Commun. Comput. Phys.
doi: 10.4208/cicp.191214.021015a

Vol. **20**, No. 1, pp. 136-155
July 2016

Eigenvalue Solver for Fluid and Kinetic Plasma Models in Arbitrary Magnetic Topology

D. A. Baver[1,*], J. R. Myra[1] and M. V. Umansky[2]

[1] *Lodestar Research Corporation, Boulder Colorado 80301, USA.*
[2] *Lawrence Livermore National Laboratory, USA.*

Received 19 December 2014; Accepted (in revised version) 2 October 2015

Abstract. ArbiTER (Arbitrary Topology Equation Reader) is a new code for solving linear eigenvalue problems arising from a broad range of physics and geometry models. The primary application area envisioned is boundary plasma physics in magnetic confinement devices; however ArbiTER should be applicable to other science and engineering fields as well. The code permits a variable numbers of dimensions, making possible application to both fluid and kinetic models. The use of specialized equation and topology parsers permits a high degree of flexibility in specifying the physics and geometry.

AMS subject classifications: 35Pxx, 76Wxx

Key words: Plasma, eigensolver, finite difference method, finite element method, SLEPc, PETSc, field-line following coordinates, snowflake divertor.

1 Introduction

Most modern computational efforts for simulating the edge region of fusion plasmas employ time advancement to capture the nonlinear and turbulent evolution of the particles and fields. Over the past several decades, many such studies of boundary turbulence have been advanced, employing varying degrees of approximation and sophistication. Some relevant examples are given in Refs. [1–18]. These simulation codes employ either fluid models, gyrofluid models, or full kinetic simulation using particle-in-cell (PIC) or Vlasov fluid approaches all of which are implemented in the time domain. The plasma simulation community needs such tools for theoretical analysis, numerical experimentation, experimental modeling, and even for the hardware component design.

On the other hand, there is a small but significant class of important edge plasma problems which are amenable to linear and/or quasilinear analysis of partial integro-differential equations. These include source-driven problems and eigenvalue problems

*Corresponding author. *Email addresses:* dabaver65@hotmail.com (D. A. Baver), jrmyra@lodestar.com (J. R. Myra), umansky1@llnl.gov (M. V. Umansky)

 ©2016 Global-Science Press

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

such as those arising in the computation of the growth rates of linear plasma instabilities. In addition, linear calculations can be used for convergence tests and to study stability regions in parameter space.

Time-evolution codes can be used for treating such problems, however there are advantages in using a non time-evolution approach. Indeed, consider a general time-evolution equation

$$df(\vec{x})/dt = \vec{F}(f(\vec{x}),\vec{x},t). \tag{1.1}$$

In the case where the problem is linear, or can be linearized, the right-hand side can be represented as a matrix M, i.e.,

$$df(\vec{x})/dt = \hat{M}(\vec{x},t)f(\vec{x}). \tag{1.2}$$

Furthermore we are assuming M(x,t)=M(x), with no explicit time dependence. In this case, for calculation of the linear instabilities in this system, one can calculate the eigenvalues and eigenmodes of the matrix M by the methods of linear algebra,

$$\hat{M}f = -\omega^2 f. \tag{1.3}$$

This approach has advantages over solving the problem by time-evolution: it is potentially much more efficient in terms of CPU-hours needed to solve a problem of given resolution, and it is capable of finding subdominant modes, i.e. modes with less than the maximum growth rate. Such modes, even if stable, can lend useful insights into the underlying physics. [19, 20] Exploiting the full potential of this approach does require that a sparse eigensolver is employed; if full matrix techniques are used instead, the computational cost will be orders of magnitude higher. In addition, it is the experience of the authors that eigenmodes of interest can be found most easily and efficiently using spectral transform techniques, which are discussed in greater detail in Ref. [29]. Of particular interest is the Cayley transform, which both greatly decreases convergence times and allows eigenmodes from nearly any part of eigenvalue space to be selected.

