cbcpost Documentation

Release 1.3.0

Martin Alnæs and Øyvind Evju

July 13, 2016

Contents

1 Installation 3
1.1 Quick Install 3
1.2 Dependencies 3

2 Features 5

3 Demos 7
3.1 A Basic Use Case 7
3.2 Restart a Problem 12
3.3 Replay a Problem 12

4 Functionality 15
4.1 The `Field`-class and subclasses 15
4.2 The postprocessor 18
4.3 Replay 21
4.4 Restart 23
4.5 Utilities 25

5 Programmer's reference 29
5.1 Postprocessor 30
5.2 Replay 30
5.3 Restart 30
5.4 Parameter system 30
5.5 FunctionSpace pooling 30
5.6 Field bases 30
5.7 Implemented cbcpost.metafields 30
5.8 Utilities 30

6 Contributing 31
6.1 Pull requests 31
6.2 Report problems 31
6.3 Contact developers 31

7 Indices and tables 33

cbcpost is developed to simplify the postprocessing of simulation results, produced by FEniCS solvers.

The framework is designed to take any given solution, and compute and save any derived data. The interface is designed to be simple, with minimal cluttering of a typical solver code. This is illustrated by the following simple example:

```
# ... problem set up ...

# Set up postprocessor
solution = SolutionField("Displacement", dict(save=True))
postprocessor = PostProcessor(dict(casedir="Results/"))
postprocessor.add_field(solution)

t = 0.0
timestep = 0
while t < T:
    timestep += 1
    # ... solve equation ...

    # Update postprocessor
    postprocessor.update_all(dict("Displacement"=lambda: u), timestep, t)

    # continue
```

cbcpost is developed at the Center for Biomedical Computing, at Simula Research Laboratory by Øyvind Evju and Martin Sandve Alnæs.

Contents:

CHAPTER 1

Installation

1.1 Quick Install

Install using git clone:

```
git clone https://bitbucket.org/simula_cbc/cbcpost.git
cd cbcpost
python setup.py install
```

Install using pip:

```
pip install git+https://bitbucket.org/simula_cbc/cbcpost.git
```

1.2 Dependencies

The installation of cbcpost requires the following environment:

- Python 2.7
- Numpy
- Scipy
- dbhash, dbm or gdbm
- FEniCS 1.4.0 or newer

To install FEniCS, please refer to the FEniCS download page. cbcpost follows the same version numbering as FEniCS, so make sure you install the correct FEniCS version. Backwards compatibility is not guaranteed (and quite unlikely).

In addition, cbcpost can utlize other libraries for added functionality

- fenicstools 1.4.0 (highly recommended, tools to inspect parts of a solution)
- mpi4py
- pytest >2.4.0 (required to run test suite)

fenicstools can be installed using pip:

```
pip install https://github.com/mikaem/fenicstools/archive/v1.4.0.zip
```

CHAPTER 2

Features

The main features of cbcpost are

- Saving in 7 different save formats (xdmf, hdf5, xml, xml.gz, pvd, shelve, txt)
- Plotting using dolfin.plot or pyplot
- Automatic planning of computations, saving and plotting
- Automatic dependency handling
- Many different fields provided, ranging from time integrals to point evaluations and norms.
- Easily expandable with custom `Field`-subclasses
- Flexible parameter system
- Small footprint on solver code
- Replay functionality
- Restart support

CHAPTER 3

Demos

To get started, we recommend starting with the demos. If you are unfamiliar with FEniCS, please refer to the FEniCS Tutorial for the FEniCS-specifics of these demos.

Documented demos:

3.1 A Basic Use Case

To demonstrate the functionality of the postprocessor, consider the 3D-case of the heat equation with variable diffusivity. The full demo can be found in `Basic.py`.

The general heat equation reads

$$\frac{\partial u}{\partial t} + \alpha(x)\Delta u = f$$

where u typically denotes the temperature and α denotes the material diffusivity.

Boundary conditions are in our example given as

$$u(x,t) = Asin(2\pi t x_0), x \in \partial\Omega$$

and initial condition

$$u(x,0) = 0.$$

We also use f=0, and solve the equations at the unit cube for $t \in (0,3]$.

3.1.1 Setting up the problem

We start by defininge a set of parameters for our problem:

```
from cbcpost import *
from cbcpost.utils import cbc_print
from dolfin import *
set_log_level(WARNING)

# Create parameters for problem
params = ParamDict(
    T = 3.0,            # End time
    dt = 0.05,          # Time step
    theta = 0.5,        # Time stepping scheme (0.5=Crank-Nicolson)
```

```
    alpha0 = 10.0,       # Outer diffusivity
    alpha1 = 1e-3,       # Inner diffusivity
    amplitude = 3.0,     # Amplitude of boundary condition
)
```

The parameters are created using the utility class `ParamDict`, which extend the built-in python dict.

