

Time Integration in BOUT++

Nick Walkden on behalf of Ben
Dudson and the BOUT++ Team

- Numerical Time Integration
- Time solvers in BOUT++
- Optimization: Physics based preconditioning
- Summary

- A general evolution equation of the form

$$\partial_t \mathbf{f} = \mathbf{F}(\mathbf{f}) \quad \square \quad \mathbf{f} = \begin{pmatrix} n \\ v \\ \text{etc} \end{pmatrix}$$

can be solved numerically by discretizing the time domain

- This can be done equivalently by

\square *Explicit methods*

$$\partial_t \mathbf{f} \approx \frac{\mathbf{f}^{n+1} - \mathbf{f}^n}{\Delta t} \approx \mathbf{F}(\mathbf{f}^n)$$

\square *Implicit methods*

$$\partial_t \mathbf{f} \approx \frac{\mathbf{f}^n - \mathbf{f}^{n-1}}{\Delta t} \approx \mathbf{F}(\mathbf{f}^n)$$

or by combining the two using

\square *ImEx methods*

$$\partial_t \mathbf{f} = \mathbf{F}_{slow}(\mathbf{f}) + \mathbf{G}_{fast}(\mathbf{f}) = \mathbf{F}_{exp}(\mathbf{f}) + \mathbf{G}_{imp}(\mathbf{f})$$

Explicit Methods

$$\square \quad \partial_t f \approx \frac{f^{n+1} - f^n}{\Delta t} \approx F(f^n)$$

Good for:

- Single timescales
- Non-linear problems

Bad for:

- Stiff problems

Advantages:

- Easy to implement
- Relatively little cost per timestep

Disadvantages:

- Stability (ie CFL condition)

Implicit Methods

$$\square \quad \partial_t f \approx \frac{f^n - f^{n-1}}{\Delta t} \approx F(f^n)$$

Good for:

- Stiff problems/multiple timescales

Bad for:

- Heavily non-linear problems

Advantages:

- Unconditionally stable
- Large timesteps

Disadvantages:

- Complexity
- Cost per timestep
- Under-relaxation

ImEx Methods

$$\partial_t f = F_{exp}(f) + G_{imp}(f)$$

Good for:

- Stiff problems/multiple timescales

Bad for:

Advantages:

- Timestep only limited by slow scales

Disadvantages:

- Complexity
- Cost per timestep
- Dependant on physics model

- Implicit Methods are often cast as a Newton iteration

$$\square \mathbf{G}(\mathbf{f}^n) = \mathbf{f}^n - \mathbf{f}^{n-1} - \gamma \mathbf{F}(\mathbf{f}^n) = 0$$

can be Taylor expanded to give

$$\square \mathbf{G}(\mathbf{f}^n) \approx \mathbf{G}(\mathbf{f}^n_m) + (\mathbf{f}^n - \mathbf{f}^n_m) \frac{\partial \mathbf{G}}{\partial \mathbf{f}^n} = 0$$

where

$$\square \frac{\partial \mathbf{G}}{\partial \mathbf{f}^n} = \mathbf{I} - \gamma \mathbf{J} \quad \square \mathbf{J} = \frac{\partial \mathbf{F}}{\partial \mathbf{f}^n} = \begin{pmatrix} \frac{\partial F_1}{\partial f_1} & \dots & \frac{\partial F_N}{\partial f_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial F_1}{\partial f_N} & \dots & \frac{\partial F_N}{\partial f_N} \end{pmatrix}$$

- Implicit Methods are often cast as a Newton iteration

So we can now solve the problem

$$(\mathbf{I} - \gamma \mathbf{J}) \delta \mathbf{f}^n = -\mathbf{G}(\mathbf{f}^n)$$

which we use to update

$$\mathbf{f}_{m+1}^n = \mathbf{f}_m^n + \delta \mathbf{f}^n$$

iteratively until

$$\mathbf{G}(\mathbf{f}^n) = 0$$

- To advance in time, time derivatives of evolving quantities are required
 - All fields store the derivatives in another field called `deriv` which can be accessed through `var.timeDeriv()`

```
Field3D var;  
Field3D *deriv = var.timeDeriv()
```