Similarly, one can consider the problem of linear response:

$$d\vec{f}/dt = \hat{M}(x)\vec{f} + \vec{g}(x,t), \tag{1.4}$$

where $\vec{g}(x,t)$ is the forcing term.

Instead of solving by time-evolution, one can take advantage of the linearity and assume a single frequency time-dependence, $\vec{g}(x,t) = \vec{g}(x)\exp(-i\omega t)$. Assuming no secular terms and assuming the homogeneous part of the solution decays in time, one can solve for the asymptotic stationary solution by solving the linear system

$$(i\omega\hat{I} + \hat{M}(x))\vec{f} = -\vec{g}(x). \tag{1.5}$$

Finally, linear codes can also play a role in understanding some classes of nonlinear problems which are amenable to quasilinear analysis. The appropriate biquadratic forms can readily be calculated once the linear eigenfunctions are known.

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

When a linear or quasilinear treatment is appropriate there are multiple advantages to using a code which exploits this feature. First, when appropriate sparse matrix techniques are used, a linear code can be much faster than its nonlinear counterpart; while the extent of this advantage is problem-dependent and therefore difficult to quantify precisely, it is the experience of the authors that this advantage can be quite large. Second, a linear code can avoid numerical instabilities in the time domain, and by using suitable tools to select regions of interest in eigenvalue space, can often circumvent spurious numerical modes even if they are still present in the simulation. For these reasons, linear codes are very well suited to parameter space studies, and can promise a high degree of confidence in their results. The latter attribute makes them especially important for use as verification tools [21,22] in comparative studies with their nonlinear time-domain counterparts. Quasilinear calculations, performed as a post-processing step, provide fast and robust results which can provide insight into some processes, such as induced forces and transport fluxes.

The Arbitrary Topology Equation Reader (ArbiTER) code is a general tool for discretizing and solving linear partial differential equations. It is an extension of the 2DX code [23], a code which was developed to solve a general class of fluid-based eigenvalue problems in the edge and scrape-off layer region of a diverted tokamak. ArbiTER inherits much of its code structure from 2DX, as well as the essential characteristics of its equation parser. It is capable of emulating 2DX, but has additional capabilities that give it significant advantages over its predecessor.

What the ArbiTER code adds is a topology parser. Topology in this case refers to the connectivity of grid nodes; this is particularly important in problems that use field-line following coordinates, since in this case grid connectivity corresponds directly to magnetic topology. The ability to reconfigure topology is useful, for instance, in the simulation of advanced divertor configurations such as snowflakes [24]. The manner in which this topological flexibility is achieved also permits the topology parser to vary the number of dimensions in a problem, even to the point of including interactions between spaces of unequal dimensionality; this is useful for solving kinetic models. Finally, the topology parser is capable of performing calculations using unstructured grids.

As with 2DX, the flexibility and modularity of ArbiTER have considerable advantages from the standpoint of code verification. By splitting the modeling of each problem into two parts, the equation and topology input files and the ArbiTER source code, it allows each part to be independently verified by different means; the source code can be verified on other models because the same basic routines are used in many different types of models, whereas the equation input file (or "structure" file) describing the physics model can be converted into algebraic form for ease of inspection. This modularity has further advantages from the standpoint of basic physics studies. By separating equation (structure) and topology features into different inputs, it is possible to mix and match equations and topologies. Thus, a model that has already been demonstrated on a simple topology can operate reliably on an entirely new topology. Likewise, once a topology has been tested, it is relatively easy to change the model equations being solved. It is from this perspec-

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

tive that the immense flexibility of the ArbiTER code becomes important; while many of its features would be more straightforward to implement as a series of separate codes, by including all of these features into a single code, a greater degree of confidence can be established regarding the reliability of its results, and greater flexibility can be achieved in matching model equations to numerical techniques.