We the use the parameters to set up the problem using FEniCS:

```
# Create mesh
mesh = UnitCubeMesh(21,21,21)

# Function spaces
V = FunctionSpace(mesh, "CG", 1)
u,v = TrialFunction(V), TestFunction(V)

# Time and time-stepping
t = 0.0
timestep = 0
dt = Constant(params.dt)

# Initial condition
U = Function(V)

# Define inner domain
def inside(x):
    return (0.5 < x[0] < 0.8) and (0.3 < x[1] < 0.6) and (0.2 < x[2] < 0.7)

class Alpha(Expression):
    "Variable conductivity expression"
    def __init__(self, alpha0, alpha1):
        self.alpha0 = alpha0
        self.alpha1 = alpha1

    def eval(self, value, x):
        if inside(x):
            value[0] = self.alpha1
        else:
            value[0] = self.alpha0

# Conductivity
alpha = project(Alpha(params.alpha0, params.alpha1), V)

# Boundary condition
u0 = Expression("ampl*sin(x[0]*2*pi*t)", t=t, ampl=params.amplitude)
bc = DirichletBC(V, u0, "on_boundary")

# Source term
f = Constant(0)

# Bilinear form
a = 1.0/dt*inner(u,v)*dx() + Constant(params.theta)*alpha*inner(grad(u), grad(v))*dx()
L = 1.0/dt*inner(U,v)*dx() + Constant(1-params.theta)*alpha*inner(grad(U), grad(v))*dx() + inner(f,v
A = assemble(a)
b = assemble(L)
bc.apply(A)
```

3.1.2 Setting up the PostProcessor

To set up the use case, we specify the case directory, and asks to clean out the case directory if there is any data remaining from a previous simulation:

```
pp = PostProcessor(dict(casedir="Results", clean_casedir=True))
```

Since we‘re solving for temperature, we add a SolutionField to the postprocessor:

```
pp.add_field(SolutionField("Temperature", dict(save=True,
                                save_as=["hdf5", "xdmf"],
                                plot=True,
                                plot_args=dict(range_min=-params.amplitude, range_max=params.amplituc
                                )))
```

Note that we pass parameters, specifying that the field is to be saved in hdf5 and xdmf formats. These formats are default for dolfin.Function-type objects. We also ask for the Field to be plotted, with plot_args specifying the plot window. These arguments are passed directly to the dolfin.plot-command.

Time derivatives and time integrals

We can compute both integrals and derivatives of other Fields. Here, we add the integral of temperature from t=1.0 to t=2.0, the time-average from t=0.0 to t=5.0 as well as the derivative of the temperature field.

```
pp.add_fields([
    TimeIntegral("Temperature", dict(save=True, start_time=1.0, end_time=2.0)),
    TimeAverage("Temperature", dict(save=True, end_time=params.T)),
    TimeDerivative("Temperature", dict(save=True)),
    ])
```

Again, we ask the fields to be saved. The save formats are decided by the datatype returned from the *compute*-functions.

Inspecting parts of a solution

We can also define fields to inspect parts of other fields. For this, we use some utilities from `cbcpost.utils`. For this problem, the domain of a different diffusivity lies entirely within the unit cube, and thus it may make sense to view some of the interior. We start by creating (sub)meshes of the domains we wish to inspect:

```
from cbcpost.utils import create_submesh, create_slice
celldomains = CellFunction("size_t", mesh)
celldomains.set_all(0)
AutoSubDomain(inside).mark(celldomains, 1)

slicemesh = create_slice(mesh, (0.7,0.5,0.5), (0.0,0.0,1.0))
submesh = create_submesh(mesh, celldomains, 1)
```

We then add instances of the fields `PointEval`, `SubFunction` and `Restrict` to the postprocessor:

```
pp.add_fields([
    PointEval("Temperature", [[0.7,0.5, 0.5]], dict(plot=True)),
    SubFunction("Temperature", slicemesh, dict(plot=True, plot_args=dict(range_min=-params.amplitude,
    Restrict("Temperature", submesh, dict(plot=True, save=True)),
    ])
```

Averages and norms

We can also compute scalars from other fields. `DomainAvg` compute the average of a specified domain (if not specified, the whole domain). Here, we compute the average temperature inside and outside the domain of different diffusivity, as specified by the variable *cell_domains*:

```
pp.add_fields([
    DomainAvg("Temperature", cell_domains=cell_domains, indicator=1, label="inner"),
    DomainAvg("Temperature", cell_domains=cell_domains, indicator=0, label="outer"),
])
```

The added parameter *label* does that these fields are now identified by *DomainAvg_Temperature-inner* and *DomainAvg_Temperature-inner*, respectively.

We can also compute the norm of any field:

```
pp.add_field(Norm("Temperature", dict(save=True)))
```

If no norm is specified, the L2-norm (or l2-norm) is computed.

Custom fields

The user may also customize fields as he wishes. In this section we demonstrate two ways to compute the difference in average temperature between the two areas of different diffusivity at any given time. First, we take an approach based solely on accessing the *Temperature*-field:

```
class TempDiff1(Field):
    def __init__(self, domains, ind1, ind2, *args, **kwargs):
        Field.__init__(self, *args, **kwargs)
        self.domains = domains
        self.ind1 = ind1
        self.ind2 = ind2

    def before_first_compute(self, get):
        self.V1 = assemble(Constant(1)*dx(self.ind1), cell_domains=self.domains, mesh=self.domains.me
        self.V2 = assemble(Constant(1)*dx(self.ind2), cell_domains=self.domains, mesh=self.domains.me

    def compute(self, get):
        u = get("Temperature")
        T1 = 1.0/self.V1*assemble(u*dx(self.ind1), cell_domains=self.domains)
        T2 = 1.0/self.V2*assemble(u*dx(self.ind2), cell_domains=self.domains)
        return T1-T2
```

