

---

# **Shenfun Documentation**

*Release 2.1.0*

**Mikael Mortensen**

**Nov 14, 2019**



# CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Spectral Galerkin . . . . .	1
1.2	Tensor products . . . . .	2
1.3	Tribute . . . . .	3
<b>2</b>	<b>Getting started</b>	<b>5</b>
2.1	Basic usage . . . . .	5
2.2	Operators . . . . .	8
2.3	Multidimensional problems . . . . .	8
2.4	Coupled problems . . . . .	10
2.5	Integrators . . . . .	12
2.6	MPI . . . . .	14
<b>3</b>	<b>Post processing</b>	<b>19</b>
3.1	ParaView . . . . .	21
<b>4</b>	<b>Installation</b>	<b>23</b>
4.1	Optimization . . . . .	24
4.2	Additional dependencies . . . . .	24
4.3	Test installation . . . . .	24
<b>5</b>	<b>Demos</b>	<b>27</b>
5.1	Demo - 1D Poisson's equation . . . . .	27
5.1.1	Model problem . . . . .	28
5.1.2	Implementation . . . . .	30
5.2	Demo - Cubic nonlinear Klein-Gordon equation . . . . .	32
5.2.1	The nonlinear Klein-Gordon equation . . . . .	32
5.2.2	Implementation . . . . .	35
5.3	Demo - 3D Poisson's equation . . . . .	40
5.3.1	Model problem . . . . .	40
5.3.2	Implementation . . . . .	43
5.4	Demo - Kuramoto-Sivashinsky equation . . . . .	49
5.4.1	The Kuramoto-Sivashinsky equation . . . . .	49
5.4.2	Implementation . . . . .	51
5.5	Demo - Stokes equations . . . . .	53
5.5.1	Model problem . . . . .	53
5.5.2	Implementation . . . . .	56
	<b>Bibliography</b>	<b>61</b>



## INTRODUCTION

### 1.1 Spectral Galerkin

The spectral Galerkin method solves partial differential equations through a special form of the [method of weighted residuals](#) (WRM). As a Galerkin method it is very similar to the [finite element method](#) (FEM). The most distinguishable feature is that it uses global shape functions, where FEM uses local. This feature leads to highly accurate results with very few shape functions, but the downside is much less flexibility when it comes to computational domain than FEM.

Consider the Poisson equation with a right hand side function  $f(\mathbf{x})$

$$-\nabla^2 u(\mathbf{x}) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega. \quad (1.1)$$

To solve this equation, we will eventually need to supplement appropriate boundary conditions. However, for now just assume that any valid boundary conditions (Dirichlet, Neumann, periodic).

With the method of weighted residuals we attempt to find  $u(\mathbf{x})$  using an approximation,  $u_N$ , to the solution

$$u(\mathbf{x}) \approx u_N(\mathbf{x}) = \sum_{k=0}^{N-1} \hat{u}_k \phi_k(\mathbf{x}). \quad (1.2)$$

Here the  $N$  expansion coefficients  $\hat{u}_k$  are unknown and  $\{\phi_k\}_{k \in \mathcal{I}^N}$ ,  $\mathcal{I}^N = 0, 1, \dots, N-1$  are *trial* functions. Inserting for  $u_N$  in Eq. (1.1) we get a residual

$$R_N(\mathbf{x}) = \nabla^2 u_N(\mathbf{x}) + f(\mathbf{x}) \neq 0. \quad (1.3)$$

With the WRM we now force this residual to zero in an average sense using *test* function  $v(\mathbf{x})$  and *weight* function  $w(\mathbf{x})$

$$(R_N, v)_w := \int_{\Omega} R_N(\mathbf{x}) \bar{v}(\mathbf{x}) w(\mathbf{x}) d\mathbf{x} = 0, \quad (1.4)$$

where  $\bar{v}$  is the complex conjugate of  $v$ . If we now choose the test functions from the same space as the trial functions, i.e.,  $V_N = \text{span}\{\phi_k\}_{k \in \mathcal{I}^N}$ , then the WRM becomes the Galerkin method, and we get  $N$  equations for  $N$  unknowns  $\{\hat{u}_k\}_{k \in \mathcal{I}^N}$

$$\sum_{j \in \mathcal{I}^N} \underbrace{(-\nabla^2 \phi_j, \phi_k)_w}_{A_{kj}} \hat{u}_j = (f, \phi_k)_w, \quad \text{for } k \in \mathcal{I}^N. \quad (1.5)$$

Note that this is a regular linear system of algebra equations

$$A_{kj} \hat{u}_j = \tilde{f}_k,$$

where the matrix  $A \in \mathbb{R}^{N \times N}$ .

The choice of basis functions  $v(\mathbf{x})$  is highly central to the method. For the Galerkin method to be *spectral*, the basis is usually chosen as linear combinations of Chebyshev, Legendre, Laguerre, Hermite, Jacobi or trigonometric functions. In one spatial dimension typical choices for  $\phi_k$  are

$$\begin{aligned}\phi_k(x) &= T_k(x) \\ \phi_k(x) &= T_k(x) - T_{k+2}(x) \\ \phi_k(x) &= L_k(x) \\ \phi_k(x) &= L_k(x) - L_{k+2}(x) \\ \phi_k(x) &= \exp(ikx)\end{aligned}$$

where  $T_k, L_k$  are the  $k$ 'th Chebyshev polynomial of the first kind and the  $k$ 'th Legendre polynomial, respectively. Note that the second and fourth functions above satisfy the homogeneous Dirichlet boundary conditions  $\phi_k(\pm 1) = 0$ , and as such these basis functions may be used to solve the Poisson equation (1.1) with homogeneous Dirichlet boundary conditions. Similarly, two basis functions that satisfy homogeneous Neumann boundary condition  $u'(\pm 1) = 0$  are

$$\begin{aligned}\phi_k &= T_k - \left(\frac{k}{k+2}\right)^2 T_{k+2} \\ \phi_k &= L_k - \frac{k(k+1)}{(k+2)(k+3)} L_{k+2}\end{aligned}$$

Shenfun contains classes for working with several such bases, to be used for different equations and boundary conditions.

Complete demonstration programs that solves the Poisson equation (1.1), and some other problems can be found by following these links

- [Demo - 1D Poisson's equation](#)
- [Demo - 3D Poisson's equation](#)
- [Demo - Cubic nonlinear Klein-Gordon equation](#)
- [Demo - Kuramoto-Sivashinsky equation](#)
- [Demo - Stokes equations](#)
- [Demo - Lid driven cavity](#)

## 1.2 Tensor products

If the problem is two-dimensional, then we need two basis functions, one per dimension. If we call the basis function along  $x$ -direction  $\mathcal{X}(x)$  and along  $y$ -direction  $\mathcal{Y}(y)$ , a test function can then be computed as

$$v(x, y) = \mathcal{X}(x)\mathcal{Y}(y).$$

If we now have a problem that has Dirichlet boundaries in the  $x$ -direction and periodic boundaries in the  $y$ -direction, then we can choose  $\mathcal{X}_k(x) = T_k - T_{k+2}$ , for  $k \in \mathcal{I}^{N-2}$  (with  $N - 2$  because  $T_{k+2}$  then equals  $T_N$  for  $k = N - 2$ ),  $\mathcal{Y}_l(y) = \exp(ily)$  for  $l \in \mathcal{I}^M$  and a tensor product test function is then

$$v_{kl}(x, y) = (T_k(x) - T_{k+2}(x)) \exp(ily), \text{ for } (k, l) \in \mathcal{I}^{N-2} \times \mathcal{I}^M. \quad (1.6)$$

In other words, we choose one test function per spatial dimension and create global basis functions by taking the outer products (or tensor products) of these individual test functions. Since global basis functions simply are the tensor products of one-dimensional basis functions, it is trivial to move to even higher-dimensional spaces. The multi-dimensional basis functions then form a basis for a multi-dimensional tensor product space. The associated

domains are similarly formed by taking Cartesian products of the one-dimensional domains. For example, if the one-dimensional domains in  $x$ - and  $y$ -directions are  $[-1, 1]$  and  $[0, 2\pi]$ , then the two-dimensional domain formed from these two are  $[-1, 1] \times [0, 2\pi]$ , where  $\times$  represents a Cartesian product.

The one-dimensional domains are discretized using the quadrature points of the chosen basis functions. If the meshes in  $x$ - and  $y$ -directions are  $x = \{x_i\}_{i \in \mathcal{I}^N}$  and  $y = \{y_j\}_{j \in \mathcal{I}^M}$ , then a Cartesian product mesh is  $x \times y$ . With index and set builder notation it is given as

$$x \times y = \{(x_i, y_j) \mid (i, j) \in \mathcal{I}^N \times \mathcal{I}^M\}. \quad (1.7)$$

With shenfun a user chooses the appropriate bases for each dimension of the problem, and may then combine these bases into tensor product spaces and Cartesian product domains. For example, to create the required spaces for the aforementioned domain, with Dirichlet in  $x$ - and periodic in  $y$ -direction, we need the following:

$$\begin{aligned} N, M &= (16, 16) \\ B^N(x) &= \text{span}\{T_k(x) - T_{k+2}(x)\}_{k \in \mathcal{I}^{N-2}} \\ B^M(y) &= \text{span}\{\exp(il y)\}_{l \in \mathcal{I}^M} \\ V(x, y) &= B^N(x) \otimes B^M(y) \end{aligned}$$

where  $\otimes$  represents a tensor product.

This can be implemented in *shenfun* as follows:

```
from shenfun import Basis, TensorProductSpace
from mpi4py import MPI
comm = MPI.COMM_WORLD
N, M = (16, 16)
BN = Basis(N, 'Chebyshev', bc=(0, 0))
BM = Basis(M, 'Fourier', dtype='d')
V = TensorProductSpace(comm, (BN, BM))
```

Note that the Chebyshev basis is created using  $N$  and not  $N - 2$ . The chosen boundary condition `bc=(0, 0)` ensures that only  $N - 2$  bases will be used. The Fourier basis `BM` has been defined for real inputs to a forward transform, which is ensured by the `dtype` keyword being set to `d` for double. `dtype` specifies the data type that is input to the `forward` method, or the data type of the solution in physical space. Setting `dtype='D'` indicates that this datatype will be complex. Note that it will not trigger an error, or even lead to wrong results, if `dtype` is by mistake set to `D`. It is merely less efficient to work with complex data arrays where double precision is sufficient. See Sec [Getting started](#) for more information on getting started with using bases.

Shenfun is parallelized with MPI through the `mpi4py-fft` package. If we store the current example in `filename.py`, then it can be run with more than one processor, e.g., like:

```
mpirun -np 4 python filename.py
```

In this case the tensor product space  $V$  will be distributed with the *slab* method (since the problem is 2D) and it can here use a maximum of 9 CPUs. The maximum is 9 since the last dimension is transformed from 16 real numbers to 9 complex, using the Hermitian symmetry of real transforms, i.e., the shape of a transformed array in the  $V$  space will be (14, 9). You can read more about MPI in the later section [MPI](#).

## 1.3 Tribute

Shenfun is named as a tribute to Prof. Jie Shen, as it contains many tools for working with his modified Chebyshev and Legendre bases, as described here:

- Jie Shen, SIAM Journal on Scientific Computing, 15 (6), 1489-1505 (1994) (JS1)

- Jie Shen, SIAM Journal on Scientific Computing, 16 (1), 74-87, (1995) (JS2)

Shenfun has implemented classes for the bases described in these papers, and within each class there are methods for fast transforms, inner products and for computing matrices arising from bilinear forms in the spectral Galerkin method.

## GETTING STARTED

### 2.1 Basic usage

Shenfun consists of classes and functions whose purpose are to make it easier to implement PDE's with spectral methods in simple tensor product domains. The most important everyday tools are

- `TensorProductSpace`
- `MixedTensorProductSpace`
- `TrialFunction`
- `TestFunction`
- `Function`
- `Array`
- `inner()`
- `div()`
- `grad()`
- `project()`
- `Basis()`

A good place to get started is by creating a `Basis()`. There are five families of bases: Fourier, Chebyshev, Legendre, Laguerre, Hermite and Jacobi. All bases are defined on a one-dimensional domain, with their own basis functions and quadrature points. For example, we have the regular Chebyshev basis  $\{T_k\}_{k=0}^{N-1}$ , where  $T_k$  is the  $k$ 'th Chebyshev polynomial of the first kind. To create such a basis with 8 quadrature points (i.e.,  $\{T_k\}_{k=0}^7$ ) do:

```
from shenfun import Basis
N = 8
T = Basis(N, 'Chebyshev', bc=None)
```

Here `bc=None` is used to indicate that there are no boundary conditions associated with this basis, which is the default, so it could just as well have been left out. To create a regular Legendre basis (i.e.,  $\{L_k\}_{k=0}^{N-1}$ , where  $L_k$  is the  $k$ 'th Legendre polynomial), just replace `Chebyshev` with `Legendre` above. And to create a Fourier basis, just use `Fourier`.

The basis  $T = \{T_k\}_{k=0}^{N-1}$  has many useful methods associated with it, and we may experiment a little. A `Function`  $u$  using basis  $T$  has expansion

$$u(x) = \sum_{k=0}^7 \hat{u}_k T_k(x) \quad (2.1)$$

and an instance of this function (initialized with  $\{\hat{u}_k\}_{k=0}^7 = 0$ ) is created in shenfun as:

```
u = Function(T)
```

Consider now for example the polynomial  $2x^2 - 1$ , which happens to be exactly equal to  $T_2(x)$ . We can create this polynomial using `sympy`

```
import sympy as sp
x = sp.Symbol('x')
u = 2*x**2 - 1 # or simply u = sp.chebyshevt(2, x)
```

The `Sympy` function `u` can now be evaluated on the quadrature points of basis `T`:

```
xj = T.mesh()
ue = Array(T)
ue[:] = [u.subs(x, xx) for xx in xj]
print(xj)
[ 0.98078528  0.83146961  0.55557023  0.19509032 -0.19509032 -0.55557023
 -0.83146961 -0.98078528]
print(ue)
[ 0.92387953  0.38268343 -0.38268343 -0.92387953 -0.92387953 -0.38268343
 0.38268343  0.92387953]
```

We see that `ue` is an `Array` on the basis `T`, and not a `Function`. The `Array` and `Function` classes are both subclasses of Numpy's `ndarray`, and represent the two arrays associated with the spectral Galerkin function, like (2.1). The `Function` represent the entire spectral Galerkin function, with array values corresponding to the expansion coefficients  $\hat{u}$ . The `Array` represent the spectral Galerkin function evaluated on the quadrature mesh of the basis `T`, i.e., here  $u(x_i), \forall i \in 0, 1, \dots, 7$ .

We now want to find the `Function` `uh` corresponding to `Array` `ue`. Considering (2.1), this corresponds to finding  $\hat{u}_k$  if the left hand side  $u(x_j)$  is known for all quadrature points  $x_j$ .

Since we already know that `ue` is equal to the second Chebyshev polynomial, we should get an array of expansion coefficients equal to  $\hat{u} = (0, 0, 1, 0, 0, 0, 0, 0)$ . We can compute `uh` either by using `project()` or a forward transform:

```
uh = Function(T)
uh = T.forward(ue, uh)
# or
# uh = ue.forward(uh)
# or
# uh = project(ue, T)
print(uh)
[-1.38777878e-17  6.72002101e-17  1.00000000e+00 -1.95146303e-16
 1.96261557e-17  1.15426347e-16 -1.11022302e-16  1.65163507e-16]
```

So we see that the projection works to machine precision.

The projection is mathematically: find  $u_h \in T$ , such that

$$(u_h - u, v)_w = 0 \quad \forall v \in T,$$

where  $v$  is a test function,  $u_h$  is a trial function and the notation  $(\cdot, \cdot)_w$  was introduced in (1.4). Using now  $v = T_k$  and  $u_h = \sum_{j=0}^7 \hat{u}_j T_j$ , we get

$$\left(\sum_{j=0}^7 \hat{u}_j T_j, T_k\right)_w = (u, T_k)_w,$$

$$\sum_{j=0}^7 (T_j, T_k)_w \hat{u}_j = (u, T_k)_w,$$

for all  $k \in 0, 1, \dots, 7$ . This can be rewritten on matrix form as

$$B_{kj} \hat{u}_j = \tilde{u}_k,$$

where  $B_{kj} = (T_j, T_k)_w$ ,  $\tilde{u}_k = (u, T_k)_w$  and summation is implied by the repeating  $j$  indices. Since the Chebyshev polynomials are orthogonal the mass matrix  $B_{kj}$  is diagonal. We can assemble both  $B_{kj}$  and  $\tilde{u}_j$  with shenfun, and at the same time introduce the `TestFunction`, `TrialFunction` classes and the `inner()` function:

```
from shenfun import TestFunction, TrialFunction, inner
u = TrialFunction(T)
v = TestFunction(T)
B = inner(u, v)
u_tilde = inner(u, v)
print(B)
{0: array([3.14159265, 1.57079633, 1.57079633, 1.57079633, 1.57079633,
          1.57079633, 1.57079633, 1.57079633])}
print(u_tilde)
[-4.35983562e-17  1.05557843e-16  1.57079633e+00 -3.06535096e-16
  3.08286933e-17  1.81311282e-16 -1.74393425e-16  2.59438230e-16]
```

The `inner()` function represents the (weighted) inner product and it expects one test function, and possibly one trial function. If, as here, it also contains a trial function, then a matrix is returned. If `inner()` contains one test, but no trial function, then an array is returned. Finally, if `inner()` contains no test nor trial function, but instead a number and an `Array`, like:

```
a = Array(T, val=1)
print(inner(1, a))
2.0
```

then `inner()` represents a non-weighted integral over the domain. Here it returns the length of the domain (2.0) since `a` is initialized to unity.

Note that the matrix  $B$  assembled above is stored using shenfun's `SpectralMatrix` class, which is a subclass of Python's dictionary, where the keys are the diagonals and the values are the diagonal entries. The matrix  $B$  is seen to have only one diagonal (the principal)  $\{B_{ii}\}_{i=0}^7$ .

With the matrix comes a `solve` method and we can solve for  $\hat{u}$  through:

```
u_hat = Function(T)
u_hat = B.solve(u_tilde, u=u_hat)
print(u_hat)
[-1.38777878e-17  6.72002101e-17  1.00000000e+00 -1.95146303e-16
  1.96261557e-17  1.15426347e-16 -1.11022302e-16  1.65163507e-16]
```

which obviously is exactly the same as we found using `project()` or the `T.forward` function.

Note that `Array` merely is a subclass of Numpy's `ndarray`, whereas `Function` is a subclass of both Numpy's `ndarray` and the `BasisFunction` class. The latter is used as a base class for arguments to bilinear and linear forms, and is as such a base class also for `TrialFunction` and `TestFunction`. An instance of the `Array` class cannot be used in forms, except from regular inner products of numbers or test function vs an `Array`. To illustrate, lets create some forms, where all except the last one is ok:

```
T = Basis(12, 'Legendre')
u = TrialFunction(T)
v = TestFunction(T)
uf = Function(T)
ua = Array(T)
A = inner(v, u) # Mass matrix
```

(continues on next page)

(continued from previous page)

```

c = inner(v, ua) # ok, a scalar product
d = inner(v, uf) # ok, a scalar product (slower than above)
e = inner(1, ua) # ok, non-weighted integral of ua over domain
df = Dx(uf, 0, 1) # ok
da = Dx(ua, 0, 1) # Not ok

AssertionError                                Traceback (most recent call last)
<ipython-input-14-3b957937279f> in <module>
----> 1 da = inner(v, Dx(ua, 0, 1))

~/MySoftware/shenfun/shenfun/forms/operators.py in Dx(test, x, k)
     82     Number of derivatives
     83     """
----> 84     assert isinstance(test, (Expr, BasisFunction))
     85
     86     if isinstance(test, BasisFunction):

AssertionError:

```

So it is not possible to perform operations that involve differentiation ( $Dx$  represents a partial derivative) on an `Array` instance. This is because the `ua` does not contain more information than its values and its `TensorProductSpace`. A `BasisFunction` instance, on the other hand, can be manipulated with operators like `div()` `grad()` in creating instances of the `Expr` class, see *Operators*.