- When we call `ddt(var)` we are actually calling the inline function (`include/field3d.hxx` line 346)

```
inline Field3D& ddt(Field3D &f){  
    return *(f.timeDeriv()); }  
}
```

- To advance in time, time derivatives of evolving quantities are required
 - All fields store the derivatives in another field called `deriv` which can be accessed through `var.timeDeriv()`

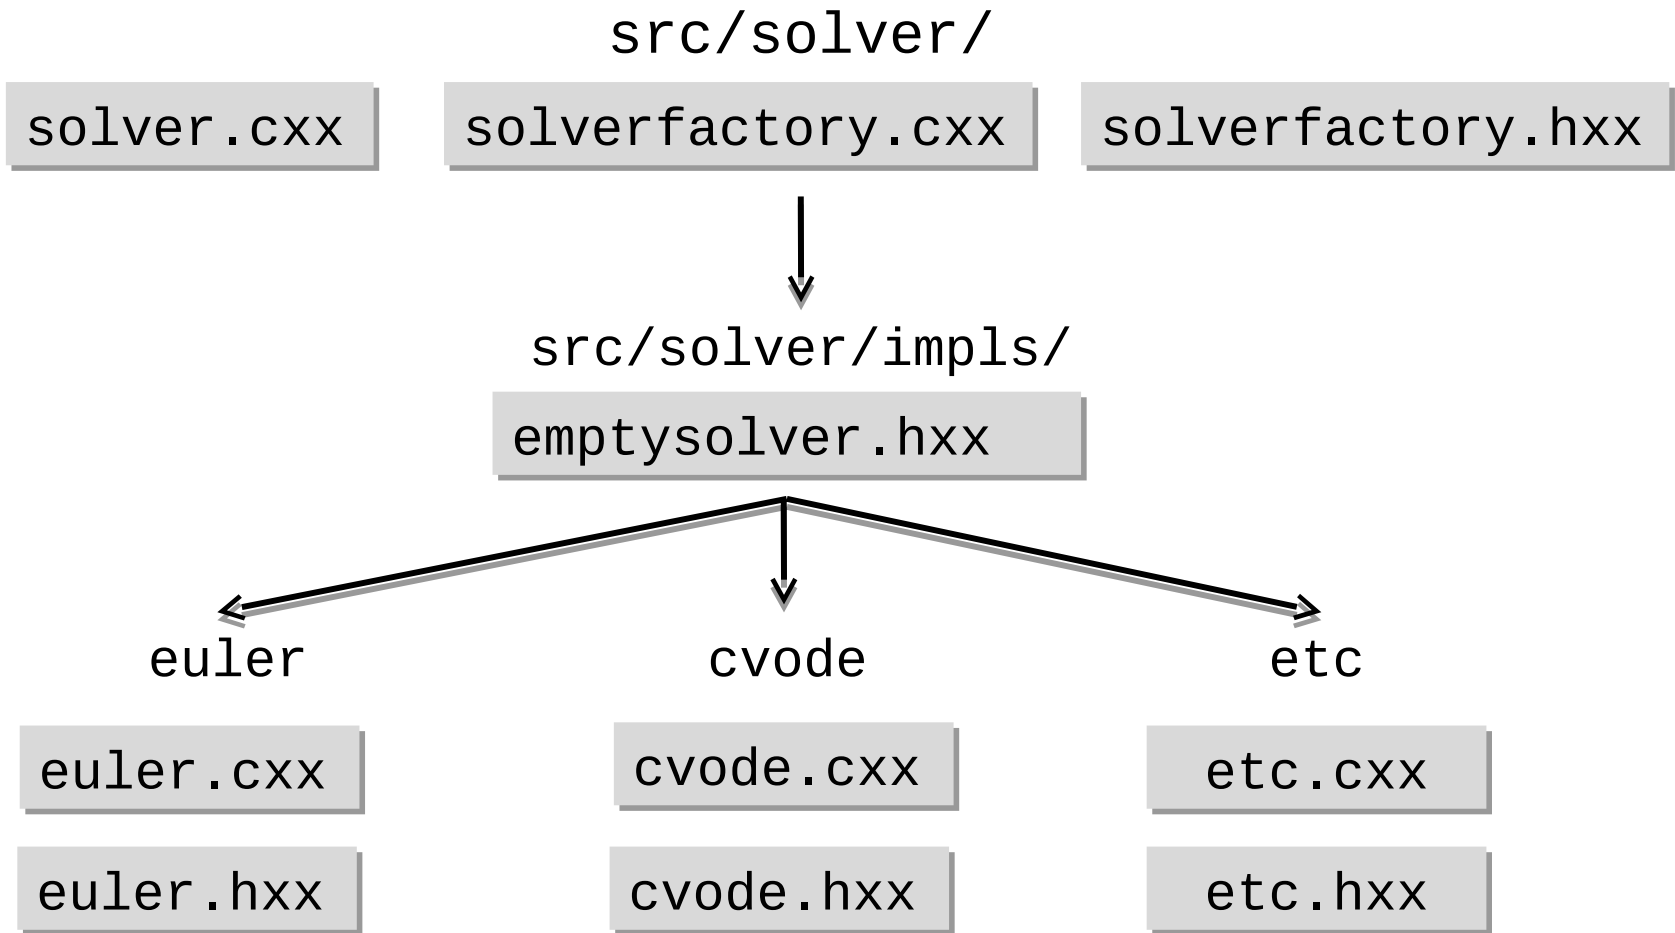
This allows us to treat `ddt(var)` as a variable in our code so that we can write statements like

