

This is a repository copy of *BOUT++:Recent and current developments*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/105679/>

Version: Accepted Version

Article:

Dudson, Benjamin Daniel orcid.org/0000-0002-0094-4867, Allen, Andrew Robert, Breyiannis, George et al. (12 more authors) (2015) *BOUT++:Recent and current developments*. *Journal of Plasma Physics*. ISSN 1469-7807

<https://doi.org/10.1017/S0022377814000816>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

BOUT++: Recent and current developments

**B. D. Dudson^{1†}, A. Allen¹, G. Breyiannis², E. Brugger³,
J. Buchanan⁴, L. Easy^{1,4}, S. Farley⁵, I. Joseph³, M. Kim⁶,
A. D. McGann¹, J. T. Omotani⁴, M. V. Umansky³,
N. R. Walkden^{1,4}, T. Xia^{3,7}, X. Q. Xu³**

¹York Plasma Institute, Department of Physics, University of York, Heslington, York, YO10 5DD, UK

²Japan Atomic Energy Agency, Rokkasho Fusion Institute, Rokkasho-mura, 039-3212, Japan

³Lawrence Livermore National Laboratory, Livermore, CA 94551, USA

⁴CCFE, Culham Science Centre, Abingdon, OX14 3DB, UK

⁵Mathematics Department, Illinois Institute of Technology, Chicago, IL 60616, USA

⁶Department of Physics, POSTECH, Pohang, Gyeongbuk 790-784, Korea

⁷Institute of Plasma Physics, Chinese Academy of Sciences, Hefei, China

(Received 19 August 2014)

BOUT++ is a 3D nonlinear finite-difference plasma simulation code, capable of solving quite general systems of Partial Differential Equations (PDEs), but targeted particularly on studies of the edge region of tokamak plasmas. BOUT++ is publicly available, and has been adopted by a growing number of researchers worldwide. Here we present improvements which have been made to the code since its original release, both in terms of structure and its capabilities. Some recent applications of these methods are reviewed, and areas of active development are discussed. We also present algorithms and tools which have been developed to enable creation of inputs from analytic expressions and experimental data, and for processing and visualisation of output results. This includes a new tool HYPNOTOAD for the creation of meshes from experimental equilibria.

Algorithms have been implemented in BOUT++ to solve a range of linear algebraic problems encountered in the simulation of reduced Magnetohydrodynamics (MHD) and gyro-fluid models: A preconditioning scheme is presented which enables the plasma potential to be calculated efficiently using iterative methods supplied by the PETSc library (the Portable, Extensible Toolkit for Scientific Computation) (Balay *et al.* 2014), without invoking the Boussinesq approximation. Scaling studies are also performed of a linear solver used as part of physics-based preconditioning to accelerate the convergence of implicit time-integration schemes.

1. Introduction

The edge region of tokamak plasmas (Wesson 1997; Stangeby 2000) is of crucial importance to their feasibility and economic viability as fusion reactors. It is the interface between the hot ($\sim 1 - 10$ keV), fully ionised core required for fusion, and the plasma facing material surfaces which must remain cool (< 1 eV) to avoid excessive damage or sputtering of impurities, which could contaminate the core plasma. The wide range of temperatures and hence collisionality regimes; nonlinear plasma dynamics; atomic ionisation, charge-exchange and recombination processes; impurities; and interaction with material surfaces, make modelling the plasma edge challenging. This is further complicated

† Email address for correspondence: benjamin.dudson@york.ac.uk

by the magnetic geometry, which is usually arranged into an 'X-point' configuration, in which the closed magnetic surfaces of the core are surrounded by open magnetic field-lines, which transport energy and particles leaving the core away into divertor regions designed to handle high heat fluxes.

Simulations of the edge of tokamaks which include many of the important atomic and impurity radiation processes are routinely performed using 1-D (Goswami *et al.* 2001; Fundamenski *et al.* 2001; Nakamura 2011) and 2-D (Rognlien *et al.* 1992; Schneider *et al.* 2006; Kawashima *et al.* 2006) codes. These however cannot correctly predict the plasma transport across the magnetic field, which is usually anomalous (turbulent) (D'Ippolito *et al.* 2011), and for which diffusion (Fick's law) is a poor approximation (Garcia *et al.* 2007; Naulin *et al.* 2007; Gendrih *et al.* 2012). The need for a first-principles understanding of cross-field transport in the edge region has received increasing attention in recent years, and several new 3-D codes have been developed to study edge turbulence (Ricci *et al.* 2012; Ricci & Rogers 2013; Tamain *et al.* 2010). As discussed above, the structure and dynamics of the edge region of tokamaks involves a complicated interaction between many physical processes, and as a result it is not clear *a priori* which model or combination of models is most appropriate. To minimise duplication of effort, there is a need for a flexible code which can be adapted to solve a range of different models, and is modular enough that it can be extended in multiple ways by a large group of users. BOUT++ (Dudson *et al.* 2009) is an open-source 3D nonlinear finite difference code which aims to fill that need.

BOUT++ was developed originally to study tokamak edge plasma physics (Dudson *et al.* 2009), taking ideas and lessons learned from the earlier BOUT code (Xu *et al.* 1999, 2000; Umansky *et al.* 2006; Xu *et al.* 2008). It is highly modular, operates in general curvilinear coordinates and complicated mesh topologies, and can be applied to the solution of quite general Partial Differential Equations (PDEs) in three dimensions + time. Recent applications of BOUT++ include the study of plasma transients (Edge Localised Modes, ELMs) (Zohm 1996; Xu *et al.* 2010; Xia *et al.* 2012; Xi *et al.* 2013), plasma turbulence (Friedman *et al.* 2012), and the dynamics of isolated 'blobs' in 3D (Angus *et al.* 2012*b,a*; Walkden *et al.* 2013).

The BOUT++ distribution is publicly available on Github†, and comes with a test suite and variety of plasma physics models and examples. Some have been used to produce results published elsewhere (e.g. ELM, turbulence in the LArge Plasma Device (Gekelman *et al.* 1991), and blob models), whilst others can be used as a starting point for new physics studies. This paper describes the 2.0 release of BOUT++, which was used as a basis for the 2013 BOUT++ workshop‡: Section 2 briefly describes modifications to the structure of BOUT++ which have been made to accommodate further development; Section 3 describes improvements and new capabilities which have been added to BOUT++ since its original release (Dudson *et al.* 2009). Section 4 details the development of pre- and post-processing tools for equilibrium input and visualisation. In section 5 we conclude and discuss future directions for development.

2. Code structure

The BOUT++ development community has expanded significantly following its release and 2011 workshop, and with that has come the need to adopt more professional software development practices: Git¶ is used for version control, along with a system of

† BOUT++ public distribution <http://github.com/boutproject>

‡ BOUT++ workshop 2013: <http://bout2013.llnl.gov>

¶ Git version control: <http://git-scm.com/>

feature branches adopted from the PETSc development group (Balay *et al.* 1997) which is described in detail on the BOUT++ development page^{||}. An important addition has been a test suite which can be run quickly before changes are committed, to check that nothing obvious has been broken. This has greatly simplified the process of checking code, and has resulted in many bugs being caught before they could affect production code. The majority of these tests are not physics simulations, as these would take too long to perform and so discourage regular use. Instead, tests are designed to check individual components independently for a range of inputs and processor configurations, so that the cause of a test failure can be quickly identified.

2.1. Factory pattern

To enable BOUT++ to be extended, and new implementations of solvers for boundary and initial-value problems to be added independently, components have been refactored and organised along the Factory pattern (Gamma *et al.* 1995; Knepley 2012), a widely used method to separate interface from implementation. Each component of BOUT++, such as file I/O or time integration solver, has a well defined interface (in C++ a base class with virtual members). Several implementations of this interface can coexist, and the user code doesn't depend on which implementation is used. To create a particular instance, the static member function “create” is called. For example, time-integration schemes implement the “Solver” interface, so creating a new solver is done by the following:

```
Solver *s = Solver::create();
```

so that `s` is a pointer (hence the `*`) to a `Solver` object. Which particular instance of `Solver` is created (RK4, PETSc etc.) and assigned to `s` is set by options stored in a tree structure, which can be set in the input file or on the command-line. By default the options section for the `Solver` class is called “solver”, so to choose the rk4 method on the command-line the user adds

```
solver:type=rk4
```

To allow multiple solvers with different settings to be used simultaneously, the option section can be passed during creation:

```
Solver *s = Solver::create(
    Options::getRoot()->getSection("mysolver")
);
```

The Options tree works in a similar manner to a file system: `Options::getRoot()` returns a pointer to the root (lowest level) of the options tree structure, analogous to the root “/” directory of a UNIX file system; `getSection()` traverses the tree by returning pointers to sub-trees, and is analogous to changing to a sub-directory using “cd mysolver”. By passing this sub-tree to `Solver::create()`, the result is that options for this solver will be taken from the “mysolver” section of the options tree, and can be changed independently.