Note that any rules for efficient use of Numpy `ndarrays`, like vectorization, also applies to `Function` and `Array` instances.

## 2.2 Operators

Operators act on any single instance of a `BasisFunction`, which can be `Function`, `TrialFunction` or `TestFunction`. The implemented operators are:

- `div()`
- `grad()`
- `curl()`
- `Dx()`

Operators are used in variational forms assembled using `inner()` or `project()`, like:

```
A = inner(grad(u), grad(v))
```

which assembles a stiffness matrix `A`. Note that the two expressions fed to `inner` must have consistent rank. Here, for example, both `grad(u)` and `grad(v)` have rank 1 of a vector.

## 2.3 Multidimensional problems

As described in the introduction, a multidimensional problem is handled using tensor product spaces, that have basis functions generated from taking the outer products of one-dimensional basis functions. We create tensor product spaces using the class `TensorProductSpace`:

```
N, M = (12, 16)
C0 = Basis(N, 'L', bc=(0, 0), scaled=True)
K0 = Basis(M, 'F', dtype='d')
T = TensorProductSpace(comm, (C0, K0))
```

Associated with this is a Cartesian mesh  $[-1, 1] \times [0, 2\pi]$ . We use classes `Function`, `TrialFunction` and `TestFunction` exactly as before:

```
u = TrialFunction(T)
v = TestFunction(T)
A = inner(grad(u), grad(v))
```

However, now `A` will be a tensor product matrix, or more correctly, the sum of two tensor product matrices. This can be seen if we look at the equations beyond the code. In this case we are using a composite Legendre basis for the first direction and Fourier exponentials for the second, and the tensor product basis function is

$$\begin{aligned} v_{kl}(x, y) &= \frac{1}{\sqrt{4k+6}}(L_k(x) - L_{k+2}(x)) \exp(iy), \\ &= \Psi_k(x)\phi_l(y), \end{aligned}$$

where  $L_k$  is the  $k$ 'th Legendre polynomial,  $\psi_k = (L_k - L_{k+2})/\sqrt{4k+6}$  and  $\phi_l = \exp(iy)$  are used for simplicity in later derivations. The trial function becomes

$$u(x, y) = \sum_k \sum_l \hat{u}_{kl} v_{kl}$$

and the inner product is

$$\begin{aligned} (\nabla u, \nabla v)_w &= \int_{-1}^1 \int_0^{2\pi} \nabla u \cdot \nabla v dx dy, \\ &= \int_{-1}^1 \int_0^{2\pi} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y} \frac{\partial v}{\partial y} dx dy, \\ &= \int_{-1}^1 \int_0^{2\pi} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} dx dy + \int_{-1}^1 \int_0^{2\pi} \frac{\partial u}{\partial y} \frac{\partial v}{\partial y} dx dy, \end{aligned} \tag{2.2}$$

showing that it is the sum of two tensor product matrices. However, each one of these two terms contains the outer product of smaller matrices. To see this we need to insert for the trial and test functions (using  $v_{mn}$  for test):

$$\begin{aligned} \int_{-1}^1 \int_0^{2\pi} \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} dx dy &= \int_{-1}^1 \int_0^{2\pi} \frac{\partial}{\partial x} \left( \sum_k \sum_l \hat{u}_{kl} \Psi_k(x) \phi_l(y) \right) \frac{\partial}{\partial x} (\Psi_m(x) \phi_n(y)) dx dy, \\ &= \sum_k \sum_l \underbrace{\int_{-1}^1 \frac{\partial \Psi_k(x)}{\partial x} \frac{\partial \Psi_m(x)}{\partial x} dx}_{A_{mk}} \underbrace{\int_0^{2\pi} \phi_l(y) \phi_n(y) dy}_{B_{nl}} \hat{u}_{kl}, \end{aligned}$$

where  $A \in \mathbb{R}^{N-2 \times N-2}$  and  $B \in \mathbb{R}^{M \times M}$ . The tensor product matrix  $A_{mk} B_{nl}$  (or in matrix notation  $A \otimes B$ ) is the first item of the two items in the list that is returned by `inner(grad(u), grad(v))`. The other item is of course the second term in the last line of (2.2):

$$\begin{aligned} \int_{-1}^1 \int_0^{2\pi} \frac{\partial u}{\partial y} \frac{\partial v}{\partial y} dx dy &= \int_{-1}^1 \int_0^{2\pi} \frac{\partial}{\partial y} \left( \sum_k \sum_l \hat{u}_{kl} \Psi_k(x) \phi_l(y) \right) \frac{\partial}{\partial y} (\Psi_m(x) \phi_n(y)) dx dy \\ &= \sum_k \sum_l \underbrace{\int_{-1}^1 \Psi_k(x) \Psi_m(x) dx}_{C_{mk}} \underbrace{\int_0^{2\pi} \frac{\partial \phi_l(y)}{\partial y} \frac{\partial \phi_n(y)}{\partial y} dy}_{D_{nl}} \hat{u}_{kl} \end{aligned}$$

The tensor product matrices  $A_{mk}B_{nl}$  and  $C_{mk}D_{nl}$  are both instances of the `TPMatrix` class. Together they lead to linear algebra systems like:

$$(A_{mk}B_{nl} + C_{mk}D_{nl})\hat{u}_{kl} = \tilde{f}_{mn}, \quad (2.3)$$

where

$$\tilde{f}_{mn} = (v_{mn}, f)_w,$$

for some right hand side  $f$ , see, e.g., (2.5). Note that an alternative formulation here is

$$A\hat{u}B^T + C\hat{u}D^T = \tilde{f}$$

where  $\hat{u}$  and  $\tilde{f}$  are treated as regular matrices ( $\hat{u} \in \mathbb{R}^{N-2 \times M}$  and  $\tilde{f} \in \mathbb{R}^{N-2 \times M}$ ). This formulation is utilized to derive efficient solvers for tensor product bases in multiple dimensions using the matrix decomposition method in [She94] and [She95].

Note that in our case the equation system (2.3) can be greatly simplified since three of the submatrices ( $A_{mk}$ ,  $B_{nl}$  and  $D_{nl}$ ) are diagonal. Even more, two of them equal the identity matrix

$$\begin{aligned} A_{mk} &= \delta_{mk}, \\ B_{nl} &= \delta_{nl}, \end{aligned}$$

whereas the last one can be written in terms of the identity (no summation on repeating indices)

$$D_{nl} = -nl\delta_{nl}.$$

Inserting for this in (2.3) and simplifying by requiring that  $l = n$  in the second step, we get

$$\begin{aligned} (\delta_{mk}\delta_{nl} - lnC_{mk}\delta_{nl})\hat{u}_{kl} &= \tilde{f}_{mn}, \\ (\delta_{mk} - l^2C_{mk})\hat{u}_{kl} &= \tilde{f}_{ml}. \end{aligned} \quad (2.4)$$

Now if we keep  $l$  fixed this latter equation is simply a regular linear algebra problem to solve for  $\hat{u}_{kl}$ , for all  $k$ . Of course, this solve needs to be carried out for all  $l$ .

Note that there is a generic solver available for the system (2.3) in `SolverGeneric2NP` that makes no assumptions on diagonality. However, this solver will, naturally, be quite a bit slower than a tailored solver that takes advantage of diagonality. For the Poisson equation such solvers are available for both Legendre and Chebyshev bases, see the extended demo [Demo - 3D Poisson's equation](#) or the demo programs `dirichlet_poisson2D.py` and `dirichlet_poisson3D.py`.

## 2.4 Coupled problems

With Shenfun it is possible to solve equations coupled and implicit using the `MixedTensorProductSpace` class for multidimensional problems and `MixedBasis` for one-dimensional problems. As an example, lets consider a mixed formulation of the Poisson equation. The Poisson equation is given as always as

$$\nabla^2 u(\mathbf{x}) = f(\mathbf{x}), \quad \text{for } \mathbf{x} \in \Omega, \quad (2.5)$$

but now we recast the problem into a mixed formulation

$$\begin{aligned} \sigma(\mathbf{x}) - \nabla u(\mathbf{x}) &= 0, \quad \text{for } \mathbf{x} \in \Omega, \\ \nabla \cdot \sigma(\mathbf{x}) &= f(\mathbf{x}), \quad \text{for } \mathbf{x} \in \Omega. \end{aligned}$$

where we solve for the vector  $\sigma$  and scalar  $u$  simultaneously. The domain  $\Omega$  is taken as a multidimensional Cartesian product  $\Omega = [-1, 1] \times [0, 2\pi]$ , but the code is more or less identical for a 3D problem. For boundary conditions we use Dirichlet in the  $x$ -direction and periodicity in the  $y$ -direction:

$$\begin{aligned} u(\pm 1, y) &= 0 \\ u(x, 2\pi) &= u(x, 0) \end{aligned}$$

Note that there is no boundary condition on  $\sigma$ , only on  $u$ . For this reason we choose a Dirichlet basis  $SD$  for  $u$  and a regular Legendre or Chebyshev  $ST$  basis for  $\sigma$ . With  $K0$  representing the function space in the periodic direction, we get the relevant 2D tensor product spaces as  $TD = SD \otimes K0$  and  $TT = ST \otimes K0$ . Since  $\sigma$  is a vector we use a `VectorTensorProductSpace`  $VT = TT \times TT$  and finally a `MixedTensorProductSpace`  $Q = VT \times TD$  for the coupled and implicit treatment of  $(\sigma, u)$ :

```
N, M = (16, 24)
family = 'Legendre'
SD = Basis(N[0], family, bc=(0, 0))
ST = Basis(N[0], family)
K0 = Basis(N[1], 'Fourier', dtype='d')
TD = TensorProductSpace(comm, (SD, K0), axes=(0, 1))
TT = TensorProductSpace(comm, (ST, K0), axes=(0, 1))
VT = VectorTensorProductSpace(TT)
Q = MixedTensorProductSpace([VT, TD])
```

In variational form the problem reads: find  $(\sigma, u) \in Q$  such that

$$\begin{aligned} (\sigma, \tau)_w - (\nabla u, \tau)_w &= 0, \quad \forall \tau \in VT, \\ (\nabla \cdot \sigma, v)_w &= (f, v)_w \quad \forall v \in TD \end{aligned} \tag{2.6}$$

To implement this we use code that is very similar to regular, uncoupled problems. We create test and trialfunction:

```
gu = TrialFunction(Q)
tv = TestFunction(Q)
sigma, u = gu
tau, v = tv
```

and use these to assemble all blocks of the variational form (2.6):

```
# Assemble equations
A00 = inner(sigma, tau)
if family.lower() == 'legendre':
    A01 = inner(u, div(tau))
else:
    A01 = inner(-grad(u), tau)
A10 = inner(div(sigma), v)
```

Note that we here can use integration by parts for Legendre, since the weight function is a constant, and as such get the term  $(-\nabla u, \tau)_w = (u, \nabla \cdot \tau)_w$  (boundary term is zero due to homogeneous Dirichlet boundary conditions).

We collect all assembled terms in a `BlockMatrix`:

```
H = BlockMatrix(A00+A01+A10)
```

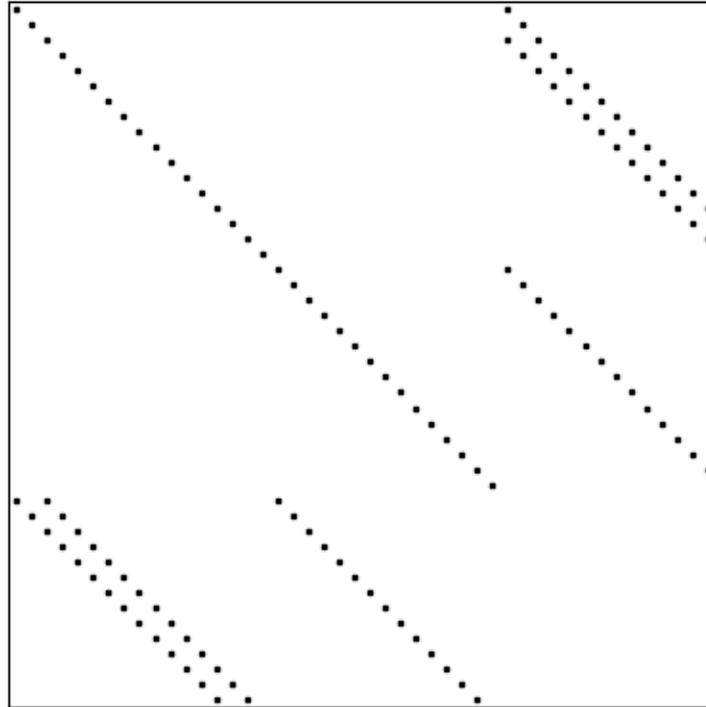
This block matrix  $H$  is then simply (for Legendre)

$$\begin{bmatrix} (\sigma, \tau)_w & (u, \nabla \cdot \tau)_w \\ (\nabla \cdot \sigma, v)_w & 0 \end{bmatrix} \tag{2.7}$$

Note that each item in (2.7) is a collection of instances of the `TPMatrix` class, and for similar reasons as given around (2.4), we get also here one regular block matrix for each Fourier wavenumber. The sparsity pattern is the same for all

matrices except for wavenumber 0. The (highly sparse) sparsity pattern for block matrix  $H$  with wavenumber  $\neq 0$  is shown in the image below

Sparsity pattern - Legendre Mixed Poisson



A complete demo for the coupled problem discussed here can be found in [MixedPoisson.py](#) and a 3D version is in [MixedPoisson3D.py](#).

## 2.5 Integrators

The `integrators` module contains some interator classes that can be used to integrate a solution forward in time. However, for now these integrators are only implemented for purely Fourier tensor product spaces. There are currently 3 different integrator classes

- RK4: Runge-Kutta fourth order
- ETD: Exponential time differencing Euler method
- ETD RK4: Exponential time differencing Runge-Kutta fourth order

See, e.g., H. Montanelli and N. Bootland “Solving periodic semilinear PDEs in 1D, 2D and 3D with exponential integrators”, <https://arxiv.org/pdf/1604.08900.pdf>

Integrators are set up to solve equations like

$$\frac{\partial u}{\partial t} = Lu + N(u) \tag{2.8}$$

where  $u$  is the solution,  $L$  is a linear operator and  $N(u)$  is the nonlinear part of the right hand side.

To illustrate, we consider the time-dependent 1-dimensional Kortveeg-de Vries equation

$$\frac{\partial u}{\partial t} + \frac{\partial^3 u}{\partial x^3} + u \frac{\partial u}{\partial x} = 0$$

which can also be written as

$$\frac{\partial u}{\partial t} + \frac{\partial^3 u}{\partial x^3} + \frac{1}{2} \frac{\partial u^2}{\partial x} = 0$$

We neglect boundary issues and choose a periodic domain  $[0, 2\pi]$  with Fourier exponentials as test functions. The initial condition is chosen as

$$u(x, t = 0) = 3A^2 / \cosh(0.5A(x - \pi + 2))^2 + 3B^2 / \cosh(0.5B(x - \pi + 1))^2 \quad (2.9)$$

where  $A$  and  $B$  are constants. For discretization in space we use the basis  $V_N = \text{span}\{\exp(ikx)\}_{k=0}^N$  and formulate the variational problem: find  $u \in V_N$  such that

$$\frac{\partial}{\partial t} (u, v) = - \left( \frac{\partial^3 u}{\partial x^3}, v \right) - \left( \frac{1}{2} \frac{\partial u^2}{\partial x}, v \right), \quad \forall v \in V_N$$

We see that the first term on the right hand side is linear in  $u$ , whereas the second term is nonlinear. To implement this problem in shenfun we start by creating the necessary basis and test and trial functions

```
import numpy as np
from shenfun import *

N = 256
T = Basis(N, 'F', dtype='d')
u = TrialFunction(T)
v = TestFunction(T)
u_ = Array(T)
u_hat = Function(T)
```

We then create two functions representing the linear and nonlinear part of (??):

```
def LinearRHS(**params):
    return -inner(Dx(u, 0, 3), v)

k = T.wavenumbers(scaled=True, eliminate_highest_freq=True)
def NonlinearRHS(u, u_hat, rhs, **params):
    rhs.fill(0)
    u_[:] = T.backward(u_hat, u_)
    rhs = T.forward(-0.5*u_**2, rhs)
    return rhs*1j*k # return inner(grad(-0.5*Up**2), v)
```

Note that we differentiate in `NonlinearRHS` by using the wavenumbers `k` directly. Alternative notation, that is given in commented out text, is slightly slower, but the results are the same.

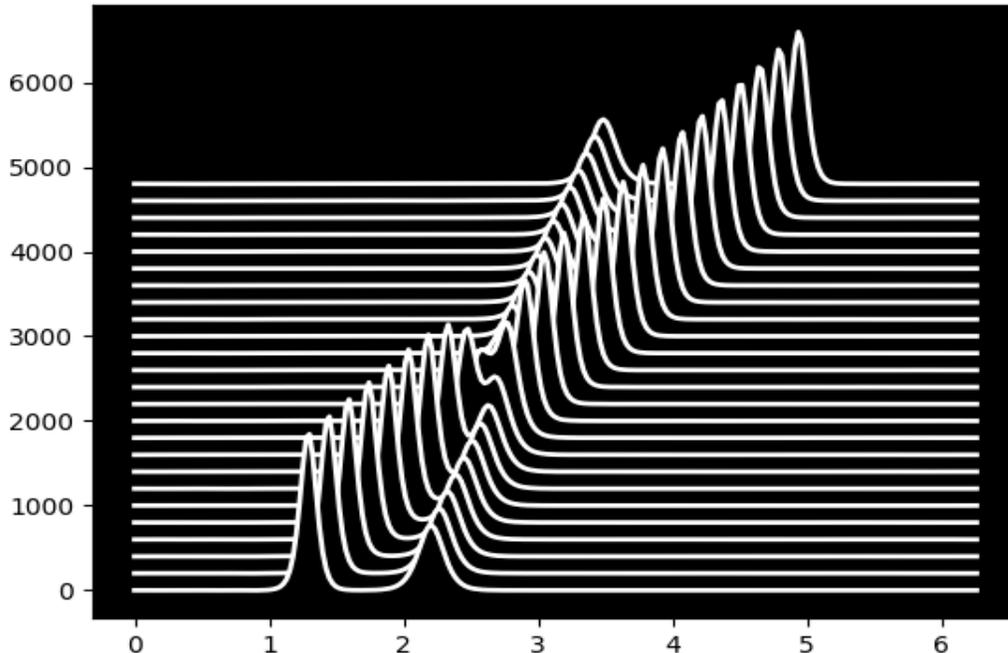
The solution vector `u_` needs also to be initialized according to (??)

```
A = 25.
B = 16.
x = T.points_and_weights()[0]
u_[:] = 3*A**2/np.cosh(0.5*A*(x-np.pi+2))**2 + 3*B**2/np.cosh(0.5*B*(x-np.pi+1))**2
u_hat = T.forward(u_, u_hat)
```

Finally we create an instance of the `ETDRK4` solver, and integrate forward with a given timestep

```
dt = 0.01/N**2
end_time = 0.006
integrator = ETDK4(T, L=LinearRHS, N=NonlinearRHS)
integrator.setup(dt)
u_hat = integrator.solve(u_, u_hat, dt, (0, end_time))
```

The solution is two waves travelling through each other, seemingly undisturbed.



## 2.6 MPI

Shenfun makes use of the Message Passing Interface (MPI) to solve problems on distributed memory architectures. OpenMP is also possible to enable for FFTs.