$$\text{ddt}(\text{var}) = \dots + \dots$$

```
inline Field3D& ddt(Field3D &f){  
    return *(f.timeDeriv()); }  

```

- In BOUT++ the solver has a factory format



- In BOUT++ the solver has a factory format

src/solver/

`solver.cxx`

`solverfactory.cxx`

`solverfactory.hxx`

The solver base class contains a set of base routines such as loading and saving variables

It also contains a set of virtual functions which must be contained in the individual solver implementations

euler

cvode

etc

`euler.cxx`

`cvode.cxx`

`etc.cxx`

`euler.hxx`

`cvode.hxx`


`etc.hxx`

- In BOUT++ the solver has a factory format



The solver factory is the only bit that knows about the implementations so the rest of the code is forced to be independent of the solver choice

It builds a solver out of whichever implementation has been chosen by the user allowing the user to change solver at runtime (ie without a recompilation)



- In BOUT++ the solver has a factory format

src/solver/

`solver.cxx`

`solverfactory.cxx`

`solverfactory.hxx`

Different time-stepping techniques are included as separate implementations of the base solver class

Adding a new time integration method should require minimal changes to the rest of the code

euler

cvode

etc

`euler.cxx`

`cvode.cxx`

`etc.cxx`

`euler.hxx`

`cvode.hxx`

`etc.hxx`

- BOUT++ has a range of solver implementations:

| Explicit Solvers | Implicit Solvers | ImEx Solvers |
|------------------|------------------|---------------|
| euler | pvode | arkode |
| rk4 | cvode | BDF multistep |
| rk3-ssp | petsc (various) | |
| karniadakis | ida | |
| power | | |

The default behaviour is to use either CVODE or IDA if present, otherwise use PVIDE

- BOUT++ has a range of solver implementations:

| Explicit Solvers | Implicit Solvers | ImEx Solvers |
|------------------|------------------|---------------|
| euler | pvode | arkode |
| rk4 | cvode | BDF multistep |
| rk3-ssp | petsc (various) | |
| karniadakis | ida | |
| power | | |

Require external libraries

The default behaviour is to use either CVODE or IDA if present, otherwise use PVODE

- BOUT++ has a range of solver implementations:
- The solver type is set in the BOUT.inp file or on the command line

In BOUT.inp:

```
[solver]  
type = ...
```

On the command line

```
./executable  
solver:type=...
```

- Usually best to try a few (ie RK4 and pvoke) to find the optimal choice for your case

- Some additional useful options are:

| Option | Function |
|--------------------------------|--|
| <code>mxstep = 500</code> | Number of timesteps to try before timestep is a failure |
| <code>atol = 1e-10</code> | Absolute tolerance used to determine error norm. Determines noise level of solution |
| <code>rtol = 1e-5</code> | Relative tolerances used to determine error norm. Indicates no of digits of relative accuracy for a single time step |
| <code>adaptive = false</code> | Use adaptive timestepping in rk4 |
| <code>use_precon = true</code> | Use preconditioning in ccode and petsc |