2 Program structure

The ArbiTER code consists of two main parts, as well as a number of tools required to set up its input files and process its output files. The relation between these parts is shown in Fig. 1. Of these, the grid generation tools and the data analysis tools are currently implemented as Mathematica notebooks, although Python scripts for grid generation are also available. Conversion between the editable and machine-readable versions of the structure and topology file are done using Python scripts. The Field/Domain variant referred to in Fig. 1 is an alternate build of the ArbiTER kernel. This program permits grid files to be generated via an iterative process, in which information stored in the topology file is used to generate an intermediate file that is then in turn used to aid in formatting data in the grid file.

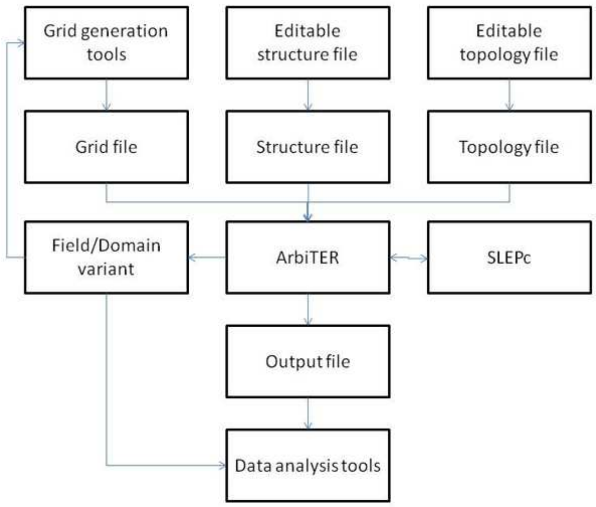


Figure 1: Relationship between components of the ArbiTER code.

The principal parts of the ArbiTER code are the ArbiTER kernel and the eigenvalue solver. The ArbiTER kernel generates a pair of sparse matrices in coordinate list format (sometimes referred to as COO format), that is, a list of entries each containing the value, row and column of each non-zero entry in the matrix. These matrices describe a generalized eigenvalue problem,

$$Ax = \lambda Bx. \tag{2.1}$$

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

The pair of matrices is passed to an eigenvalue solver, which returns the eigenvalues λ and eigenvectors x. The eigenvalue solver used by ArbiTER is SLEPc [29]. In addition, there exist variants of the ArbiTER code that are discussed in greater detail in Section 2.3 One of these uses the same matrices A and B generated by the ArbiTER kernel but solves them as ordinary matrix equations,

$$Ax = Bb. \tag{2.2}$$

Problems of this "source-driven" type are solved using PETSc [30], which is already included in the process of installing SLEPc. In either case, a variety of command-line options are available with the respective solver packages. In the case of SLEPc, the use of these options is discussed in the reference on the 2DX code [23].

The major distinguishing features of the ArbiTER code are related to how it builds the matrices A and B that define the eigenvalue problem. Like its predecessor 2DX, ArbiTER uses elementary matrix operations to build these matrices from simpler matrices. Unlike 2DX, rather than rely on built-in elementary operators as basic building blocks, ArbiTER allows the user to define their own elementary operators using a topology parser. This allows the user to define the relations in coordinate space between nodes in their computational domain, and from that information define differential operators that take the connectivity of those nodes into account. This ability to define both computational domains and differential operators gives the code its immense flexibility, making it suitable for both higher-dimensional models such as kinetic models as well as problems involving unstructured grids.

2.1 Equation language

The equation language is a file input format for defining sets of partial differential equations to discretize and solve. The resulting file, henceforth referred to as the structure file, contains all of the information the ArbiTER code needs to convert profile functions (i.e. input data such as temperature, density, magnetic geometry, etc. needed to define the matrix elements) into sparse matrices to solve. The equation language used in ArbiTER is based on that used in 2DX. This format is discussed in more detail in the paper on 2DX [23]. An example is also present in Appendix B. For clarity some general characteristics of this feature are reviewed here.

The equation language comes in two forms, a machine-readable form used by the ArbiTER code, and a human-readable form for editing. The human-readable form of the ArbiTER equation language is nearly reverse-compatible with 2DX, whereas the machine-readable form is not. The input file contains a number of sections, of which three are of special importance: the input language, the formula language, and the element language.