Dataarrays in Shenfun are distributed using a [new and completely generic method](#), that allows for any index of a multidimensional array to be distributed. To illustrate, let's consider a `TensorProductSpace` of three dimensions, such that the arrays living in this space will be 3-dimensional. We create two spaces that are identical, except from the MPI decomposition, and we use 4 CPUs (`mpirun -np 4 python mpitest.py`, if we store the code in this section as `mpitest.py`):

```
from shenfun import *
from mpi4py import MPI
from mpi4py_fft import generate_xdmf
comm = MPI.COMM_WORLD
N = (20, 40, 60)
K0 = Basis(N[0], 'F', dtype='D', domain=(0, 1))
K1 = Basis(N[1], 'F', dtype='D', domain=(0, 2))
K2 = Basis(N[2], 'F', dtype='d', domain=(0, 3))
```

(continues on next page)

(continued from previous page)

```
T0 = TensorProductSpace(comm, (K0, K1, K2), axes=(0, 1, 2), slab=True)
T1 = TensorProductSpace(comm, (K0, K1, K2), axes=(1, 0, 2), slab=True)
```

Here the keyword `slab` determines that only *one* index set of the 3-dimensional arrays living in `T0` or `T1` should be distributed. The default is to use two, which corresponds to a so-called pencil decomposition. The `axes`-keyword determines the order of which transforms are conducted, starting from last to first in the given tuple. Note that `T0` now will give arrays in real physical space that are distributed in the first index, whereas `T1` will give arrays that are distributed in the second. This is because 0 and 1 are the first items in the tuples given to `axes`.

We can now create some Arrays on these spaces:

```
u0 = Array(T0, val=comm.Get_rank())
u1 = Array(T1, val=comm.Get_rank())
```

such that `u0` and `u1` have values corresponding to their communicating processors rank in the `COMM_WORLD` group (the group of all CPUs).

Note that both the `TensorProductSpaces` have functions with expansion

$$u(x, y, z) = \sum_{n=-N/2}^{N/2-1} \sum_{m=-N/2}^{N/2-1} \sum_{l=-N/2}^{N/2-1} \hat{u}_{l,m,n} e^{i(lx+my+nz)}. \quad (2.10)$$

where  $u(x, y, z)$  is the continuous solution in real physical space, and  $\hat{u}$  are the spectral expansion coefficients. If we evaluate expansion (??) on the real physical mesh, then we get

$$u(x_i, y_j, z_k) = \sum_{n=-N/2}^{N/2-1} \sum_{m=-N/2}^{N/2-1} \sum_{l=-N/2}^{N/2-1} \hat{u}_{l,m,n} e^{i(lx_i+my_j+nz_k)}. \quad (2.11)$$

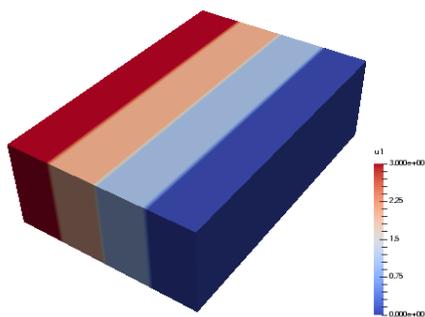
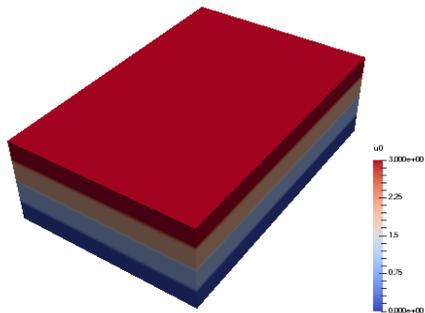
The function  $u(x_i, y_j, z_k)$  corresponds to the arrays `u0`, `u1`, whereas we have not yet computed the array  $\hat{u}$ . We could get  $\hat{u}$  as:

```
u0_hat = Function(T0)
u0_hat = T0.forward(u0, u0_hat)
```

Now, `u0` and `u1` have been created on the same mesh, which is a structured mesh of shape (20, 40, 60). However, since they have different MPI decomposition, the values used to fill them on creation will differ. We can visualize the arrays in Paraview using some postprocessing tools, to be further described in Sec Postprocessing:

```
u0.write('myfile.h5', 'u0', 0, domain=T0.mesh())
u1.write('myfile.h5', 'u1', 0, domain=T1.mesh())
if comm.Get_rank() == 0:
    generate_xdmf('myfile.h5')
```

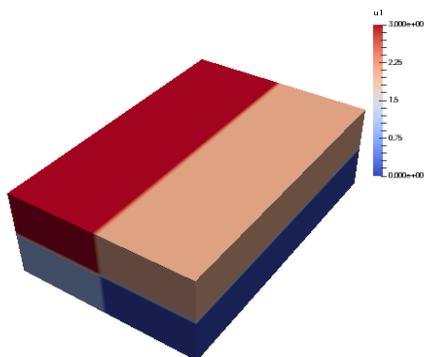
And when the generated `myfile.xdmf` is opened in Paraview, we can see the different distributions. The function `u0` is shown first, and we see that it has different values along the short first dimension. The second figure is evidently distributed along the second dimension. Both arrays are non-distributed in the third and final dimension, which is fortunate, because this axis will be the first to be transformed in, e.g., `u0_hat = T0.forward(u0, u0_hat)`.



We can now decide to distribute not just one, but the first two axes using a pencil decomposition instead. This is achieved simply by dropping the slab keyword:

```
T2 = TensorProductSpace(comm, (K0, K1, K2), axes=(0, 1, 2))
u2 = Array(T2, val=comm.Get_rank())
u2.write('pencilfile.h5', 'u2', 0)
if comm.Get_rank() == 0:
    generate_xdmf('pencilfile.h5')
```

Running again with 4 CPUs the array u2 will look like:



The local slices into the global array may be obtained through:

```
>>> print(comm.Get_rank(), T2.local_slice(False))
0 [slice(0, 10, None), slice(0, 20, None), slice(0, 60, None)]
```

(continues on next page)

(continued from previous page)

```
1 [slice(0, 10, None), slice(20, 40, None), slice(0, 60, None)]
2 [slice(10, 20, None), slice(0, 20, None), slice(0, 60, None)]
3 [slice(10, 20, None), slice(20, 40, None), slice(0, 60, None)]
```

In spectral space the distribution will be different. This is because the discrete Fourier transforms are performed one axis at the time, and for this to happen the dataarrays need to be realigned to get entire axis available for each processor. Naturally, for the array in the pencil example (see image), we can only perform an FFT over the third and longest axis, because only this axis is locally available to all processors. To do the other directions, the dataarray must be realigned and this is done internally by the `TensorProductSpace` class. The shape of the datastructure in spectral space, that is the shape of  $\hat{u}$ , can be obtained as:

```
>>> print(comm.Get_rank(), T2.local_slice(True))
0 [slice(0, 20, None), slice(0, 20, None), slice(0, 16, None)]
1 [slice(0, 20, None), slice(0, 20, None), slice(16, 31, None)]
2 [slice(0, 20, None), slice(20, 40, None), slice(0, 16, None)]
3 [slice(0, 20, None), slice(20, 40, None), slice(16, 31, None)]
```

Evidently, the spectral space is distributed in the last two axes, whereas the first axis is locally available to all processors. The dataarray is said to be aligned in the first dimension.



## POST PROCESSING

MPI is great because it means that you can run Shenfun on pretty much as many CPUs as you can get your hands on. However, MPI makes it more challenging to do visualization, in particular with Python and Matplotlib. For this reason there is a `utilities` module with helper classes for dumping dataarrays to [HDF5](#) or [NetCDF](#)

Most of the IO has already been implemented in `mpi4py-fft`. The classes `HDF5File` and `NCFFile` are used exactly as they are implemented in `mpi4py-fft`. As a common interface we provide

- `ShenfunFile()`

where `ShenfunFile()` returns an instance of either `HDF5File` or `NCFFile`, depending on choice of backend.

For example, to create an HDF5 writer for a 3D `TensorProductSpace` with Fourier bases in all directions:

```
from shenfun import *
from mpi4py import MPI
N = (24, 25, 26)
K0 = Basis(N[0], 'F', dtype='D')
K1 = Basis(N[1], 'F', dtype='D')
K2 = Basis(N[2], 'F', dtype='d')
T = TensorProductSpace(MPI.COMM_WORLD, (K0, K1, K2))
fl = ShenfunFile('myh5file', T, backend='hdf5', mode='w')
```

The file instance `fl` will now have two methods that can be used to either write dataarrays to file, or read them back again.

- `fl.write`
- `fl.read`

With the HDF5 backend we can write both arrays from physical space (`Array`), as well as spectral space (`Function`). However, the `NetCDF4` backend cannot handle complex dataarrays, and as such it can only be used for real physical dataarrays.

In addition to storing complete dataarrays, we can also store any slices of the arrays. To illustrate, this is how to store three snapshots of the `u` array, along with some *global* 2D and 1D slices:

```
u = Array(T)
u[:] = np.random.random(u.shape)
d = {'u': [u, (u, np.s_[4, :, :]), (u, np.s_[4, 4, :])]}
fl.write(0, d)
u[:] = 2
fl.write(1, d)
```

The `ShenfunFile` may also be used for the `MixedTensorProductSpace`, or `VectorTensorProductSpace`, that are collections of the scalar `TensorProductSpace`. We can create a `MixedTensorProductSpace` consisting of two `TensorProductSpaces`, and an accompanying writer class as:

```
TT = MixedTensorProductSpace([T, T])
fl_m = ShenfunFile('mixed', TT, backend='hdf5', mode='w')
```

Let's now consider a transient problem where we step a solution forward in time. We create a solution array from the Array class, and update the array inside a while loop:

```
TT = VectorTensorProductSpace(T)
fl_m = ShenfunFile('mixed', TT, backend='hdf5', mode='w')
u = Array(TT)
tstep = 0
du = {'uv': (u,
             (u, [4, slice(None), slice(None)]),
             (u, [slice(None), 10, 10]))}
while tstep < 3:
    fl_m.write(tstep, du, forward_output=False)
    tstep += 1
```

Note that on each time step the arrays `u`, `(u, [4, slice(None), slice(None)])` and `(u, [slice(None), 10, 10])` are vectors, and as such of global shape `(3, 24, 25, 26)`, `(3, 25, 26)` and `(3, 25)`, respectively. However, they are stored in the hdf5 file under their spatial dimensions 1D, 2D and 3D, respectively.

Note that the slices in the above dictionaries are *global* views of the global arrays, that may or may not be distributed over any number of processors. Also note that these routines work with any number of CPUs, and the number of CPUs does not need to be the same when storing or retrieving the data.

After running the above, the different arrays will be found in groups stored in `myyfile.h5` with directory tree structure as:

```
myh5file.h5/
├─ u/
│  ├─ 1D/
│  │  └─ 4_4_slice/
│  │     └─ 0
│  │     └─ 1
│  ├─ 2D/
│  │  └─ 4_slice_slice/
│  │     └─ 0
│  │     └─ 1
│  └─ 3D/
│     └─ 0
│     └─ 1
└─ mesh/
   └─ x0
   └─ x1
   └─ x2
```

Likewise, the `mixed.h5` file will at the end of the loop look like:

```
mixed.h5/
├─ uv/
│  ├─ 1D/
│  │  └─ slice_10_10/
│  │     └─ 0
│  │     └─ 1
│  │     └─ 3
│  └─ 2D/
```

(continues on next page)

(continued from previous page)

```

|   └─ 4_slice_slice/
|       └─ 0
|       └─ 1
|       └─ 3
├─ 3D/
|   └─ 0
|   └─ 1
|   └─ 3
└─ mesh/
    └─ x0
    └─ x1
    └─ x2

```

Note that the mesh is stored as well as the results. The three mesh arrays are all 1D arrays, representing the domain for each basis in the `TensorProductSpace`.

With `NetCDF4` the layout is somewhat different. For `mixed` above, if we were using backend `netcdf` instead of `hdf5`, we would get a datafile where `ncdump -h mixed.nc` would result in:

```

netcdf mixed {
dimensions:
    time = UNLIMITED ; // (3 currently)
    i = 3 ;
    x = 24 ;
    y = 25 ;
    z = 26 ;
variables:
    double time(time) ;
    double i(i) ;
    double x(x) ;
    double y(y) ;
    double z(z) ;
    double uv(time, i, x, y, z) ;
    double uv_4_slice_slice(time, i, y, z) ;
    double uv_slice_10_10(time, i, x) ;
}

```

Note that it is also possible to store vector arrays as scalars. For `NetCDF4` this is necessary for direct visualization using `Visit`. To store vectors as scalars, simply use:

```
fl_m.write(tstep, du, forward_output=False, as_scalar=True)
```

## 3.1 ParaView

The stored datafiles can be visualized in `ParaView`. However, `ParaView` cannot understand the content of these `HDF5`-files without a little bit of help. We have to explain that these data-files contain structured arrays of such and such shape. The way to do this is through the simple XML descriptor `XDMF`. To this end there is a function imported from `mpi4py-fft` called `generate_xdmf` that can be called with any one of the generated `hdf5`-files:

```
generate_xdmf('myh5file.h5')
generate_xdmf('mixed.h5')
```

This results in some light `xdmf`-files being generated for the 2D and 3D arrays in the `hdf5`-file:

- `myh5file.xdmf`

- myh5file\_4\_slice\_slice.xdmf
- mixed.xdmf
- mixed\_4\_slice\_slice.xdmf

These xdmf-files can be opened and inspected by ParaView. Note that 1D arrays are not wrapped, and neither are 4D.

An annoying feature of Paraview is that it views a three-dimensional array of shape  $(N_0, N_1, N_2)$  as transposed compared to shenfun. That is, for Paraview the *last* axis represents the  $x$ -axis, whereas shenfun (like most others) considers the first axis to be the  $x$ -axis. So when opening a three-dimensional array in Paraview one needs to be aware. Especially when plotting vectors. Assume that we are working with a Navier-Stokes solver and have a three-dimensional `VectorTensorProductSpace` to represent the fluid velocity:

```
from mpi4py import MPI
from shenfun import *

comm = MPI.COMM_WORLD
N = (32, 64, 128)
V0 = Basis(N[0], 'F', dtype='D')
V1 = Basis(N[1], 'F', dtype='D')
V2 = Basis(N[2], 'F', dtype='d')
T = TensorProductSpace(comm, (V0, V1, V2))
TV = VectorTensorProductSpace(T)
U = Array(TV)
U[0] = 0
U[1] = 1
U[2] = 2
```

To store the resulting Array `U` we can create an instance of the `HDF5File` class, and store using keyword `as_scalar=True`:

```
hdf5file = ShenfunFile("NS", TV, backend='hdf5', mode='w')
...
file.write(0, {'u': [U]}, as_scalar=True)
file.write(1, {'u': [U]}, as_scalar=True)
```

Alternatively, one may store the arrays directly as:

```
U.write('U.h5', 'u', 0, domain=T.mesh(), as_scalar=True)
U.write('U.h5', 'u', 1, domain=T.mesh(), as_scalar=True)
```

Generate an xdmf file through:

```
generate_xdmf('NS.h5')
```

and open the generated `NS.xdmf` file in Paraview. You will then see three scalar arrays `u0`, `u1`, `u2`, each one of shape  $(32, 64, 128)$ , for the vector component in what Paraview considers the  $z$ ,  $y$  and  $x$  directions, respectively. Other than the swapped coordinate axes there is no difference. But be careful if creating vectors in Paraview with the Calculator. The vector should be created as:

```
u0*kHat+u1*jHat+u2*iHat
```

## INSTALLATION

Shenfun has a few dependencies

- `mpi4py`
- `FFTW`
- `mpi4py-fft`
- `cython`
- `numpy`
- `sympy`
- `scipy`
- `h5py`

that are mostly straight-forward to install, or already installed in most Python environments. The first two are usually most troublesome. Basically, for `mpi4py` you need to have a working MPI installation, whereas `FFTW` is available on most high performance computer systems. If you are using `conda`, then all you need to install a fully functional shenfun, with all the above dependencies, is

```
conda install -c conda-forge shenfun
```

You probably want to install into a fresh environment, though, which can be achieved with

```
conda create --name shenfun -c conda-forge shenfun
conda activate shenfun
```

Note that this gives you shenfun with default settings. This means that you will probably get the `openmpi` backend, and it is also likely that `conda-forge` chooses `numpy` with the `mkl` backend. Unfortunately, the `mkl` python package makes adjustments to the `FFTW` library and hard to resolve bugs may arise. For this reason it is advisable to make sure that `mkl` is not installed. This can be achieved with, e.g.,

```
conda create --name shenfun -c conda-forge shenfun mpich nomkl
```

Note that the `nomkl` package makes sure that `numpy` is installed without `mkl`, whereas `mpich` here chooses this backend over `openmpi`.

If you do not use `conda`, then you need to make sure that MPI and `FFTW` are installed by some other means. You can then install any version of shenfun hosted on `pypi` using `pip`

```
pip install shenfun
```

whereas the following will install the latest version from github

```
pip install git+https://github.com/spectralDNS/shenfun.git@master
```

You can also build shenfun yourselves from the top directory, after cloning or forking

```
pip install .
```

or using `conda-build` with the recipes in folder `conf/conda`

```
conda build -c conda-forge -c spectralDNS conf/conda
conda create --name shenfun -c conda-forge -c spectralDNS shenfun --use-local
conda activate shenfun
```

## 4.1 Optimization

Shenfun contains a few routines (essentially linear algebra solvers and matrix vector products) that are difficult to vectorize with numpy, and for this reason they have been implemented in either (or both of) [Numba](#) or [Cython](#). The user may choose which implementation to use through the environment variable `SHENFUN_OPTIMIZATION`. The default is to use cython, but it is possible to enable either one by making the appropriate choice in the active terminal

```
export SHENFUN_OPTIMIZATION={CYTHON,NUMBA}
```

## 4.2 Additional dependencies

For storing and retrieving data you need either [HDF5](#) or [netCDF4](#), compiled with support for MPI (see Postprocessing). [HDF5](#) is already available with parallel support on [conda-forge](#) and, if it was not installed at the same time as shenfun, it can be installed (with the mpich backend for MPI) as

```
conda install -c conda-forge h5py=*mpi_mpich_*
```

A parallel version of [netCDF4](#) cannot be found on the conda-forge channel, but a precompiled version has been made available for python 2.7, 3.6 and 3.7 on the [spectralDNS channel](#), for both osx and linux

```
conda install -c spectralDNS netcdf4-parallel
```

Note that parallel HDF5 and NetCDF4 often are available as modules on supercomputers. Otherwise, see the respective packages for how to install with support for MPI.

## 4.3 Test installation

After installing (from source) it may be a good idea to run all the tests located in the `tests` folder. The tests are run with `pytest` from the main directory of the source code

```
python -m pytest tests/
```

However, note that for conda you need to install pytest into the correct environment as well. A common mistake is to run a version of pytest that has already been installed in a different conda environment, perhaps using a different Python version.

The tests are run automatically on every commit to github, see





## 5.1 Demo - 1D Poisson's equation

**Authors** Mikael Mortensen (mikaem at math.uio.no)

**Date** Nov 14, 2019

*Summary.* This is a demonstration of how the Python module `shenfun` can be used to solve Poisson's equation with Dirichlet boundary conditions in one dimension. Spectral convergence, as shown in Figure *Convergence of 1D Poisson solvers for both Legendre and Chebyshev modified basis function*, is demonstrated. The demo is implemented in a single Python file `dirichlet_poisson1D.py`, and the numerical method is described in more detail by J. Shen [She94] and [She95].

