Quickstart Tutorial* and perhaps the *Worked Examples* sections, and then you should be good to go.

The main news when switching to XMDS2 is the long list of new things you can do. If it's an initial value problem, XMDS2 has a good chance of being able to solve it.

We have made the decision to call the executables “xm2s2” and “xsil2graphics2” so that you can keep using your old installation in parallel with the new version.

OPTIMISATION HINTS

There are a variety of things you can do to make your simulations run faster.

9.1 Geometry and transform-based tricks

9.1.1 Simpler simulation geometries

Consider symmetry, can you use `dct` transforms or `bessel` transforms? Do you really need that many points? How big does your grid need to be? Could absorbing boundary conditions help?

9.1.2 Tricks for Bessel and Hermite-Gauss transforms

Dimensions using matrix transforms should be first for performance reasons. Unless you're using MPI, in which case XMDS can work it out for the first two dimensions. Ideally, XMDS would sort it out in all cases, but it's not that smart yet.

9.2 Reduce code complexity

Avoid transcendental functions like $\sin(x)$ or $\exp(x)$ in inner loops. Not all operations are made equal, use multiplication over division.

9.2.1 Optimising with the Interaction Picture (IP) operator

You should use the IP operator when you can. Only use the EX operator when you have to. If you must use the EX operator, consider making it `constant="no"`. It uses less memory. When you use the IP operator, make sure you know what it's doing. Do not pre- or post-multiply that term in your equations, XMDS will do a fairly thorough check to see you aren't using the IP operator improperly, but it is possible to confuse XMDS's check.

If your simulation uses two or more dimensions, check to see if your IP operator is separable, i.e. can be written in the form $f(kx) + g(ky)$ (this is frequently possible in atom-optics simulations). If your IP operator is separable, create separate IP operators for each dimension. This provides a significant speedup (~30%) for simulations using an adaptive integrator. For example, instead of using the IP operator:

```
<operator kind="ip">
  <operator_names>T</operator_names>
  <![CDATA[
    T = -i*0.5*hbar/M*(kx*kx + ky*ky);
  ]]>
</operator>
<![CDATA[
  dpsi_dt = T[psi] + /* other terms */
]>
```

replace it with the pair of IP operators:

```
<operator kind="ip" dimensions="x">
  <operator_names>Tx</operator_names>
  <![CDATA[
    Tx = -i*0.5*hbar/M*kx*kx;
  ]]>
</operator>
<operator kind="ip" dimensions="y">
  <operator_names>Ty</operator_names>
  <![CDATA[
    Ty = -i*0.5*hbar/M*ky*ky;
  ]]>
</operator>
<![CDATA[
  dpsi_dt = Tx[psi] + Ty[psi] + /* other terms */
]]>
```

When using the IP operator, check if your operator is purely real or purely imaginary. If real, (e.g. $L = -0.5 \cdot k_x \cdot k_x$), then add the attribute `type="real"` to the `<operator kind="ip">` tag. If purely imaginary, use `type="imaginary"`. This optimisation saves performing the part of the complex exponential that is unnecessary.

9.2.2 Consider writing the evolution in spectral basis

Evolution equations do not need to be written in the position basis. If your equations are diagonal in the spectral basis, then it makes more sense to compute the time derivative terms in that basis. For example, if you have the system

$$\begin{aligned}\frac{d\psi_1(x)}{dt} &= i \frac{\hbar}{2M} \frac{d^2\psi_1(x)}{dx^2} - i\Omega\psi_2(x) \\ \frac{d\psi_2(x)}{dt} &= i \frac{\hbar}{2M} \frac{d^2\psi_2(x)}{dx^2} - i\Omega\psi_1(x)\end{aligned}$$

then this is diagonal in the Fourier basis where it takes the form

$$\begin{aligned}\frac{d\psi_1(k_x)}{dt} &= -i \frac{\hbar k_x^2}{2M} \psi_1(k_x) - i\Omega\psi_2(k_x) \\ \frac{d\psi_2(k_x)}{dt} &= -i \frac{\hbar k_x^2}{2M} \psi_2(k_x) - i\Omega\psi_1(k_x)\end{aligned}$$

The first term in each evolution equation can be solved exactly with an IP operator, and the second term is diagonal in Fourier space. This can be written in XMD S as:

```
<operators>
  <integration_vectors basis="kx">wavefunction</integration_vectors>
  <operator kind="ip" type="imaginary" >
    <operator_names>Lxx</operator_names>
    <![CDATA[
      Lxx = -i*0.5*hbar_M*(kx*kx);
    ]]>
  </operator>
  <![CDATA[

    dpsi0_dt = Lxx[psi0] - i*Omega*psi1;
    dpsi1_dt = Lxx[psi1] - i*Omega*psi0;

  ]]>
</operators>
```

Although the `dpsi0_dt` code reads the same in position and Fourier space, it is the `basis=kx` attribute on `<integration_vectors>` that causes the evolution code to be executed in Fourier space.

A final optimisation is to cause the integration code itself to operate in Fourier space. By default, all time stepping (i.e. $f(t + \Delta t) = f(t) + f'(t)\Delta t$ for forward-Euler integration) occurs in the position space. As the derivative terms can be computed in Fourier space, it is faster to also to the time stepping in Fourier space too. This then means that no Fourier transforms will be needed at all during this integrate block (except as needed by sampling). To cause time-stepping to happen in Fourier space, we add the `home_space="k"` attribute to the surrounding `<integrate>` block. By default, `home_space` has the value `"x"` which means position space, even if you don't have an `x` dimension.

The fully optimised code then reads:

```
<integrate algorithm="ARK45" interval="1" tolerance="1e-6" home_space="k">
  <samples> 10 </samples>
  <operators>
    <integration_vectors basis="kx">wavefunction</integration_vectors>
    <operator kind="ip" type="imaginary" >
      <operator_names>Lxx</operator_names>
      <![CDATA[
        Lxx = -i*0.5*hbar_M*(kx*kx);
      ]]>
    </operator>
    <![CDATA[

      dpsi0_dt = Lxx[psi0] - i*Omega*psi1;
      dpsi1_dt = Lxx[psi1] - i*Omega*psi0;

    ]]>
  </operators>
</integrate>
```

This code will not use any Fourier transforms during an ordinary time-stepping, and will be much faster than if the code were written without the `home_space` and `basis` attributes.

9.2.3 Don't recalculate things you don't have to

Use `computed_vectors` appropriately.

9.3 Compiler and library tricks

9.3.1 Faster compiler

If you're using an Intel CPU, then you should consider using their compiler, `icc`. They made the silicon, and they also made a compiler that understands how their chips work significantly better than the more-portable `GCC`.

9.3.2 Faster libraries

Intel MKL is faster than ATLAS, which is faster than GSL CBLAS. If you have a Mac, then Apple's `vecLib` is plenty fast.

Note for Linux users

If you used the linux installer, and are using Ubuntu or Debian, the version of the ATLAS package that gets installed is the generic version in the repositories. This version lacks architecture and CPU-specific optimizations.

Creating an ATLAS package locally tuned for your specific machine will result in a faster linear algebra implementation, which can significantly speed up problems utilizing matrix based transforms (bessel, hermite-gauss etc). Some simple tests using a cylindrically symmetric problem with one bessel transform dimension and one

FFT transform dimension showed speed increases from 5% - 100% over the default ATLAS package, depending on the number of grid points.

To create and install an ATLAS package optimized for your machine, carry out the following procedure:

Using your favourite package manager (e.g. Synaptic) to remove any current ATLAS libraries (probably `libatlas3gf-base`, `libatlas-dev`, `libatlas-base-dev`). Then create an empty directory whose path doesn't include any spaces. In this directory, do

```
apt-get source atlas
apt-get build-dep atlas
apt-get install devscripts dpkg-dev

cd atlas-*
sudo fakeroot debian/rules custom
cd ..
ls libatlas*.deb
```

Then, for each of the `.deb` packages listed by the `ls` command, install via:

```
sudo dpkg -i <filename here>.deb
```

This procedure was tested on Ubuntu 12.04 LTS, but an identical or very similar procedure should work for other Ubuntu/Debian versions.

Finally, note that the “`sudo fakeroot debian/rules custom`” package creation step carries out an exhaustive series of tests to optimize for your architecture, SSE support, cache hierarchy and so on, and can take a long time. Be patient.

9.3.3 Auto-vectorisation

Auto-vectorisation is a compiler feature that makes compilers generate more efficient code that can execute the same operation on multiple pieces of data simultaneously. To use this feature, you need to add the following to the `<features>` block at the start of your simulation:

```
<auto_vectorise />
```

This will make `xpdeint` generate code that is more friendly to compiler's auto-vectorisation features so that more code can be vectorised. It will also add the appropriate compiler options to turn on your compiler's auto-vectorisation features. For auto-vectorisation to increase the speed of your simulations, you will need a compiler that supports it such as `gcc 4.2` or later, or Intel's C compiler, `icc`.

9.3.4 OpenMP

OpenMP is a set of compiler directives to make it easier to use threads (different execution contexts) in programs. Using threads in your simulation does incur some overhead, so for the speedup to outweigh the overhead, you must have a reasonably large simulation grid. To add these compiler directives to the generated simulations, add the tag `<openmp />` in the `<features>` block. This can be used in combination with the auto-vectorisation feature above. Note that if you are using `gcc`, make sure you check that your simulations are faster by using this as `gcc`'s OpenMP implementation isn't as good as `icc`'s.