In this implementation we have to specify the domains, as well as compute the respective averages directly each time. However, since we already added fields to compute the averages in both domains, there is another, much less code-demanding way to do this:

```
class TempDiff2(Field):
    def compute(self, get):
        T1 = get("DomainAvg_Temperature-inner")
        T2 = get("DomainAvg_Temperature-outer")
        return T1-T2
```

Here, we use the provided *get*-function to access the fields named as above, and compute the difference. We add an instance of both to the potsprocessor:

```
pp.add_fields([
    TempDiff1(cell_domains, 1, 0, dict(plot=True)),
```

```
    TempDiff2(dict(plot=True)),
])
```

Since both these should be the same, we can check this with `ErrorNorm`:

```
pp.add_field(
    ErrorNorm("TempDiff1", "TempDiff2", dict(plot=True), name="error"),
)
```

We ask for the error to be plotted. Since this is a scalar, this will be done using matplotlibs *pyplot*-module. We also pass the keyword argument *name*, which overrides the default naming (which would have been ErrorNorm_TempDiff1_TempDiff2) with *error*.

Combining fields

Finally, we can also add combination of fields, provided all dependencies have already been added to the postprocessor. For example, we can compute the space average of a time-average of our field *Restrict_Temperature* the following way:

```
pp.add_fields([
    TimeAverage("Restrict_Temperature"),
    DomainAvg("TimeAverage_Restrict_Temperature", params=dict(save=True)),
])
```

If *TimeAverage("Restrict_Temperature")* is not added first, adding the `DomainAvg`-field would fail with a `DependencyException`, since the postprocessor would have no knowledge of the field *TimeAverage_Restrict_Temperature*.

Saving mesh and parameters

We choose to store the mesh, domains and parameters associated with the problem:

```
pp.store_mesh(mesh, cell_domains=cell_domains)
pp.store_params(params)
```

These will be stored to *mesh.hdf5*, *params.pickle* and *params.txt* in the case directory.

3.1.3 Solving the problem

Solving the problem is done very simply here using simple FEniCS-commands:

```
solver = KrylovSolver(A, "cg", "hypre_amg")
while t <= params.T+DOLFIN_EPS:
    cbc_print("Time: "+str(t))
    u0.t = float(t)

    assemble(L, tensor=b)
    bc.apply(b)
    solver.solve(U.vector(), b)

    # Update the postprocessor
    pp.update_all({"Temperature": lambda: U}, t, timestep)

    # Update time
    t += float(dt)
    timestep += 1
```

Note the single call to the postprocessor, *pp.update_all*, which will then execute the logic for the postprocessor. The solution *Temperature* is passed in a dict as a lambda-function. This lambda-function gives the user flexibility to process the solution in any way before it is used in the postprocessor. This can for example be a scaling to physical units or joining scalar functions to a vector function.

Finally, at the end of the time-loop we finalize the postprocessor through

```
pp.finalize_all()
```

This command will finalize and return values for fields such as for example time integrals.

3.2 Restart a Problem

Say we wish to run our simulation further than t=3.0, to see how it develops. To restart a problem, all you need is to use the computed solution as initial conditions in a simiular problem setup.

Restarting the heat equation solved as in *A Basic Use Case*, can be done really simple with cbcpost. Starting with the python-file in *A Basic Use Case*, we only have to make a couple of minor changes.

We change the parameters T0 and T to look at the interval $t \in [3, 6]$:

```
params.T0 = 3.0
params.T = 6.0
```

and we replace the initial condition, using the `Restart`-class:

```
# Get restart data
restart = Restart(dict(casedir="../Basic/Results/"))
restart_data = restart.get_restart_conditions()

# Initial condition
U = restart_data.values()[0]["Temperature"]
```

Note that we point `Restart` to the case directory where the solution is stored. We could also choose to write our restart data to the same directory when setting up the postprocessor:

```
pp = PostProcessor(dict(casedir="../Basic/Results"))
```

3.3 Replay a Problem

Once a simulation is completed, one might want to compute other fields of the solution. This can be done with cbcposts `Replay`-functionality. The process can be done in very few lines of code.

In the following, we initialize a replay of the heat equation solved in *A Basic Use Case* and restarted in *Restart a Problem*. First, we set up a postprocessor with the fields we wish to compute:

```
from cbcpost import *
from dolfin import set_log_level, WARNING, interactive
set_log_level(WARNING)

pp = PostProcessor(dict(casedir="../Basic/Results"))

pp.add_fields([
    SolutionField("Temperature", dict(plot=True)),
    Norm("Temperature", dict(save=True, plot=True)),
```

```
    TimeIntegral("Norm_Temperature", dict(save=True, start_time=0.0, end_time=6.0)),
])
```

To *replay* the simulation, we do:

```
replayer = Replay(pp)
replayer.replay()
interactive()
```

CHAPTER 4

Functionality

The main functionality is handled with a `PostProcessor`-instance, populated with several `Field`-items.

The `Field`-items added to the `PostProcessor` can represent *meta* computations (`MetaField`, `MetaField2`) such as time integrals or time derivatives, restrictions or subfunction, or norms. They can also represent custom computations, such as stress, strain, stream functions etc. All subclasses of the `Field`-class inherits a set of parameters used to specify computation logic, and has a set of parameters related to saving, plotting, and computation intervals.

The `Planner`, instantiated by the PostProcessor, handles planning of computations based on Field-parameters. It also handles the dependency, and plans ahead for computations at a later time.