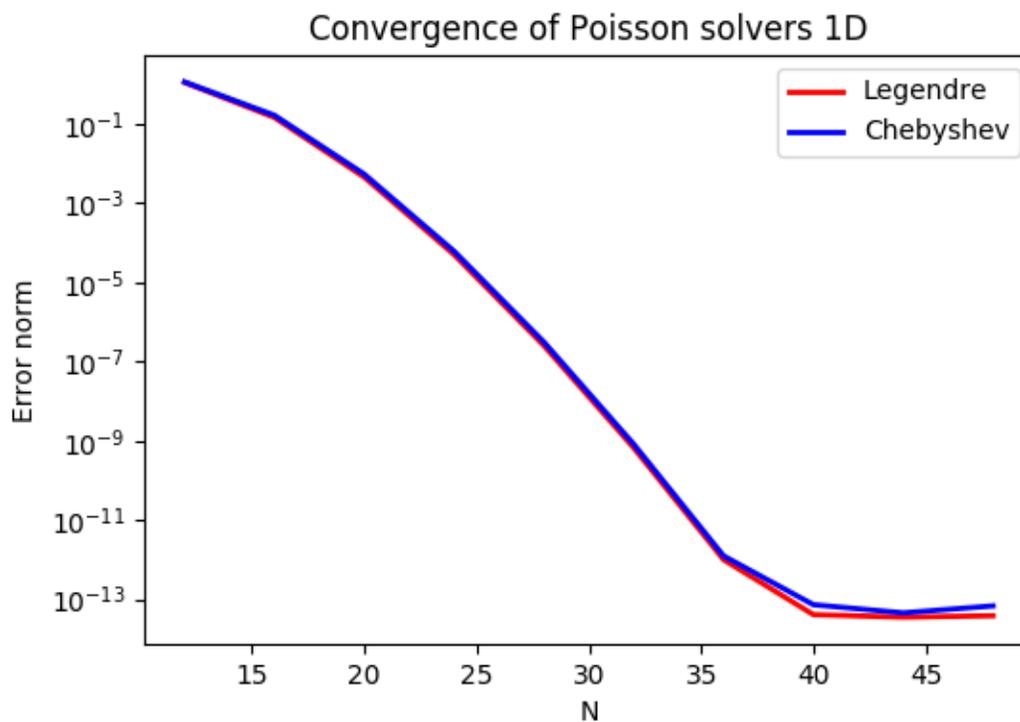


Fig. 1: *Convergence of 1D Poisson solvers for both Legendre and Chebyshev modified basis function*

## 5.1.1 Model problem

### Poisson's equation

Poisson's equation is given as

$$\nabla^2 u(x) = f(x) \quad \text{for } x \in [-1, 1], \quad (5.1)$$

$$u(-1) = a, u(1) = b,$$

where  $u(x)$  is the solution,  $f(x)$  is a function and  $a, b$  are two possibly non-zero constants.

To solve Eq. (5.1) with the Galerkin method we need smooth continuously differentiable basis functions,  $v_k$ , that satisfy the given boundary conditions. And then we look for solutions like

$$u(x) = \sum_{k=0}^{N-1} \hat{u}_k v_k(x), \quad (5.2)$$

where  $N$  is the size of the discretized problem,  $\hat{u} = \{\hat{u}_k\}_{k=0}^{N-1}$  are the unknown expansion coefficients, and the basis is  $\text{span}\{v_k\}_{k=0}^{N-1}$ .

The basis functions can, for example, be constructed from [Chebyshev](#),  $T_k(x)$ , or [Legendre](#),  $L_k(x)$ , polynomials and we use the common notation  $\phi_k(x)$  to represent either one of them. It turns out that it is easiest to use basis functions with homogeneous Dirichlet boundary conditions

$$v_k(x) = \phi_k(x) - \phi_{k+2}(x), \quad (5.3)$$

for  $k = 0, 1, \dots, N-3$ . This gives the basis  $V_0^N = \text{span}\{v_k(x)\}_{k=0}^{N-3}$ . We can then add two more linear basis functions (that belong to the kernel of Poisson's equation)

$$v_{N-2} = \frac{1}{2}(\phi_0 + \phi_1), \quad (5.4)$$

$$v_{N-1} = \frac{1}{2}(\phi_0 - \phi_1). \quad (5.5)$$

which gives the inhomogeneous basis  $V^N = \text{span}\{v_k\}_{k=0}^{N-1}$ . With the two linear basis functions it is easy to see that the last two degrees of freedom,  $\hat{u}_{N-2}$  and  $\hat{u}_{N-1}$ , now are given as

$$u(-1) = \sum_{k=0}^{N-1} \hat{u}_k v_k(-1) = \hat{u}_{N-1} = a, \quad (5.6)$$

$$u(+1) = \sum_{k=0}^{N-1} \hat{u}_k v_k(+1) = \hat{u}_{N-2} = b, \quad (5.7)$$

and, as such, we only have to solve for  $\{\hat{u}_k\}_{k=0}^{N-3}$ , just like for a problem with homogeneous boundary conditions (for homogeneous boundary condition we simply have  $\hat{u}_{N-2} = \hat{u}_{N-1} = 0$ ). We now formulate a variational problem using the Galerkin method: Find  $u \in V^N$  such that

$$\int_{-1}^1 \nabla^2 u v w dx = \int_{-1}^1 f v w dx \quad \forall v \in V_0^N. \quad (5.8)$$

Note that since we only have  $N-3$  unknowns we are only using the homogeneous test functions from  $V_0^N$ .

The weighted integrals, weighted by  $w(x)$ , are called inner products, and a common notation is

$$\int_{-1}^1 u v w dx = (u, v)_w. \quad (5.9)$$

The integral can either be computed exactly, or with quadrature. The advantage of the latter is that it is generally faster, and that non-linear terms may be computed just as quickly as linear. For a linear problem, it does not make much of a difference, if any at all. Approximating the integral with quadrature, we obtain

$$\begin{aligned} \int_{-1}^1 u v w dx &\approx (u, v)_w^N, \\ &\approx \sum_{j=0}^{N-1} u(x_j) v(x_j) w(x_j), \end{aligned}$$

where  $\{w(x_j)\}_{j=0}^{N-1}$  are quadrature weights. The quadrature points  $\{x_j\}_{j=0}^{N-1}$  are specific to the chosen basis, and even within basis there are two different choices based on which quadrature rule is selected, either Gauss or Gauss-Lobatto.

Inserting for test and trial functions, we get the following bilinear form and matrix  $A \in \mathbb{R}^{N-3 \times N-3}$  for the Laplacian (using the summation convention in step 2)

$$\begin{aligned} (\nabla^2 u, v)_w^N &= \left( \nabla^2 \sum_{k=0}^{N-3} \hat{u}_k v_k, v_j \right)_w^N, \quad j = 0, 1, \dots, N-3 \\ &= (\nabla^2 v_k, v_j)_w^N \hat{u}_k, \\ &= a_{jk} \hat{u}_k. \end{aligned}$$

Note that the sum in  $a_{jk} \hat{u}_k$  runs over  $k = 0, 1, \dots, N-3$  since the second derivatives of  $v_{N-1}$  and  $v_N$  are zero. The right hand side linear form and vector is computed as  $\tilde{f}_j = (f, v_j)_w^N$ , for  $j = 0, 1, \dots, N-3$ , where a tilde is used because this is not a complete transform of the function  $f$ , but only an inner product.

The linear system of equations to solve for the expansion coefficients of  $u(x)$  is given as

$$A \hat{\mathbf{u}} = \tilde{\mathbf{f}}. \quad (5.10)$$

Now, when the expansion coefficients  $\hat{\mathbf{u}}$  are found by solving this linear system, they may be transformed to real space  $u(x)$  using (5.2), and here the contributions from  $\hat{u}_{N-2}$  and  $\hat{u}_{N-1}$  must be accounted for. Note that the matrix  $A$  (different for Legendre or Chebyshev) has a very special structure that allows for a solution to be found very efficiently in order of  $\mathcal{O}(N)$  operations, see [She94] and [She95]. These solvers are implemented in shenfun for both bases.

## Method of manufactured solutions

In this demo we will use the method of manufactured solutions to demonstrate spectral accuracy of the shenfun Dirichlet bases. To this end we choose an analytical function that satisfies the given boundary conditions:

$$u_e(x) = \sin(k\pi x)(1-x^2) + a(1+x)/2 + b(1-x)/2, \quad (5.11)$$

where  $k$  is an integer and  $a$  and  $b$  are constants. Now, feeding  $u_e$  through the Laplace operator, we see that the last two linear terms disappear, whereas the first term results in

$$\nabla^2 u_e(x) = \frac{d^2 u_e}{dx^2}, \quad (5.12)$$

$$= -4k\pi x \cos(k\pi x) - 2 \sin(k\pi x) - k^2 \pi^2 (1-x^2) \sin(k\pi x). \quad (5.13)$$

Now, setting  $f_e(x) = \nabla^2 u_e(x)$  and solving for  $\nabla^2 u(x) = f_e(x)$ , we can compare the numerical solution  $u(x)$  with the analytical solution  $u_e(x)$  and compute error norms.

## 5.1.2 Implementation

### Preamble

We will solve Poisson's equation using the `shenfun` Python module. The first thing needed is then to import some of this module's functionality plus some other helper modules, like `Numpy` and `Sympy`:

```
from shenfun import inner, div, grad, TestFunction, TrialFunction, Function, \
    project, Dx, Array, Basis
import numpy as np
from sympy import symbols, cos, sin, exp, lambdify
```

We use `Sympy` for the manufactured solution and `Numpy` for testing.

### Manufactured solution

The exact solution  $u_e(x)$  and the right hand side  $f_e(x)$  are created using `Sympy` as follows

```
a = -1
b = 1
k = 4
x = symbols("x")
ue = sin(k*np.pi*x)*(1-x**2) + a*(1 + x)/2. + b*(1 - x)/2.
fe = ue.diff(x, 2)

# Lambdify for faster evaluation
ul = lambdify(x, ue, 'numpy')
fl = lambdify(x, fe, 'numpy')
```

These solutions are now valid for a continuous domain. The next step is thus to discretize, using a discrete mesh  $\{x_j\}_{j=0}^{N-1}$  and a finite number of basis functions.

Note that it is not mandatory to use `Sympy` for the manufactured solution. Since the solution is known (5.13), we could just as well simply use `Numpy` to compute  $f_e$  at  $\{x_j\}_{j=0}^{N-1}$ . However, with `Sympy` it is much easier to experiment and quickly change the solution.

### Discretization

We create a basis with a given number of basis functions, and extract the computational mesh from the basis itself

```
N = 32
SD = Basis(N, 'Chebyshev', bc=(a, b))
#SD = Basis(N, 'Legendre', bc=(a, b))
X = SD.mesh(N)
```

Note that we can either choose a Legendre or a Chebyshev basis.

### Variational formulation

The variational problem (5.97) can be assembled using `shenfun`'s `TrialFunction`, `TestFunction` and `inner()` functions.

```

u = TrialFunction(SD)
v = TestFunction(SD)
# Assemble left hand side matrix
A = inner(v, div(grad(u)))
# Assemble right hand side
fj = Array(SD, buffer=f1(X))
f_hat = Function(SD)
f_hat = inner(v, fj, output_array=f_hat)

```

Note that `f1(X)` returns a Numpy array of the correct shape and type of the left hand side of (5.2), evaluated on all quadrature points `X`. We wrap this Numpy array in an `Array` class (`fj = Array(SD, buffer=f1(X))`), because an `Array` is required as input to the `inner()` function.

## Solve linear equations

Finally, solve linear equation system and transform solution from spectral  $\{\hat{u}_k\}_{k=0}^{N-1}$  vector to the real space  $\{u(x_j)\}_{j=0}^{N-1}$  and then check how the solution corresponds with the exact solution  $u_e$ .

```

u_hat = A.solve(f_hat)
uj = SD.backward(u_hat)
ue = ul(X)
print("Error=%2.16e" % (np.linalg.norm(uj-ue)))
assert np.allclose(uj, ue)

```

## Convergence test

A complete solver is given in Sec. *Complete solver*. This solver is created such that it takes in two commandline arguments and prints out the  $l_2$ -errornorm of the solution in the end. We can use this to write a short script that performs a convergence test. The solver is run like

```

>>> python dirichlet_poisson1D.py 32 legendre
Error=6.5955040031498912e-10

```

for a discretization of size  $N = 32$  and for the Legendre basis. Alternatively, change `legendre` to `chebyshev` for the Chebyshev basis.

We set up the solver to run for a list of  $N = [12, 16, \dots, 48]$ , and collect the errornorms in arrays to be plotted. Such a script can be easily created with the `subprocess` module

```

import subprocess

N = range(12, 50, 4)
error = {}
for basis in ('legendre', 'chebyshev'):
    error[basis] = []
    for i in range(len(N)):
        output = subprocess.check_output("python dirichlet_poisson1D.py {} {}".
    ↪format(N[i], basis), shell=True)
        exec(output) # Error is printed as "Error=%2.16e"%(np.linalg.norm(uj-ua))
        error[basis].append(Error)

```

The error can be plotted using `matplotlib`, and the generated figure is shown in the summary's Fig. *Convergence of 1D Poisson solvers for both Legendre and Chebyshev modified basis function*. The spectral convergence is evident and we can see that after  $N = 40$  roundoff errors dominate as the errornorm trails off around  $10^{-14}$ .

```
import matplotlib.pyplot as plt
plt.figure(figsize=(6, 4))
for basis, col in zip(('legendre', 'chebyshev'), ('r', 'b')):
    plt.semilogy(N, error[basis], col, linewidth=2)
plt.title('Convergence of Poisson solvers 1D')
plt.xlabel('N')
plt.ylabel('Error norm')
plt.savefig('poisson1D_errornorm.png')
plt.legend(('Legendre', 'Chebyshev'))
plt.show()
```

## Complete solver

A complete solver, that can use either Legendre or Chebyshev bases, chosen as a command-line argument, can be found [here](#).

## 5.2 Demo - Cubic nonlinear Klein-Gordon equation

**Authors** Mikael Mortensen (mikaem at math.uio.no)

**Date** Nov 14, 2019

*Summary.* This is a demonstration of how the Python module `shenfun` can be used to solve the time-dependent, nonlinear Klein-Gordon equation, in a triply periodic domain. The demo is implemented in a single Python file `Klein-Gordon.py`, and it may be run in parallel using MPI. The Klein-Gordon equation is solved using a mixed formulation. The discretization, and some background on the spectral Galerkin method is given first, before we turn to the actual details of the `shenfun` implementation.

### 5.2.1 The nonlinear Klein-Gordon equation

Movie showing the evolution of the solution  $u$  from Eq. (5.14), in a slice through the center of the domain, computed with the code described in this demo.

#### Model equation

The cubic nonlinear Klein-Gordon equation is a wave equation important for many scientific applications such as solid state physics, nonlinear optics and quantum field theory [Waz08]. The equation is given as

$$\frac{\partial^2 u}{\partial t^2} = \nabla^2 u - \gamma(u - |u|^2) \quad \text{for } u \in \Omega, \quad (5.14)$$

with initial conditions

$$u(\mathbf{x}, t = 0) = u^0 \quad \text{and} \quad \frac{\partial u(\mathbf{x}, t = 0)}{\partial t} = u_t^0. \quad (5.15)$$

The spatial coordinates are here denoted as  $\mathbf{x} = (x, y, z)$ , and  $t$  is time. The parameter  $\gamma = \pm 1$  determines whether the equations are focusing (+1) or defocusing (-1) (in the movie we have used  $\gamma = 1$ ). The domain  $\Omega = [-2\pi, 2\pi]^3$  is triply periodic and initial conditions will here be set as

$$u^0 = 0.1 \exp(-\mathbf{x} \cdot \mathbf{x}), \quad (5.16)$$

$$u_t^0 = 0. \quad (5.17)$$

We will solve these equations using a mixed formulation and a spectral Galerkin method. The mixed formulation reads

$$\frac{\partial f}{\partial t} = \nabla^2 u - \gamma(u - u|u|^2), \quad (5.18)$$

$$\frac{\partial u}{\partial t} = f. \quad (5.19)$$

The energy of the solution can be computed as

$$E(u) = \int_{\Omega} \left( \frac{1}{2} f^2 + \frac{1}{2} |\nabla u|^2 + \gamma \left( \frac{1}{2} u^2 - \frac{1}{4} u^4 \right) \right) dx \quad (5.20)$$

and it is crucial that this energy remains constant in time.

The movie (*The nonlinear Klein-Gordon equation*) is showing the solution  $u$ , computed with the code shown in the bottom of Sec. *Complete solver*.

### Spectral Galerkin formulation

The PDEs in (5.18) and (5.19) can be solved with many different numerical methods. We will here use the `shenfun` software and this software makes use of the spectral Galerkin method. Being a Galerkin method, we need to reshape the governing equations into proper variational forms, and this is done by multiplying (5.18) and (5.19) with the complex conjugate of proper test functions and then integrating over the domain. To this end we use testfunctions  $g \in V(\Omega)$  with Eq. (5.18) and  $v \in V(\Omega)$  with Eq. (5.19), where  $V(\omega)$  is a suitable function space, and obtain

$$\frac{\partial}{\partial t} \int_{\Omega} f \bar{g} w dx = \int_{\Omega} (\nabla^2 u - \gamma(u - u|u|^2)) \bar{g} w dx, \quad (5.21)$$

$$\frac{\partial}{\partial t} \int_{\Omega} u \bar{v} w dx = \int_{\Omega} f \bar{v} w dx. \quad (5.22)$$

Note that the overline is used to indicate a complex conjugate, and  $w$  is a weight function associated with the test functions. The functions  $f$  and  $u$  are now to be considered as trial functions, and the integrals over the domain are often referred to as inner products. With inner product notation

$$(u, v) = \int_{\Omega} u \bar{v} w dx.$$

and an integration by parts on the Laplacian, the variational problem can be formulated as:

$$\frac{\partial}{\partial t} (f, g) = -(\nabla u, \nabla g) - \gamma (u - u|u|^2, g), \quad (5.23)$$

$$\frac{\partial}{\partial t} (u, v) = (f, v). \quad (5.24)$$

The time and space discretizations are still left open. There are numerous different approaches that one could take for discretizing in time, and the first two terms on the right hand side of (5.23) can easily be treated implicitly as well as explicitly. However, the approach we will follow in Sec. (*Runge-Kutta integrator*) is a fully explicit 4th order *Runge-Kutta* method.

### Discretization

To find a numerical solution we need to discretize the continuous problem (5.23) and (5.24) in space as well as time. Since the problem is triply periodic, Fourier exponentials are normally the best choice for trial and test functions, and as such we use basis functions

$$\phi_l(x) = e^{ilx}, \quad -\infty < l < \infty, \quad (5.25)$$

where  $l$  is the wavenumber, and  $\underline{l} = \frac{2\pi}{L}l$  is the scaled wavenumber, scaled with domain length  $L$  (here  $4\pi$ ). Since we want to solve these equations on a computer, we need to choose a finite number of test functions. A function space  $V^N$  can be defined as

$$V^N(x) = \text{span}\{\phi_l(x)\}_{l \in \mathbf{l}}, \quad (5.26)$$

where  $N$  is chosen as an even positive integer and  $\mathbf{l} = (-N/2, -N/2 + 1, \dots, N/2 - 1)$ . And now, since  $\Omega$  is a three-dimensional domain, we can create tensor products of such bases to get, e.g., for three dimensions

$$W^N(x, y, z) = V^N(x) \otimes V^N(y) \otimes V^N(z), \quad (5.27)$$

where  $\mathbf{N} = (N, N, N)$ . Obviously, it is not necessary to use the same number ( $N$ ) of basis functions for each direction, but it is done here for simplicity. A 3D tensor product basis function is now defined as

$$\Phi_{lmn}(x, y, z) = e^{i\underline{l}x} e^{i\underline{m}y} e^{i\underline{n}z} = e^{i(\underline{l}x + \underline{m}y + \underline{n}z)} \quad (5.28)$$

where the indices for  $y$ - and  $z$ -direction are  $\underline{m} = \frac{2\pi}{L}m$ ,  $\underline{n} = \frac{2\pi}{L}n$ , and  $\mathbf{m}$  and  $\mathbf{n}$  are the same as  $\mathbf{l}$  due to using the same number of basis functions for each direction. One distinction, though, is that for the  $z$ -direction expansion coefficients are only stored for  $n = (0, 1, \dots, N/2)$  due to Hermitian symmetry (real input data).