This pattern enables users to experiment with different numerical methods with minimal changes to the program inputs. As new capabilities are added to BOUT++, such as new PDE and ODE solvers, existing models can take advantage of them without needing to modify any code, only the input settings. To the extent possible, this separation of

^{||} BOUT++ development model <http://boutproject.github.io/devel.html>

interface and implementation reduces the number of dependencies between parts of the code and allows researchers to benefit from each other’s work on separate components. A similar system is used by the Trilinos project (Heroux *et al.* 2005), in which parameter lists are passed to objects on creation in order to allow changing of options at run-time.

2.2. Library interface

The interface between the core BOUT++ code and problem-specific “user” code has also been modified since the original publication. The original BOUT++ was structured as a framework so that the `main()` function was defined internally, and the user supplied two functions: one for initialisation of the desired physics model, and one which calculated the time-derivative of each evolving variable given the system state. Further callback functions were later added for optional preconditioning (section 3.3) and system Jacobian calculations. This method was simple to implement, and familiar to those with a background in C programming, but caused complications when combining BOUT++ with other libraries and frameworks. In addition to this original style, an object-oriented style is now supported: Rather than callback functions, users implement a class which inherits from `PhysicsModel`, overriding the default functions as needed. A simple example is a diffusion equation in 1-D, which could be implemented in the following code:

```
class Diffusion : public PhysicsModel {
private:
    Field3D T;
protected:
    void init(bool restarting) {
        SOLVE_FOR(T);
    }
    void rhs(double time) {
        ddt(T) = Laplace_par(T);
    }
};
```

This defines a 3D scalar field (`Field3D`) T ; specifies that T should be evolved in the initialisation function `init`; and then calculates the time-derivative as $\frac{\partial T}{\partial t} = \nabla \cdot (\mathbf{b}_0 \mathbf{b}_0 \cdot \nabla T)$ in the function `rhs`, in which `ddt()` should be read as $\partial/\partial t$. Users can use a macro `BOUTMAIN` to define a standard `main()` function, or define their own to enable BOUT++ to be combined with other libraries. Separating physics models into classes allows the possibility of combining several models into a single simulation, for example a model for neutral gas with a plasma model, and could be exploited for multiscale simulations.

The above example illustrates some other minor improvements which have been made to BOUT++: `ddt()` and `SOLVE_FOR` are preprocessor macros, which are used sparingly wherever the resulting improvement in readability outweighs their potential for causing hard-to-find bugs.

3. Solvers and capabilities

As increasingly sophisticated plasma simulation models are studied with BOUT++, in particular gyro-fluid extensions (Xu *et al.* 2013), the range of differential operators which needs to be solved has expanded, and the computational cost of the simulations has increased. To address this, new capabilities have been added to BOUT++: More general

elliptic solvers for calculating the electrostatic potential ϕ are presented in section 3.1, and an algorithm to solve parabolic equations along magnetic field lines is presented in section 3.2. This latter solver has been used as part of a physics-based preconditioning strategy (Mousseau *et al.* 2000; Knoll *et al.* 2001) to improve convergence for large time steps, described in section 3.3, and to calculate closure terms for gyro-fluid operators (Dimitis *et al.* 2013). Wherever possible, these new solvers have been implemented using the factory pattern (section 2.1) and a generic interface, so that they can be reused in many plasma models and geometries.

3.1. Calculation of potential ϕ from vorticity

Reduced MHD models solved in BOUT++ are commonly formulated in terms of a vorticity equation, from which the electrostatic potential is calculated. In reduced MHD, this can be derived from either the momentum equation or charge conservation (current continuity) (Hazeltine & Meiss 2003; Catto & Simakov 2004). Gyro-fluid models, which are an area of current research in BOUT++ (Xu *et al.* 2013), can also be cast in a vorticity formulation (Ottaviani & Manfredi 1999), or the potential can be calculated from a polarisation equation coupling electron and ion gyro-centre densities (Beer *et al.* 1997; Scott 2005a). In either case, the electrostatic potential ϕ is calculated by solving an equation of the form:

$$\nabla \cdot \left(\frac{m_i n}{B^2} \nabla_{\perp} \phi \right) = \omega \quad (3.1)$$

with ion mass m_i , plasma density n , and magnetic field strength B . The time evolution of the right hand side ω (vorticity) depends on the particular model. The coefficient n/B^2 arises from the ion polarisation, and in general will vary in 3D as the density n is an evolving quantity.

Elliptic equations of the same form as equation 3.1 arise in many fields, and so numerical methods for their solution have been extensively studied. There are therefore many different methods available in the literature (e.g. (Iserles 2009)); the challenge is in finding one which is efficient enough for practical applications. Solving for ϕ requires the solution of a linear (matrix) problem for every evaluation of the time-derivatives of the system, which will usually be several times per time step. A typical turbulence or ELM simulation might require $10^4 - 10^6$ time steps, and so efficient solution to equation 3.1 is critical for the overall run-time of the simulation.

A common approximation in plasma simulations is to neglect the variation of n/B^2 in space and/or time, referred to as the Boussinesq approximation (Yu *et al.* 2006). BOUT++ simulations have usually replaced the full density n in equation 3.1 with the axisymmetric (constant in toroidal angle) equilibrium density n_0 . Since B is also axisymmetric, the left hand side of equation 3.1 can then be Fourier transformed in toroidal angle, decoupling the toroidal harmonics. Each toroidal harmonic can then be solved efficiently as a 1D tridiagonal system of complex equations in the radial coordinate. This scheme will be referred to here as the FFT or Boussinesq method.

The effect of the Boussinesq approximation on simulation results can be subtle, and its importance depends on the problem being considered. As stressed elsewhere (Scott 2002, 2003, 2005b) the energetics of a model are important for long time simulations of turbulence. Unless the Boussinesq approximation is introduced carefully, it can result in an unphysical source of energy which can grow and eventually dominate the simulation. In other cases, the Boussinesq approximation has been found to have little effect, for example in BOUT++ simulations of blobs (Angus & Umansky 2014) it was found that the Boussinesq approximation made only a small difference, and that in 3D the drift-wave dynamics made a greater difference to the result.

To remove the Boussinesq approximation, and allow the solution to the full vorticity equation, BOUT++ has been coupled to the PETSc library (Balay *et al.* 1997, 2010, 2014). Here we present details of the numerical scheme, and leave exploration of the impact on simulations of plasma phenomena to a future publication.

3.1.1. Iterative solution with PETSc

By discretising equation 3.1, the calculation of ϕ from ω can be cast as a linear algebra problem of the form

$$\mathcal{A}x = b \quad (3.2)$$

In BOUT++ this discretisation is done by Finite Differences, but other choices such as Finite Element are used elsewhere. The resulting matrix can then be solved using the Krylov subspace (KSP) iterative solvers available in PETSc, such as GMRES. Iterative methods are attractive because they have smaller memory requirements than direct solvers, as \mathcal{A} need never be explicitly stored, and in principle iterative methods can be parallelised more efficiently.

In general, equation 3.1 will couple all points in the domain, so the $N \times N$ sparse matrix has a size $N \simeq 10^6$. By neglecting derivatives parallel to the magnetic field, which are assumed small in drift-ordered fluid (Hazeltine & Meiss 2003) and gyro-fluid (Ottaviani & Manfredi 1999) models, this can be simplified to solving multiple independent $N \simeq 10^4$ matrices. Since the density n is evolving in time, the coefficients in this matrix change every time step. This makes direct solution methods based on matrix factorisation inefficient, as the matrix must be frequently re-calculated and re-factored. It is for this reason that iterative methods have been implemented in BOUT++, as these do not require the explicit calculation of the matrix elements or costly matrix factorisations.