If you are using the OpenMP feature and are using **FFTW**-based transforms (Discrete Fourier/Cosine/Sine Transforms), you should consider using threads with your FFT's by adding the following to the `<features>` block at the start of your simulation:

```
<fftw threads="2" />
```

Replace the number of threads in the above code by the number of threads that you want to use.

9.3.5 Parallelisation with MPI

Some simulations are so large or take so much time that it is not reasonable to run them on a single CPU on a single machine. Fortunately, the [Message Passing Interface](#) was developed to enable different computers working on the same program to exchange data. You will need a MPI package installed to be able to use this feature with your simulations. One popular implementation of MPI is [OpenMPI](#).

9.4 Atom-optics-specific hints

9.4.1 Separate out imaginary-time calculation code

When doing simulations that require the calculation of the groundstate (typically via the imaginary time algorithm), typically the groundstate itself does not need to be changed frequently as it is usually the dynamics of the simulation that have the interesting physics. In this case, you can save having to re-calculate groundstate every time by having one script (call it `groundstate.xmgs`) that saves the calculated groundstate to a file using a breakpoint, and a second simulation that loads this calculated groundstate and then performs the evolution. More often than not, you won't need to re-run the groundstate finder.

The file format used in this example is [HDF5](#), and you will need the HDF5 libraries installed to use this example. The alternative is to use the deprecated `binary` format, however to load `binary` format data `xmgs`, the predecessor to `xpdeint` must be installed. Anyone who has done this before will tell you that installing it isn't a pleasant experience, and so HDF5 is the recommended file format.

If your wavefunction vector is called `'wavefunction'`, then to save the groundstate to the file `groundstate_break.h5` in the HDF5 format, put the following code immediately after the `integrate` block that calculates your groundstate:

```
<breakpoint filename="groundstate_break" format="hdf5">
  <dependencies>wavefunction</dependencies>
</breakpoint>
```

In addition to the `groundstate_break.h5` file, an XSIL wrapper `groundstate_break.xsil` will also be created for use with `xsil2graphics2`.

To load this groundstate into your evolution script, the declaration of your `'wavefunction'` vector in your evolution script should look something like

```
<vector name="wavefunction">
  <components>phi1 phi2</components>
  <initialisation kind="hdf5">
    <filename>groundstate_break.h5</filename>
  </initialisation>
</vector>
```

Note that the groundstate-finder doesn't need to have all of the components that the evolution script needs. For example, if you are considering the evolution of a two-component BEC where only one component has a population in the groundstate, then your groundstate script can contain only the `phi1` component, while your evolution script can contain both the `phi1` component and the `phi2` component. Note that the geometry of the script generating the groundstate and the evolution script must be the same.

9.4.2 Use an energy or momentum offset

This is just the interaction picture with a constant term in the Hamiltonian. If your state is going to rotate like $e^{i(\omega+\delta\omega)t}$, then transform your equations to remove the $e^{i\omega t}$ term. Likewise for spatial rotations, if one mode will be moving on average with momentum $\hbar k$, then transform your equations to remove that term. This way, you may be able to reduce the density of points you need in that dimension. Warning: don't forget to consider this when looking at your results. I (Graham Dennis) have been tripped up on multiple occasions when making this optimisation.

XSIL2GRAPHICS2

xsil2graphics2 is a way of converting ".xsil" files to formats that other programs can read. The syntax is described in the [Quickstart Tutorial](#), and by using the `xsil2graphics2 --help` option. It currently can convert any output format for use by Mathematica, MATLAB and Octave.

We recommend HDF5 format instead of the binary format for output and input, as many visualisation tools can already read/write to this format directly.

DEVELOPER DOCUMENTATION

Developers need to know more than users. For example, they need to know about the test suite, and writing test cases. They need to know how to perform a developer installation. They need to know how to edit and compile this documentation. They need a step-by-step release process.

11.1 Test scripts

Every time you add a new feature and/or fix a new and exciting bug, it is a great idea to make sure that the new feature works and/or the bug stays fixed. Fortunately, it is pleasantly easy to add a test case to the testing suite.

1. Write normal XMDS script that behaves as you expect.
2. Add a `<testing>` element to your script. You can read the description of this element and its contents below, and have a look at other testcases for examples, but the basic structure is simple:.

```
<testing>
  <command_line> </command_line>
  <arguments>
    <argument />
    <argument />
    ...
  </arguments>
  <input_xsil_file />
  <xsil_file>
    <moment_group />
    <moment_group />
    ...
  </xsil_file>
</testing>
```

3. Put into the appropriate `testsuite/` directory.
4. run `./run_tests.py` This will automatically generate your `_expected` files.
5. Commit the `.xmids`, `*_expected.xsil` file and any `*_expected*` data files.

11.1.1 Testing element

11.1.2 command_line element

11.1.3 input_xsil_file element

11.1.4 xsil_file element

11.1.5 moment_group element

11.2 XMDS Documentation

Documentation in XMDS is written as reStructuredText files (.rst), which are then parsed into HTML files to be displayed on their website.

You can find the user documentation folder located `admin/userdoc-source`. This is where all of the .rst files are kept. If you're wanting to add documentation to the site, you'll need to create your own .rst file, with the name of the webpage as the filename.

RST is a relatively simple language, which is basically simplified HTML markup. For documentation on how to make Lists, Href Links, Embed images etc, you should check here;

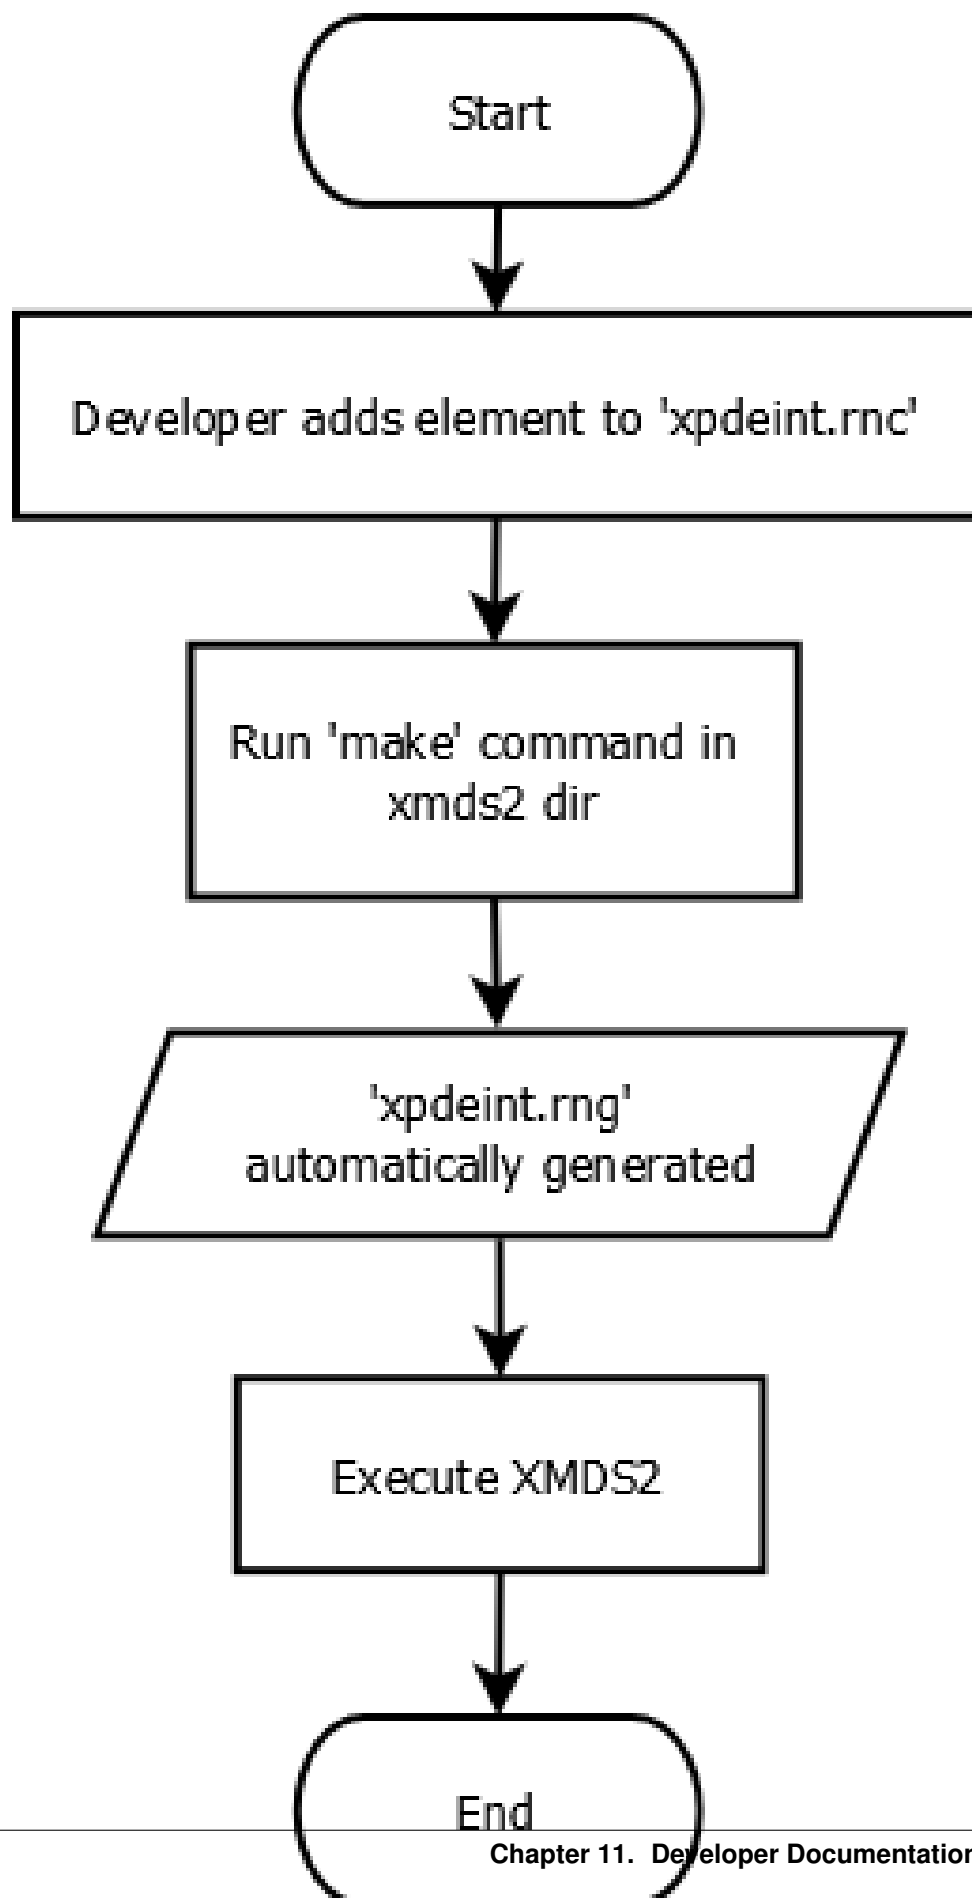
<http://docutils.sourceforge.net/docs/user/rst/quickref.html> <http://docutils.sourceforge.net/docs/ref/rst/restructuredtext.html>