We now look for solutions of the form

$$u(x, y, z, t) = \sum_{n=-N/2}^{N/2-1} \sum_{m=-N/2}^{N/2-1} \sum_{l=-N/2}^{N/2-1} \hat{u}_{lmn}(t) \Phi_{lmn}(x, y, z). \quad (5.29)$$

The expansion coefficients  $\hat{\mathbf{u}} = \{\hat{u}_{lmn}(t)\}_{(l,m,n) \in \mathbf{l} \times \mathbf{m} \times \mathbf{n}}$  can be related directly to the solution  $u(x, y, z, t)$  using Fast Fourier Transforms (FFTs) if we are satisfied with obtaining the solution in quadrature points corresponding to

$$x_i = \frac{4\pi i}{N} - 2\pi \quad \forall i \in \mathbf{i}, \text{ where } \mathbf{i} = (0, 1, \dots, N-1), \quad (5.30)$$

$$y_j = \frac{4\pi j}{N} - 2\pi \quad \forall j \in \mathbf{j}, \text{ where } \mathbf{j} = (0, 1, \dots, N-1), \quad (5.31)$$

$$z_k = \frac{4\pi k}{N} - 2\pi \quad \forall k \in \mathbf{k}, \text{ where } \mathbf{k} = (0, 1, \dots, N-1). \quad (5.32)$$

Note that these points are different from the standard (like  $2\pi j/N$ ) since the domain is set to  $[-2\pi, 2\pi]^3$  and not the more common  $[0, 2\pi]^3$ . We have

$$\mathbf{u} = \mathcal{F}_k^{-1} (\mathcal{F}_j^{-1} (\mathcal{F}_i^{-1} (\hat{\mathbf{u}}))) \quad (5.33)$$

with  $\mathbf{u} = \{u(x_i, y_j, z_k)\}_{(i,j,k) \in \mathbf{i} \times \mathbf{j} \times \mathbf{k}}$  and where  $\mathcal{F}_i^{-1}$  is the inverse Fourier transform along the direction of index  $i$ , for all  $(j, k) \in \mathbf{j} \times \mathbf{k}$ . Note that the three inverse FFTs are performed sequentially, one direction at the time, and that there is no scaling factor due to the definition used for the inverse [Fourier transform](#)

$$u(x_j) = \sum_{l=-N/2}^{N/2-1} \hat{u}_l e^{i\underline{l}x_j}, \quad \forall j \in \mathbf{j}. \quad (5.34)$$

Note that this differs from the definition used by, e.g., [Numpy](#).

The inner products used in Eqs. (5.23), (5.24) may be computed using forward FFTs. However, there is a tiny detail that deserves a comment. The regular Fourier inner product is given as

$$\int_0^L e^{i\underline{k}x} e^{-i\underline{l}x} dx = L \delta_{kl}$$

where a weight function is chosen as  $w(x) = 1$  and  $\delta_{kl}$  equals unity for  $k = l$  and zero otherwise. In Shenfun we choose instead to use a weight function  $w(x) = 1/L$ , such that the weighted inner product integrates to unity:

$$\int_0^L e^{ikx} e^{-ilx} \frac{1}{L} dx = \delta_{kl}.$$

With this weight function the scalar product and the forward transform are the same and we obtain:

$$(u, \Phi_{lmn}) = \hat{u}_{lmn} = \left(\frac{1}{N}\right)^3 \mathcal{F}_l(\mathcal{F}_m(\mathcal{F}_n(u))) \quad \forall (l, m, n) \in \mathbf{l} \times \mathbf{m} \times \mathbf{n}, \quad (5.35)$$

From this we see that the variational forms (5.23) and (5.24) may be written in terms of the Fourier transformed quantities  $\hat{u}$  and  $\hat{f}$ . Expanding the exact derivatives of the nabla operator, we have

$$(\nabla u, \nabla v) = (\underline{l}^2 + \underline{m}^2 + \underline{n}^2) \hat{u}_{lmn}, \quad (5.36)$$

$$(u, v) = \hat{u}_{lmn}, \quad (5.37)$$

$$(u|u|^2, v) = \widehat{u|u|^2}_{lmn} \quad (5.38)$$

and as such the equations to be solved for each wavenumber can be found directly as

$$\frac{\partial \hat{f}_{lmn}}{\partial t} = \left( -(\underline{l}^2 + \underline{m}^2 + \underline{n}^2 + \gamma) \hat{u}_{lmn} + \gamma \widehat{u|u|^2}_{lmn} \right), \quad (5.39)$$

$$\frac{\partial \hat{u}_{lmn}}{\partial t} = \hat{f}_{lmn}. \quad (5.40)$$

There is more than one way to arrive at these equations. Taking the 3D Fourier transform of both equations (5.18) and (5.19) is one obvious way. With the Python module `shenfun`, one can work with the inner products as seen in (5.23) and (5.24), or the Fourier transforms directly. See for example Sec. *Runge-Kutta integrator* for how  $(\nabla u, \nabla v)$  can be implemented. In short, `shenfun` contains all the tools required to work with the spectral Galerkin method, and we will now see how `shenfun` can be used to solve the Klein-Gordon equation.

For completion, we note that the discretized problem to solve can be formulated with the Galerkin method as: for all  $t > 0$ , find  $(f, u) \in W^N \times W^N$  such that

$$\frac{\partial}{\partial t}(f, g) = -(\nabla u, \nabla g) - \gamma (u - u|u|^2, g), \quad (5.41)$$

$$\frac{\partial}{\partial t}(u, v) = (f, v) \quad \forall (g, v) \in W^N \times W^N. \quad (5.42)$$

where  $u(x, y, z, 0)$  and  $f(x, y, z, 0)$  are given as the initial conditions according to Eq. (5.15).

## 5.2.2 Implementation

To solve the Klein-Gordon equations we need to make use of the Fourier bases in `shenfun`, and these base are found in submodule `shenfun.fourier.bases`. The triply periodic domain allows for Fourier in all three directions, and we can as such create one instance of this base class using `Basis()` with family `Fourier` for each direction. However, since the initial data are real, we can take advantage of Hermitian symmetries and thus make use of a real to complex class for one (but only one) of the directions, by specifying `dtype='d'`. We can only make use of the real-to-complex class for the direction that we choose to transform first with the forward FFT, and the reason is obviously that the output from a forward transform of real data is now complex. We may start implementing the solver as follows

```

from shenfun import *
from mpi4py import MPI
import numpy as np

# Set size of discretization
N = (32, 32, 32)

# Create bases
K0 = Basis(N[0], 'F', domain=(-2*np.pi, 2*np.pi), dtype='D')
K1 = Basis(N[1], 'F', domain=(-2*np.pi, 2*np.pi), dtype='D')
K2 = Basis(N[2], 'F', domain=(-2*np.pi, 2*np.pi), dtype='d')

```

We now have three instances `K0`, `K1` and `K2`, corresponding to the space (5.26), that each can be used to solve one-dimensional problems. However, we want to solve a 3D problem, and for this we need a tensor product space, like (5.27), created as a tensor product of these three spaces

```

# Create communicator
comm = MPI.COMM_WORLD
T = TensorProductSpace(comm, (K0, K1, K2), **{'planner_effort':
                                             'FFTW_MEASURE'})

```

Here the `planner_effort`, which is a flag used by `FFTW`, is optional. Possible choices are from the list (`FFTW_ESTIMATE`, `FFTW_MEASURE`, `FFTW_PATIENT`, `FFTW_EXHAUSTIVE`), and the flag determines how much effort `FFTW` puts in looking for an optimal algorithm for the current platform. Note that it is also possible to use `FFTW wisdom` with `shenfun`, and as such, for production, one may perform exhaustive planning once and then simply import the result of that planning later, as `wisdom`.

The `TensorProductSpace` instance `T` contains pretty much all we need for computing inner products or fast transforms between real and wavenumber space. However, since we are going to solve for a mixed system, it is convenient to also use the `MixedTensorProductSpace` class

```

TT = MixedTensorProductSpace([T, T])

```

We need containers for the solution as well as intermediate work arrays for, e.g., the Runge-Kutta method. Arrays are created as

```

uf = Array(TT)           # Solution array in physical space
u, f = uf               # Split solution array by creating two views u and f
duf = Function(TT)     # Array for right hand sides
du, df = duf           # Split into views
uf_hat = Function(TT)  # Solution in spectral space
uf_hat0 = Function(TT) # Work array 1
uf_hat1 = Function(TT) # Work array 2
u_hat, f_hat = uf_hat  # Split into views

```

The `Array` class is a subclass of Numpy's `ndarray`, without much more functionality than constructors that return arrays of the correct shape according to the basis used in the construction. The `Array` represents the left hand side of (5.29), evaluated on the quadrature mesh. A different type of array is returned by the `Function` class, that subclasses both Numpy's `ndarray` as well as an internal `BasisFunction` class. An instance of the `Function` represents the entire spectral Galerkin function (5.29). As such, it can be used in complex variational linear forms. For example, if you want to compute the partial derivative  $\partial u / \partial x$ , then this may be achieved by projection, i.e., find  $u_x \in V^N$  such that  $(u_x - \partial u / \partial x, v) = 0$ , for all  $v \in V^N$ . This projection may be easily computed in `shenfun` using

```

ux = project(Dx(u_hat, 0, 1), T)

```

The following code, on the other hand, will raise an error since you cannot take the derivative of an interpolated `Array` `u`, only a `Function`

```

try:
    project(Dx(u, 0, 1), T)
except AssertionError:
    print("AssertionError: Dx not for Arrays")

```

## Initialization

The solution array `uf` and its transform `uf_hat` need to be initialized according to Eq. (5.15). To this end it is convenient (but not required, we could just as easily use Numpy for this as well) to use `Sympy`, which is a Python library for symbolic mathematics.

```

from sympy import symbols, exp, lambdify

x, y, z = symbols("x,y,z")
ue = 0.1*exp(-(x**2 + y**2 + z**2))
ul = lambdify((x, y, z), ue, 'numpy')
X = T.local_mesh(True)
u[:] = Array(T, buffer=ul(*X))
u_hat = T.forward(u, u_hat)

```

Here `X` is a list of the three mesh coordinates (`x`, `y`, `z`) local to the current processor. Each processor has its own part of the computational mesh, and the distribution is handled during the creation of the `TensorProductSpace` class instance `T`. There is no need to do anything about the `f/f_hat` arrays since they are already initialized by default to zero. Note that calling the `ul` function with the argument `*X` is the same as calling with `X[0]`, `X[1]`, `X[2]`.

## Runge-Kutta integrator

A fourth order explicit Runge-Kutta integrator requires only a function that returns the right hand sides of (5.39) and (5.40). Such a function can be implemented as

```

# focusing (+1) or defocusing (-1)
gamma = 1
uh = TrialFunction(T)
vh = TestFunction(T)
k2 = -(inner(grad(vh), grad(uh)).scale + gamma)

def compute_rhs(duf_hat, uf_hat, up, Tp, w0):
    duf_hat.fill(0)
    u_hat, f_hat = uf_hat
    du_hat, df_hat = duf_hat
    df_hat[:] = k2*u_hat
    up = Tp.backward(u_hat, up)
    df_hat += Tp.forward(gamma*up**3, w0)
    du_hat[:] = f_hat
    return duf_hat

```

The code is fairly self-explanatory. `k2` represents the coefficients in front of the linear  $\hat{u}$  in (5.39). The output array is `duf_hat`, and the input array is `uf_hat`, whereas `up` and `w0` are work arrays. The array `duf_hat` contains the right hand sides of both (5.39) and (5.40), where the linear and nonlinear terms are recognized in the code as comments (1) and (2). The array `uf_hat` contains the solution at initial and intermediate Runge-Kutta steps.

With a function that returns the right hand side in place, the actual integrator can be implemented as

```

w0 = Function(T)
a = [1./6., 1./3., 1./3., 1./6.]           # Runge-Kutta parameter
b = [0.5, 0.5, 1.]                       # Runge-Kutta parameter
t = 0
dt = 0.01
end_time = 1.0
while t < end_time-1e-8:
    t += dt
    uf_hat1[:] = uf_hat0[:] = uf_hat
    for rk in range(4):
        duf = compute_rhs(duf, uf_hat, u, T, w0)
        if rk < 3:
            uf_hat[:] = uf_hat0 + b[rk]*dt*duf
        uf_hat1 += a[rk]*dt*duf
    uf_hat[:] = uf_hat1

```

### Complete solver

A complete solver is given below, with intermediate plotting of the solution and intermediate computation of the total energy. Note that the total energy is unchanged to 8 decimal points at  $t = 100$ .

```

from sympy import symbols, exp, lambdify
import numpy as np
import matplotlib.pyplot as plt
from mpi4py import MPI
from time import time
from shenfun import *

comm = MPI.COMM_WORLD
rank = comm.Get_rank()

# Use sympy to set up initial condition
x, y, z = symbols("x,y,z")
ue = 0.1*exp(-(x**2 + y**2 + z**2))
ul = lambdify((x, y, z), ue, 'numpy')

# Size of discretization
N = (64, 64, 64)

# Defocusing or focusing
gamma = 1

K0 = Basis(N[0], 'F', domain=(-2*np.pi, 2*np.pi), dtype='D')
K1 = Basis(N[1], 'F', domain=(-2*np.pi, 2*np.pi), dtype='D')
K2 = Basis(N[2], 'F', domain=(-2*np.pi, 2*np.pi), dtype='d')
T = TensorProductSpace(comm, (K0, K1, K2), slab=False,
                        **{'planner_effort': 'FFTW_MEASURE'})

TT = MixedTensorProductSpace([T, T])

X = T.local_mesh(True)
uf = Array(TT)
u, f = uf[:]
up = Array(T)
duf = Function(TT)
du, df = duf[:]

```

(continues on next page)

(continued from previous page)

```

uf_hat = Function(TT)
uf_hat0 = Function(TT)
uf_hat1 = Function(TT)
w0 = Function(T)
u_hat, f_hat = uf_hat[:]

# initialize (f initialized to zero, so all set)
u[:] = ul(*X)
u_hat = T.forward(u, u_hat)

uh = TrialFunction(T)
vh = TestFunction(T)
k2 = -inner(grad(vh), grad(uh)).scale - gamma

count = 0
def compute_rhs(duf_hat, uf_hat, up, T, w0):
    global count
    count += 1
    duf_hat.fill(0)
    u_hat, f_hat = uf_hat[:]
    du_hat, df_hat = duf_hat[:]
    df_hat[:] = k2*u_hat
    up = T.backward(u_hat, up)
    df_hat += T.forward(gamma*up**3, w0)
    du_hat[:] = f_hat
    return duf_hat

def energy_fourier(comm, a):
    result = 2*np.sum(abs(a[... , 1:-1])**2) + np.sum(abs(a[... , 0])**2) + np.
    ↪sum(abs(a[... , -1])**2)
    result = comm.allreduce(result)
    return result

# Integrate using a 4th order Rung-Kutta method
a = [1./6., 1./3., 1./3., 1./6.] # Runge-Kutta parameter
b = [0.5, 0.5, 1.] # Runge-Kutta parameter
t = 0.0
dt = 0.005
end_time = 1.
tstep = 0
if rank == 0:
    plt.figure()
    image = plt.contourf(X[1][... , 0], X[0][... , 0], u[... , 16], 100)
    plt.draw()
    plt.pause(1e-4)
t0 = time()
K = np.array(T.local_wavenumbers(True, True, True))
TV = VectorTensorProductSpace([T, T, T])
gradu = Array(TV)
while t < end_time-1e-8:
    t += dt
    tstep += 1
    uf_hat1[:] = uf_hat0[:] = uf_hat
    for rk in range(4):
        duf = compute_rhs(duf, uf_hat, up, T, w0)
        if rk < 3:

```

(continues on next page)

(continued from previous page)

```

        uf_hat[:] = uf_hat0 + b[rk]*dt*duf
        uf_hat1 += a[rk]*dt*duf
    uf_hat[:] = uf_hat1

    if timestep % 100 == 0:
        uf = TT.backward(uf_hat, uf)
        ekin = 0.5*energy_fourier(T.comm, f_hat)
        es = 0.5*energy_fourier(T.comm, 1j*K*u_hat)
        eg = gamma*np.sum(0.5*u**2 - 0.25*u**4)/np.prod(np.array(N))
        eg = comm.allreduce(eg)
        gradu = TV.backward(1j*K*u_hat, gradu)
        ep = comm.allreduce(np.sum(f*gradu)/np.prod(np.array(N)))
        ea = comm.allreduce(np.sum(np.array(X)*(0.5*f**2 + 0.5*gradu**2
            - (0.5*u**2 - 0.25*u**4)*f))/np.prod(np.array(N)))

        if rank == 0:
            image.ax.clear()
            image.ax.contourf(X[1][..., 0], X[0][..., 0], u[..., 16], 100)
            plt.pause(1e-6)
            plt.savefig('Klein_Gordon_{}_real_{}.png'.format(N[0], timestep))
            print("Time = %2.2f Total energy = %2.8e Linear momentum %2.8e Angular_
↪momentum %2.8e" %(t, ekin+es+eg, ep, ea))
            comm.barrier()

print("Time ", time()-t0)

```

## 5.3 Demo - 3D Poisson's equation

**Authors** Mikael Mortensen (mikaem at math.uio.no)

**Date** Nov 14, 2019

*Summary.* This is a demonstration of how the Python module `shenfun` can be used to solve a 3D Poisson equation in a 3D tensor product domain that has homogeneous Dirichlet boundary conditions in one direction and periodicity in the remaining two. The solver described runs with MPI without any further considerations required from the user. Spectral convergence, as shown in Figure *Convergence of 3D Poisson solvers for both Legendre and Chebyshev modified basis function*, is demonstrated. The demo is implemented in a single Python file `dirichlet_poisson3D.py`, and the numerical method is described in more detail by J. Shen [She94] and [She95].

### 5.3.1 Model problem

#### Poisson equation

The Poisson equation is given as

$$\nabla^2 u(\mathbf{x}) = f(\mathbf{x}) \quad \text{for } \mathbf{x} = (x, y, z) \in \Omega, \quad (5.43)$$

$$u(\pm 1, y, z) = 0, \quad (5.44)$$

$$u(x, 2\pi, z) = u(x, 0, z), \quad (5.45)$$

$$u(x, y, 2\pi) = u(x, y, 0), \quad (5.46)$$

where  $u(\mathbf{x})$  is the solution and  $f(\mathbf{x})$  is a function. The domain  $\Omega = [-1, 1] \times [0, 2\pi]^2$ .