- To tell the time-solver to solve for a particular field we use the routine

```
bout_solve(n, "density");
```

or the macro `SOLVE_FOR(n);`

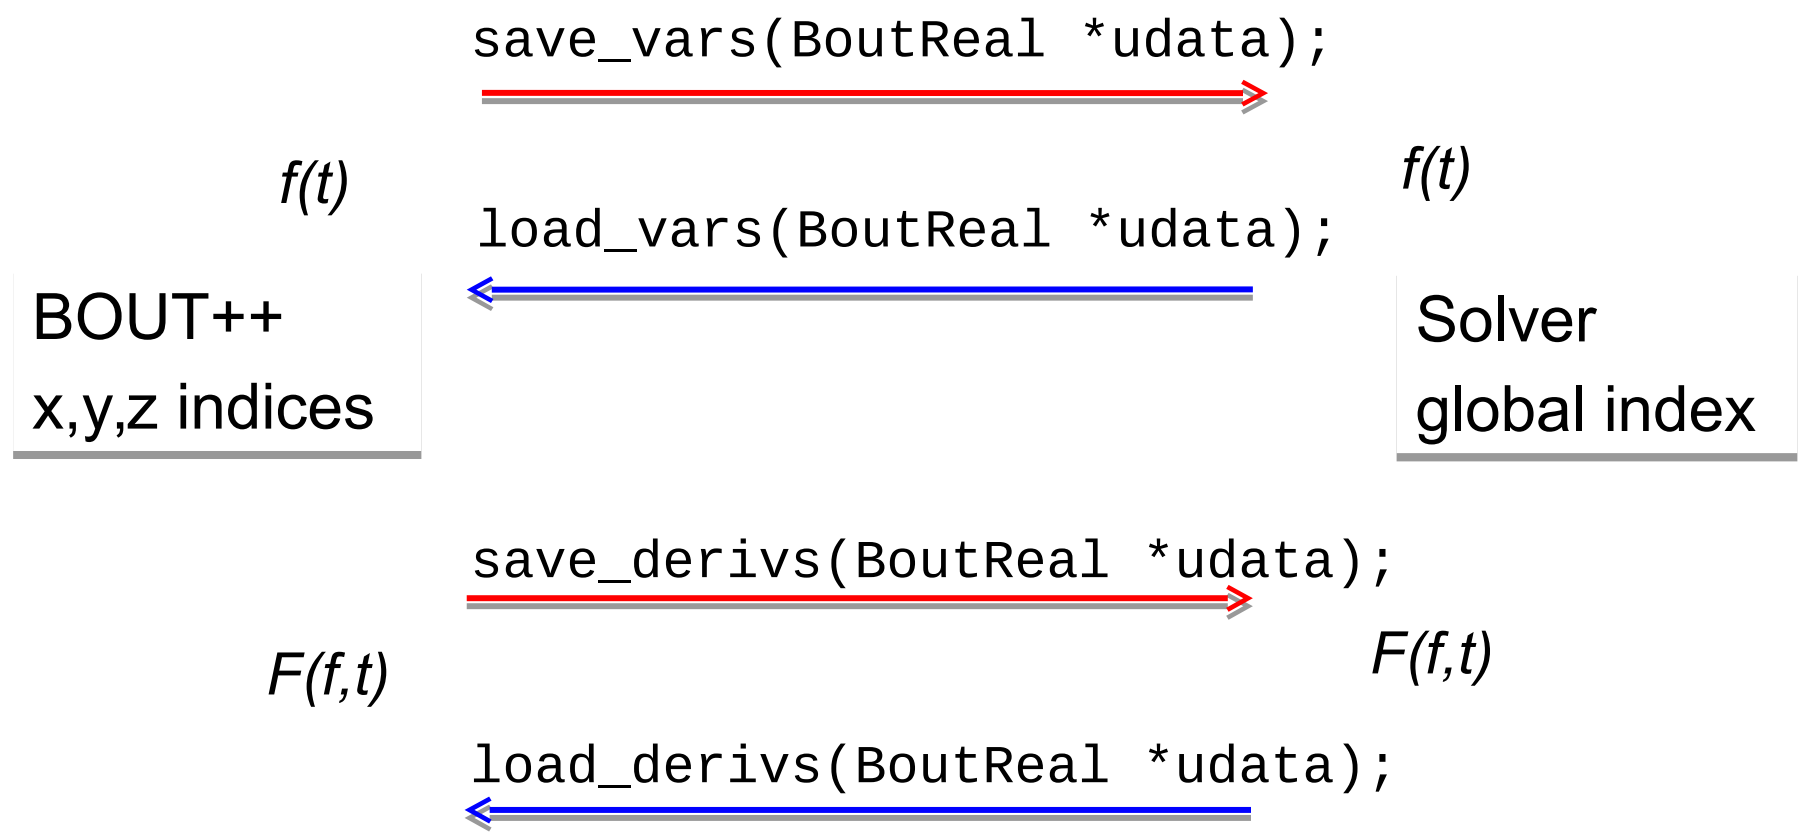
or similarly for $2 \leq n \leq 6$ `SOLVE_FORn(n, ...);`