3.1.2. Preconditioning of iterative solver

When solving large and/or ill-conditioned problems, iterative solvers can fail to converge, or converge very slowly after a small number of iterations (referred to as stalling). To accelerate convergence, an approximate solver is often used to “precondition” the problem, improving the condition of the matrix which the iterative solver is inverting. See for example (Golub & Van Loan 2013). A preconditioner \mathcal{P} is an approximate inverse of \mathcal{A} , which needs to be calculated quickly for the overall scheme to be efficient. The equation above can be multiplied through by \mathcal{P} as a *left preconditioner*:

$$(\mathcal{P}\mathcal{A})x = \mathcal{P}b \quad (3.3)$$

or as a *right preconditioner*:

$$(\mathcal{A}\mathcal{P})(\mathcal{P}^{-1}x) = b \quad (3.4)$$

and this modified system is solved using the iterative method. In the limit that the preconditioner \mathcal{P} is the inverse of \mathcal{A} , $\mathcal{P}\mathcal{A}$ is the identity, and no iterations should be required.

To test preconditioning methods, a 2D (x, z) test case was used with a density profile of the form:

$$n = \sin(x) e^{-x^2} (1 + p \cos(z)) \quad (3.5)$$

where the radial coordinate x goes between 0 and 1, and the constant p in the above equation is adjusted to change the variation of density with toroidal angle z . When $p = 0$ density does not vary with toroidal angle, and so the Boussinesq approximation is exact, but as p is increased the approximation will break down. Results are shown in table 1 for a small 40×32 mesh. For three values of p the time taken in seconds is given, followed by

TABLE 1. Timing for solution on a 40×32 mesh. Shown are the wall clock times in seconds, and the iteration counts in brackets

| Preconditioner | Density perturbation | | |
|----------------|----------------------|-------------|-------------|
| | 10% | 50% | 90% |
| None | 0.157 (319) | 0.110 (226) | 0.142 (299) |
| Jacobi | 0.162 (318) | 0.113 (224) | 0.137 (299) |
| SOR | 0.022 (30) | 0.048 (35) | 0.048 (40) |
| FFT | 0.013 (4) | 0.015 (5) | 0.024 (9) |

the iteration count in brackets for a single solve, as the perturbation size is increased from 10% to 90% (i.e. p varies from 0.1 to 0.9). Without preconditioning the iterative method requires $\simeq 300$ iterations and $\simeq 1.5$ seconds to converge, compared to a time of $\simeq 2$ ms for a single Boussinesq solve. To improve on this, several “black box” preconditioning schemes are available in PETSc, such as Jacobi iteration or Successive Over Relaxation (SOR) methods (Iserles 2009). These methods are not problem specific, and so require only a run-time switch to enable and configure. For the problems tested, the SOR method reduced the run-time (see table 1), but the number of iterations remained prohibitive. To improve on these, a problem-specific preconditioner has been implemented.

As discussed above, the purpose of a preconditioner is to quickly find an approximate inverse to the linear problem (matrix \mathcal{A} above). The Boussinesq FFT-based solver is just such a solver, as it finds a fast solution by simplifying the coefficients. We therefore use the FFT solver as a preconditioner for the full problem, by wrapping the FFT solver in a PETSc PCShell preconditioner object which is then passed to PETSc to be used in the iterative solver. This preconditioner is extremely good when the density perturbation is small, but we should expect it to become less effective as the size of the density perturbation becomes large. This is what is observed in Table 1: For small perturbations (10%), the run-time is a little over half that of the SOR method, but as the density perturbation amplitude increases so does the iteration count and run-time for all methods. Even at 90% density perturbation, however, this preconditioner is still highly effective. Based on this small test, a larger study was performed to compare the SOR and FFT (Boussinesq) preconditioners.

For a 516×256 mesh more typical of ELM and turbulence calculations, the iterative solver will typically stall without good preconditioning, as shown in figure 1. Using a 90% density perturbation (reducing the accuracy of the FFT preconditioner), and using the FFT solver result as the starting point for the iterative solver gives the convergence shown in figure 1. Without preconditioning, or using Jacobi or SOR preconditioners, the residual is only reduced by a factor of 10 in nearly 10^4 iterations; using the FFT-based preconditioner the residual is reduced by a factor of $\simeq 10^5$ in 10 iterations. This shows that for large meshes the FFT-based solver provides a good rate of convergence for realistic problem sizes, even when the deviation from axisymmetry is large.




FIGURE 1. Iterative KSP solver residual, as a function of iteration number. Solving elliptic equation 3.1 on a 516×256 mesh.

3.2. Parabolic solver along magnetic field-lines

In the reduced MHD and gyro-fluid models which BOUT++ specialises in solving, the magnetosonic fast wave is removed analytically, and so the fastest physical processes are usually the shear Alfvén wave and heat conduction along (parallel to) magnetic field-lines. Preconditioning of either of these processes for implicit time integration (see section 3.3) requires the solution to a parabolic equation of the form (Chacón *et al.* 2002)

$$\left(\mathcal{A} + \mathcal{B}\partial_{\parallel 0}^2\right)x = b \quad (3.6)$$

where $\partial_{\parallel 0} = \mathbf{b}_0 \cdot \nabla$ is the derivative along the equilibrium magnetic field \mathbf{b}_0 . Even though equation 3.6 appears to be a one-dimensional problem, due to the structure of the equilibrium magnetic field, it is in general a two-dimensional problem: The equilibrium magnetic field in a tokamak is helical, and lies on nested toroidal surfaces. If the pitch angle of the magnetic field is such that it makes an irrational number of poloidal to toroidal turns, then a single field-line will fill the 2D surface.

If magnetic perturbations were included, so that $\partial_{\parallel 0}^2$ became $\partial_{\parallel}^2 = (\mathbf{b} \cdot \nabla)^2$ in equation 3.6, then the magnetic field lines in general no longer lie on magnetic flux surfaces, but fill a volume. This case is of significant interest, for example in studying the transport of heat in ELM crashes and in the presence of externally applied magnetic perturbations, but solving this more general problem is left to future work.

To solve equations of the form (3.6), a solver has been implemented in BOUT++ using a partitioning scheme described in (Austin *et al.* 2004), which has similarities to cyclic reduction (Hockney 1965; Temperton 1975): Two interface equations are calculated on each processor, one for each boundary. The reduced system of interface equations is gathered onto one processor and solved using the serial Thomas algorithm. The resulting solution for the boundary values on each processor are then scattered back and used to solve for the remainder of the solution. As with many schemes for the solution of tridi-

agonal systems, this method is suitable for diagonally dominant systems. This condition is satisfied for equation 3.6 when coefficient \mathcal{B} is negative, and coefficient \mathcal{A} is positive definite, which is the case for all cases of interest here.

In the following we assume that the magnetic flux surface has N points in toroidal angle, and M points in poloidal angle. Firstly we exploit the toroidal symmetry of the equilibrium to decompose the problem into Fourier harmonics in toroidal angle ϕ . This then decomposes the problem into N complex tridiagonal systems, one for each toroidal mode, each of size M . If the domain is a closed magnetic surface, then the tridiagonal systems are cyclic, with a complex phase shift between the first and last row which is determined by the pitch of the magnetic field-lines. Each of these N systems of equations therefore has the form:

$$\left(\begin{array}{ccccccc} b_0 & c_0 & & & & & a_0 \\ a_1 & b_1 & c_1 & & & & \\ & \ddots & \ddots & \ddots & & & \\ & & a_{m-1} & b_{m-1} & c_{m-1} & & \\ \hline & & & a_m & b_m & c_m & \\ & & & & \ddots & \ddots & \ddots \\ c_{M-1} & & & & & a_{M-1} & b_{M-1} \end{array} \right) x = b \quad (3.7)$$

The domain is split between processors in the poloidal direction θ , with m rows per processor, illustrated by a horizontal line in the above equation. Within each processor the rows are eliminated, reducing the problem to two boundary equations for each processor, which forms a smaller (cyclic) tridiagonal matrix:

$$\left(\begin{array}{cccc} \beta_0 & \gamma_0 & & \alpha_0 \\ \alpha_1 & \beta_1 & \gamma_1 & \\ \hline & \alpha_2 & \beta_2 & \gamma_2 \\ \gamma_3 & & \alpha_3 & \gamma_3 \end{array} \right) \chi = \xi \quad (3.8)$$

where ξ and χ are the boundary values of b and x respectively. When solving N independent systems of equations (one for each toroidal mode), they are divided between processors, and all boundary equations for a given system are gathered onto a single processor. For example if N systems are split between 2 processors, then $N/2$ sets of boundary equations are gathered onto each processor. Once on a single processor, the serial Thomas algorithm (with Sherman-Morrison-Woodbury formula for cyclic tridiagonal systems (Temperton 1975)) is used to solve for the boundary values. These are then scattered back, and substituted into the original equation to obtain the solution inside each processor's domain.