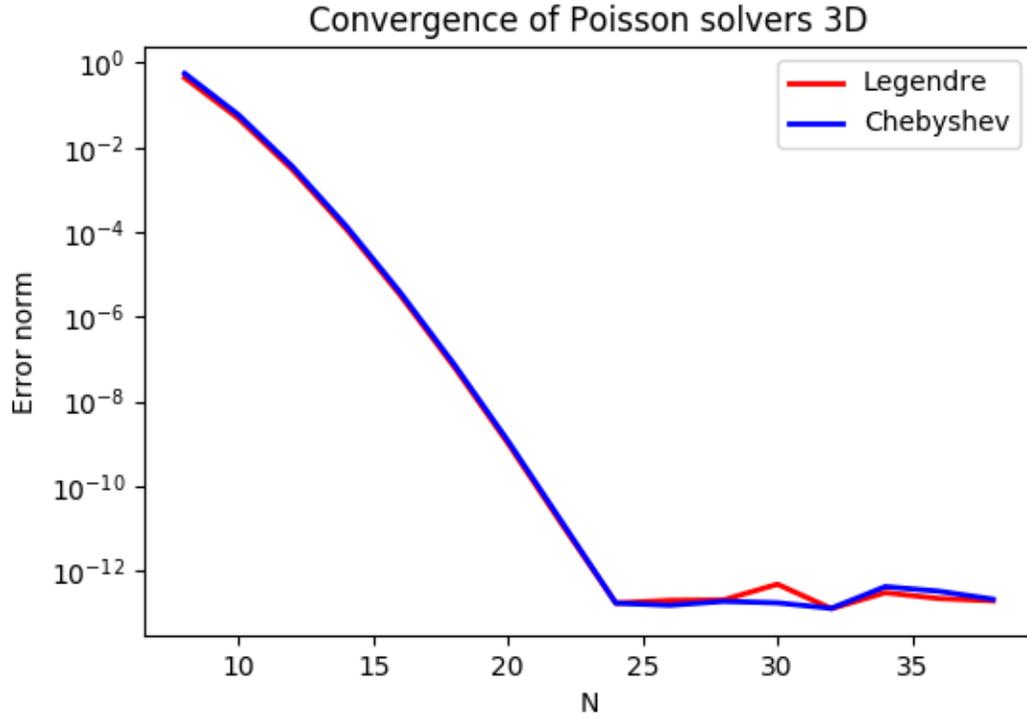


Fig. 2: Convergence of 3D Poisson solvers for both Legendre and Chebyshev modified basis function

To solve Eq. (5.43) with the Galerkin method we need smooth basis functions,  $v(x)$ , that live in the Hilbert space  $H^1(\Omega)$  and that satisfy the given boundary conditions. To this end we will use one basis function for the  $x$ -direction,  $\mathcal{X}(x)$ , one for the  $y$ -direction,  $\mathcal{Y}(y)$ , and one for the  $z$ -direction,  $\mathcal{Z}(z)$ . And then we create three-dimensional basis functions like

$$v(x, y, z) = \mathcal{X}(x)\mathcal{Y}(y)\mathcal{Z}(z).$$

The basis functions  $\mathcal{Y}(y)$  and  $\mathcal{Z}(z)$  are chosen as Fourier exponentials, since these functions are periodic. Likewise, the basis functions  $\mathcal{X}(x)$  are chosen as modified Legendre or Chebyshev polynomials, using  $\phi_l(x)$  to refer to either one

$$\mathcal{X}_l(x) = \phi_l(x) - \phi_{l+2}(x), \forall l \in \mathbf{l}^{N_0}, \quad (5.47)$$

$$\mathcal{Y}_m(y) = e^{imy}, \forall m \in \mathbf{m}^{N_1}, \quad (5.48)$$

$$\mathcal{Z}_n(z) = e^{inz}, \forall n \in \mathbf{n}^{N_2}, \quad (5.49)$$

where the size of the discretized problem is  $\mathbf{N} = (N_0, N_1, N_2)$ ,  $\mathbf{l}^{N_0} = (0, 1, \dots, N_0-3)$ ,  $\mathbf{m}^{N_1} = (-N_1/2, -N_1/2+1, \dots, N_1/2-1)$  and  $\mathbf{n}^{N_2} = (-N_2/2, -N_2/2+1, \dots, N_2/2-1)$ . However, due to [Hermitian symmetry](#), we only store  $N_2/2+1$  wavenumbers in the  $z$ -direction, such that  $\mathbf{n}^{N_2} = (0, 1, \dots, N_2/2)$ . We refer to the Cartesian wavenumber mesh on vector form as  $\mathbf{k}$ :

$$\mathbf{k} = \{(l, m, n) \mid (l, m, n) \in \mathbf{l}^{N_0} \times \mathbf{m}^{N_1} \times \mathbf{n}^{N_2}\}.$$

We have the one-dimensional spaces

$$V^{N_0} = \text{span}\{\mathcal{X}_l\}_{l \in \mathbf{l}^{N_0}}, \quad (5.50)$$

$$V^{N_1} = \text{span}\{\mathcal{Y}_m\}_{m \in \mathbf{m}^{N_1}}, \quad (5.51)$$

$$V^{N_2} = \text{span}\{\mathcal{Z}_n\}_{n \in \mathbf{n}^{N_2}}, \quad (5.52)$$

and from these we create a tensor product space  $W^N(\mathbf{x})$

$$W^N(\mathbf{x}) = V^{N_0}(x) \otimes V^{N_1}(y) \otimes V^{N_2}(z). \quad (5.53)$$

And then we look for discrete solutions  $u \in W^N$  like

$$u(\mathbf{x}) = \sum_{l \in \mathbf{l}^{N_0}} \sum_{m \in \mathbf{m}^{N_1}} \sum_{n \in \mathbf{n}^{N_2}} \hat{u}_{lmn} \mathcal{X}_l(x) \mathcal{Y}_m(y) \mathcal{Z}_n(z), \quad (5.54)$$

$$= \sum_{\mathbf{k} \in \mathbf{k}} \hat{u}_{\mathbf{k}} v_{\mathbf{k}}(\mathbf{x}), \quad (5.55)$$

where  $\hat{u}_{lmn}$  are components of the expansion coefficients for  $u$  and the second form,  $\{\hat{u}_{\mathbf{k}}\}_{\mathbf{k} \in \mathbf{k}}$ , is a shorter, simplified notation, with sans-serif  $\mathbf{k} = (l, m, n)$ . The expansion coefficients are the unknowns in the spectral Galerkin method.

We now formulate a variational problem using the Galerkin method: Find  $u \in W^N$  such that

$$\int_{\Omega} \nabla^2 u \bar{v} w \, d\mathbf{x} = \int_{\Omega} f \bar{v} w \, d\mathbf{x} \quad \forall v \in W^N. \quad (5.56)$$

Here  $d\mathbf{x} = dx dy dz$ , and the overline represents a complex conjugate, which is needed here because the Fourier exponentials are complex functions. The weighted integrals, weighted by  $w(\mathbf{x})$ , are called inner products, and a common notation is

$$\int_{\Omega} u \bar{v} w \, d\mathbf{x} = \langle u, v \rangle_w. \quad (5.57)$$

The integral can either be computed exactly, or with quadrature. The advantage of the latter is that it is generally faster, and that non-linear terms may be computed just as quickly as linear. For a linear problem, it does not make much of a difference, if any at all. Approximating the integral with quadrature, we obtain

$$\int_{\Omega} u \bar{v} w \, d\mathbf{x} \approx \langle u, v \rangle_w^N, \quad (5.58)$$

$$\approx \sum_{i=0}^{N_0-1} \sum_{j=0}^{N_1-1} \sum_{k=0}^{N_2-1} u(x_i, y_j, z_k) \bar{v}(x_i, y_j, z_k) w(x_i, y_j, z_k), \quad (5.59)$$

where  $w(\mathbf{x})$  now are the quadrature weights. The quadrature points  $\{x_i\}_{i=0}^{N_0-1}$  are specific to the chosen basis, and even within basis there are two different choices based on which quadrature rule is selected, either Gauss or Gauss-Lobatto. The quadrature points for the Fourier bases are the uniform  $\{y_j\}_{j=0}^{N_1-1} = 2\pi j/N_1$  and  $\{z_k\}_{k=0}^{N_2-1} = 2\pi k/N_2$ .

Inserting for test function (5.54) and trialfunction  $v_{pqr} = \mathcal{X}_p \mathcal{Y}_q \mathcal{Z}_r$  on the left hand side of (5.56), we get

$$\begin{aligned} \langle \nabla^2 u, v \rangle_w^N &= \left\langle \nabla^2 \sum_{l \in \mathbf{l}^{N_0}} \sum_{m \in \mathbf{m}^{N_1}} \sum_{n \in \mathbf{n}^{N_2}} \hat{u}_{lmn} \mathcal{X}_l \mathcal{Y}_m \mathcal{Z}_n, \mathcal{X}_p \mathcal{Y}_q \mathcal{Z}_r \right\rangle_w^N, \\ &= \left[ \left( \mathcal{X}_l'' \right)_w^N - (m^2 + n^2) \left( \mathcal{X}_l \right)_w^N \right] \delta_{mq} \delta_{nr} \hat{u}_{lmn}, \\ &= (a_{pl} - (m^2 + n^2) b_{pl}) \hat{u}_{lqr}, \end{aligned}$$

where the notation  $(\cdot, \cdot)_w^{N_0}$

$$b_{pl} = \left( \mathcal{X}_l, \mathcal{X}_p \right)_w^{N_0} = \sum_{i=0}^{N_0-1} \mathcal{X}_l(x_i) \mathcal{X}_p(x_i) w(x_i), \quad (5.60)$$

is used to represent an  $L_2$  inner product along only the first, nonperiodic, direction. The delta functions above come from integrating over the two periodic directions, where we use constant weight functions  $w = 1/(2\pi)$  in the inner products

$$\int_0^{2\pi} \mathcal{Y}_m(y) \overline{\mathcal{Y}_q(y)} \frac{1}{2\pi} dy = \delta_{mq}, \quad (5.61)$$

$$\int_0^{2\pi} \mathcal{Z}_n(z) \overline{\mathcal{Z}_r(z)} \frac{1}{2\pi} dz = \delta_{nr}, \quad (5.62)$$

The Kronecker delta-function  $\delta_{ij}$  is one for  $i = j$  and zero otherwise.

The right hand side of Eq. (5.56) is computed as

$$\tilde{f}_{pqr} = \langle f, \mathcal{X}_p \mathcal{Y}_q \mathcal{Z}_r \rangle_w^N, \quad (5.63)$$

where a tilde is used because this is not a complete transform of the function  $f$ , but only an inner product.

The linear system of equations to solve for the expansion coefficients can now be found as follows

$$(a_{lj} - (m^2 + n^2)b_{lj}) \hat{u}_{jmn} = \tilde{f}_{lmn} \quad \forall (l, m, n) \in \mathbf{k}. \quad (5.64)$$

Now, when  $\hat{u} = \{\hat{u}_k\}_{k \in \mathbf{k}}$  is found by solving this linear system over the entire computational mesh, it may be transformed to real space  $u(\mathbf{x})$  using (5.54). Note that the matrices  $A \in \mathbb{R}^{N_0-3 \times N_0-3}$  and  $B \in \mathbb{R}^{N_0-3 \times N_0-3}$  differ for Legendre or Chebyshev bases, but for either case they have a special structure that allows for a solution to be found very efficiently in the order of  $\mathcal{O}(N_0 - 3)$  operations given  $m$  and  $n$ , see [She94] and [She95]. Fast solvers for (5.64) are implemented in `shenfun` for both bases.

## Method of manufactured solutions

In this demo we will use the method of manufactured solutions to demonstrate spectral accuracy of the `shenfun` bases. To this end we choose a smooth analytical function that satisfies the given boundary conditions:

$$u_e(x, y, z) = (\cos(4x) + \sin(2y) + \sin(4z)) (1 - x^2). \quad (5.65)$$

Sending  $u_e$  through the Laplace operator, we obtain the right hand side

$$\nabla^2 u_e(x, y, z) = -16(1 - x^2) \cos(4x) + 16x \sin(4x) - 2 \cos(4x) - (1 - x^2)(4 \sin(2y) + 16 \sin(4z)). \quad (5.66)$$

Now, setting  $f_e(\mathbf{x}) = \nabla^2 u_e(\mathbf{x})$  and solving for  $\nabla^2 u(\mathbf{x}) = f_e(\mathbf{x})$ , we can compare the numerical solution  $u(\mathbf{x})$  with the analytical solution  $u_e(\mathbf{x})$  and compute error norms.

## 5.3.2 Implementation

### Preamble

We will solve the Poisson problem using the `shenfun` Python module. The first thing needed is then to import some of this module's functionality plus some other helper modules, like `Numpy` and `Sympy`:

```
from sympy import symbols, cos, sin, exp, lambdify
import numpy as np
from shenfun.tensorproductspace import TensorProductSpace
from shenfun import inner, div, grad, TestFunction, TrialFunction, Function, \
    project, Dx, Basis
from mpi4py import MPI
```

We use `Sympy` for the manufactured solution and `Numpy` for testing. `MPI` for Python (`mpi4py`) is required for running the solver with `MPI`.

## Manufactured solution

The exact solution  $u_e(x, y, z)$  and the right hand side  $f_e(x, y, z)$  are created using Sympy as follows

```
x, y, z = symbols("x,y,z")
ue = (cos(4*x) + sin(2*y) + sin(4*z))*(1-x**2)
fe = ue.diff(x, 2) + ue.diff(y, 2) + ue.diff(z, 2)

# Lambdify for faster evaluation
ul = lambdify((x, y, z), ue, 'numpy')
fl = lambdify((x, y, z), fe, 'numpy')
```

These solutions are now valid for a continuous domain. The next step is thus to discretize, using the computational mesh

$$(x_i, y_j, z_k) \forall (i, j, k) \in [0, 1, \dots, N_0 - 1] \times [0, 1, \dots, N_1 - 1] \times [0, 1, \dots, N_2 - 1]$$

and a finite number of basis functions.

Note that it is not mandatory to use Sympy for the manufactured solution. Since the solution is known (5.66), we could just as well simply use Numpy to compute  $f_e$ . However, with Sympy it is much easier to experiment and quickly change the solution.

## Discretization and MPI

We create three bases with given size, one for each dimension of the problem. From these three bases a TensorProductSpace is created.

```
# Size of discretization
N = [14, 15, 16]

SD = Basis(N[0], 'Chebyshev', bc=(0, 0))
#SD = Basis(N[0], 'Legendre', bc=(0, 0))
K1 = Basis(N[1], 'Fourier', dtype='D')
K2 = Basis(N[2], 'Fourier', dtype='d')
T = TensorProductSpace(comm, (SD, K1, K2), axes=(0, 1, 2))
X = T.local_mesh()
```

Note that we can either choose a Legendre or a Chebyshev basis for the nonperiodic direction. The TensorProductSpace class takes an MPI communicator as first argument and the computational mesh is distributed internally using the pencil method. The T.local\_mesh method returns the mesh local to each processor. The axes keyword determines the order of transforms going back and forth between real and spectral space. With axes=(0, 1, 2) and a forward transform (from real space to spectral, i.e., from  $u$  to  $\hat{u}$ ) axis 2 is transformed first and then 1 and 0, respectively.

The manufactured solution is created with Dirichlet boundary conditions in the  $x$ -direction, and for this reason SD is the first basis in T. We could just as well have put the nonperiodic direction along either  $y$ - or  $z$ -direction, though, but this would then require that the order of the transformed axes be changed as well. For example, putting the Dirichlet direction along  $y$ , we would need to create the tensorproductspace as

```
T = TensorProductSpace(comm, (K1, SD, K2), axes=(1, 0, 2))
```

such that the Dirichlet direction is the last to be transformed. The reason for this is that only the Dirichlet direction leads to matrices that need to be inverted (or solved). And for this we need the entire data array along the Dirichlet direction to be local to the processor. If the SD basis is the last to be transformed, then the data will be aligned in this direction, whereas the other two directions may both, or just one of them, be distributed.

Note that  $X$  is a list containing local values of the arrays  $\{x_i\}_{i=0}^{N_0-1}$ ,  $\{y_j\}_{j=0}^{N_1-1}$  and  $\{z_k\}_{k=0}^{N_2-1}$ . For example, using 4 processors and a processor mesh of shape  $2 \times 2$ , then the local slices for each processor in spectral space are

```
>>> print(comm.Get_rank(), T.local_slice())
3 [slice(0, 14, None), slice(8, 15, None), slice(5, 9, None)]
1 [slice(0, 14, None), slice(0, 8, None), slice(5, 9, None)]
2 [slice(0, 14, None), slice(8, 15, None), slice(0, 5, None)]
0 [slice(0, 14, None), slice(0, 8, None), slice(0, 5, None)]
```

where the global shape is  $N = (14, 15, 9)$  after taking advantage of Hermitian symmetry in the  $z$ -direction. So, all processors have the complete first dimension available locally, as they should. Furthermore, processor three owns the slices from  $8 : 15$  and  $5 : 9$  along axes  $y$  and  $z$ , respectively. Processor 2 owns slices  $0 : 8$  and  $0 : 5$  etc. In real space the mesh is distributed differently. First of all the global mesh shape is  $N = (14, 15, 16)$ , and it is distributed along the first two dimensions. The local slices can be inspected as

```
>>> print(comm.Get_rank(), T.local_slice(False))
0 [slice(0, 7, None), slice(0, 8, None), slice(0, 16, None)]
1 [slice(0, 7, None), slice(8, 15, None), slice(0, 16, None)]
2 [slice(7, 14, None), slice(0, 8, None), slice(0, 16, None)]
3 [slice(7, 14, None), slice(8, 15, None), slice(0, 16, None)]
```

Since two directions are distributed, both in spectral and real space, we say that we have a two-dimensional decomposition (here a  $2 \times 2$  shaped processor mesh) and the MPI distribution is of type *pencil*. It is also possible to choose a *slab* decomposition, where only one dimension of the array is distributed. This choice needs to be made when creating the tensorproductspace as

```
T = TensorProductSpace(comm, (SD, K1, K2), axes=(0, 1, 2), slab=True)
```

which will lead to a mesh that is distributed along  $x$ -direction in real space and  $y$ -direction in spectral space. The local slices are

```
>>> print(comm.Get_rank(), T.local_slice()) # spectral space
1 [slice(0, 14, None), slice(4, 8, None), slice(0, 9, None)]
2 [slice(0, 14, None), slice(8, 12, None), slice(0, 9, None)]
0 [slice(0, 14, None), slice(0, 4, None), slice(0, 9, None)]
3 [slice(0, 14, None), slice(12, 15, None), slice(0, 9, None)]
>>> print(comm.Get_rank(), T.local_slice(False)) # real space
3 [slice(11, 14, None), slice(0, 15, None), slice(0, 16, None)]
0 [slice(0, 4, None), slice(0, 15, None), slice(0, 16, None)]
2 [slice(8, 11, None), slice(0, 15, None), slice(0, 16, None)]
1 [slice(4, 8, None), slice(0, 15, None), slice(0, 16, None)]
```

Note that the *slab* decomposition is usually the fastest choice. However, the maximum number of processors with *slab* is  $\min\{N_0, N_1\}$ , whereas a *pencil* approach can be used with up to  $\min\{N_1(N_2/2 + 1), N_0N_1\}$  processors.

## Variational formulation

The variational problem (5.56) can be assembled using shenfun's form language, which is perhaps surprisingly similar to FEniCS.

```
u = TrialFunction(T)
v = TestFunction(T)
K = T.local_wavenumbers()
# Get f on quad points
fj = Array(T, buffer=fl(*X))
# Compute right hand side of Poisson equation
```

(continues on next page)

(continued from previous page)

```
f_hat = inner(v, fj)
# Get left hand side of Poisson equation
matrices = inner(v, div(grad(u)))
```

The Laplacian operator is recognized as `div(grad)`. The `matrices` object is a dictionary representing the left hand side of (5.64), and there are two keys: (`ADDmat`, `BDDmat`). The value of `matrices["ADDmat"]` is an object of type `SpectralMatrix`, which is shenfun's type for a matrix. This matrix represents  $A_{lj}$ , see (5.64), and it has an attribute `scale` that is equal to  $(2\pi)^2$  (also see (5.64)). The other key in `matrices` is `BDDmat`, and the value here is a `SpectralMatrix` representing  $B_{lj}$  from (5.64). This matrix has an attribute `scale` that is equal to  $m^2 + n^2$ . This `scale` is stored as a numpy array of shape  $(1, 15, 9)$ , representing the set  $\{m^2 + n^2 : (m, n) \in \mathbf{m}^{N_1} \times \mathbf{n}^{N_2}\}$ . Note that  $\mathbf{n}^{N_2}$  is stored simply as an array of length  $N_2/2 + 1$  (here 9), since the transform in direction  $z$  takes a real signal and transforms it taking advantage of Hermitian symmetry, see `rfft`.