The input language is used to interpret the main data input file, known as the grid file. Each entry in the input language declares a variable name and assigns it a variable type (i.e. integer, real, function, etc.). In the case of functions the variable is assigned a computational domain, thus defining which entries in the function are assigned to which coordinate positions. The code then searches the input file (either ASCII or HDF5) for

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

the appropriate variable name and loads the associated data block into an appropriate structure in memory.

The formula language constructs the partial differential equation to be solved from some combination of functions, differential operators, and constants. These elements may be input as profile functions, constructed as elementary operators using the topology language, or built from simpler elements using the element language.

The element language builds elements (functions, constants, or operators) from simpler elements. This is particularly useful when using specialized coordinate systems, such as field-line following coordinates, as the definition of simple gradient operators depends on functions defining magnetic geometry and is quite complicated. Unlike 2DX, the ArbiTER element language can output constants, which is useful for applying physical formulas to the data.

2.2 Topology language

The main feature that distinguishes ArbiTER from 2DX is the topology language. This is an input file format for defining all features related to the spatial connectivity of grid points. This is primarily concerned with defining topological regions in the computational domain, but additionally is concerned with defining elementary differential operators used by the equation language. The capabilities of this language are fairly general, and among other things allow the ArbiTER code to emulate the capabilities of 2DX. The topology used to emulate 2DX is shown in Fig. 2, namely single-null x-point geometry with three topological regions: a closed flux surface region, the main scrape-off layer and the private flux region.

The topology language consists of a list of definitions for four different types of objects: blocks, linkages, domains, and operators. Blocks are simple Cartesian sub-domains, used as a building block for domains. Linkages are sparse matrices that describe how to connect two blocks or domains. Domains are sets of grid points that can be assigned to a profile (input) function or a dependent variable. Operators are sparse matrices that describe how grid points are related to one another, e.g. for the purpose of calculating derivatives. In addition, the topology language permits renumbering of domains to ensure that grid points are sequenced in an intuitive order, and the topology language file contains an integer language (integer analog of the element language) to aid in calculating the sizes of blocks and offsets of linkages.

Much of the diversity of the topology language stems from the wide variety of linkage and domain subtypes. Linkages can simply tie together blocks, or they can be multiplied by complex functions to produce phase-shift periodic boundary conditions (relevant to periodicity in toroidal domains described by field-line following coordinates), or they can even be defined by using an integer array to specify individual grid cells to be connected. Domains likewise can be simply assembled from blocks and linkages, or they can be constructed from existing domains in a number of different ways, such as convolution (combining two domains to create an outer product space) or indentation (removing cells

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

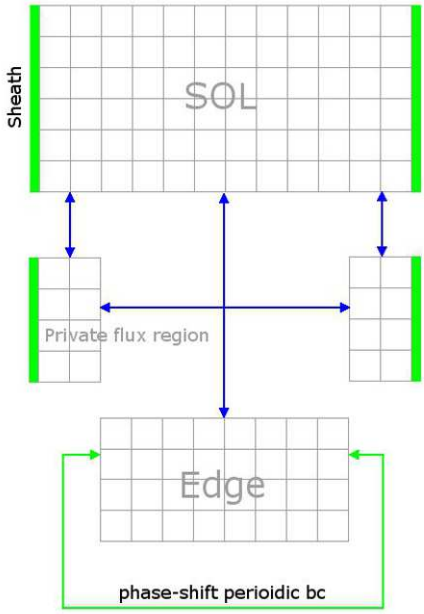


Figure 2: Block structure of 2DX emulation topology for ArbiTER.

from selected boundaries, relevant to creating staggered grids). These features provide significant advantages for kinetic models, in which operators must be defined that link domains of different dimensionalities.