For saving purposes the PostProcessor also creates a `Saver`-instance. This will save Fields as specified by the Field-parameters and computed fields. It saves in a structured manner within a specified case directory.

In addition, there is support for plotting in the `Plotter`-class, also created within the PostProcessor. It uses either dolfin.plot or pyplot.plot to plot data, based on data format.

4.1 The `Field`-class and subclasses

To understand how cbcpost works, one first needs to understand the role of *Fields*. All desired postprocessing must be added to the PostProcessor as subclasses of `Field`. The class itself is to be considered as an abstract base class, and must be subclassed to make sense.

All subclasses are expected to implement (at minimum) the `Field.compute()`-method. This takes a single argument which can be used to retrieve dependencies from other fields.

An important property of the `Field`-class, is the parameters. Through the `Parameterized`-interface, it implements a set of default parameters that is used by the PostProcessor when determining how to handle any given Field, with respect to computation frequency, saving and plotting.

4.1.1 Subclassing the `Field`-class

To compute any quantity of interest, one needs to either use one of the provided metafields or subclass `Field`. In the following, we will first demonstrate the simplicity of the interface, before demonstrating the flexibility of it.

A viscous stress tensor

The viscous stress tensor for a Newtonian fluid is computed as

$$\sigma(\mathbf{u}, p) = -p\mathbb{I} + \mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$$

where μ is the dynamic viscosity, $\mathbf{u}$ is the fluid velocity and p is the pressure. A Field to compute this might be specified as the following:

```
from dolfin import *
from cbcpost import Field
from cbcpost.spacepool import get_grad_space
class Stress(Field):
    def __init__(self, mu, params=None, name="default", label=None):
        Field.__init__(self, params, name, label)
        self.mu = mu

    def before_first_compute(self, get):
        u = get("Velocity")

        # Create Function container on space of velocity gradient
        V = get_grad_space(u)
        self._function = Function(V, name=self.name)

    def compute(self, get):
        u = get("Velocity")
        p = get("Pressure")
        mu = self.mu

        expr = - p*Identity(u.cell().d) + mu*(grad(u)+grad(u)^T)

        return self.expr2function(expr, self._function)
```

Note that we have overridden three methods defined in `Field`:

- __init__
- before_first_compute
- compute

The __init__ method is only used to pass any additional arguments to our Field, in this case the viscosity. The keyword arguments *params*, *name* and *label* are passed directly to `Field.__init__()`.

before_first_compute is used to do any costly computations or allocations that are only required once. This is called from the postprocessor before any calls to compute is made. In this case we create a container (*_function*) that we can later use to store our computations. We use the *get*-argument to fetch the field named *Velocity*, and the helper function `get_grad_space()` to get the gradient space of the Velocity (a TensorFunctionSpace).

The compute method is responsible for computing our quantity. This is called from the postprocessor every time the `Planner` determines that this field needs to be computed. Here we use the *get*-argument to fetch the *Velocity* and *Pressure* required to compute the stress. We formulate the stress, and converts to a function using the helper function `Field.expr2function()`.

Computing the maximum pressure drop

In this next section, we demonstrate some more functionality one can take advantage of when subclassing the `Field`-class. In a flow, the maximum pressure drop gives an indication of the forces involved in the flow. It can be written as

$$\tilde{p} := \max_{t\in[0,T]} \left(\max_{\mathbf{x}\in\Omega} p(\mathbf{x},t) - \min_{\mathbf{x}\in\Omega} p(\mathbf{x},t)\right)$$

A `Field`-class to compute this can be implemented as

```
from dolfin import *
from cbcpost import Field
from cbcpost.spacepool import get_grad_space
class PTilde(Field):
    def add_fields(self):
        return [ Maximum("Pressure"), Minimum("Pressure") ]

    def before_first_compute(self, get):
        self._ptilde = 0.0
        self._tmax = 0.0

    def compute(self, get):
        pmax = get("Maximum_Pressure")
        pmin = get("Minimum_Pressure")
        t = get("t")

        if pmax-pmin > self._ptilde:
            self._ptilde = pmax-pmin
            self._tmax = t

        return None

    def after_last_compute(self, get):
        return (self._ptilde, self._tmax)
```

Here, we implement two more `Field`-methods:

- add_fields
- after_last_compute

The add_fields method is a convenience function to make sure that dependent Fields are added to the postprocessor. This can also be handled manually, but this makes for a cleaner code. Here we add two fields to compute the (spatial) `Maximum` and `Minimum` of the pressure.

The method after_last_compute is called when the compution is finished. This is determined by the time parameters (see *Parameters*), and handled within the postprocessors `Planner`-instance.

4.1.2 Field names

The internal communication of fields is based on the name of the `Field`-instances. The default name is

```
[class name]-[optional label]
```

The label can be specified in the *__init__*-method (through the *label*-keyword), or a specific name can be set using the *name*-keyword.

When subclassing the `Field`-class, the default naming convention can overloaded in the `Field.name`-property.

4.1.3 The *get*-argument

In the three methods *before_first_compute*, *compute* and *after_last_compute* a single argument (in addition to *self*) is passed from the postprocessor, namely the *get*-argument. This argument is used to fetch the computed value from other fields, through the postprocessor. The argument itself points to the `PostProcessor.get()`-method, and is typically used with these two arguments:

- Field name

- Relative timestep

A call using the *get*-function will trigger a computation of the field with the given name, and cache it in the postprocessor. Therefore, a second call with the same arguments, will return the cached value and not trigger a new computation.