The number of boundary equations for each system of equations, and number of communications is independent of the size of the problem M , and proportional to the number of processors. This means that the algorithm scales well with problem size, but uses gather and scatter operations which reduces performance on large numbers of processors. Scaling of the solver with processor number and problem size has been performed on HECToR, with 32 cores per node, and up to 2048 cores in total. Soft scaling, in which the problem size is increased proportionally with the number of processors, is shown in figure 2. The factor of $\simeq 10$ increase in wall clock time above 32 processors in figure 2 is because above this point communications occur between nodes, rather than solely within a single node. Waiting for global gather operations from across nodes is significantly slower than within




FIGURE 2. Soft scaling of parallel solver. m is the number of rows per processor, and N is the number of separate problems which are inverted simultaneously, corresponding to the number of toroidal Fourier modes

a node, and so becomes a bottleneck: for more than 64 processors the run time becomes almost independent of problem size m .

For fewer than 32 processors, it can be seen that doubling the number of independent systems N has a similar effect to doubling the size of each system M , and the wall time is approximately proportional to the problem size. This can also be seen in figure 3: on single core the algorithm scales approximately linearly with problem size, for larger numbers of processors there is a constant offset which depends on the number of processors, but only weakly on the number of rows per processor. The result is that as the number of rows per processor m is increased, the scaling with processor number appears more favourable: In the limit that each processor only has a two rows, the algorithm becomes equivalent to gathering all rows onto one processor, and using a serial algorithm, which would result in a linear scaling with number of processors N_p . In figure 2, when $m = 5$, the wall time scales with $N_p^{0.77}$, whilst for $m = 40$ the exponent is $\simeq 0.20$.

This scaling analysis demonstrates that the use of boundary equations to reduce the problem within each domain to two rows, results in an approximately linear scaling with problem size on a fixed number of processors. The global gather and scatter operations used to solve these boundary equations is a bottleneck for large numbers of processors, and leads to poor parallel scaling. Improving this will be the subject of future investigation. A possible solution is to use the cyclic reduction algorithm to solve the boundary value equations, which would eliminate global gather/scatter operations in favour of more point-to-point communications. In practice, the current solver has been found to be sufficiently fast for simulations to-date: the domain is decomposed in both radial and poloidal directions, so the poloidal direction is typically not divided into more than 32 processors. In addition, the total run-time spent in solving parallel parabolic equations is a small fraction of the total. As a result, this method has enabled the use of algorithms which have overall good scaling (Dudson *et al.* 2012; Dimits *et al.* 2013).




FIGURE 3. Soft scaling of parallel solver showing same data as figure 2 for $N = 64$. Shows variation of wall-clock time with problem size per processor m for different numbers of processors N_p .

This tridiagonal solver has been used to implement physics-based preconditioning (Mousseau *et al.* 2000; Knoll *et al.* 2001) in BOUT++ (Dudson *et al.* 2012), which is discussed in section 3.3. It has also been used to implement gyro-fluid parallel closures such as those due to (Hammett & Perkins 1990) by approximating to high accuracy the $k/|k|$ using a small number of Lorentzians, which take the form of equation 3.6. This has enabled these closures to be applied in situations where Fourier methods cannot be used, such as global simulations in realistic toroidal geometry. Further details can be found in (Dimitis *et al.* 2013; Xu *et al.* 2013).

3.3. Time integration

The time integration schemes available in BOUT++ have been expanded to include explicit schemes (RK4, Euler (Iserles 2009)), a scheme derived by Karniadakis *et al.* (Karniadakis *et al.* 1991; Scott 2005b), and a range of implicit schemes through coupling to the SUNDIALS (SUite of Nonlinear and Differential/ALgebraic equation Solvers) (Hindmarsh *et al.* 2005) and PETSc (Portable, Extensible Toolkit for Scientific Computation) (Balay *et al.* 1997, 2010) libraries.

The need to solve increasingly complicated models at increasingly high resolution has made efficiency and parallel scaling of algorithms used important. Plasma models are typically stiff, meaning that explicit time-integration methods are limited to small time-steps relative to the time-scales of interest. Implicit time integration methods overcome this restriction, but require the solution of a large linear system ($N \times N$ where N is the number of evolving variables, typically 1 – 10 million) at every timestep. Unless this can be done efficiently, the overall time taken by an implicit solver may not be less than an explicit solver.

As discussed in section 3.1.2, a preconditioner is an approximate inverse of the large matrix being solved, which should be efficient to evaluate. At each timestep an implicit

time integration method is solving a nonlinear problem to find the state at the next time point. Using a scheme such as Newton's method, this nonlinear problem is reduced to one or more steps which require the solution to a linear problem of the form $\mathcal{A}\mathbf{x} = \mathbf{b}$ where \mathbf{b} is known, and depends on current and previous state and their time-derivatives; \mathbf{x} is related to the unknown state at the next timestep; and \mathcal{A} is a large matrix. Advancing a single time-step therefore involves an outer loop to solve the nonlinear problem, which contains an inner loop to find the solution to a series of linear problems. Preconditioning targets this inner loop, improving convergence of the linear solve in order to reduce the overall run time.

Physics-based preconditioning describes a family of approaches to deriving preconditioners, which use knowledge of the physical system to simplify the model equations (Mousseau *et al.* 2000; Knoll *et al.* 2001). The aim is to retain in the preconditioner only those processes (oscillatory or diffusive) which are limiting the timestep. This then improves the condition number of the matrix which the iterative (usually Krylov subspace) method has to solve. Because the iterative method converges towards the solution to the full system, approximations can be made in the preconditioner without affecting the final solution, only the convergence rate towards the solution. The approach we have followed is based on work by Chacon and others (Chacón *et al.* 2002). Preconditioning of the implicit solvers in BOUT++ (using SUNDIALS and PETSc libraries) has been implemented in BOUT++, described in (Dudson *et al.* 2012).

An example is the shear Alfvén wave which is present in all 3D turbulence models. In the simplest form of the reduced MHD equations, this wave can be described by the following coupled equations for vorticity ω (equation 3.1), and magnetic potential A_{\parallel} which describes the perturbed magnetic field $\delta\mathbf{B} = \nabla \times (\mathbf{b}_0 A_{\parallel})$:

$$\begin{aligned} \frac{\partial\omega}{\partial t} &= \nabla_{\parallel 0} (j_{\parallel}) & \frac{\partial A_{\parallel}}{\partial t} &= -\partial_{\parallel 0} \phi \\ \omega &= \nabla \cdot \left(\frac{m_i n}{B^2} \nabla_{\perp} \phi \right) & \nabla_{\perp}^2 A_{\parallel} &= -\mu_0 j_{\parallel} \end{aligned} \quad (3.9)$$

where $\nabla_{\parallel 0} f = \nabla \cdot (\mathbf{b}_0 f)$, $\partial_{\parallel 0} = \mathbf{b}_0 \cdot \nabla$, $\nabla_{\perp} f = \nabla f - \mathbf{b}_0 (\mathbf{b}_0 \cdot \nabla f)$, and $\nabla_{\perp}^2 f = \nabla \cdot (\nabla_{\perp} f)$. These equations can be combined to give a wave equation, which in the case that n/B^2 is a constant reduces to:

$$\frac{\partial^2 \omega}{\partial t^2} = \nabla_{\parallel 0} \frac{\partial j_{\parallel}}{\partial t} = \nabla_{\parallel 0} \left(\frac{1}{\mu_0} \nabla_{\perp}^2 \partial_{\parallel 0} \frac{B^2}{m_i n} \nabla_{\perp}^{-2} \omega \right) \quad (3.10)$$