2.3 Variant builds

While the core ArbiTER code was designed to solve eigenvalue problems, a number of variant builds have been developed for more specialized purposes. Some are intended as tools for debugging equation and topology input files, some are intended as tools for extracting geometric and topological information for purposes of formatting profile function inputs, while some introduce entirely new capabilities to the code.

The full list of variant builds is shown in Table 1. In addition, many of the variants on this list can be built with or without HDF5 file capability. This allows the program to be compiled on machines that do not have HDF5 installed. Some variants do not use SLEPc, and can therefore be compiled on machines that do not have SLEPc installed. In this case, matrices can be generated for debugging or for use in external solvers.

Some of these variants merit special attention. In particular, the field/domain variant (arbiterfd) can be used in the process of gridfile generation as shown in Fig. 1. In this case, the field/domain variant is run using the same topology file as the intended model equations, but a simplified structure file. The purpose of this is to generate a file containing a list of nodes and their coordinate positions for each domain. This is important because ArbiTER inputs profile functions as 1-D arrays. In order to input arrays

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

Table 1: Variant builds of the ArbiTER code. A and B represent eigenvalue matrices, λ represents eigenvalues, x represents eigenvectors, np is number of processors.

Variant	Outputs	Uses
arbiter	λ and x	Solving eigenvalue problems
arbiterp	λ and x in np files	Solving eigenvalue problems in parallel
arbiterm	A and B	Debugging, interfacing with stand-alone eigensolver
arbiterog	all data except above	Debugging
arbitersd	x from $Ax=Bb$	Source-driven problems
arbiterdi	x from $x=Ab$	Measuring effect of specific terms in equation
arbiterfd	functions, constants, domains	Iterative generation of grid files

with a larger numbers of dimensions, data values must be placed in a specific sequence, which is determined when the domain is generated internally. If the domain is particularly simple, external routines can be written to sequence the data correctly. However, for more complicated domains, this means the topology must be coded twice: once in the topology file and once in the grid setup routine. By using the field/domain variant to calculate the correct data sequence, one only needs to generate the topology once, and this same topology is used both to create the grid file and to generate the full eigenvalue problem. In addition to creating domains, the field/domain variant can also be used to process constants (for instance to calculate basic plasma parameters) and to process profile functions. The latter case assumes that data sequenced according to the correct input topology is already available from somewhere, but this can still reduce the complexity of grid file setup in some cases, particularly if the profile functions are to be interpolated or truncated after being calculated.

3 Verification and test cases

3.1 Resistive ballooning in snowflake geometry

One of the most basic advantages of the ArbiTER code over its predecessor 2DX is the ability to change topologies easily. This is particularly important in models that use field-line following coordinates, as any change in the magnetic topology will correspondingly change the structure of the computational grid in non-trivial ways. As it turns out, field-line following coordinates are commonly used in both fluid and kinetic models for studying plasma edge instabilities, since the large anisotropies found in magnetized plasmas tend to result in eigenmodes that are highly localized in the radial direction but extended along field lines.

A particularly relevant demonstration of this capability is the application of a simple plasma model in a snowflake divertor [24]. The snowflake divertor is based on a higher-order null than an x-point divertor, resulting in weaker poloidal fields near the null. In practice, a snowflake divertor is better described as containing two separate x-points in close proximity to one another. Depending on the relative positions of these x-points, the

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

snowflake may be described as snowflake-plus or snowflake-minus. The case used in this study is of the snowflake-minus variety.

For purposes of simplicity and ease of comparison to preexisting benchmark cases, a resistive ballooning model [25–27] was used for this test. The model equations for this are as follows:

$$\gamma\nabla_\perp^2\delta\Phi = +\frac{2B}{n}C_r\delta p - \frac{B^2}{n}\partial_\parallel\nabla_\perp^2\delta A_\parallel, \tag{3.1}$$

$$\gamma\delta n = -\delta v_E\cdot\nabla n, \tag{3.2}$$

$$-\gamma\nabla_\perp^2\delta A_\parallel = \nu_e\nabla_\perp^2\delta A_\parallel - \mu n\nabla_\parallel\delta\Phi, \tag{3.3}$$