The calls to the *get*-function also determines the dependencies of a Field (see *Dependency handling*).

4.1.4 Parameters

The logic of the postprocessor relies on a set of parameters defined on each Field. For explanation of the common parameters and their default, see `Field.default_params()`.

4.1.5 SolutionField

The `SolutionField`-class is a convenience class, for specifying Field(s) that will be provded as solution variables. It requires a single argument as the name of the Field. Since it is a solution field, it does not implement it does not implement a *compute*-method, but relies on data passed to the `PostProcessor.update_all()` for its associatied data. It is used to be able to build dependencies in the postprocessor.

4.1.6 MetaField and MetaField2

Two additional base classes are also available. These are designed to allow for computations that are not specific (such as PTilde or Stress), but where you need to specify the Field(s) to compute on.

Subclasses of the `MetaField`-class include for example `Maximum`, `Norm` and `TimeIntegral`, and takes a single name (or Field) argument to specify which Field to do the computation on.

Subclasses of the `MetaField2` include `ErrorNorm`, and takes two name (or Field) arguments to specify which Fields to compute with.

4.1.7 Provided fields

Several meta fields are provided in cbcpost, for general computations. These are summarized in the following table:

Time dependent	Spatially restricted	Norms and averages	Other
TimeDerivative	SubFunction	DomainAvg	Magnitude
TimeIntegral	Restrict	Norm	
TimeAverage	Boundary PointEval	ErrorNorm Maximum Minimum	

For more details of each field, refer to *Implemented cbcpost.metafields*.

4.2 The postprocessor

The `PostProcessor`-class is responsible for all the logic behind the scenes. This includes logic related to:

- Dependency handling
- Planning and caching of computation
- Saving
- Plotting

The planning, saving and plotting is delegated to dedicated classes (`Planner`, `Saver` and `Plotter`), but is called from within a `PostProcessor`-instance.

4.2.1 The *update_all*-function

The main interface to the user is through the `PostProcessor.update_all()`-method. This takes three arguments: a *dict* representing the solution, the solution time and the solution timestep.

The time and timestep is used for saving logic, and stored in a *play log* and metadata of the saved data. This is necessary for the replay and restart functionality, as well as order both the saved and plotted fields.

The solution argument should be of the format:

```
solution = dict(
    "Velocity": lambda: u
    "Pressure": lambda: p
)
```

Note that we pass a lambda function as values in the dict. This is done to give the user the flexibility for special solvers, and can be replaced with any callable to do for example a conversion. This can be useful when there are discrepancies between the solver solution, and the desired *physical* solution. This could be for example a simple scaling, or it could be that a mixed or segregated approach is used in the solver.

Because this function might be non-negligible in cost, it will be treated in the same manner as the `Field.compute()`-method, and not called unless required.

4.2.2 Dependency handling

When a field is added to the postprocessor, a dependency tree is built. These dependencies represent the required fields (or time parameters) required to succesfully execute the *compute*-method.

The source code of the *compute*-function is inspected with the *inspect*-module, by looking for calls through the *get*-argument, and build a dependency tree from that.

Assume that the following code is executed:

```
pp = PostProcessor()
pp.add_field(SolutionField("F"))
pp.add_field(TimeDerivative("F"))
```

In that case, when the `TimeDerivative`-field is added to the postprocessor, the following code is inspected:

```
class TimeDerivative(MetaField):
    def compute(self, get):
        u1 = get(self.valuename)
        u0 = get(self.valuename, -1)

        t1 = get("t")
        t0 = get("t", -1)

        # ... [snip] ...
```

By evaluating the *get*-calls here, we are able to build the following dependency tree:

If we extend the above example to add the time derivative of the viscous stress tensor (see *A viscous stress tensor*) like the following:

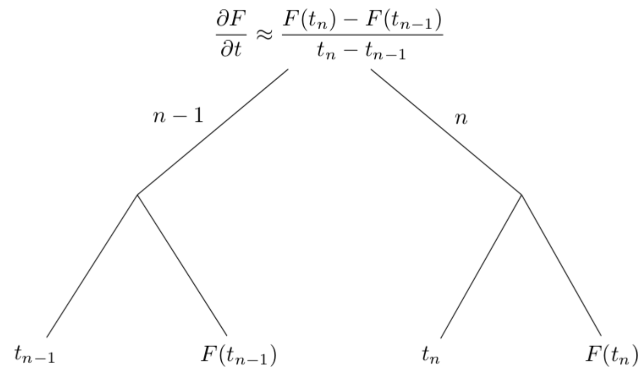


```
pp = PostProcessor()
pp.add_fields([SolutionField("Velocity"), SolutionField("Pressure")])
pp.add_field(Stress())
pp.add_field(TimeDerivative("Stress"))
```

The first emphasized line will trigger building of the dependency tree for the stress:

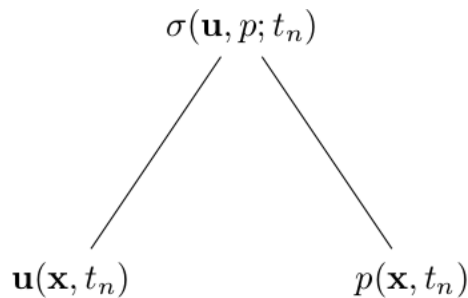

while the second emphasized line will use this dependency tree, and trigger the building of the larger dependency tree

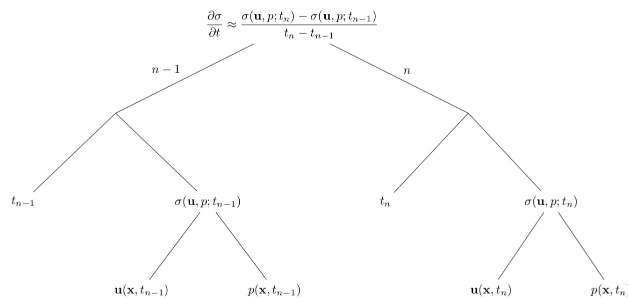

4.2.3 Planner

The `Planner`-class will set up a plan of the computations for the coming timesteps. This algorithm will inspect the dependencies of each field, and compute the necessary fields at the required time.