However, you should easily be able to use some of the pre-existing .rst files in the project as a template to create yours.

Once your documentation is in this folder, it should be deployed along with the project to their website when you run `create_release_version.sh`, which can be found in the `/Trunk/xpdeint/admin` folder. If you would like to test to see what your rst file generates without running this shell script, you can use the Makefile in the `userdoc-source` folder, by running "make html".

NOTE: Before you can run the `create_release_version.sh` file, there are a few packages you will need. This command uses latex to generate the XMDS2 pdf, so you'll be needing the following packages; `texlive-fonts-recommended`, `texlive-lang-cjk`, `texlive-latex-base`.

11.3 How to update XMDS2 script validator (XML schema)



This is a short guide to adding an element to XMDS2, so that it can be validated by the XMDS2 script validator. In this guide, the example being used will be the addition of a matrix element to the validator. The matrix will have a 'name' and a 'type' (so it can be called later, and the type is known for future reference). Each matrix will also need a 'row' component, and possibly an initialisation value.

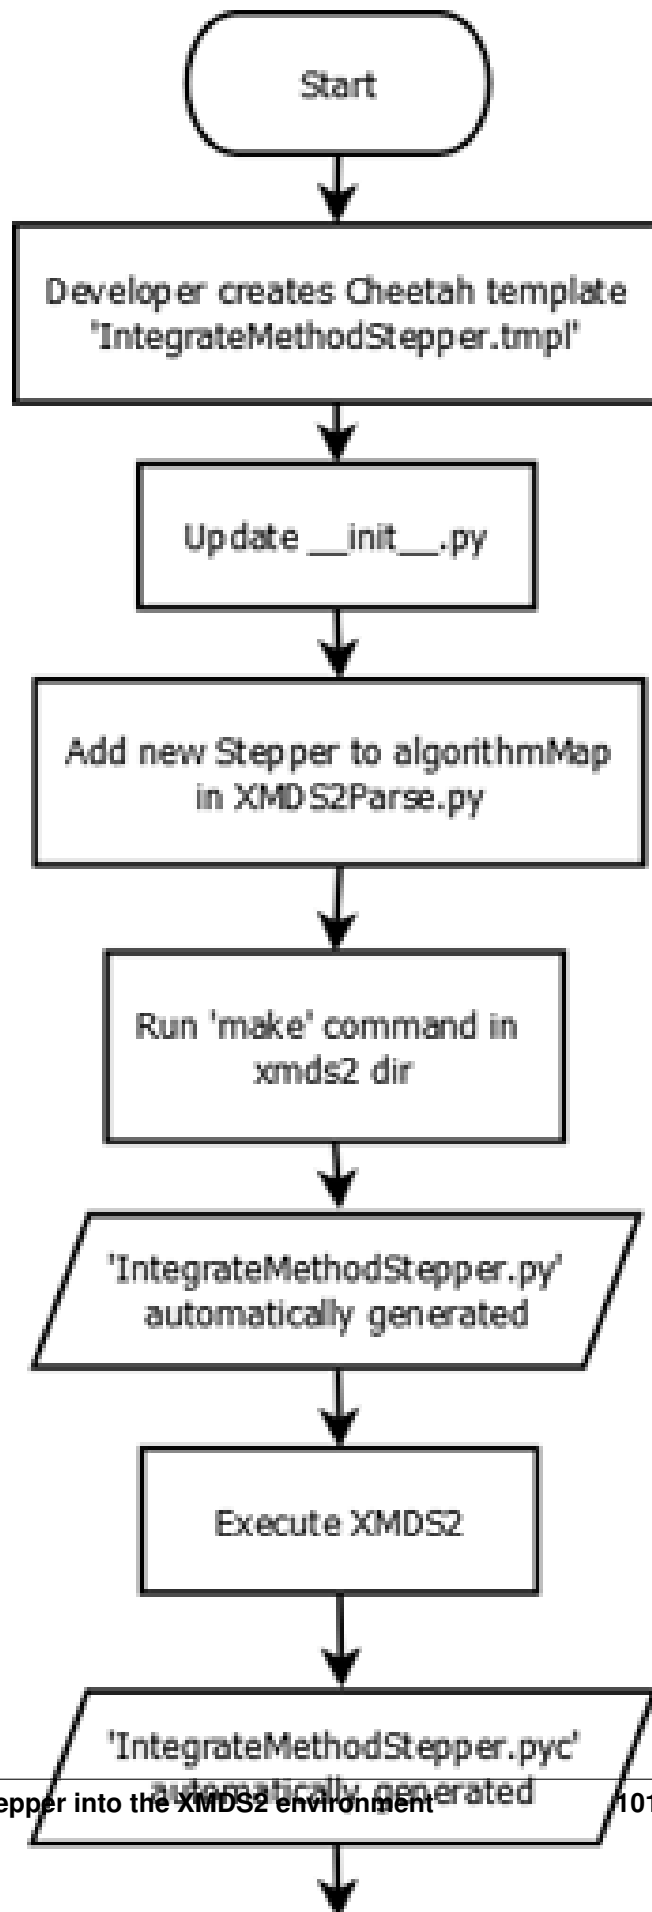
Navigate to `xpdeint/support/xpdeint.rnc`. This is a RelaxNG compact file, which specifies the XML schema which is only used for issuing warnings to users about missing or extraneous XML tags / attributes. Add the following lines to the end of the file (so that it is outside all other brackets in the file):

```
Matrix = element matrix {  
    attribute name { text }  
    , attribute type { text }?  
    , element components { text }  
    , element initialisation {  
        attribute kind { text }?  
    }?  
}
```

Save this file, and then in the terminal navigate to the folder `xpdeint/support/` and run `make`. This updates the XML based file `xpdeint/support/xpdeint.rng`, which is the file the parser uses to validate elements in XMDS2. This file which is used is in RelaxNG format, but RelaxNG compact is easier to read and edit.

Commit both `xpdeint/support/xpdeint.rnc` and `xpdeint/support/xpdeint.rng` to the code repository.

11.4 How to introduce a new integrator Stepper into the XMD S2 environment



This is a short guide to adding a new stepper containing a new mathematical technique to XMDS2, which can then be used by to integrate equations. This guide describes the logistics of introducing a new stepper and as such, the code inside the stepper template is outside the scope of this document. The new stepper which will be used in this guide will be called 'IntegrateMethodStepper'.

Navigate to the `xpdeint/Segments/Integrators` directory. Create a file called `IntegrateMethodStepper.tmpl` in this directory. In this file, implement the new integration algorithm (follow the convention of existing steppers in that folder). In this same folder, open the file named `__init__.py` and add the following line to the bottom of the file and save it:

```
import IntegrateMethodStepper
```

Navigate up until you are in the `xpdeint` directory. Open the file `XMDS2Parser.py`, and 'find' the algorithm map (Ctrl+F > `algorithmMap` works for most text editors). The mnemonic 'IM' will be used for our Stepper. If the stepper uses fixed step sizes, then add the following line to the algorithm map:

```
'IM': (Integrators.FixedStep.FixedStep, Integrators.IntegrateMethodStepper.IntegrateMethodStepper)
```

Otherwise, if your stepper is an adaptive Stepper, add the following line:

```
'IM': (Integrators.AdaptiveStep.AdaptiveStep, Integrators.IntegrateMethodStepper.IntegrateMethodStepper)
```

In the terminal, navigate to the `xpdeint` directory, and run `make` over the entire directory. 'IM' can now be used to specify the new Stepper as your integration algorithm inside your `.xmds` files, e.g.