## Solve linear equations

Finally, solve linear equation system and transform solution from spectral  $\hat{u}_k$  vector to the real space  $u(\mathbf{x})$  and then check how the solution corresponds with the exact solution  $u_e$ .

```
# Create Helmholtz linear algebra solver
H = Solver(*matrices)

# Solve and transform to real space
u_hat = Function(T)           # Solution spectral space
u_hat = H(u_hat, f_hat)      # Solve
uq = T.backward(u_hat)

# Compare with analytical solution
uj = ul(*X)
error = comm.reduce(np.linalg.norm(uj-uq)**2)
if comm.Get_rank() == 0:
    print("Error=%2.16e" % (np.sqrt(error)))
```

## Convergence test

A complete solver is given in Sec. *Complete solver*. This solver is created such that it takes in two commandline arguments and prints out the  $L_2$ -errornorm of the solution in the end. We can use this to write a short script that performs a convergence test. The solver is run like

```
>>> python dirichlet_poisson3D.py 32 legendre
Error=6.5955040031498912e-10
```

for a discretization of size  $N = N^3 = 32^3$  and for the Legendre basis. Alternatively, change `legendre` to `chebyshev` for the Chebyshev basis.

We set up the solver to run for a list of  $N = [8, 10, \dots, 38]$ , and collect the errornorms in arrays to be plotted. Such a script can be easily created with the `subprocess` module

```
import subprocess
from numpy import log, array
from matplotlib import pyplot as plt

N = range(8, 40, 2)
error = {}
```

(continues on next page)

(continued from previous page)

```

for basis in ('legendre', 'chebyshev'):
    error[basis] = []
    for i in range(len(N)):
        output = subprocess.check_output("python dirichlet_poisson3D.py {} {}".
→format(N[i], basis), shell=True)
        exec(output) # Error is printed as "Error=%2.16e"%(np.linalg.norm(uj-ua))
        error[basis].append(Error)
        if i == 0:
            print("Error          hmin          r          ")
            print("%2.8e %2.8e %2.8f"%(error[basis][-1], 1./N[i], 0))
        if i > 0:
            print("%2.8e %2.8e %2.8f"%(error[basis][-1], 1./N[i], log(error[basis][-
→1]/error[basis][-2])/log(N[i-1]/N[i])))

```

The error can be plotted using matplotlib, and the generated figure is shown in the summary's Fig. *Convergence of 3D Poisson solvers for both Legendre and Chebyshev modified basis function*. The spectral convergence is evident and we can see that after  $N = 25$  roundoff errors dominate as the error norm trails off around  $10^{-13}$ .

```

plt.figure(figsize=(6, 4))
for basis, col in zip(('legendre', 'chebyshev'), ('r', 'b')):
    plt.semilogy(N, error[basis], col, linewidth=2)
plt.title('Convergence of Poisson solvers 3D')
plt.xlabel('N')
plt.ylabel('Error norm')
plt.legend(('Legendre', 'Chebyshev'))
plt.savefig('poisson3D_errornorm.png')
plt.show()

```

## Complete solver

A complete solver, that can use either Legendre or Chebyshev bases, and any quadrature size chosen as a command-line argument, is shown below.

```
>>> python dirichlet_poisson3D.py 36 legendre
```

or similarly with chebyshev instead of legendre.

```

import sys, os
import importlib
from sympy import symbols, cos, sin, lambdify
import numpy as np
from shenfun import inner, div, grad, TestFunction, TrialFunction, Array, \
    Function, Basis, TensorProductSpace
import time
from mpi4py import MPI
try:
    import matplotlib.pyplot as plt
except ImportError:
    plt = None

comm = MPI.COMM_WORLD

assert len(sys.argv) == 3
assert sys.argv[-1].lower() in ('legendre', 'chebyshev')
assert isinstance(int(sys.argv[-2]), int)

```

(continues on next page)

(continued from previous page)

```

# Collect basis and solver from either Chebyshev or Legendre submodules
family = sys.argv[-1].lower()
base = importlib.import_module('.'.join(('shenfun', family)))
Solver = base.la.Helmholtz

# Use sympy to compute a rhs, given an analytical solution
a = -0
b = 0
x, y, z = symbols("x,y,z")
ue = (cos(4*x) + sin(2*y) + sin(4*z))*(1-z**2) + a*(1 + z)/2. + b*(1 - z)/2.
fe = ue.diff(x, 2) + ue.diff(y, 2) + ue.diff(z, 2)

# Lambdify for faster evaluation
ul = lambdify((x, y, z), ue, 'numpy')
fl = lambdify((x, y, z), fe, 'numpy')

# Size of discretization
N = int(sys.argv[-2])
N = [N, N, N]

SD = Basis(N[0], family=family, bc=(a, b))
K1 = Basis(N[1], family='F', dtype='D')
K2 = Basis(N[2], family='F', dtype='d')
T = TensorProductSpace(comm, (K1, K2, SD), axes=(0, 1, 2), slab=True)
X = T.local_mesh()
u = TrialFunction(T)
v = TestFunction(T)

K = T.local_wavenumbers()

# Get f on quad points
fj = Array(T, buffer=fl(*X))

# Compute right hand side of Poisson equation
f_hat = inner(v, fj)
if family == 'legendre':
    f_hat *= -1.

# Get left hand side of Poisson equation
if family == 'chebyshev':
    matrices = inner(v, div(grad(u)))
else:
    matrices = inner(grad(v), grad(u))

# Create Helmholtz linear algebra solver
H = Solver(*matrices)

# Solve and transform to real space
u_hat = Function(T) # Solution spectral space
t0 = time.time()
u_hat = H(u_hat, f_hat) # Solve
uq = T.backward(u_hat, fast_transform=False)

# Compare with analytical solution
uj = ul(*X)
error = comm.reduce(np.linalg.norm(uj-uq)**2)

```

(continues on next page)

(continued from previous page)

```

if comm.Get_rank() == 0:
    print("Error=%2.16e" % (np.sqrt(error)))

```

## 5.4 Demo - Kuramoto-Sivashinsky equation

**Authors** Mikael Mortensen (mikaem at math.uio.no)

**Date** Nov 14, 2019

*Summary.* This is a demonstration of how the Python module `shenfun` can be used to solve the time-dependent, nonlinear Kuramoto-Sivashinsky equation, in a doubly periodic domain. The demo is implemented in a single Python file `KuramotoSivashinsky.py`, and it may be run in parallel using MPI.

### 5.4.1 The Kuramoto-Sivashinsky equation

Movie showing the evolution of the solution  $u$  from Eq. (5.67).

#### Model equation

The Kuramoto-Sivashinsky (KS) equation is known for its chaotic behaviour, and it is often used in study of turbulence or turbulent combustion. We will here solve the KS equation in a doubly periodic domain  $[-30\pi, 30\pi]^2$ , starting from a single Gaussian pulse

$$\frac{\partial u(\mathbf{x}, t)}{\partial t} + \nabla^2 u(\mathbf{x}, t) + \nabla^4 u(\mathbf{x}, t) + |\nabla u(\mathbf{x}, t)|^2 = 0 \quad \text{for } \mathbf{x} \in \Omega = [-30\pi, 30\pi]^2 \quad (5.67)$$

$$u(\mathbf{x}, 0) = \exp(-0.01\mathbf{x} \cdot \mathbf{x})$$

#### Spectral Galerkin method

The PDE in (5.67) can be solved with many different numerical methods. We will here use the `shenfun` software and this software makes use of the spectral Galerkin method. Being a Galerkin method, we need to reshape the governing equations into proper variational forms, and this is done by multiplying (5.67) with the complex conjugate of a proper test function and then integrating over the domain. To this end we use testfunction  $v \in V(\Omega)$ , where  $V(\Omega)$  is some suitable function space, and obtain

$$\frac{\partial}{\partial t} \int_{\Omega} u \bar{v} w \, dx = - \int_{\Omega} (\nabla^2 u + \nabla^4 u + |\nabla u|^2) \bar{v} w \, dx. \quad (5.68)$$

Note that the overline is used to indicate a complex conjugate, whereas  $w$  is a weight function. The function  $u$  is now to be considered a trial function, and the integrals over the domain are often referred to as inner products. With inner product notation

$$(u, v) = \int_{\Omega} u \bar{v} w \, dx.$$

the variational problem can be formulated as

$$\frac{\partial}{\partial t} (u, v) = - (\nabla^2 u + \nabla^4 u + |\nabla u|^2, v). \quad (5.69)$$

The space and time discretizations are still left open. There are numerous different approaches that one could take for discretizing in time. Here we will use a fourth order exponential Runge-Kutta method.

## Discretization

We discretize the model equation in space using continuously differentiable Fourier basis functions

$$\phi_l(x) = e^{ilx}, \quad -\infty < l < \infty, \quad (5.70)$$

where  $l$  is the wavenumber, and  $\underline{l} = \frac{2\pi}{L}l$  is the scaled wavenumber, scaled with domain length  $L$  (here  $60\pi$ ). Since we want to solve these equations on a computer, we need to choose a finite number of test functions. A discrete function space  $V^N$  can be defined as

$$V^N(x) = \text{span}\{\phi_l(x)\}_{l \in \mathbf{l}}, \quad (5.71)$$

where  $N$  is chosen as an even positive integer and  $\mathbf{l} = (-N/2, -N/2 + 1, \dots, N/2 - 1)$ . And now, since  $\Omega$  is a two-dimensional domain, we can create a tensor product of two such one-dimensional spaces:

$$W^N(x, y) = V^N(x) \otimes V^N(y), \quad (5.72)$$

where  $\mathbf{N} = (N, N)$ . Obviously, it is not necessary to use the same number ( $N$ ) of basis functions for each direction, but it is done here for simplicity. A 2D tensor product basis function is now defined as

$$\Phi_{lm}(x, y) = e^{ilx} e^{imy} = e^{i(\underline{l}x + \underline{m}y)}, \quad (5.73)$$

where the indices for  $y$ -direction are  $\underline{m} = \frac{2\pi}{L}m$ , and  $\mathbf{m}$  is the same set as  $\mathbf{l}$  due to using the same number of basis functions for each direction. One distinction, though, is that for the  $y$ -direction expansion coefficients are only stored for  $m = (0, 1, \dots, N/2)$  due to Hermitian symmetry (real input data).

We now look for solutions of the form

$$u(x, y) = \sum_{l=-N/2}^{N/2-1} \sum_{m=-N/2}^{N/2-1} \hat{u}_{lm} \Phi_{lm}(x, y). \quad (5.74)$$

The expansion coefficients  $\hat{u}_{lm}$  can be related directly to the solution  $u(x, y)$  using Fast Fourier Transforms (FFTs) if we are satisfied with obtaining the solution in quadrature points corresponding to

$$x_i = \frac{60\pi i}{N} - 30\pi \quad \forall i \in \mathbf{i}, \text{ where } \mathbf{i} = (0, 1, \dots, N-1), \quad (5.75)$$

$$y_j = \frac{60\pi j}{N} - 30\pi \quad \forall j \in \mathbf{j}, \text{ where } \mathbf{j} = (0, 1, \dots, N-1). \quad (5.76)$$

Note that these points are different from the standard (like  $2\pi j/N$ ) since the domain is set to  $[-30\pi, 30\pi]^2$  and not the more common  $[0, 2\pi]^2$ . We now have

$$u(x_i, y_j) = \mathcal{F}_y^{-1}(\mathcal{F}_x^{-1}(\hat{u})) \quad \forall (i, j) \in \mathbf{i} \times \mathbf{j}, \quad (5.77)$$

where  $\mathcal{F}_x^{-1}$  is the inverse Fourier transform along direction  $x$ , for all  $j \in \mathbf{j}$ . Note that the two inverse FFTs are performed sequentially, one direction at the time, and that there is no scaling factor due the definition used for the inverse Fourier transform:

$$u(x_j) = \sum_{l=-N/2}^{N/2-1} \hat{u}_l e^{ilx_j}, \quad \forall j \in \mathbf{j}. \quad (5.78)$$

Note that this differs from the definition used by, e.g., `Numpy`.

The inner products used in Eq. (5.69) may be computed using forward FFTs (using weight functions  $w = 1/L$ ):

$$(u, \Phi_{lm}) = \hat{u}_{lm} = \frac{1}{N^2} \mathcal{F}_l(\mathcal{F}_m(u)) \quad \forall (l, m) \in \mathbf{l} \times \mathbf{m}, \quad (5.79)$$

From this we see that the variational forms may be written in terms of the Fourier transformed  $\hat{u}$ . Expanding the exact derivatives of the nabla operator, we have

$$(\nabla^2 u, v) = -(\underline{l}^2 + \underline{m}^2) \hat{u}_{lm}, \quad (5.80)$$

$$(\nabla^4 u, v) = (\underline{l}^2 + \underline{m}^2)^2 \hat{u}_{lm}, \quad (5.81)$$

$$(|\nabla u|^2, v) = \widehat{|\nabla u|^2}_{lm} \quad (5.82)$$

and as such the equation to be solved for each wavenumber can be found directly as

$$\frac{\partial \hat{u}_{lm}}{\partial t} = (\underline{l}^2 + \underline{m}^2 - (\underline{l}^2 + \underline{m}^2)^2) \hat{u}_{lm} - \widehat{|\nabla u|^2}_{lm}, \quad (5.83)$$

## 5.4.2 Implementation

The model equation (5.67) is implemented in shenfun using Fourier basis functions for both  $x$  and  $y$  directions. We start the solver by implementing necessary functionality from required modules like `Numpy`, `Sympy` `matplotlib` and `mpi4py`, in addition to `shenfun`:

```
from sympy import symbols, exp, lambdify
import numpy as np
import matplotlib.pyplot as plt
from mpi4py import MPI
from shenfun import *
```

The size of the problem (in real space) is then specified, before creating the `TensorProductSpace`, which is using a tensor product of two Fourier bases as basis functions. We also create a `VectorTensorProductSpace`, since this is required for computing the gradient of the scalar field  $u$ . The gradient is required for the nonlinear term.

```
# Size of discretization
N = (128, 128)

comm = MPI.COMM_WORLD
K0 = Basis(N[0], 'F', domain=(-30*np.pi, 30*np.pi), dtype='D')
K1 = Basis(N[1], 'F', domain=(-30*np.pi, 30*np.pi), dtype='d')
T = TensorProductSpace(comm, (K0, K1), **{'planner_effort': 'FFTW_MEASURE'})
TV = VectorTensorProductSpace([T, T])
```

Test and trialfunctions are required for assembling the variational forms:

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

and some arrays are required to hold the solution. We also create an array `gradu`, that will be used to compute the gradient in the nonlinear term. Finally, the wavenumbers are collected in list  $\mathcal{K}$ . Here one feature is worth mentioning. The gradient in spectral space can be computed as `1j*K*U_hat`. However, since this is an odd derivative, and we are using an even number  $N$  for the size of the domain, the highest wavenumber must be set to zero. This is the purpose of the last keyword argument to `local_wavenumbers` below.

```
U = Array(T)
U_hat = Function(T)
gradu = Array(TV)
K = np.array(T.local_wavenumbers(True, True, eliminate_highest_freq=True))
```

Note that using this  $K$  in computing derivatives has the same effect as achieved by symmetrizing the Fourier series by replacing the first sum below with the second when computing odd derivatives.

$$u = \sum_{k=-N/2}^{N/2-1} \hat{u} e^{ikx} \quad (5.84)$$

$$u = \sum'_{k=-N/2}^{N/2} \hat{u} e^{ikx} \quad (5.85)$$

Here  $\sum'$  means that the first and last items in the sum are divided by two. Note that the two sums are equal as they stand (due to aliasing), but only the latter (known as the Fourier interpolant) gives the correct (zero) derivative of the basis with the highest wavenumber.

Sympy is used to generate an initial condition, as stated in Eq (5.67)

```
# Use sympy to set up initial condition
x, y = symbols("x,y")
ue = exp(-0.01*(x**2+y**2))
ul = lambdify((x, y), ue, 'numpy')
```

Shenfun has a few integrators implemented in the `integrators` submodule. Two such integrators are the 4th order explicit Runge-Kutta method `RK4`, and the exponential 4th order Runge-Kutta method `ETDRK4`. Both these integrators need two methods provided by the problem being solved, representing the linear and nonlinear terms in the problem equation. We define two methods below, called `LinearRHS` and `NonlinearRHS`

```
def LinearRHS(self):
    # Assemble diagonal bilinear forms
    L = -(inner(div(grad(u))+div(grad(div(grad(u))))), v))
    return L

def NonlinearRHS(self, U, U_hat, dU):
    # Assemble nonlinear term
    global gradu
    gradu = TV.backward(1j*K*U_hat, gradu)
    dU = T.forward(0.5*(gradu[0]*gradu[0]+gradu[1]*gradu[1]), dU)
    return -dU
```

The code should, hopefully, be self-explanatory.

All that remains now is to initialize the solution arrays and to setup the integrator plus some plotting functionality for visualizing the results. Note that visualization is only nice when running the code in serial. For parallel, it is recommended to use `HDF5File`, to store intermediate results to the HDF5 format, for later viewing in, e.g., Paraview.

The solution is initialized as

```
#initialize
X = T.local_mesh(True)
U[:] = ul(*X)
U_hat = T.forward(U, U_hat)
```

And we also create an update function for plotting intermediate results with a cool colormap:

```
# Integrate using an exponential time integrator
plt.figure()
cm = plt.get_cmap('hot')
image = plt.contourf(X[0], X[1], U, 256, cmap=cm)
```

(continues on next page)

(continued from previous page)

```

plt.draw()
plt.pause(1e-6)
count = 0
def update(u, u_hat, t, tstep, **params):
    global count
    if tstep % params['plot_step'] == 0 and params['plot_step'] > 0:
        u = T.backward(u_hat, u)
        image.ax.clear()
        image.ax.contourf(X[0], X[1], U, 256, cmap=cm)
        plt.pause(1e-6)
        count += 1
        plt.savefig('Kuramoto_Sivashinsky_N_{0}_{1}.png'.format(N[0], count))

```

Now all that remains is to create the integrator and call it

```

if __name__ == '__main__':
    par = {'plot_step': 100}
    dt = 0.01
    end_time = 100
    integrator = ETDRK4(T, L=LinearRHS, N=NonlinearRHS, update=update, **par)
    #integrator = RK4(T, L=LinearRHS, N=NonlinearRHS, update=update, **par)
    integrator.setup(dt)
    U_hat = integrator.solve(U, U_hat, dt, (0, end_time))

```

## 5.5 Demo - Stokes equations

**Authors** Mikael Mortensen (mikaem at math.uio.no)

**Date** Nov 14, 2019

*Summary.* The Stokes equations describe the flow of highly viscous fluids. This is a demonstration of how the Python module `shenfun` can be used to solve Stokes equations using a mixed (coupled) basis in a 3D tensor product domain. We assume homogeneous Dirichlet boundary conditions in one direction and periodicity in the remaining two. The solver described runs with MPI without any further considerations required from the user. The solver assembles a block matrix with sparsity pattern as shown below for the Legendre basis.