- This calls a routine in the solver class

```
solver->add(n, "density");
```

which adds the variable to the state and residual vectors for input to the time-solver

- BOUT++ data is then be passed to the solver through a few protected functions



- There are two mandatory functions that a time-solver implementation must contain

```
int mysolver::init(bool restarting, int nout, BoutReal  
                  tstep)
```

Initialization of the solver. Calls a generic solver initialization as well as implementation specific solver options.

- There are two mandatory functions that a time-solver implementation must contain

```
int mysolver::init(bool restarting, int nout, BoutReal  
                  tstep)  
int mysolver::run()
```

Running the solver. This function integrates in time until nout is reached and the simulation is over.

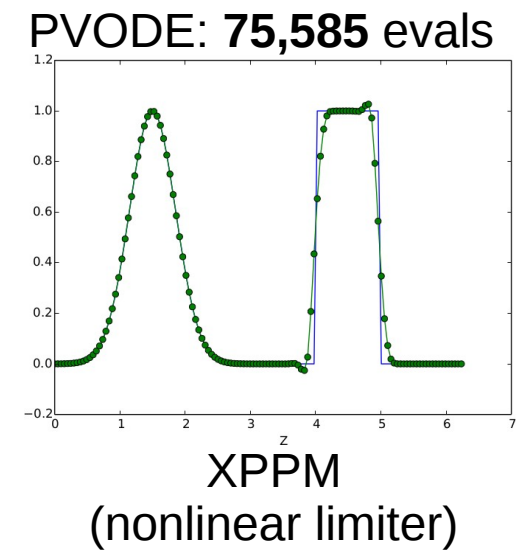
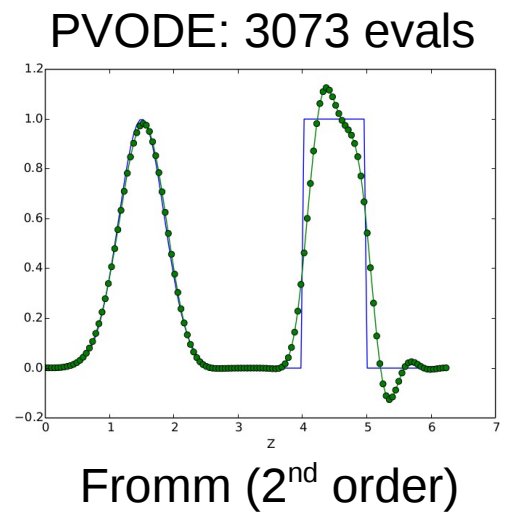
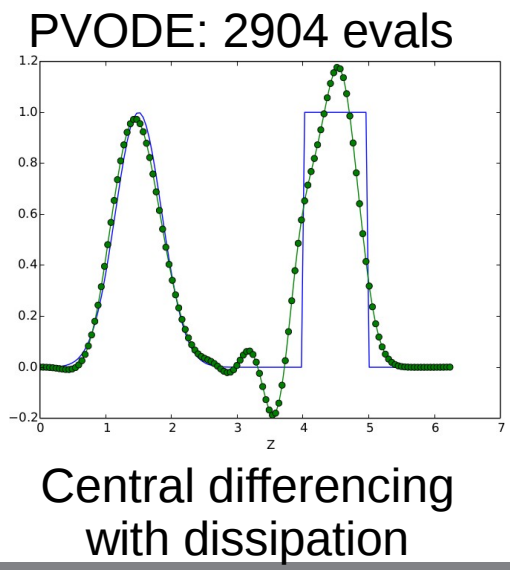
- There are two mandatory functions that a time-solver implementation must contain
- A call to `int bout_run(Solver *solver, rhsfunc physics_run)` (bout++.cxx, L 287) is made in the main function which runs the solver
- That's how the magic happens...