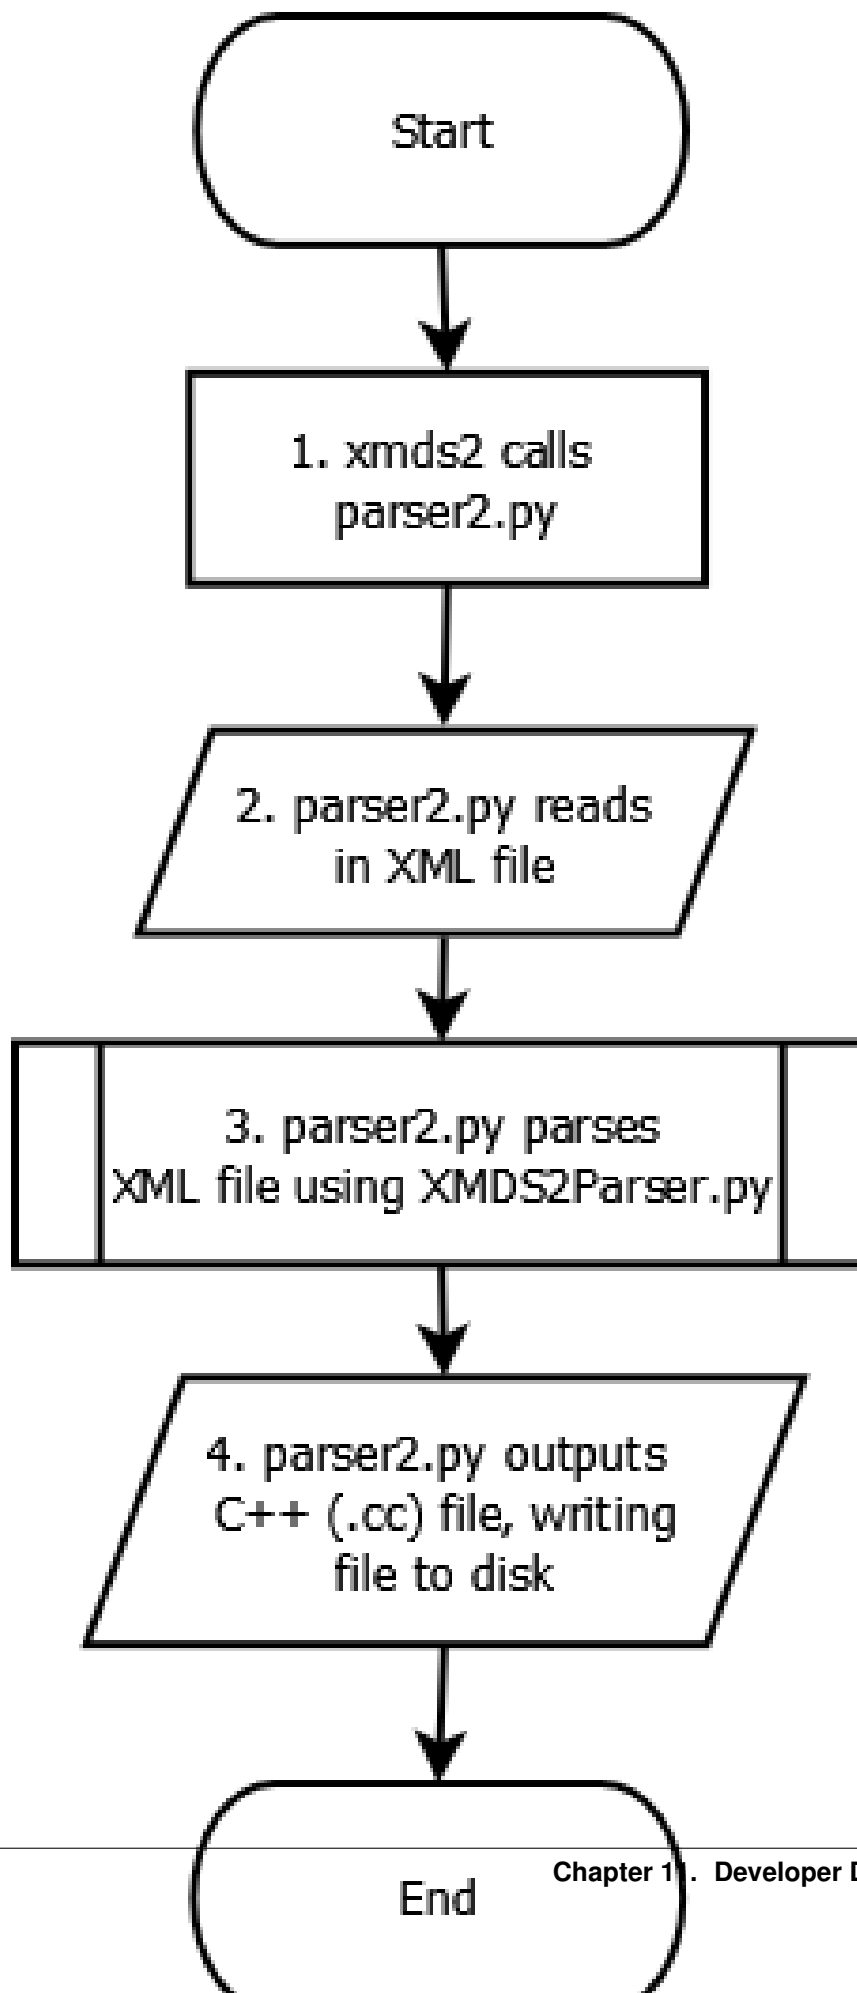
```
<integrate algorithm="IM" interval="5.0" steps="2000">
...
</integrate>
```

11.5 Logical breakdown of XMDS2 Parsing Process

The following information is intended to assist developers in understanding the logical process undertaken by the XMDS2 system when parsing an `.xmds` file. The documentation was not designed to be exhaustive, but rather to help paint a picture of part of the way XMDS2 works.