where:

$$\delta p = (T_e+T_i)\delta n, \tag{3.4}$$

$$C_r = \mathbf{b}\times\kappa\cdot\nabla = -\kappa_g RB_p\partial_x + i(\kappa_n k_z - \kappa_g k_\psi), \tag{3.5}$$

$$\nabla_\perp^2 = -k_z^2 - \jmath B(k_\psi - i\partial_x RB_p)(1/\jmath B)(k_\psi - iRB_p\partial_x), \tag{3.6}$$

$$\partial_\parallel Q = B\nabla_\parallel(Q/B), \tag{3.7}$$

$$\nabla_\parallel = \jmath\partial_y, \tag{3.8}$$

$$\delta v_E\cdot\nabla Q = -i\frac{k_z(RB_p\partial_x Q)}{B}\delta\Phi, \tag{3.9}$$

$$\nu_e = .51\nu_r n/T_e^{3/2}. \tag{3.10}$$

Here the growth rate γ is the eigenvalue for a three-field model coupling perturbed electrostatic potential $\delta\Phi$, density δn, and parallel vector potential δA. C_r is the curvature operator, κ_n and κ_g are the normal and geodesic curvatures, respectively, ν_e is the electron-ion collision frequency, μ is the mass ratio, and $\jmath$ is the Jacobian. In Eqs. (3.7) and (3.9) Q is an arbitrary function, i.e. $\partial_\parallel\nabla_\perp^2\delta A = B\nabla_\parallel(\nabla_\perp^2\delta A/B)$. The normal and geodesic curvature refer to the components of the curvature vector perpendicular to and tangential to the flux surface on which a field line resides, respectively.

These model equations were applied to a snowflake profile based on DIII-D data. Ad-hoc illustrative density and temperature profiles were generated during pre-processing. The test cases used a flat temperature profile with density defined by a tanh function with respect to poloidal flux. By adjusting the position of the inflexion point of the tanh function, it was possible to control the location of the most unstable eigenmode. This in turn permitted study of the effect of magnetic topology on the structure and growth rate of a resistive ballooning mode.

The results of this are shown in Fig. 3. As the peak gradient region is moved outward, the eigenmode shifts outward accordingly. In the case of the middle plot, one can see an eigenmode localizing between the two x-points. The second from the right shows an eigenmode that is localized near an x-point, and shows forking with the eigenmode having a significant amplitude in two different divertor legs.

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

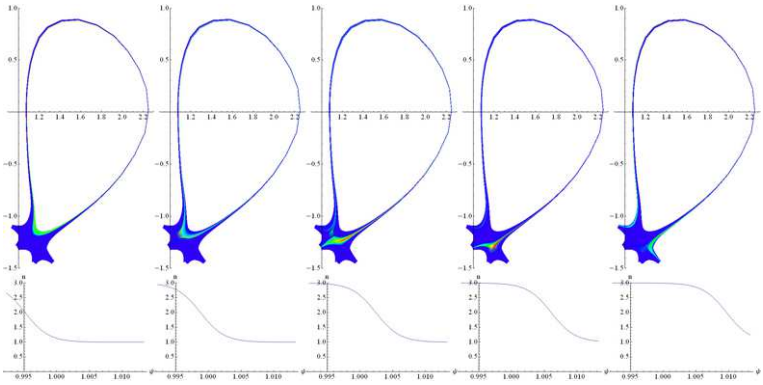


Figure 3: Results from simulation of resistive ballooning modes in snowflake geometry. The upper plots show the distribution of mode amplitude for Φ. The lower plots show the corresponding density profiles used in the above simulations. Upper plot horizontal and vertical axes are radius and elevation, lower plot horizontal axis is flux coordinate.