By further assuming that the magnetic field \mathbf{B} varies slowly, and so neglecting derivative of B terms,

$$\nabla_{\parallel 0} \nabla_{\perp}^2 \partial_{\parallel 0} \nabla_{\perp}^{-2} \simeq \partial_{\parallel 0}^2 \quad (3.11)$$

and the shear Alfvén wave propagates only along the (equilibrium) magnetic field:

$$\frac{\partial^2 \omega}{\partial t^2} = \frac{B^2}{\mu_0 m_i n} \partial_{\parallel 0}^2 \omega = V_A^2 \partial_{\parallel 0}^2 \omega \quad (3.12)$$

where V_A is the Alfvén speed. In tokamak simulations, this speed can be $V_A \simeq 10^7$ m/s, severely restricting the time step for explicit time integration schemes. If hyperbolic equation 3.12, or the original equations 3.9 are solved implicitly, then this requires the solution to a parabolic equation. For example using a backwards Euler method:

$$\begin{pmatrix} \omega \\ \omega' \end{pmatrix}^{t+1} = \begin{pmatrix} \omega \\ \omega' \end{pmatrix}^t + \delta t \begin{pmatrix} 0 & 1 \\ V_A^2 \partial_{\parallel 0}^2 & 0 \end{pmatrix} \begin{pmatrix} \omega \\ \omega' \end{pmatrix}^{t+1} \quad (3.13)$$

the equation to be solved at each timestep is parabolic:

$$\left(1 - \delta t^2 V_A^2 \partial_{||0}^2\right) \begin{pmatrix} \omega \\ \omega' \end{pmatrix}^{t+1} = \begin{pmatrix} 1 & \delta t \\ \delta t V_A^2 \partial_{||0}^2 & 1 \end{pmatrix} \begin{pmatrix} \omega \\ \omega' \end{pmatrix}^t \quad (3.14)$$

and of the same form as equation 3.6. To precondition waves of this type, the parabolic solver discussed in section 3.2 can therefore be used.

The above procedure has converted a wave equation 3.12 into a parabolic equation, so that waves which are not resolved by the time integration, for which the Courant number is greater than 1, are slowed and damped. This is an inevitable consequence of implicit time integration (Chacón *et al.* 2002), rather than a property of this preconditioning scheme. Implicit time integration with preconditioning (physics-based or otherwise) is only effective if the timescales of interest are long compared with the smallest timescales in the system. This is almost always the case in simulations of tokamak plasmas, for which drift or resistive timescales of interest are usually slow compared to dynamics such as electron heat conduction and shear Alfvén wave propagation along magnetic field-lines.

The assumptions made to reduce the full set of equations 3.9 to wave equation 3.12 cannot be made in the calculation of time-derivatives for the full problem, as this would affect the solution, but they can be made in the preconditioner since this is used to find an approximate solution to accelerate convergence to the solution of the full set of equations. The full procedure to derive a preconditioner using Schur factorisation is described in (Dudson *et al.* 2012), and is somewhat more involved than outlined here, but it makes the same assumptions and so results in the same form of equations. The resulting preconditioner using the solver presented in section 3.2 has been found to result in significant speed-ups for sets of equations where the timestep was limited by the shear Alfvén wave and parallel heat conductivity (Dudson *et al.* 2012), reducing overall wall-clock time by an order of magnitude in some cases.

4. Pre- and Post-processing tools

A simulation code is of little use without the tools to create high quality inputs such as meshes, and to analyse and present the results of the simulations. For post-processing the most important requirement is to be able to read the simulation output data in the user’s language of choice. Routines to do this are now available for IDL[†], Python, Matlab, Octave, and Mathematica (Wolfram Research, Inc. 2014) as part of the public BOUT++ repository.

For most publications, 1D plots and 2D contours are sufficient, but there are occasions when the ability to visualise data in three dimensions is useful. In the early stages of a scientific investigation, seeing the entire simulation domain rather than slices through it, can help spot anomalies or unexpected features. When presenting results, particularly at conferences, 3-D visualisation can quickly convey a large amount of information. Wrappers have been developed to enable two scientific data visualisation packages to be used with BOUT++: Mayavi[‡] and VisIT[¶].

Tools have also been developed to enable more convenient input of initial profiles and sources (section 4.1), and processing of experimental equilibria into input meshes (section 4.2).

[†] IDL: <http://www.exelisvis.co.uk/ProductsServices/IDL.aspx>

[‡] Mayavi project, <http://code.enthought.com/projects/mayavi/>

[¶] VisIT tool <https://wci.llnl.gov/codes/visit>




FIGURE 4. A tree of generator objects (Abstract Syntax Tree) to evaluate the expression $(\dots)*\sin(3*y-z)$

4.1. Input expression parser

Many simulations do not require complex geometry, but are intended to study basic physical mechanisms in slab or cylindrical geometries. For these cases the initial conditions and parameters often follow analytic expressions. If these expressions can be stored in the input files rather than preprocessing scripts, then inputs can be modified more quickly, and a clearer record of simulation inputs is kept for later reference. Examples include 2D blob simulations, where the initial density profile is a Gaussian in x specified using:

```
function = 1 + 0.2*gauss(x-0.25, 0.1)
```

In slab simulations of forced reconnection, a helical external magnetic potential is applied to an initial sheared magnetic field. This external field can be specified in the input file using:

```
function = (1-4*x*(1-x))*sin(3*y - z)
```

A recursive descent parser with operator precedence (see e.g.(Aho *et al.* 2006; LLVM Project 2014)) is used to build an Abstract Syntax Tree (AST) of generator objects from the input text: A constant generator like '3' or 'pi' always returns the same value; a coordinate generator like 'x' returns a value depending on the cell location; and a binary operation generator like '+' or 'sin' depends on the value of its children generators. Part of the AST for the above example is shown in figure 4. This tree is then evaluated for each cell in the domain to obtain the required initial conditions. This method is not computationally efficient, but is only used during initialisation and provides the flexibility to adapt the code in future.

Initially a learning exercise in compilers, the capability to construct ASTs from inputs, then manipulate and execute them at runtime has proved to be useful for scientific applications, eliminating the need to write input preprocessing scripts in many cases. This is proving particularly useful for verification using the Method of Manufactured Solutions (Roache 1998; Salari & Knupp 2000), in which analytic expressions for sources

and boundary conditions need to be specified. This verification activity is ongoing, and will be presented in a future publication.

As demonstrated by libraries such as SciPy (Jones *et al.* 2001–), the combination of an efficient statically compiled library with the flexibility of a run-time interpreter can be very powerful. A possibility for future exploration is to use a scripting language such as Python, Lua, or Scheme to provide input and output to BOUT++, or even to implement parts of physics modules. These languages provide a complete programming environment, but would require significant work to interface with the C++ classes in BOUT++.

4.2. Mesh generation

The accuracy and robustness of plasma simulations is strongly dependent on the quality of the input mesh: noise in the metric components, or large variations in the grid spacing leads to noise in the solution, restrictions in the timestep, and occasionally numerical instabilities.

Generating meshes for tokamak equilibria with X-points is challenging due to the change in topology across the separatrix, the shape of the boundary around the plasma, and the variation in geometry between machines. The original BOUT code used UEDGE (Rognlien *et al.* 1992, 2002) to generate grids, and BOUT++ can also use these input files with a little preprocessing. There are several other tools available for generating these meshes such as CARRE (Marchand & Dumberry 1996), but none are open-source licensed and suitable for distribution with BOUT++.

To generate X-point meshes for BOUT++ from experimental free boundary equilibria, a new code named HYPNOTOAD has been developed, and is available in the BOUT++ public repository. The main features of this grid generator are that it:

(a) Was originally written entirely in IDL. This allows Hypnotoad to be run anywhere where IDL is available, without the need for a compilation step and complicated dependencies. IDL is widely available and used in fusion research institutions, and comes with a large library of built-in functions. Work is currently ongoing to port the algorithms which have been developed into Python, and indeed many of the figures shown here will be from the Python version, due to the superior graphical capabilities of the Python library Matplotlib[†].

(b) Automatically adjusts settings when needed. The grid produced can be customised, but the minimum number of inputs is very small (number of grid points and a range of poloidal flux ψ). The ψ range asked for is adjusted to fit within the boundary, with configurable levels of strictness.

(c) Can handle an arbitrary number of X-points. Whilst not of obvious benefit since most tokamak equilibria are single or double-null, this means that HYPNOTOAD is quite generalised and can cope with unusual configurations. It has been applied to Snowflake-like configurations (Ma *et al.* 2014), but only in snowflake-plus configurations where the second X-point was not included in the mesh.