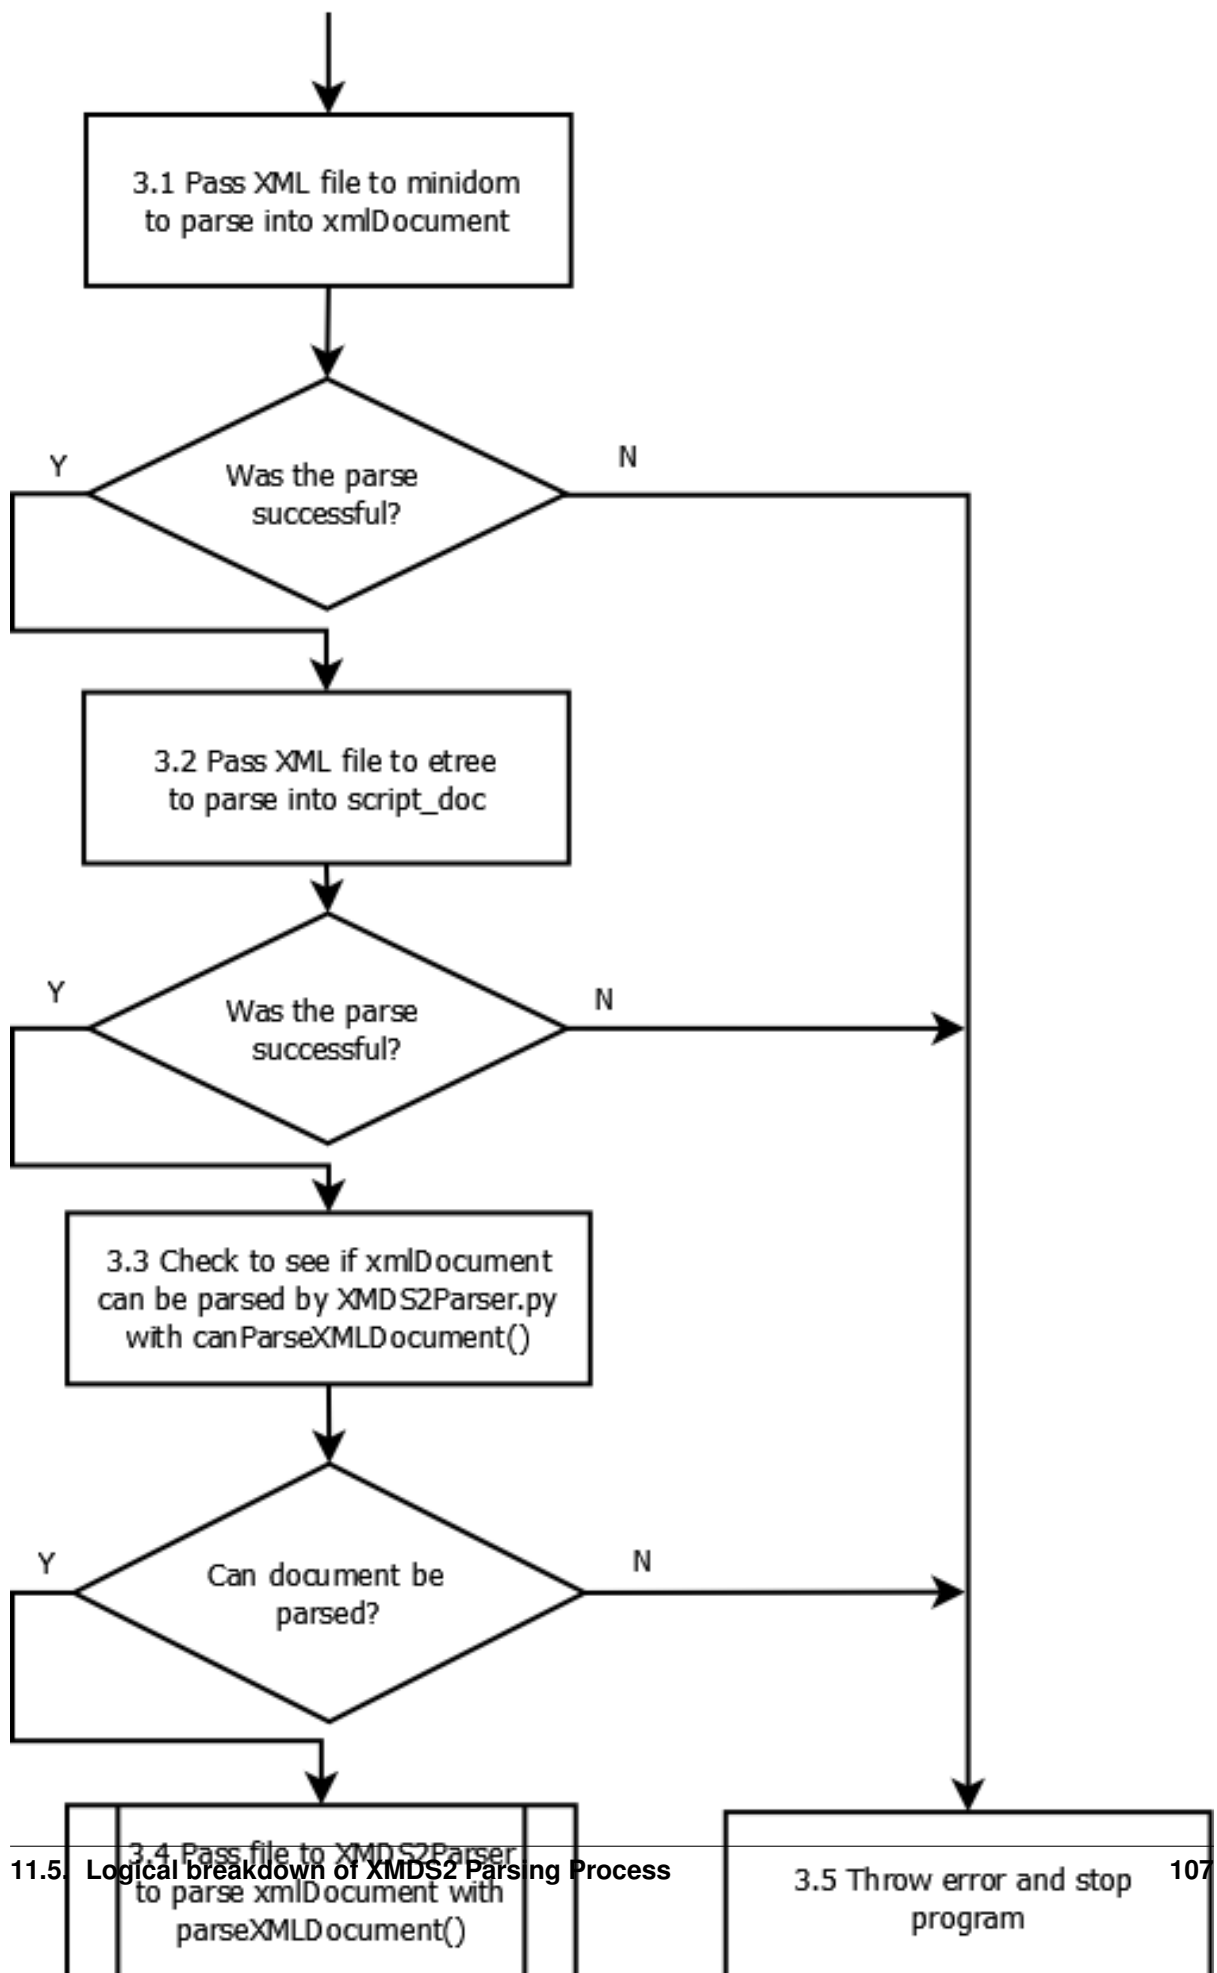
The flowcharts have been created in open source diagram drawing program Dia, and compiled into `.png` files which are displayed below. This page contains links to the original `.dia` files, so if you find any error in the information below (or you'd like to extend it, by adding in more information), please update the `.dia` files and commit them (and their compiled versions) to svn.

11.5.1 Overall process for parsing XML file in XMDs2



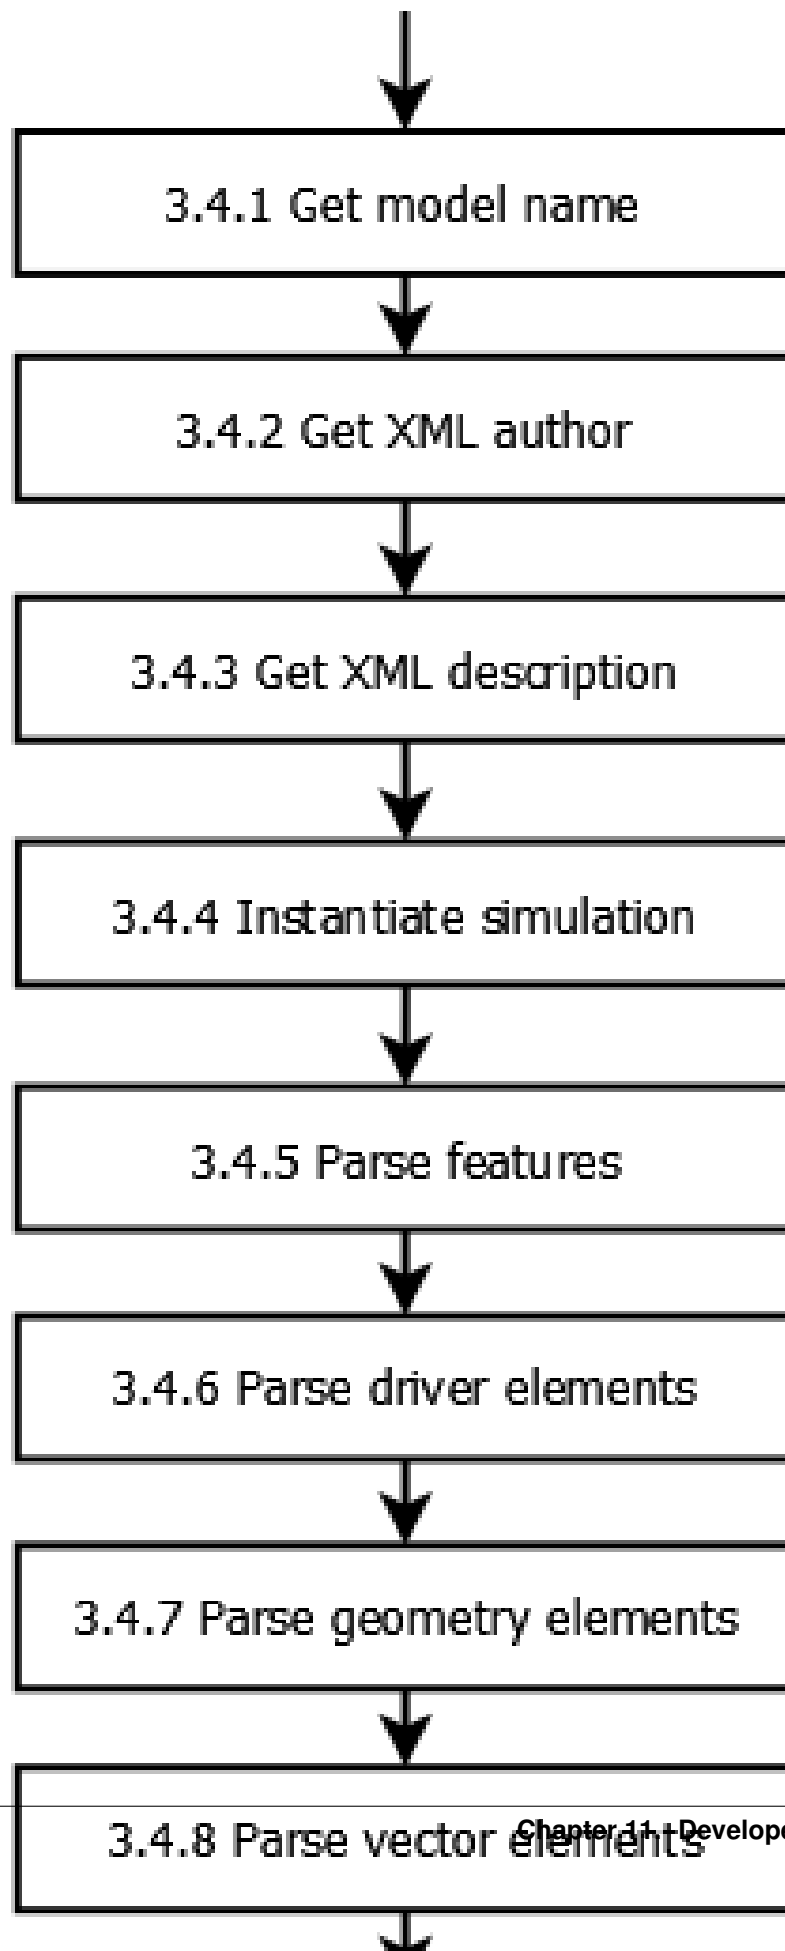
The original .dia file can be downloaded [here](#).