In addition, it determines how long each computation should be kept in cache.

Note: This does not yet support variable timestepping.

4.2.4 Saver

The `Saver`-class handles all the saving operations in cbcpost. It will determine if and how to save based on Field-parameters. In addition, there are helper methods in `PostProcessor` for saving mesh and parameters.

For fields, several saveformats are available:

Replay/restart-compatible	Visualization	Plain text
hdf5	xdmf	txt
xml	pvd	
xml.gz		
shelve		

The default save formats are:

- *hdf5* and *xdmf* if data is dolfin.Function
- *txt* and *shelve* if data is float, int, list, tuple or dict

The saving is done in a structured manner below the postprocessors case director. Consider the following example:

```
pp = PostProcessor(dict(casedir="Results/"))
pp.add_fields([
    SolutionField("Pressure", save=True),
    Norm("Pressure", save=True),
])
pp.store_mesh(mesh, facet_domains=my_facet_domains, cell_domains=my_cell_domains)
pp.store_params(
    ParamDict(
        mu = 1.5,
        case = "A",
        bc = "p(0)=1",
    )
)
```

Here, we ask the postprocessor to save the Pressure and the (L2-)norm of the pressure, we store the mesh with associated cell- and facet domains, and we save some (arbitrary) parameters. (Note the use of `ParamDict`).

This will result in the following structure of the *Results*-folder:

4.2.5 Plotter

Two types of data are supported for plotting:

- dolfin.Function-type objects
- Scalars (int, float, etc)

The `Plotter`-class plots using dolfin.plot or pyplot.plot depending on the input data. The plotting is updated each timestep the Field is directly triggered for recomputation, and rescaled if necessary. For dolfin plotting, arguments can be passed to the dolfin.plot-command through the parameter *plot_args*.

4.3 Replay

One of the key functionalities of the cbcpost framework is the ability to replay problem. Consider the case where one wants to extract additional information from a simulation. Simulations are typically costly, and redoing simulations are not generally desired (or even feasible). This motivates the functionality to *replay* the simulation by loading the computed solution back into memory and compute additional fields.