- For many problems the implicit PVIDE/CVIDE solvers work well
- In some cases however they can fail

Example: Advection of a pulse in 1D $\frac{\partial f}{\partial t} = -v \frac{\partial f}{\partial x}$

PVIDE integrator, absolute tolerance 1e-12, relative tolerance 1e-5

CFL condition limits explicit methods to 128 steps, or 512 evaluations for RK4
 The RK3-SSP method requires $dt < 0.2 dt(\text{CFL})$, or 1920 evaluations



- For ImEx schemes the split operators must be defined

C style interface

```
examples/split-operator/  
  
int physics_init(bool restart) {  
    Solver->setSplitOperator  
        (physics_run, reaction);  
    ...  
}  
  
int physics_run(BoutReal time) {  
    // Explicit part  
}  
  
int reaction(BoutReal time) {  
    // Implicit part  
}
```

C++ style interface

```
examples/test-drift/  
  
class DriftWave : public  
PhysicsModel {  
    int init(bool restart) {  
        SetSplitOperator();  
        ...  
    }  
    int convective(BoutReal time) {  
        // Explicit part  
    }  
    int diffusive(BoutReal time) {  
        // Implicit part  
    }  
}
```

"We should forget about small efficiencies, say about 97% of the time: premature optimization is the root of all evil. Yet we should not pass up our opportunities in that critical 3%." D. Knuth

- Optimization in BOUT++ is handled almost entirely internally and should rarely be tinkered with
- One exception to this is preconditioning

- Recall that the implicit time-solver(s) solves a newton iteration of the form

$$(I - \gamma J)\delta f = -G(f)$$

which requires the inversion of $(I - \gamma J)$

- It may be possible to define a new operator, P , such that

$$P(I - \gamma J)\delta f = -PG(f)$$

If $P(I - \gamma J)$ is easier to invert then we may speed up our solver!!

- We can use our knowledge of the physics system to help us here
- If we have some stiff physics on some timescale then we can try and construct P such that

$$P \approx (I - \gamma J)^{-1}$$

over

- This then means that over the time-scale the inversion is trivial

$$P(I - \gamma J)\delta f = -PG(f) \xrightarrow{\tau} I\delta f = -(I - \gamma J)^{-1}G(f)$$

- Problem: How do we find P ?
 1. Reduce equations
 2. Calculate analytical Jacobian matrix
 3. Factorize matrix using schur factorization
 4. Simplify the problem (decouple perpendicular and parallel dynamics)

1. Reduce equations

$$\square \quad \partial_t n = f(n, T) \quad \leftarrow \text{Slow (non-stiff) physics}$$

$$\partial_t T = g(n, T) + \partial_{||} \chi_{||} \partial_{||} T \quad \leftarrow \text{Fast (stiff) physics}$$

$$\chi_{||} = \chi_0 T^{5/2}$$

We want to isolate the fast physics, so consider the reduced system

$$\square \quad \partial_t n = 0$$

$$\partial_t T = \partial_{||} \chi_{||} \partial_{||} T$$

2. Calculate the Jacobian

$$\square J = \frac{\partial \mathbf{F}}{\partial \mathbf{f}} = \begin{pmatrix} \partial_n \partial_t n & \partial_n \partial_t T \\ \partial_T \partial_t n & \partial_T \partial_t T \end{pmatrix}$$

For our reduced system

$$\square J = \begin{pmatrix} 0 & 0 \\ 0 & \chi_0 T^{5/2} \partial_{||}^2 \end{pmatrix}$$