Because this grid generator is intended to be used for many different tokamaks, robust algorithms have been developed which can handle the large number of possible configurations encountered. To date, this grid generator has been successfully used to generate meshes from C-MOD, DIII-D, EAST, ITER, JET, MAST and NSTX tokamak equilibria (Wesson 1997), without requiring any machine-specific alterations or inputs beyond the EFIT generated G-EQDSK or 'g' file (Lao *et al.* 1985, 2005; General Atomics 2014).

The production of a BOUT++ input mesh consists of three stages: Analysis of the equilibrium to determine X-point locations, construction of mesh points, and calculation

[†] Matplotlib plotting library <http://matplotlib.org/>




FIGURE 5. Automated identification of O- and X-points in a MAST double-null configuration

of metric tensor and equilibrium quantities. The key features and algorithms used in each of these stages are described in the following sections.

4.2.1. Finding X-points

The first task in generating a mesh is to determine the number and location of the O- and X-points of the plasma. These correspond to critical points (maxima, minima, or saddle points) of the poloidal flux function $\psi(R, Z)$, which for tokamak equilibria is a 2D function of major radius R and height Z . The most robust technique for finding these has been found empirically to be to produce contour lines of $\frac{d\psi}{dR} = 0$ and $\frac{d\psi}{dZ} = 0$. Intersections of these lines then give locations of critical points. By comparing second derivatives of ψ at these critical points it can be determined whether they are O-points (minima/maxima) or X-points (saddle points). An example of a double-null equilibrium from the Mega-Amp Spherical Tokamak (MAST) is shown in figure 5. Contours of $\frac{d\psi}{dZ} = 0$ and $\frac{d\psi}{dR} = 0$ are plotted, and their intersections identified and categorised. Some additional heuristics are needed to eliminate false positives or duplicate critical points, which can occur if the input data contains grid-scale features or noise. The primary O-point can be reliably identified as the one closest to the middle of the grid, but identifying the plasma X-points is more prone to error, as a typical equilibrium will contain several X-points.




FIGURE 6. Mesh produced for MAST double-null equilibrium (shot 14220)

When the boundary shape is specified, critical points close to the poloidal field coils can be discarded as being outside the boundary, though ripples in the solution due to the central solenoid at major radius $R = 0$ can lead to spurious O- and X-points on the inboard side: In figure 5 two O-points (labelled '0' and '2'), and one X-point (labelled '1') can be seen close to the centre column. These can usually be excluded through choice of ψ range or boundary contour.

4.2.2. Meshing

Transport of heat and particles in magnetised plasmas are strongly anisotropic, and as a result fluid simulations commonly use meshes aligned with magnetic flux surfaces, in order to minimise mixing the directions perpendicular and parallel to the magnetic field. The coordinates currently used by BOUT++ for tokamak simulations are orthogonal in the poloidal (R-Z) plane, illustrated in figure 6 for a typical MAST double-null discharge. A region close to the X-point is shown enlarged, showing coordinate lines passing around the X-point, but leaving a hole at the null itself. The mesh has a branch-cut around the magnetic X-point, which must be treated carefully to avoid creating large variations in mesh spacing, which could lead to poor numerical behavior.

A starting contour line is created, aligned with a magnetic flux surface. In the plasma core this line is just inside the separatrix, whilst for the divertor legs the separatrix is used. From this starting line the locations of the X-points are found, and the regions between X-points are then meshed independently, adjusting the distance to the X-point to obtain a smoothly varying mesh spacing. From these starting locations the gradient of ψ is found using either a Discrete Cosine transform (DCT) method, or 2D cubic convolution interpolation (Park & Schowengerdt 1983) provided in the IDL standard library. The DCT method provides smooth interpolation, but is slow for large input meshes and sharp features in the input can lead to ringing (Gibbs phenomena). The interpolation method is therefore generally preferred. The gradient in ψ is followed using the LSODE (Livermore

Solver for Ordinary Differential Equations) algorithm (Hindmarsh 1983; Radhakrishnan & Hindmarsh 1993) through IDL or SciPy, to construct the mesh points for each value of θ . This is then repeated for each θ coordinate to construct a 2D mesh. Once the mesh points have been found, the metric tensor and equilibrium components are calculated.

4.2.3. Calculating metric components

Experimental equilibria are often of low resolution (65×65 is the standard EFIT output for many tokamaks), and once the pressure P and magnetic field \mathbf{B} is interpolated onto a new mesh there is no guarantee that the new values will still obey ideal MHD force balance ($\nabla \times \mathbf{B}) \times \mathbf{B} = \mu_0 \nabla P$. Ideal MHD force balance may not be an equilibrium solution to the plasma model being simulated, but is almost invariably a good approximation, and it is important that the mesh generation process does not lead to artefacts or sources of numerical noise and instability. Care is therefore taken to ensure the accuracy and smoothness of the interpolated solution, and quantities such as parallel current density J_{\parallel} are calculated multiple ways and compared as a consistency check. The interpolation method used can have a significant impact on the quality of the results: many terms such as the curvature and parallel current involve second derivatives of equilibrium flux, and these quantities should themselves be smoothly varying inputs to the simulation.

As discussed elsewhere (Xu *et al.* 2008; Dudson *et al.* 2009), in order to efficiently simulate structures (predominantly) aligned to the magnetic field, BOUT++ uses grid-points placed in a field-aligned coordinate system. From the standard, orthogonal, toroidal coordinate system (ψ, θ, ζ) new coordinates (x, y, z) :

$$\begin{aligned} x &= \sigma_{B\theta} (\psi - \psi_0) & y &= \theta \\ z &= \sigma_{B\theta} \left(\zeta - \int_{\theta_0}^{\theta} \nu(\psi, \theta) d\theta \right) \end{aligned} \quad (4.1)$$

where $\sigma_{B\theta} \equiv B_{\theta}/|B_{\theta}|$ is the sign of the poloidal field, and ν is the local field-line pitch given by

$$\nu(\psi, \theta) = \frac{\mathbf{B} \cdot \nabla \zeta}{\mathbf{B} \cdot \nabla \theta} = \frac{B_{\zeta} h_{\theta}}{B_{\theta} R} \quad (4.2)$$

where B_{ζ} is the toroidal magnetic field, B_{θ} the poloidal magnetic field, R is the major radius, and h_{θ} is poloidal arc length divided by 2π . In the limit of a circular cross-section equilibrium, h_{θ} becomes the minor radius r . The coordinate system is chosen so that x increases radially outwards, from plasma to the wall. The sign of the toroidal field B_{ζ} can then be either positive or negative.

By equating contravariant x components of $\mathbf{J} \times \mathbf{B} = \nabla P$, radial force balance in field-aligned coordinates can be written as:

$$\frac{\partial}{\partial x} \left(\frac{B^2 h_{\theta}}{B_{\theta}} \right) - B_{\zeta} R \frac{\partial}{\partial x} \left(\frac{B_{\zeta} h_{\theta}}{R B_{\theta}} \right) + \frac{\mu_0 h_{\theta}}{B_{\theta}} \frac{\partial P}{\partial x} = 0 \quad (4.3)$$

Close to the X-points, $B_{\theta} \rightarrow 0$ and the above expression becomes singular, so a better way to write this is:

$$\begin{aligned} \frac{\partial}{\partial x} (B^2 h_{\theta}) - h_{\theta} B_{\theta} \frac{\partial B_{\theta}}{\partial x} \\ - B_{\zeta} R \frac{\partial}{\partial x} \left(\frac{B_{\zeta} h_{\theta}}{R} \right) + \mu_0 h_{\theta} \frac{\partial P}{\partial x} = 0 \end{aligned} \quad (4.4)$$

This expression is used to calculate the pressure, by integrating $\frac{\partial P}{\partial x}$, and compared with the input pressure profile. By using the input pressure profile for P , h_{θ} is also calculated




FIGURE 7. Comparison of pressure and h_θ profiles calculated using radial force balance (equation 4.4) with a reference calculated independently.

and compared with the values calculated from geometric arc-length. An example result of this comparison is shown in figure 7. In figure 7(a) the pressure at the outermost radial point has been set to the value from the EFIT input, and then integrated inwards according to equation 4.4. In figure 7(b) the pressure gradient from EFIT has been used, and equation 4.4 solved for h_θ . Both show good agreement, indicating that the mesh generation is sufficiently accurate and smooth to retain radial force balance. This is not always observed, and where large discrepancies are observed the equilibrium should not be used.