11.5.2 parser2.py parses XML file (Sub process 3)



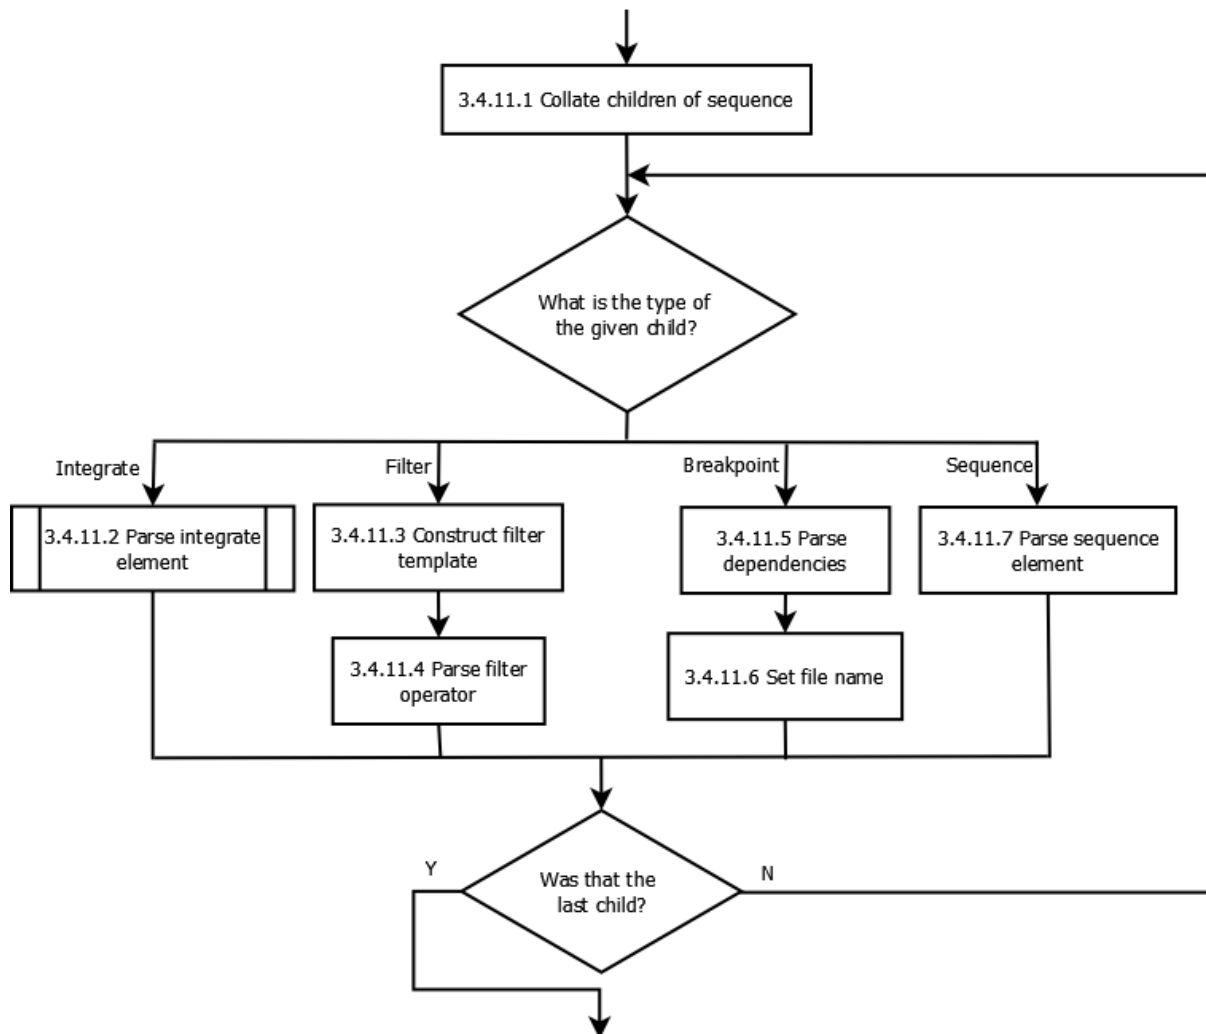
You can download the original dia file [here](#).

11.5.3 Pass file to XMDS2Parser to parse xmlDocument with parseXMLDocument() (Sub process 3.4)



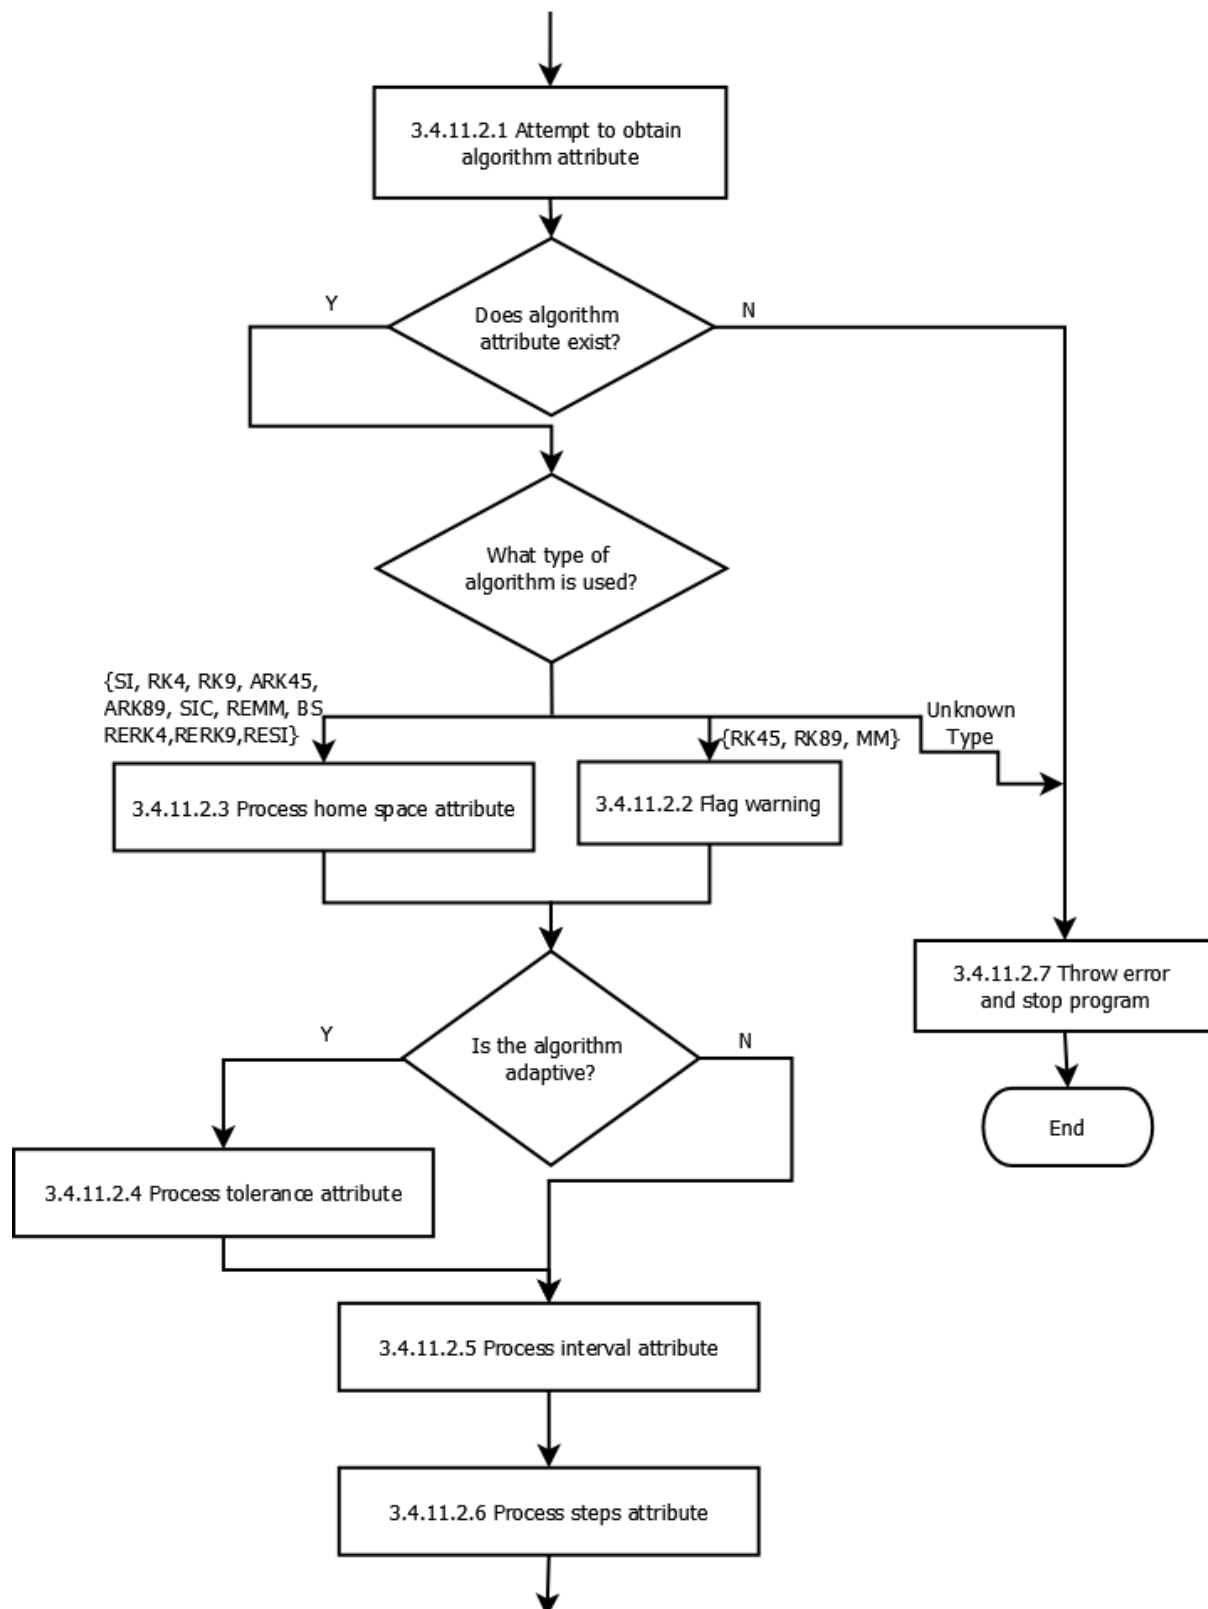
You can download the original dia file [here](#).