### 5.5.1 Model problem

#### Stokes equations

The Stokes equations are given in strong form as

$$\begin{aligned}
 \nabla^2 \mathbf{u} - \nabla p &= \mathbf{f} & \text{in } \Omega, \\
 \nabla \cdot \mathbf{u} &= h & \text{in } \Omega \\
 \int_{\Omega} p dx &= 0
 \end{aligned}$$

where  $\mathbf{u}$  and  $p$  are, respectively, the fluid velocity vector and pressure, and the domain  $\Omega = [0, 2\pi]^2 \times [-1, 1]$ . The flow is assumed periodic in  $x$  and  $y$ -directions, whereas there is a no-slip homogeneous Dirichlet boundary condition on  $\mathbf{u}$  on the boundaries of the  $z$ -direction, i.e.,  $\mathbf{u}(x, y, \pm 1) = (0, 0, 0)$ . (Note that we can configure `shenfun` with non-periodicity in any of the three directions. However, since we are to solve linear algebraic systems in the non-periodic direction, there is a speed benefit from having the nonperiodic direction last. This has to do with Numpy

Block matrix:  $l, m = (5, 5)$

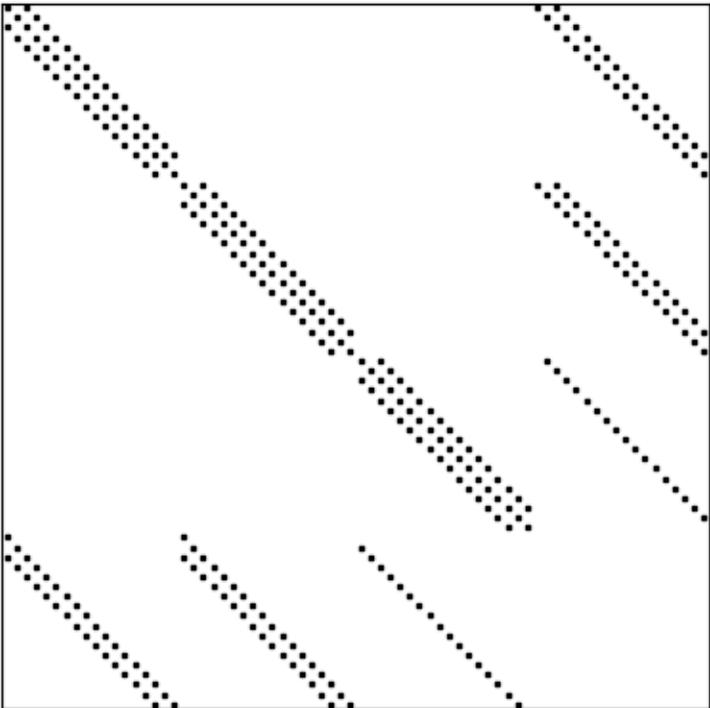


Fig. 3: *Coupled block matrix for Stokes equations*

using a C-style row-major storage of arrays by default.) The right hand side vector  $\mathbf{f}(\mathbf{x})$  is an external applied body force. The right hand side  $h$  is usually zero in the regular Stokes equations. Here we include it because it will be nonzero in the verification, which is using the method of manufactured solutions. Note that the final  $\int_{\Omega} p dx = 0$  is there because there is no Dirichlet boundary condition on the pressure and the system of equations would otherwise be ill conditioned.

To solve Stokes equations with the Galerkin method we need basis functions for both velocity and pressure. A Dirichlet basis will be used for velocity, whereas there is no boundary restriction on the pressure basis. For both three-dimensional bases we will use one basis function for the  $x$ -direction,  $\mathcal{X}(x)$ , one for the  $y$ -direction,  $\mathcal{Y}(y)$ , and one for the  $z$ -direction,  $\mathcal{Z}(z)$ . And then we create three-dimensional basis functions like

$$v(x, y, z) = \mathcal{X}(x)\mathcal{Y}(y)\mathcal{Z}(z). \quad (5.86)$$

The basis functions  $\mathcal{X}(x)$  and  $\mathcal{Y}(y)$  are chosen as Fourier exponentials, since these functions are periodic:

$$\mathcal{X}_l(x) = e^{ilx}, \forall l \in \mathbf{I}^{N_0}, \quad (5.87)$$

$$\mathcal{Y}_m(y) = e^{imy}, \forall m \in \mathbf{m}^{N_1}, \quad (5.88)$$

where  $\mathbf{I}^{N_0} = (-N_0/2, -N_0/2 + 1, \dots, N_0/2 - 1)$  and  $\mathbf{m}^{N_1} = (-N_1/2, -N_1/2 + 1, \dots, N_1/2 - 1)$ . The size of the discretized problem in real physical space is  $\mathbf{N} = (N_0, N_1, N_2)$ , i.e., there are  $N_0 \cdot N_1 \cdot N_2$  quadrature points in total.

The basis functions for  $\mathcal{Z}(z)$  remain to be decided. For the velocity we need homogeneous Dirichlet boundary conditions, and for this we use composite Legendre or Chebyshev polynomials

$$\mathcal{Z}_n^0(z) = \phi_n(z) - \phi_{n+2}(z), \forall n \in \mathbf{n}^{N_2-2}, \quad (5.89)$$

where  $\phi_n$  is the  $n$ 'th Legendre or Chebyshev polynomial of the first kind.  $\mathbf{n}^{N_2-2} = (0, 1, \dots, N_2 - 3)$ , and the zero on  $\mathcal{Z}^0$  is there to indicate the zero value on the boundary.

The pressure basis that comes with no restrictions for the boundary is a little trickier. The reason for this has to do with inf-sup stability. The obvious choice of basis is the regular Legendre or Chebyshev basis, which is denoted as

$$\mathcal{Z}_n(z) = \phi_n(z), \forall n \in \mathbf{n}^{N_2}. \quad (5.90)$$

The problem is that for the natural choice of  $n \in (0, 1, \dots, N_2 - 1)$  there is a nullspace and one degree of freedom remains unresolved. It turns out that the proper choice for the pressure basis is simply (5.90) for  $n \in \mathbf{n}^{N_2-2}$ . (Also remember that we have to fix  $\int_{\Omega} p dx = 0$ .)

With given basis functions we obtain the spaces

$$V^{N_0} = \text{span}\{\mathcal{X}_l\}_{l \in \mathbf{I}^{N_0}}, \quad (5.91)$$

$$V^{N_1} = \text{span}\{\mathcal{Y}_m\}_{m \in \mathbf{m}^{N_1}}, \quad (5.92)$$

$$V^{N_2} = \text{span}\{\mathcal{Z}_n\}_{n \in \mathbf{n}^{N_2-2}}, \quad (5.93)$$

$$V_0^{N_2} = \text{span}\{\mathcal{Z}_n^0\}_{n \in \mathbf{n}^{N_2-2}}, \quad (5.94)$$

and from these we create two different tensor product spaces

$$W_0^{\mathbf{N}}(\mathbf{x}) = V^{N_0}(x) \otimes V^{N_1}(y) \otimes V_0^{N_2}(z), \quad (5.95)$$

$$W^{\mathbf{N}}(\mathbf{x}) = V^{N_0}(x) \otimes V^{N_1}(y) \otimes V^{N_2}(z). \quad (5.96)$$

The velocity vector is using a mixed basis, such that we will look for solutions  $\mathbf{u} \in [W_0^{\mathbf{N}}]^3 (= W_0^{\mathbf{N}} \times W_0^{\mathbf{N}} \times W_0^{\mathbf{N}})$ , whereas we look for the pressure  $p \in W^{\mathbf{N}}$ . We now formulate a variational problem using the Galerkin method: Find  $\mathbf{u} \in [W_0^{\mathbf{N}}]^3$  and  $p \in W^{\mathbf{N}}$  such that

$$\int_{\Omega} (\nabla^2 \mathbf{u} - \nabla p) \cdot \bar{\mathbf{v}} dx_w = \int_{\Omega} \mathbf{f} \cdot \bar{\mathbf{v}} dx_w \quad \forall \mathbf{v} \in [W_0^{\mathbf{N}}]^3, \quad (5.97)$$

$$\int_{\Omega} \nabla \cdot \mathbf{u} \bar{q} dx_w = \int_{\Omega} h \bar{q} dx_w \quad \forall q \in W^{\mathbf{N}}. \quad (5.98)$$

Here  $dx_w = w_x dx w_y dy w_z dz$  represents a weighted measure, with weights  $w_x(x), w_y(y), w_z(z)$ . Note that it is only Chebyshev polynomials that make use of a non-constant weight  $w_x = 1/\sqrt{1-x^2}$ . The Fourier weights are  $w_y = w_z = 1/(2\pi)$  and the Legendre weight is  $w_x = 1$ . The overline in  $\bar{\mathbf{v}}$  and  $\bar{q}$  represents a complex conjugate, which is needed here because the Fourier exponentials are complex functions.

### Mixed variational form

Since we are to solve for  $\mathbf{u}$  and  $p$  at the same time, we formulate a mixed (coupled) problem: find  $(\mathbf{u}, p) \in [W_0^{\mathbf{N}}]^3 \times W^{\mathbf{N}}$  such that

$$a((\mathbf{u}, p), (\mathbf{v}, q)) = L((\mathbf{v}, q)) \quad \forall (\mathbf{v}, q) \in [W_0^{\mathbf{N}}]^3 \times W^{\mathbf{N}}, \quad (5.99)$$

where bilinear ( $a$ ) and linear ( $L$ ) forms are given as

$$a((\mathbf{u}, p), (\mathbf{v}, q)) = \int_{\Omega} (\nabla^2 \mathbf{u} - \nabla p) \cdot \bar{\mathbf{v}} dx_w + \int_{\Omega} \nabla \cdot \mathbf{u} \bar{q} dx_w, \quad (5.100)$$

$$L((\mathbf{v}, q)) = \int_{\Omega} \mathbf{f} \cdot \bar{\mathbf{v}} dx_w + \int_{\Omega} h \bar{q} dx_w. \quad (5.101)$$

Note that the bilinear form will assemble to block matrices, whereas the right hand side linear form will assemble to block vectors.

## 5.5.2 Implementation

### Preamble

We will solve the Stokes equations using the `shenfun` Python module. The first thing needed is then to import some of this module's functionality plus some other helper modules, like `Numpy` and `Sympy`:

```
import os
import sys
import numpy as np
from mpi4py import MPI
from sympy import symbols, sin, cos, lambdify
from shenfun import *
```

We use `Sympy` for the manufactured solution and `Numpy` for testing. `MPI` for Python (`mpi4py`) is required for running the solver with `MPI`.

### Manufactured solution

The exact solutions  $\mathbf{u}_e(\mathbf{x})$  and  $p(\mathbf{x})$  are chosen to satisfy boundary conditions, and the right hand sides  $\mathbf{f}(\mathbf{x})$  and  $h(\mathbf{x})$  are then computed exactly using `Sympy`. These exact right hand sides will then be used to compute a numerical solution that can be verified against the manufactured solution. The chosen solution with computed right hand sides are:

```
x, y, z = symbols('x, y, z')
uex = sin(2*y) * (1-z**2)
uey = sin(2*x) * (1-z**2)
uez = sin(2*z) * (1-z**2)
```

(continues on next page)

(continued from previous page)

```

pe = -0.1*sin(2*x)*cos(4*y)
fx = uex.diff(x, 2) + uex.diff(y, 2) + uex.diff(z, 2) - pe.diff(x, 1)
fy = uey.diff(x, 2) + uey.diff(y, 2) + uey.diff(z, 2) - pe.diff(y, 1)
fz = uez.diff(x, 2) + uez.diff(y, 2) + uez.diff(z, 2) - pe.diff(z, 1)
h = uex.diff(x, 1) + uey.diff(y, 1) + uez.diff(z, 1)

# Lambdify for faster evaluation
ulx = lambdify((x, y, z), uex, 'numpy')
uly = lambdify((x, y, z), uey, 'numpy')
ulz = lambdify((x, y, z), uez, 'numpy')
flx = lambdify((x, y, z), fx, 'numpy')
fly = lambdify((x, y, z), fy, 'numpy')
flz = lambdify((x, y, z), fz, 'numpy')
hl = lambdify((x, y, z), h, 'numpy')
pl = lambdify((x, y, z), pe, 'numpy')

```

## Bases and tensor product spaces

Bases are created using the `Basis()` function. A choice of polynomials between Legendre or Chebyshev can be made, and the size of the domain is given

```

N = (40, 40, 40)
family = 'Legendre'
K0 = Basis(N[0], 'Fourier', dtype='D', domain=(0, 2*np.pi))
K1 = Basis(N[1], 'Fourier', dtype='d', domain=(0, 2*np.pi))
SD = Basis(N[2], family, bc=(0, 0))
ST = Basis(N[2], family)
ST.slice = lambda: slice(0, ST.N-2)

```

Note that the last line of code is there to ensure that only the first  $N_2 - 2$  coefficients are used, see discussion around Eq. (5.90). At the same time, we ensure that we are still using  $N_2$  quadrature points, the same as for the Dirichlet basis.

Next the one-dimensional spaces are used to create two tensor product spaces  $Q = W^N$  and  $TD = W_0^N$ , one vector  $V = [W_0^N]^3$  and one mixed space  $VQ = V \times Q$ .

```

TD = TensorProductSpace(comm, (K0, K1, SD), axes=(2, 0, 1))
Q = TensorProductSpace(comm, (K0, K1, ST), axes=(2, 0, 1))
V = VectorTensorProductSpace(TD)
VQ = MixedTensorProductSpace([V, Q])

```

Note that we choose to transform axes in the order 1, 0, 2. This is to ensure that the fully transformed arrays are aligned in the non-periodic direction 2. And we need the arrays aligned in this direction, because this is the only direction where there are tensor product matrices that are non-diagonal. All Fourier matrices are, naturally, diagonal.

Test- and trialfunctions are created much like in a regular, non-mixed, formulation. However, one has to create one test- and trialfunction for the mixed space, and then split them up afterwards

```

up = TrialFunction(VQ)
vq = TestFunction(VQ)
u, p = up
v, q = vq

```

With the basisfunctions in place we may assemble the different blocks of the final coefficient matrix. Since Legendre is using a constant weight function, the equations may also be integrated by parts to obtain a symmetric system:

```

if family.lower() == 'chebyshev':
    A = inner(v, div(grad(u)))
    G = inner(v, -grad(p))
else:
    A = inner(grad(v), -grad(u))
    G = inner(div(v), p)
D = inner(q, div(u))

```

**Note:** The inner products may also be assembled with one single line, as

```
AA = inner(v, div(grad(u))) + inner(v, -grad(u)) + inner(q, div(u))
```

However, this requires addition, not subtraction, of inner products and it is not possible to move the negation to `-inner(v, grad(u))`

The assembled subsystems A, G and D are lists containing the different blocks of the complete, coupled matrix. A actually contains 6 tensor product matrices of type `TPMatrix`. The first two matrices are for vector component zero of the test function `v[0]` and trial function `u[0]`, the matrices 2 and 3 are for components 1 and the last two are for components 2. The first two matrices are as such for

```
A[0:2] = inner(v[0], div(grad(u[0])))
```

Breaking it down the inner product is mathematically

$$\int_{\Omega} \mathbf{v}[0] \left( \frac{\partial^2 \mathbf{u}[0]}{\partial x^2} + \frac{\partial^2 \mathbf{u}[0]}{\partial y^2} + \frac{\partial^2 \mathbf{u}[0]}{\partial z^2} \right) w_x dx w_y dy w_z dz. \quad (5.102)$$

If we now use test function `v[0]`

$$\mathbf{v}[0]_{lmn} = \mathcal{X}_l \mathcal{Y}_m \mathcal{Z}_n, \quad (5.103)$$

and trialfunction

$$\mathbf{u}[0]_{pqr} = \sum_p \sum_q \sum_r \hat{\mathbf{u}}[0]_{pqr} \mathcal{X}_p \mathcal{Y}_q \mathcal{Z}_r, \quad (5.104)$$

where  $\hat{\mathbf{u}}$  are the unknown degrees of freedom, and then insert these functions into (5.102), then we obtain after performing some exact evaluations over the periodic directions

$$\underbrace{\left( - (l^2 \delta_{lp} + m^2 \delta_{mq}) \int_{-1}^1 \mathcal{Z}_r(z) \mathcal{Z}_n(z) w_z dz \right)}_{A[0]} + \underbrace{\left( \delta_{lp} \delta_{mq} \int_{-1}^1 \frac{\partial^2 \mathcal{Z}_r(z)}{\partial z^2} \mathcal{Z}_n(z) w_z dz \right)}_{A[1]} \hat{\mathbf{u}}[0]_{pqr}, \quad (5.105)$$

Similarly for components 1 and 2 of the test and trial vectors, leading to 6 tensor product matrices in total for A. Similarly, we get three components of G and three of D.

Eliminating the Fourier diagonal matrices, we are left with block matrices like

$$H(l, m) = \begin{bmatrix} A[0] + A[1] & 0 & 0 & G[0] \\ 0 & A[2] + A[3] & 0 & G[1] \\ 0 & 0 & A[4] + A[5] & G[2] \\ D[0] & D[1] & D[2] & 0 \end{bmatrix} \quad (5.106)$$

Note that there will be one large block matrix  $H(l, m)$  for each Fourier wavenumber combination  $(l, m)$ . To solve the problem in the end we will need to loop over these wavenumbers and solve the assembled linear systems one by one. An example of the block matrix, for  $l = m = 5$  and  $\mathbf{N} = (20, 20, 20)$  is given in Fig. [Coupled block matrix for Stokes equations](#).

In the end we create a block matrix through

```
M = BlockMatrix(A+G+D)
```

The right hand side can easily be assembled since we have already defined the functions  $f$  and  $h$ , see Sec. *Manufactured solution*

```
# Get mesh (quadrature points)
X = TD.local_mesh(True)

# Get f and h on quad points
fh = Array(VQ)
f_, h_ = fh
f_[0] = flx(*X)
f_[1] = fly(*X)
f_[2] = flz(*X)
h_[:] = hl(*X)

# Compute inner products
fh_hat = Function(VQ)
f_hat, h_hat = fh_hat
f_hat = inner(v, f_, output_array=f_hat)
h_hat = inner(q, h_, output_array=h_hat)
```

In the end all that is left is to solve and compare with the exact solution.

```
# Solve problem
up_hat = M.solve(fh_hat, constraints=((3, 0, 0),))
up = up_hat.backward()
u_, p_ = up

# Exact solution
ux = ulx(*X)
uy = uly(*X)
uz = ulz(*X)
pe = pl(*X)

error = [comm.reduce(np.linalg.norm(ux-u_[0])),
         comm.reduce(np.linalg.norm(uy-u_[1])),
         comm.reduce(np.linalg.norm(uz-u_[2])),
         comm.reduce(np.linalg.norm(pe-p_))]
print(error)
```

Note that solve has a keyword argument `constraints=((3, 0, 0),)` that takes care of the restriction  $\int_{\Omega} p dx = 0$  by indenting the row in  $M$  corresponding to the first degree of freedom for the pressure. The value  $(3, 0, 0)$  indicates that pressure is in block 3 of the block vector solution (the velocity vector holds positions 0, 1 and 2), whereas the two zeros ensures that the first dof (dof 0) should obtain value 0.

## Complete solver

A complete solver can be found in demo [Stokes3D.py](#).



## BIBLIOGRAPHY

- [PSS18] Ambrish Pandey, Janet D. Scheel, and Jörg Schumacher. Turbulent superstructures in rayleigh-bénard convection. *Nature Communications*, 9(1):2118, 2018. doi:10.1038/s41467-018-04478-0.
- [PSS18] Ambrish Pandey, Janet D. Scheel, and Jörg Schumacher. Turbulent superstructures in rayleigh-bénard convection. *Nature Communications*, 9(1):2118, 2018. doi:10.1038/s41467-018-04478-0.
- [PSS18] Ambrish Pandey, Janet D. Scheel, and Jörg Schumacher. Turbulent superstructures in rayleigh-bénard convection. *Nature Communications*, 9(1):2118, 2018. doi:10.1038/s41467-018-04478-0.
- [She94] Jie Shen. Efficient spectral-galerkin method i. direct solvers of second- and fourth-order equations using legendre polynomials. *SIAM Journal on Scientific Computing*, 15(6):1489–1505, 1994. URL: <https://doi.org/10.1137/0915089>, arXiv:<https://doi.org/10.1137/0915089>, doi:10.1137/0915089.
- [She95] Jie Shen. Efficient spectral-galerkin method ii. direct solvers of second- and fourth-order equations using chebyshev polynomials. *SIAM Journal on Scientific Computing*, 16(1):74–87, 1995. URL: <https://doi.org/10.1137/0916006>, arXiv:<https://doi.org/10.1137/0916006>, doi:10.1137/0916006.
- [Waz08] Abdul-Majid Wazwaz. New travelling wave solutions to the boussinesq and the klein-gordon equations. *Communications in Nonlinear Science and Numerical Simulation*, 13(5):889–901, 2008. URL: <http://www.sciencedirect.com/science/article/pii/S1007570406001729>, doi:<http://dx.doi.org/10.1016/j.cnsns.2006.08.005>.