Many reduced MHD models make use of the quantity

$$\frac{B}{2} \nabla \times \left(\frac{\mathbf{b}}{B} \right) \simeq \mathbf{b} \times \kappa \quad (4.5)$$

where $\kappa = \mathbf{b} \cdot \nabla \mathbf{b}$ is the curvature vector, which arises from magnetic particle drifts. Calculation of this quantity involves second derivatives of the input poloidal flux ψ , and so must be calculated carefully to avoid introducing noise. Several methods have been tried, including:

(a) Calculate curvature on the original R-Z mesh supplied as input, then interpolate onto the new field-aligned mesh. This has been found to usually produce the smoothest result and so is the default. Using DCTs to calculate differentials of the magnetic field components was found to produce oscillatory results, so a 3-point Lagrangian interpolation is used instead.

(b) Calculate curvature on the field-aligned mesh in toroidal coordinates, using nearest neighbours and least-square fitting.

(c) Calculate curvature in field-aligned coordinates using

$$\begin{aligned} \nabla \times \left(\frac{\mathbf{b}}{B} \right) &= \frac{B_\theta}{h_\theta} \left[\left(\frac{\partial}{\partial x} \left(\frac{h_\theta}{B_\theta} \right) - \frac{\partial}{\partial y} \left(\frac{\sigma_{B\theta} B_\zeta I R}{B^2} \right) \right) \mathbf{e}_z \right. \\ &\quad \left. + \frac{\partial}{\partial y} \left(\frac{\sigma_{B\theta} B_\zeta R}{B^2} \right) \mathbf{e}_x \right. \\ &\quad \left. + \frac{\partial}{\partial x} \left(\frac{\sigma_{B\theta} B_\zeta R}{B^2} \right) \mathbf{e}_y \right] \end{aligned} \quad (4.6)$$




FIGURE 8. Comparison of j_{\parallel} from input profiles (equation 4.8) and divergence of current (equation 4.7)

Methods (b) and (c) are based on calculating differentials on the output (field-aligned) mesh. They work well when the input is of high resolution, but become noisy once the resolution of the output grid significantly exceeds that of the input. Since high resolutions are required for BOUT++ simulations, this is nearly always the case, and so method (a) is preferred. Methods (b) and (c) are retained for cross-comparison.

Finally, a cross-check is made between the parallel current and the curvature. In a tokamak equilibrium the divergence of the parallel current balances the divergence of diamagnetic current, so that the total current is divergence free. In reduced MHD models this appears through the vorticity equation. The following relationship should therefore be satisfied:

$$\nabla_{\parallel 0} j_{\parallel 0} + \nabla \times \left(\frac{\mathbf{b}}{B} \right) \cdot \nabla P = 0 \quad (4.7)$$

From this the parallel current can be calculated, and compared with the parallel current calculated from the input $f = RB_{\zeta}$ and pressure P profiles:

$$\mu_0 j_{\parallel} = -B \frac{\partial f}{\partial \psi} - \mu_0 \frac{f}{B} \frac{\partial P}{\partial \psi} \quad (4.8)$$

For the MAST equilibrium shown in figure 6 and figure 7 the parallel current calculated from the curvature and pressure gradient (equation 4.7) is shown in figure 8, and compared to the current given by equation 4.8.

These comparisons between quantities such as J_{\parallel} and curvature calculated in multiple ways allow an assessment of the quality of the input equilibrium. If the values obtained are inconsistent, then a higher accuracy input can be generated using a free boundary equilibrium solver such as EFIT or CORSICA (Tarditi *et al.* 1996). The result is a robust process for generating BOUT++ input grids from tokamak experimental equilibria, which has been successfully used for a range of devices.

5. Conclusions

Recent improvements to the BOUT++ simulation framework have been summarised. Since its original release (Dudson *et al.* 2009), BOUT++ has been adopted by a growing number of users, who have extended its capabilities in a number of ways: the structure of the code has been improved; more complex elliptic and parabolic equations can be solved, through coupling to the PETSc library and built-in implementations; the input and output tools have been improved to enable experimental or theoretical equilibria and profiles to be imported into BOUT++.

A preconditioning scheme has been presented which enables simulations to be performed without invoking the Boussinesq approximation in the vorticity. This method can now be used to perform more accurate simulations of tokamak edge turbulence and Edge Localised Modes (ELMs) (Zohm 1996), the exploration of which is the subject of future work. More general elliptic and parabolic solvers for preconditioning and vorticity inversion problems in 3D are currently under development, and will be presented elsewhere.

Soft scaling studies have been performed for a parabolic solver along equilibrium field-lines. We find that although the scaling with problem size is good (being based on the $O(n)$ serial Thomas algorithm), the scaling with processor number is quite poor due to the global gather and scatter operations used. Improving this will be a subject of future work.

The use of software development practices such as version control with Git and automated testing has been described. A more rigorous set of tests using the Method of Manufactured Solutions (Roache 1998; Salari & Knupp 2000) for code verification are currently under development, and will be described elsewhere.

In addition to the improvements described herein, significant progress has also been made in improving the calculation of heat fluxes along magnetic field-lines, which are often poorly described by the Spitzer formula (Wesson 1997) or limiter schemes (Schneider *et al.* 2006; Tskhakaya 2012; Tskhakaya *et al.* 2008): An extension of the Hammett-Perkins model to non-Fourier methods (Dimits *et al.* 2013), and a method based on solving a 1-D time-independent kinetic equation along magnetic field-lines (Ji *et al.* 2009; Omotani & Dudson 2013), have both been implemented in BOUT++. Testing of these methods against kinetic simulations, primarily for simulation of transients such as ELMs, is the subject of ongoing work.

Acknowledgements

This work was funded by EPSRC grant EP/K006940/1 using HECToR computing resources through the Plasma HEC consortium grant EP/L000237/1. EURATOM Mobility support is gratefully acknowledged. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

REFERENCES

- AHO, A V, LAM, M S, SETHI, R & ULLMAN, J D 2006 *Compilers: Principles, Techniques, and Tools*. Addison-Wesley.
- ANGUS, J R & UMANSKY, M V 2014 *Physics of Plasmas* **21**, 012514.
- ANGUS, J R, UMANSKY, M V & KRASHENINNIKOV, S I 2012a 3d blob modelling with bout++. *Contrib. Plasma Phys.* **52**, 348–352.
- ANGUS, J R, UMANSKY, M V & KRASHENINNIKOV, S I 2012b Effect of drift waves on plasma blob dynamics. *Phys. Rev. Lett.* **108**, 215002.