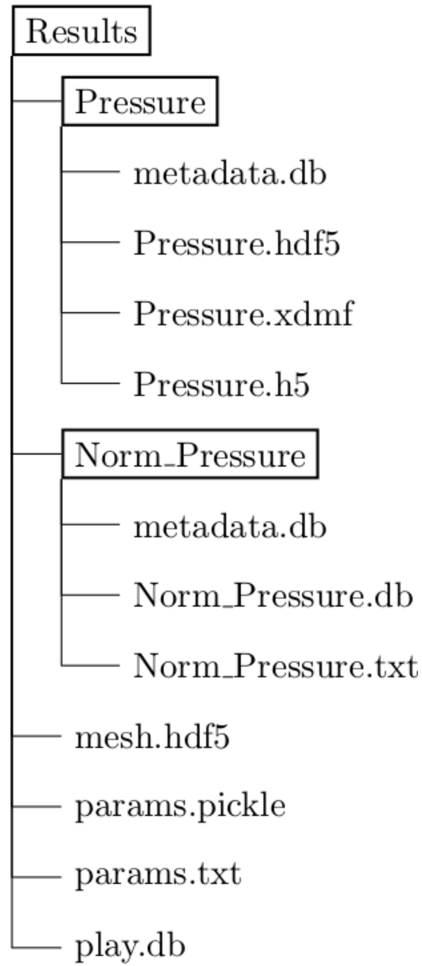

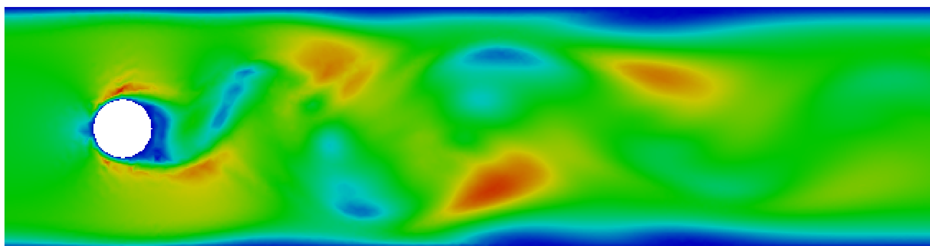

Fig. 4.1: dolfin.Function objects are plotted with dolfin.plot

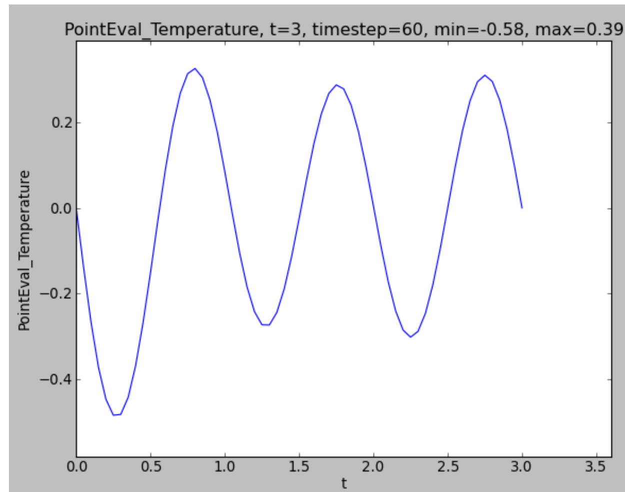

Fig. 4.2: Simple scalars are plotted with pyplot

This has several major benefits:

- Compute additional quantities
- Limit memory consumption of initial computation
- Compute quantities unsupported in parallel
- Compute costly, conditional quantities (e.g. not to be performed if simulation was unable to complete)
- Create visualization data

The interface to the replay module is minimal:

```
from cbcpost import PostProcessor, Replay

pp = PostProcessor(dict(casedir="ExistingResults/"))
pp.add_field(MyCustomField(), dict(save=True))

replayer = Replay(pp)
replayer.replay()
```

In the replay module, all fields that are stored in a reloadable format will be treated as a solution. They will be passed to a postprocessor as instances of the `Loadable`-class. This makes sure that no unnecessary I/O-operations occur, as the stored data are only loaded when they are triggered in the postprocessor.

4.4 Restart

The restart functionality lets the user set up a problem for restart. This functionality is based on the idea that a restart of a simulation is nothing more than changing the initial conditions of the problem in question. Therefore, the `Restart`-class is used to extract the solution at any given time(s) in a format that may be used as intiial conditions.

If we want to restart any problem, where a solution has been stored by cbcpost, we can simply point to the case directory:

```
from cbcpost import *
restart = Restart(dict(casedir='Results/'))
restart_data = restart.get_restart_conditions()
```

If you for instance try to restart the simple case of the heat equation, *restart_data* will be a *dict* of the format {t0: {"Temperature": U0}}. If you try to restart for example a (Navier-)Stokes-problem, it will take a format of {t0: {"Velocity": U0, "Pressure": P0}}.

There are several options for fetching the restart conditions.

4.4.1 Specify restart time

You can easily specify the restart time to fetch the solution from:

```
t0 = 2.5
restart = Restart(dict(casedir='Results/', restart_times=t0))
restart_data = restart.get_restart_conditions()
```

If the restart time does not match a solution time, it will do a linear interpolation between the closest existing solution times.

4.4.2 Fetch multiple restart times

For many problems, initial conditions are required at several time points prior to the desired restart time. This can be handled through:

```
dt = 0.01
t1 = 2.5
t0 = t1-dt
restart = Restart(dict(casedir='Results/', restart_times=[t0,t1]))
restart_data = restart.get_restart_conditions()
```

4.4.3 Rollback case directory for restart

If you wish to write the restarted solution to the same case directory, you will need to clean up the case directory to avoid write errors. This is done by setting the parameter *rollback_casedir*:

```
t0 = 2.5
restart = Restart(dict(casedir='Results/', restart_times=t0, rollback_casedir=True))
restart_data = restart.get_restart_conditions()
```

4.4.4 Specifying solution names to fetch

By default, the Restart-module will search through the case directory for all data stored as a `SolutionField`. However, you can also specify other fields to fetch as restart data:

```
solution_names = ["MyField", "MyField2"]
restart = Restart(dict(casedir='Results/', solution_names=solution_names))
restart_data = restart.get_restart_conditions()
```

In this case, all `SolutionField`-names will be ignored, and only restart conditions from fields named *MyField* and *MyField2* will be returned.

4.4.5 Changing function spaces

If you wish to restart the simulation using different function spaces, you can pass the function spaces to *get_restart_conditions*:

```
V = FunctionSpace(mesh, "CG", 3)
restart = Restart(dict(casedir='Results/'))
restart_data = restart.get_restart_conditions(spaces={"Temperature": V})
```

Note: This does not currently work for function spaces defined on a different mesh.

4.5 Utilities

A set of utilities are provided with cbcpost. Below are just a few of them. For a more complete set of utilities, refer to the *Programmer's reference*.

4.5.1 The `ParamDict`-class

The `ParamDict`-class extends to the stadard python *dict*. It supports dot-notation (*mydict["key"] == mydict.key*), and nested parameters.

Todo

Extend this documentation.

4.5.2 The `Parameterized`-class

The `Parameterized`-class is sued for classes that are associated with a set of parameters. All subclasses must implement the method `Parameterized.default_params()`, which return a ParamDict/dict with default values for the parameters.

When initialized, it takes a *params*-option where specific parameters are set, and overwriting the associated parameters returned from *default_params*. This is then stored in an attribute *params* attached to the object.

When initializing a Parameterized-object, no new keys are allowed. This means that all parameters of a Parameterized-instance must be defined with default values in *default_params*.

The class is subclasses several places within cbcpost:

- `Field`
- `PostProcessor`
- `Restart`
- `Replay`

4.5.3 Pooling of function spaces

When using many different functions across a large function, it may be useful to reuse FunctionSpace definitions. This has two basic advantages:

- Reduced memory consumption
- Reduced computational cost

Space pools are grouped according to mesh, with *Mesh.id()* used as keys in a *weakref.WeakValueDictionary*. Once a mesh is out of focus in the program, the related SpacePool is removed.