11.5.4 Parse Top Level Sequence elements (Sub process 3.4.11)



You can download the original dia file [here](#).

11.5.5 Parse Integrate Element (Sub process 3.4.11.2)



You can download the original dia file [here](#).

11.6 Directory layout

11.6.1 XMDs2's code and templates

All `.tmpl` files are Cheetah template files. These are used to generate C++ code. These templates are compiled as part of the XMDs2 build process to `.py` files of the same name. Do not edit the generated `.py` files, always edit the `.tmpl` files and regenerate the corresponding `.py` files with `make`.

- **xpdeint/:**

- **Features/:** Code for all `<feature>` elements, such as `<globals>` and `<auto_vectorise>`

- * **Transforms/:** Code for the Fourier and matrix-based transforms (including MPI variants).

- **Geometry/:** Code for describing the geometry of simulation dimensions and domains. Includes code for `Geometry`, `Field` and all `DimensionRepresentations`.

- **Operators/:** Code for all `<operator>` elements, including `IP`, `EX` and the temporal derivative operator `DeltaA`.

- **Segments/:** Code for all elements that can appear in a `<segments>` tag. This includes `<integrate>`, `<`

- * **Integrators:** Code for fixed and adaptive integration schemes, and all steppers (e.g. `RK4`, `RK45`, `RK9`, etc.)

- **Stochastic/:** Code for all random number generators and the random variables derived from them.

- * **Generators/:** Code for random number generators, includes `dSFMT`, `POSIX`, `Solirte`.

- * **RandomVariables/:** Code for the random variables derived from the random number generators. These are the gaussian, poissonian and uniform random variables.

- **SimulationDrivers/:** Code for all `<driver>` elements. In particular, this is where the location of MPI and multi-path code.

- **Vectors/:** Code for all `<vector>` elements, and their initialisation. This includes normal `<vector>` elements as well as `<computed_vector>` and `<noise_vector>` elements.

- **includes/:** C++ header and sources files used by the generated simulations.

- **support/:** Support files

- * **wscript:** waf build script for configuring and compiling generated simulations

- * **xpdeint.rnc:** Compact RelaxNG XML validation for XMDs scripts. This is the source file for the XML RelaxNG file `xpdeint.rng`

- * **xpdeint.rng:** RelaxNG XML validation for XMDs scripts. To regenerate this file from `xpdeint.rnc`, just run `make` in this directory.

- **waf/:** Our included version of the Python configuration and build tool `waf`.

- **waf_extensions/:** waf tool for compiling Cheetah templates.

- **xsil2graphics2/:** Templates for the output formats supported by `xsil2graphics2`.

- **wscript:** waf build script for XMDs2 itself.

- **CodeParser.py:** Minimally parses included C++ code for handling nonlocal dimension access, `IP/EX` operators and `IP` operator validation.

- **Configuration.py:** Manages configuration and building of generated simulations.

- `FriendlyPlusStyle.py`: Sphinx plug-in to improve formatting of XMDS scripts in user documentation.
- This directory also contains code for the input script parser, code blocks, code indentation, and the root `_ScriptElement` class.

11.6.2 Support files

- **admin/**: Documentation source, Linux installer and release scripts.
 - `developer-doc-source/`: source for epydoc python class documentation (generated from python code).
 - `userdoc-source/`: source for the user documentation (results visible at www.xmds.org and xmds2.readthedocs.org).
 - `xpdeint.tmbundle/`: TextMate support bundle for Cheetah templates and XMDS scripts
- `bin/`: Executable scripts to be installed as part of XMDS2 (includes `xmds2` and `xsil2graphics2`).
- `examples/`: Example XMDS2 input scripts demonstrating most of XMDS2's features.
- `testsuite/`: Testsuite of XMDS2 scripts. Run the testsuite by executing `./run_tests.py`

LICENSING

The XMDS2 codebase is licensed under the GPL version 2 license, which can be found in the COPYING file in the root directory of your XMDS install.

The XMDS2 documentation is licensed under the GNU Free Documentation License.

We encourage people to submit patches to us that extend XMDS2's functionality and fix bugs. If you do send us a patch, we do not require copyright assignment, but please include a statement saying you agree to license your code under the GPL v2 license, or under the GNU FDL license if your patch contributes to the documentation.

Both licenses can be found at <http://www.gnu.org/licenses/>, but we also include them below for convenience.

12.1 GNU General Public License

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - (a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - (b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - (c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - (a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - (b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - (c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distri-

bution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

12.2 GNU Free Documentation License

GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any

title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

1. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
2. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.

5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
10. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
11. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
12. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections

Entitled “Dedications”. You must delete all sections Entitled “Endorsements”. 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document. 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate. 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title. 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it. 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document. 11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing. **ADDENDUM:** How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (C) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with . . . Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and
with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

13.1 XMDS 2.2.1 “XMDS2 is a game of two halves” (September 30, 2014)

XMDS 2.2.1 contains minor bugfixes and updates. This includes documentation improvements, superior handling of external packages and more informative errors.

13.2 XMDS 2.2.0 “Out of cheese error” (January 13, 2014)

XMDS 2.2.0 contains a number of new features, as well as bugfixes and updates. Specifically

- Separated IP operators. This is a significant performance optimisation (~30%) for problems with two or more dimensions. It requires separating IP operators of the form “ $f(kx) + g(ky)$ ” (e.g. kinetic energy for quantum physics) into *two* IP operators and explicitly setting the dimensions=”x” and dimensions=”y” attributes on each. See *Optimisation hints* for details.
- Significant speed optimisations for adaptive integrators with IP operators (past IP operator calculations are re-used if the time-step hasn’t changed).
- The “constant” attribute for IP/EX operators is now unnecessary and considered advanced usage. If you don’t know whether to specify constant=”yes” or constant=”no”, don’t specify either.
- The xsil2graphics2 data exporter now supports Matlab, Octave, Mathematica and Python in all output formats, as well as R (HDF5 only). The Matlab/Octave scripts are now identical. A script generated for one will work for the other.
- Bessel-Neumann transforms have been implemented. Set transform=”bessel-neumann” if you want a Bessel (Hankel) transform but have zero derivative at the boundary (Neumann boundary conditions) instead of zero function value (Dirichlet boundary conditions). If you don’t care about your boundary condition, stick with the “bessel” transform.
- A Bulirsch-Stoer integrator. This can be useful for problems which are very smooth as you can use an arbitrarily high order algorithm. Specify algorithm=”RE” and extrapolations=”5” to have a 10th order integrator. Currently this is fixed-step only.
- “adaptive-mpi-multipath” driver. This implements a load scheduler that better spreads the work across different CPUs when different paths can take very different amounts of time. Also useful in heterogeneous clusters.
- XMDS2 is currently undergoing acceptance into Debian linux and will soon be able to be installed via the package manager. In the meantime you can find it in the private APT repository at <http://xmds.laboissiere.net>.
- A number of bug fixes.
- Expanded and improved documentation.

Many thanks to all who contributed to this release!

13.3 XMDs 2.1.4 “Well if this isn’t nice, I don’t know what is” (September 27, 2013)

The XMDs 2.1.4 update contains many new improvements and bugfixes:

- *xsil2graphics2* now supports all output formats for MATLAB, Octave and Python. The scripts generated for MATLAB/Octave are compatible with both.
- Fix a bug when *nonlocally* referencing a *dimension alias* with subsampling in *sampling_group* blocks or in some situations when MPI is used. This bug caused incorrect elements of the vector to be accessed.
- Correct the Fourier basis for dimensions using Hermite-Gauss transforms. Previously ‘kx’ was effectively behaving as ‘-kx’.
- Improve the performance of ‘nx’ <-> ‘kx’ Hermite-Gauss transforms.
- Stochastic error checking with runtime noise generation now works correctly. Previously different random numbers were generated for the full-step paths and the half-step paths.
- Documentation updates.