3.2 Kinetic resistive ballooning in x-point geometry

Because the ArbiTER topology language treats dimensionality as a free parameter to be defined as a computational domain is defined, it is possible to apply the code to higher-dimensional gyrokinetic models [20,28]. However, such models are typically quite complicated, and hence not natural first choices for developing and benchmarking a new code. Also, kinetic models may develop numerical issues that are not typically found in corresponding fluid models. For this reason, it was decided to test the kinetic capabilities of the ArbiTER code on a simple kinetic model. For this purpose, the model chosen was a resistive ballooning model with kinetic parallel electrons.

The model equations in this case are:

$$\gamma\nabla_\perp^2\delta\phi = \frac{2B}{n}C_r T\delta n + \frac{B^2}{n}\int_v dv v\partial_\parallel \delta f, \tag{3.11}$$

$$\gamma\delta n = -\delta v_E\cdot\nabla n, \tag{3.12}$$

$$\gamma\delta f = -\mu\nabla_\parallel\delta\phi f_0' - v\partial_\parallel\delta f + \nu\mu T_e\partial_v\delta f. \tag{3.13}$$

Again the growth or decay rate γ is the sought-after eigenvalue, f_0 and δf are the equilibrium and perturbed distribution functions, which depend on the spatial coordinate x and the velocity coordinate v.

In this equation set, there are two computational domains: a $nx\times ny$ domain for real space, and a $nx\times ny\times nv$ domain for phase space. Here f is defined as residing in phase space, whereas ϕ is defined as residing in real space. This takes advantage of features in the ArbiTER topology language: not only is the dimensionality of a computational grid selectable, but computational grids with different dimensionalities can coexist in the same model.

The results of this case are shown in Fig. 4. In order to verify correct construction of the structure and topology files, a test case was run in which the velocity advection

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

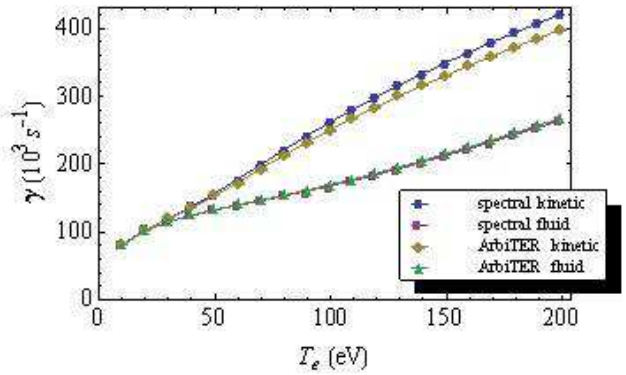


Figure 4: Results from simulation of a kinetic ballooning mode. Results are compared to fluid ballooning models and also to an iterative semi-fluid method used with the 2DX code.

terms in the model equation were turned off. This results in a model that is mathematically identical to a fluid model. Results from that are listed as ArbiTER fluid. The results from the fluid and kinetic (full model) cases were then compared to results from a spectral calculation performed using Mathematica, in one of which cases kinetic effects were modeled using the plasma response function. As expected, for the fluid case there is close agreement between the spectral model and the ArbiTER result. In the kinetic case, the ArbiTER code underestimates the growth rate by a few percent. This is adequate agreement given the differences between the treatment of collisions in Eq. (3.13) and the spectral code which employed a Krook-model collision operator.

3.3 Heat diffusion using first-order finite element analysis

A significant feature of the ArbiTER code is the ability to input connectivity matrices as integer input from the grid file. The intended use of this feature is to permit the definition of unstructured grids for finite element analysis. While this feature has so far not been applied in the context of edge plasma physics, it is nonetheless an interesting feature worthy of a demonstration case.

The computational domains for this problem begin with two domains, one for nodes and one for cells. The domain for cells is convolved with a domain containing three elements to create a $ncells\times 3$ domain and a $ncells\times 3\times 3$ domain. A linkage is constructed from integer data in the grid file linking nodes to the $ncells\times 3$ domain. This defines which nodes are included in which cells. Once this is done, differential operators are defined on the $ncells\times 3\times 3$ domain. By multiplying these operators by operators describing linkages between the $ncells\times 3\times 3$ and $ncells\times 3$ domains and between the $ncells\times 3$ and $nnodes$ domains, one can map these operators back to node space. Once this is done, the equations can then be defined with dynamical variables existing on the node space.