So

$$\square I - \gamma J = \begin{pmatrix} 1 & 0 \\ 0 & 1 - \gamma \chi_0 T^{5/2} \partial_{||}^2 \end{pmatrix}$$

3. Factorize the matrix

Using a technique called Schur factorization

$$\square \begin{pmatrix} E & U \\ L & D \end{pmatrix}^{-1} = \begin{pmatrix} I & -E^{-1}U \\ 0 & I \end{pmatrix} \begin{pmatrix} E^{-1} & 0 \\ 0 & S^{-1} \end{pmatrix} \begin{pmatrix} I & 0 \\ -LE^{-1} & I \end{pmatrix}$$

Where the Schur complement is

$$\square S = D - LE^{-1}U$$

□ For our case $E = 1, L = U = 0$ and $D = 1 - \gamma\chi_0 T^{\frac{5}{2}} \partial_{||}^2$ so

$$\square S^{-1} = \left(1 - \gamma\chi_0 T^{\frac{5}{2}} \partial_{||}^2 \right)^{-1}$$

- Problem: How do we find P ?
 1. Reduce equations
 2. Calculate analytical Jacobian matrix
 3. Factorize matrix using schur factorization
 4. Simplify the problem (decouple perpendicular and parallel dynamics)

We now have the preconditioner operator P

$$P = \begin{pmatrix} 1 & 0 \\ 0 & S^{-1} \end{pmatrix} \text{ where } S^{-1} = \left(1 - \gamma \chi_0 T^{\frac{5}{2}} \partial_{\parallel}^2 \right)^{-1}$$

- Problem: How do we apply the preconditioner?
 5. Create precon function
 6. Implement preconditioning operator on time derivatives
 7. Tell the solver to use a preconditioner
 8. Tinker with the preconditioner until you get some speedup

5. Write a precon function

At the end of your physics module you now need a function called

```
int precon(BoutReal t, BoutReal gamma, BoutReal delta)
```

This function must apply the preconditioning operator to the time derivatives of the evolving variables

`t` is the current simulation time, `gamma` is the methods timestep and `delta` may be used to apply constraints (but rarely used, so don't worry about it)

6. Apply preconditioner to time derivatives

To apply the preconditioner we will need to invert a matrix in the parallel direction. BOUT++ has a class to do this, called `InvertPar` which solves

```
InvertPar *precon_inv;  
int physics_init(bool restarting){  
...  
precon_inv = InvertPar::create();  
precon_inv->setCoefA(1.0);  
...  
}
```

6. Apply preconditioner to time derivatives

Now in the precon function we apply the operator

```
int precon(BoutReal t, BoutReal gamma, BoutReal
delta){

mesh->communicate(ddt(T));
Field 2D B;
B = -gamma*chi*(T.DC())^(5./2.);
precon_inv->setCoefB(B);
ddt(T) = precon_inv->solve(ddt(T));
ddt(T).applyBoundary("neumann_o2");

}
```

7. Tell the solver to use the preconditioner

In `BOUT.inp` we need the lines:

```
[solver]  
type = ccode #or petsc  
use_precon = true  
rightprec = false
```

In `physics_init` we need the line:

```
solver->setPrecon(precon);
```

7. Tell the solver to use the preconditioner

In a 1D high powered SOL slab equilibrium calculation on 8 cores

| Solver | Setup | Wall time (s) | ~Iteration count |
|--------|-----------------------|---------------|------------------|
| CVODE | Isothermal | 2 | 100 |
| CVODE | Conduction removed | 5 | 300 |
| PVODE | None | 221 | 37000 |
| CVODE | None | 276 | 45000 |
| CVODE | BBD Preconditioner | 58 | 17000 |
| CVODE | Custom Preconditioner | 9 | 800 |

Approximately 30x speedup with preconditioner

8. Tinker

Since the preconditioner affects convergence, but **not** the solution, you can tinker with it to find the best setup

WARNING: Because the preconditioner is problem dependant it will not always be beneficial

- BOUT++ contains a suite of time-solvers including explicit and implicit options
- The choice of solver is problem dependant but can be interchanged at runtime
- Physics based preconditioning can be used to optimize (implicit) solvers, but costs (some) blood