13.4 XMDs 2.1.3 “Happy Mollusc” (June 7, 2013)

The XMDs 2.1.3 update is a bugfix release that includes the following improvements:

- XMDs will work when MPI isn’t installed (but only for single-process simulations).
- Support for GCC 4.8
- The number of paths used by the multi-path driver can now be specified at run-time (using *<validation kind=“run-time”>*)
- Other bug fixes

13.5 XMDs 2.1.2 “Happy Mollusc” (October 15, 2012)

The XMDs 2.1.2 update has many improvements:

- Named filters. You can now specify a name for a filter block and call it like a function if you need to execute it conditionally. See the documentation for the *<filter>* block for more information.
- New *chunked_output* feature. XMDs can now output simulation results as it goes, reducing the memory requirement for simulations that generate significant amounts of output. See the documentation for more details.
- Improved OpenMP support
- The EX operator is now faster in the common case (but you should still prefer IP when possible)
- If seeds are not provided for a *noise_vector*, they are now generated at simulation run-time, so different executions will give different results. The generated noises can still be found in the output .xsil files enabling results to be reproduced if desired.
- Advanced feature: Dimensions can be accessed non-locally with the index of the lattice point. This removes the need in hacks to manually access XMDs’s underlying C arrays. This is an advanced feature and requires a little knowledge of XMDs’s internal grid representation. See the advanced topics documentation for further details.
- Fixed adaptive integrator order when noises were used in vector initialisation
- Fix the Spherical Bessel basis. There were errors in the definition of this basis which made it previously unreliable.

- Other bug fixes

13.6 XMDs 2.1 “Happy Mollusc” (June 14, 2012)

XMDs 2.1 is a significant upgrade with many improvements and bug fixes since 2.0. We also now have installers for Linux and Mac OS X, so you no longer have to build XMDs from source! See [here](#) for details about the installers.

Existing users should note that this release introduces a more concise syntax for moment groups. You can now use:

```
<sampling_group initial_sample="yes" basis="x y z">
...
</sampling_group>
```

Instead of:

```
<group>
  <sampling initial_sample="yes" basis="x y z">
    ...
  </sampling>
</group>
```

Another syntax change is that the initial basis of a vector should be specified with *initial_basis* instead of *initial_space*.

In both cases, although the old syntax is not described in the documentation, it is still supported, so existing scripts will work without any changes.

Other changes in XMDs 2.1 include:

- The *lattice* attribute for dimensions can now be specified at run-time. Previously only the minimum and maximum values of the domain could be specified at run-time. See [here](#) for details.
- *noise_vectors* can now be used in non-uniform dimensions (e.g. dimensions using the Bessel transform for cylindrical symmetry).
- “loose” *geometry_matching_mode* for HDF5 vector initialisation. This enables extending the simulation grid from one simulation to the next, or coarsening or refining a grid when importing.
- *vectors* can now be initialised by integrating over dimensions of other vectors. *computed_vectors* always supported this, now *vectors* do too.
- Update to latest version of waf, which is used for compiling simulations and detecting FFTW, HDF5, etc. This should lead to fewer waf-related problems.
- Bug fixes.

13.7 XMDs 2.0 “Shiny!” (September 13, 2010)

XMDs 2.0 is a major upgrade which has been rewritten from the ground up to make it easier for us to apply new features. And there are many. XMDs 2.0 is faster and far more versatile than previous versions, allowing the efficient integration of almost any initial value problem on regular domains.

The feature list includes:

- Quantities of different dimensionalities. So you can have a 1D potential and a 3D wavefunction.
- Integrate more than one vector (in more than one geometry), so you can now simultaneously integrate a PDE and a coupled ODE (or coupled PDEs of different dimensions).
- Non-Fourier transformations including the Bessel basis, Spherical Bessel basis and the Hermite-Gauss (harmonic oscillator) basis.

- The ability to have more than one kind of noise (gaussian, poissonian, etc) in a simulation.
- Integer-valued dimensions with non-local access. You can have an array of variables and access different elements of that array.
- Significantly better error reporting. When errors are found when compiling the script they will almost always be reported with the corresponding line of your script, instead of the generated source.
- *IP/EX* operators are separate from the integration algorithm, so you can have both *IP* and *EX* operators in a single integrate block. Also, *EX* operators can act on arbitrary code, not just vector components. (e.g. *L[phi*phi]*).
- Cross propagation in the increasing direction of a given dimension or in the decreasing dimension. And you can have more than one cross-propagator in a given integrator (going in different directions or dimensions).
- Faster Gaussian noises.
- The ability to calculate spatial correlation functions.
- OpenMP support.
- MPI support.
- Output moment groups use less memory when there isn't a *post_processing* element.
- Generated source is indented correctly.
- An *xmids1*-like script file format.
- *xmids1*-like generated source.
- All of the integrators from *xmids1* (*SI*, *RK4*, *ARK45*, *RK9*, *ARK89*).

Symbols

`_SAMPLE_COMPLEX`, 31
`_dx`, 73
`_lattice`, 73
`_max`, 73
`_min`, 73

A

Adding an integrator, 99
 Advanced topics, 76
 Aliases, 41, 53, 54

B

Boundary conditions, 57, 69
 Boundary conditions (Bessel), 57
 Boundary conditions (DCT), 55
 Boundary conditions (DFT), 54
 Boundary conditions (DST), 56
 Bug reports, 83
 Bulirsch-Stoer algorithm, 74
 Bulirsch-Stoer error scaling, 75

C

Citing XMDS2, 83
 Command line arguments, 30
 Computed vectors, 34
 Convolutions and Fourier transforms, 79
 Cross-propagation, 68

D

Dimension aliases, 82
 Distributed MPI, 30
 Documentation format, 96
 Double precision, 51
 Driver examples
 distributed-mpi, 30
 mpi-multi-path, 24
 multi-path, 20
 Drivers
 adaptive-mpi-multi-pathhyperpage, 53
 distributed-mpihyperpage, 53
 mpi-multi-pathhyperpage, 53
 multi-pathhyperpage, 53

E

EX operators, 68

Examples, 14
 Exporting data, 11, 91

F

FAQ, 82
 Fibre noise, 21

G

Geometry matching mode - loose, 81
 GNU Free Documentation License, 118
 GNU General Public License, 115
 GPL, 115
 Grid specification at run-time, 84

I

Importing data, 77
 Installers - linux, 3
 Installers - Max OS X, 4
 Installing via source, 4
 Installing XMDS2, 2
 Integer dimensions, 25
 Integration algorithms
 adaptive Runge-Kutta (ARK), 70
 all, 66
 Richardson extrapolation, 71
 Runge-Kutta (RK), 70
 semi-implicit (SI), 69
 IP operators, 68
 IP versus EX operators, 84

K

Kubo oscillator, 19

L

Licensing, 114

M

Modified midpoint method, 73
 MPI, 30
 non-local vector access, 60
 Multi-path MPI, 24

N

News, 123
 Noise
 gaussian, 64

- jump, 65
- poissonian, 65
- uniform, 63
- wiener, 64

Noise vector example, 20

Nonlinear Schroedinger equation, 15

O

Operators, 67

- EX, 68
- IP, 68

P

Precision, 51

R

Reference, 46

S

Sampling, 72

Script elements, 46

Single precision, 51

T

Test scripts, 95

Time dependence in vectors, 83

Transforms

- bessel, 56
- dct, 55
- dft, 54
- dst, 56
- hermite-gauss, 58
- spherical-bessel, 58

Tutorial, 7

V

Vectors

- non-local access, 53, 60, 67

X

XMDs2 parsing process, 102

XMDs2 project directory layout, 112

XMDs2 XML schema, 44

XML element attributes

- algorithm, 66
- aliases, 53
- basis (dependencies element), 61
- basis (sampling_group element), 72, 73
- constant, 68
- cross-propagation, 68
- cycles, 65
- default_value, 49
- dimensions (computed vector), 62
- dimensions (noise vector), 63
- domain, 53
- filename (breakpoint element), 72
- filename (output element), 72

- format (breakpoint element), 72
- format (output element), 72
- geometry_matching_mode, 60
- initial_basis (noise vector), 63
- initial_basis (vector), 60
- initial_sample, 72
- interval, 66
- kind, 52
- kind (boundary_condition element), 69
- kind (initialisation element), 60
- kind (noise vector), 63
- kind (operator), 68
- lattice, 53
- length_scale, 59
- method, 63
- name, 49
- name (computed vector), 62
- name (dimension), 53
- name (driver), 53
- name (filter), 66
- name (noise vector), 63
- name (vector), 59, 60
- plan, 51
- seed, 63
- steps, 66
- tolerance, 66
- transform, 54
- type, 49
- type (computed vector), 62
- type (dimension), 53
- type (noise vector), 63
- type (vector), 59
- volume_prefactor, 54
- where, 71

XML Elements, 46

XML elements

- argument, 49
- arguments, 48
- author, 47
- autovectorise, 49
- benchmark, 49
- bing, 49
- boundary_conditions, 69
- breakpoint, 72
- cflags, 49
- chunked_output, 50
- components, 59
- computed_vector, 62
- dependencies, 61
- description, 47
- diagnostics, 50
- dimension, 53
- driver, 53
- error_check, 50
- evaluation, 62
- features, 47
- fftw, 51
- filename, 60

- filter, 66
- filters, 71
- geometry, 53
- globals, 51
- halt_non_finite, 50
- initialisation, 60
- integrate, 66
- integration_vectors, 67
- name (of simulation), 47
- noise_vector, 63
- openmp, 51
- operator, 68
- operators, 67
- output, 72
- precision, 51
- propagation_dimension, 53
- samples, 66
- sampling_group, 72
- sequence, 65
- simulation, 47
- transverse_dimensions, 53
- validation, 52
- vector, 59

XML schema, 44

XML schema updates, 96

xsil2graphics2, 91