4.5.4 Submesh creation

The `SubMesh`-class in dolfin is not currently supported in dolfin. In cbcpost, the function `create_submesh()` is the equivalent functionality, but with parallel support.

This allows for arbitrary submeshes in parallel based by providing a MeshFunction and marker.

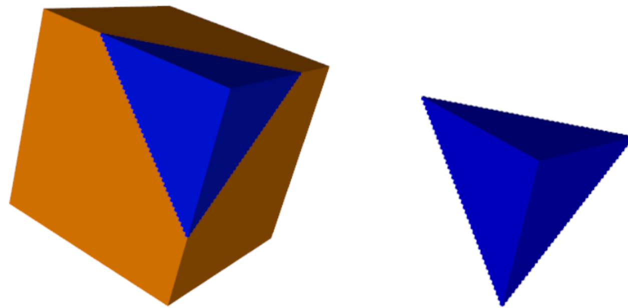

Fig. 4.3: Submesh created with *create_submesh* in cbcpost.

4.5.5 Mesh slicing

Three-dimensional meshes can be sliced in cbcpost with the Slice-class. The `Slice`-class takes basemesh, togetther with a point and normal defining the slicing plane, to create a slicemesh.

The `Slice`-class is a subclass of dolfin.Mesh.

Warning: Slice-instances are intended for visualization only, and may produce erronous results if used for computations.

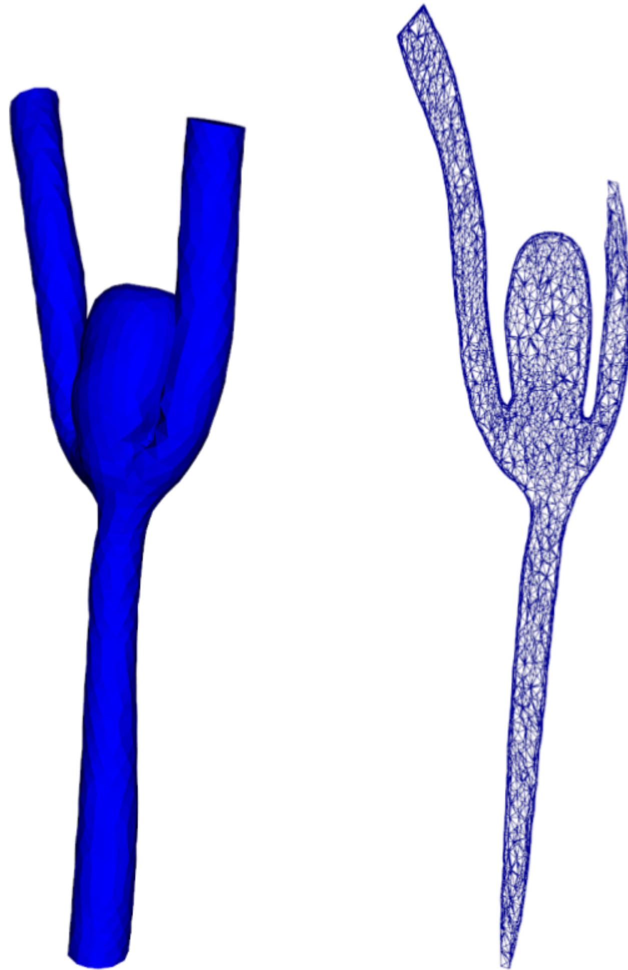

Fig. 4.4: A complex 3D-mesh, with an associated slicemesh.

CHAPTER 5

Programmer's reference

5.1 Postprocessor

5.1.1 Saver

5.1.2 Planner

5.1.3 Plotter

5.2 Replay

5.3 Restart

5.4 Parameter system

5.4.1 ParamDict

5.4.2 Parameterized

5.5 FunctionSpace pooling

5.6 Field bases

5.7 Implemented cbcpost.metafields

5.7.1 `Boundary`

5.7.2 `DomainAvg`

5.7.3 `ErrorNorm`

5.7.4 `Magnitude`

5.7.5 `Maximum`

5.7.6 `Minimum`

5.7.7 `Norm`

CHAPTER 6

Contributing

cbcpost is located at

> https://bitbucket.org/simula_cbc/cbcpost

A *git* workflow is used, and the language is python. There are no strict guidelines followed, but as a rule of thumb, consider the following:

- No '*' imports
- No unused imports
- Whitespace instead of tabs
- All modules, public functions and classes should have reasonable docstrings
- 4 space indentation levels
- Unit tests should cover the majority of the code
- Look at existing code, and use common sense

cbcpost will follow the development of FEniCS, and will likely require development versions of FEniCS-components between releases.

6.1 Pull requests

If you wish to fix issues or add features, please check out the code using git, make your changes in a clean branch and create a pull request.

Before making a pull request, make sure that all unit tests pass, and that you add sufficient unit tests for new code.

To avoid unnecessary work, please contact the developers in advance.

6.2 Report problems

Please report any bugs or feature requests to the issue tracker on Bitbucket.

6.3 Contact developers

You can contact the developers directly:

- Øyvind Evju
- Martin Sandve Alnæs.

CHAPTER 7

Indices and tables

- genindex
- modindex
- search