- AUSTIN, T M, BERNDT, M & MOULTON, J D 2004 Tech.Rep. LA-UR 03-4149. LANL, USA.
- BALAY, SATISH, GROPP, WILLIAM D., MCINNES, LOIS CURFMAN & SMITH, BARRY F. 1997 In *Modern Software Tools in Scientific Computing* (ed. E. Arge, A. M. Bruaset & H. P. Langtangen), pp. 163–202. Birkhauser Press.
- BALAY, S *et al.* 2010 *Tech. Rep.* ANL-95/11 - Revision 3.1. Argonne National Laboratory.
- BALAY, SATISH *et al.* 2014 PETSc Web page. <http://www.mcs.anl.gov/petsc>.
- BEER, M A *et al.* 1997 *Physics of Plasmas* **4** (5), 1792–1799.
- CATTO, P J & SIMAKOV, A N 2004 A drift ordered short mean free path description for magnetized plasma allowing strong spatial anisotropy. *Physics of Plasmas* **11** (1), pp. 90–102.
- CHACÓN, L., KNOLL, D. A. & FINN, J. M. 2002 *J. Comput. Phys.* **178** (1), 15–36.
- DIMITS, A M, JOSEPH, I & UMANSKY, M V 2013 A fast non-Fourier method for Landau-fluid operators. *Physics of Plasmas* **21** (5), 055907.
- D’IPPOLITO, D A, MYRA, J R & ZWEBEN, S J 2011 *Physics of Plasmas* **18**, 060501.
- DUDSON, B, FARLEY, S & CURFMANN MCINNES, L 2012 *arXiv:1209.2054* .
- DUDSON, B D *et al.* 2009 *Comp. Phys. Comm.* **180**, 1467–1480.
- FRIEDMAN, B, CARTER, T A, UMANSKY, M V, SCHAFFNER, D & DUDSON, D 2012 *Physics of Plasmas* **19**, 102307.
- FUNDAMENSKI, W *et al.* 2001 *J. Nucl. Materials* **290-293**, 593.
- GAMMA, ERICH, HELM, RICHARD, JOHNSON, RALPH & VLISSIDES, JOHN 1995 *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley.
- GARCIA, O E *et al.* 2007 *J. Nucl. Materials* **363**, 575.
- GEKELMAN, W *et al.* 1991 *Rev. Sci. Instr.* **62**, 2875.
- GENDRIH, PH *et al.* 2012 *J. Phys.: Conf. Ser.* **401**, 012007.
- GENERAL ATOMICS 2014 G EQDSK file format. <https://fusion.gat.com/theory/Efitgeqds>.
- GOLUB, G H & VAN LOAN, C F 2013 *Matrix Computations*. The Johns Hopkins University Press.
- GOSWAMI, R *et al.* 2001 *Physics of Plasmas* **8**, 857.
- HAMMETT, G W & PERKINS, F W 1990 *Phys. Rev. Lett.* **64**, 3019.
- HAZELTINE, R D & MEISS, J D 2003 *Plasma Confinement*. Dover publications.
- HEROUX, M A *et al.* 2005 *ACM Trans. Math. Softw.* **31** (3), 397–423.
- HINDMARSH, A C 1983 In *Scientific Computing* (ed. R. S. Stepleman *et al.*), *IMACS Transactions on Scientific Computation*, vol. 1, pp. 55–64. North-Holland, Amsterdam.
- HINDMARSH, A C *et al.* 2005 SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software* **31** (3), 363–396.
- HOCKNEY, R W 1965 *J. Assoc. Comput. Mach.* **12**, 95.
- ISERLES, AREIH 2009 *A First Course in the Numerical Analysis of Differential Equations*. Cambridge University Press, ISBN: 978-0-521-73490-5.
- JI, J Y, HELD, E D & SOVINEC, C R 2009 *Physics of Plasmas* **16** (2), 022312.
- JONES, ERIC, OLIPHANT, TRAVIS, PETERSON, PEARU *et al.* 2001– SciPy: Open source scientific tools for Python.
- KARNIADAKIS, G E, ISRAELI, M & ORSZAG, S A 1991 *J. Comput. Phys.* **97**, 414.
- KAWASHIMA *et al.* 2006 *Plasma Fusion Res.* **1**, 31.
- KNEPLEY, M G 2012 *arXiv:1209.1711* .
- KNOLL, D A *et al.* 2001 *SIAM J.Sci.Comput.* **23** (2), 381.
- LAO, L L, ST. JOHN, H, STAMBAUGH, R D, KELLMAN, A G & PFEIFFER, W 1985 Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nucl. Fusion* **25**, 1611–22.
- LAO, L L *et al.* 2005 MHD equilibrium reconstruction in the DIII-D tokamak. *Fusion Science and technology* **48**, p968.
- LLVM PROJECT 2014 Implementing a language with LLVM. <http://llvm.org/docs/tutorial/>.
- MA, J F, XU, X Q & DUDSON, B D 2014 **54**, 033011.
- MARCHAND, R & DUMBERRY, M 1996 CARRE: a quasi-orthogonal mesh generator for 2D edge plasma modelling. *Comp. Phys. Comm.* **96**, 232–246.
- MOUSSEAU, V A, KNOLL, D A & RIDER, W J 2000 *J. Comput. Phys.* **160**, 743.

- NAKAMURA, M 2011 *J. Nucl. Materials* **415**, S553.
- NAULIN, V *et al.* 2007 *J. Nucl. Materials* **24**, 363–365.
- OMOTANI, J T & DUDSON, B D 2013 *Plasma Phys. Control. Fusion* **55**, 055009.
- OTTAVIANI, M & MANFREDI, G 1999 *Physics of Plasmas* **6**, 3267.
- PARK, S & SCHOWENGERDT, R 1983 *Computer Vision, Graphics & Image Processing* **23**, 256.
- RADHAKRISHNAN, K & HINDMARSH, A C 1993 *Tech. Rep.*. LLNL, <http://computation.llnl.gov/casc/nsde/pubs/u113855.pdf>.
- RICCI, P & ROGERS, B N 2013 *Physics of Plasmas* **20**, 010702.
- RICCI, P *et al.* 2012 *Plasma Phys. Control. Fusion* **54** (124047).
- ROACHE, P J 1998 *Verification and Validation in Computational Science and Engineering*. Hermosa Publishers, Albuquerque NM.
- ROGNLIEN, T D, XU, X Q & HINDMARSH, A C 2002 Application of Parallel Implicit Methods to Edge-Plasma Numerical Simulations. *J. Comput. Phys.* **175**, 249–268.
- ROGNLIEN, T D *et al.* 1992 *J. Nucl. Materials* **196/198**, 347.
- SALARI, K & KNUPP, P 2000 Code verification by the method of manufactured solutions. *Tech. Rep.* SAND2000-1444. Sandia National Laboratories.
- SCHNEIDER, R *et al.* 2006 *Contrib. Plasma Phys.* **46**, 3–191.
- SCOTT, B 2003 *Plasma Phys. Control. Fusion* **45**, A385–398.
- SCOTT, B 2005a *Physics of Plasmas* **12**, 102307.
- SCOTT, B D 2002 *New J. Physics* **4**, 52.1–52.30.
- SCOTT, B D 2005b GEM - an energy conserving electromagnetic gyrofluid model. *arXiv:physics/0501124*.
- STANGEBY, P C 2000 *The Plasma Boundary of Magnetic Fusion Devices*. IoP.
- TAMAIN, P *et al.* 2010 TOKAM-3D: A 3D fluid code for transport and turbulence in the edge plasma of Tokamaks. *J. Comput. Phys.* **229** (2), 361–378.
- TARDITI, A *et al.* 1996 *Contrib. Plasma Phys.* **36**, 132.
- TEMPERTON, C 1975 Algorithms for the solution of cyclic tridiagonal systems. *J. Comput. Phys.* **19** (3), 317–323.
- TSKHAKAYA, D 2012 *Contrib. Plasma Phys.* **52** (5-6), 490–499.
- TSKHAKAYA, D., SUBBA, F., BONNIN, X., COSTER, D. P., FUNDAMENSKI, W., PITTS, R. A. & CONTRIBUTORS, JET EFDA 2008 On kinetic effects during parallel transport in the sol. *Contrib. Plasma Phys.* **48** (1-3), 89–93.
- UMANSKY, M V, ROGNLIEN, T D, XU, X Q, DUDSON, B D & KIRK, A 2006 Modelling of edge plasma turbulence in a spherical tokamak.
- WALKDEN, N R, DUDSON, B D & FISHPOOL, G 2013 Characterization of 3d filament dynamics in a mast sol flux tube geometry. *Plasma Phys. Control. Fusion* **55**, 105005.
- WESSON, J A, ed. 1997 *Tokamaks*, 2nd edn. Clarendon Press.
- WOLFRAM RESEARCH, INC. 2014 Mathematica. Champaign, Illinois.
- XI, P W, XU, X Q, XIA, T Y, NEVINS, W M & KIM, S S 2013 **53** (11), 113020.
- XIA, T Y, XU, X Q, DUDSON, B D & LI, J 2012 *Contrib. Plasma Phys.* **52**, 353–359.
- XU, X *et al.* 2013 *Physics of Plasmas* **20**, 056113.
- XU, X Q, COHEN, R H, PORTER, G D, MYRA, J R, D’IPPOLITO, D A & MOYER, R 1999 Turbulence in Boundary Plasmas. *J. Nucl. Materials* **266-269**, 993–996.
- XU, X Q, COHEN, R H, PORTER, G D, ROGNLIEN, T D, RYUTOV, D D, MYRA, J R, D’IPPOLITO, D A, MOYER, R A & GROEBNER, R J 2000 Turbulence studies in tokamak boundary plasmas with realistic divertor geometry. *Nucl. Fusion* **40** (3Y), 731–736.
- XU, X Q, UMANSKY, M V, DUDSON, B & SNYDER, P B 2008 Boundary plasma turbulence simulations for tokamaks. *Comm. in Comput. Phys.* **4** (5), pp. 949–979.
- XU, X Q *et al.* 2010 *Phys. Rev. Lett.* **105**, 175005.
- YU, G Q, KRASHENINNIKOV, S I & GUZDAR, P N 2006 *Physics of Plasmas* **13**, 042508.
- ZOHM, H 1996 *Plasma Phys. Control. Fusion* **38**, 105–128.