Once differential operators have been defined on the finite element grid, one can proceed with construction of a suitable grid file. This was done by using a Python script to

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

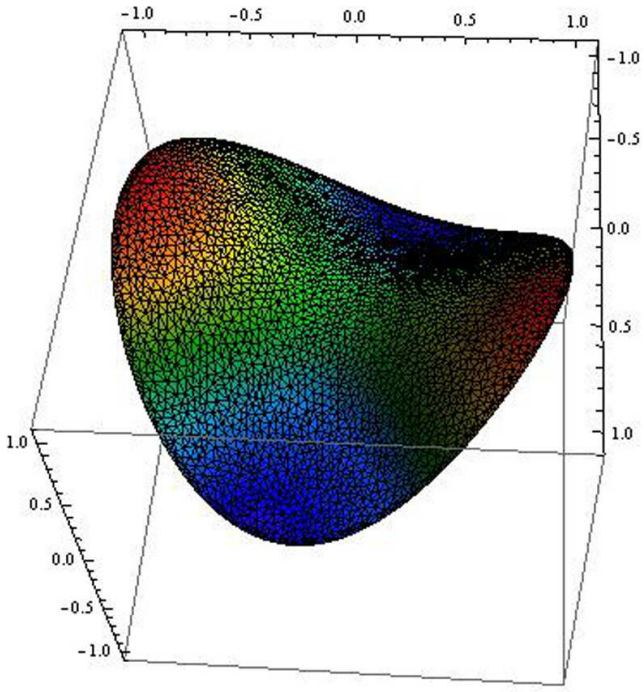


Figure 5: Fourth eigenmode from a finite-element simulation of a diffusion/wave equation with zero-derivative boundary conditions. Color and height correspond to function value. Mesh used is superimposed onto solution. All axes are in dimensionless units.

convert output from a Delaunay triangulation program called Triangle [31]. The resulting final grid contains 3007 nodes. These differential equations were used to solve the following equation in a circular domain with zero-derivative boundary conditions:

$$\lambda T = \nabla^2 T. \tag{3.14}$$

Results from this test are shown in Fig. 5. The fourth slowest decaying eigenmode was shown for clarity, as lower eigenmodes displayed few distinctive features. The resulting solution is smooth given that only a first order method was used, and has the expected form for a solution to the model equations.

3.4 Test of source-driven capability on a magnetic dipole

To test the source-driven capability of the ArbiTER code, it was used to calculate the fields of a point magnetic dipole inside a superconducting cylinder. This problem was chosen because an analytic solution had already been calculated for another purpose, otherwise the choice of test problem is largely arbitrary.

In this problem, we place a magnetic dipole along the axis of a cylinder, oriented perpendicular to that axis. We then define a potential, much as one would in the case

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

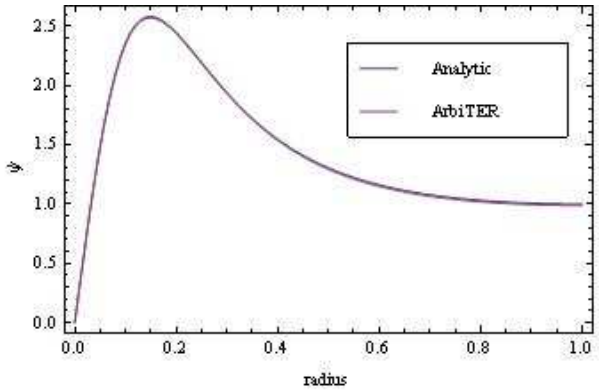


Figure 6: Comparison of analytic and numerical solutions of a point magnetic dipole in a superconducting cylinder along the cylinder midplane.

of an electric dipole; the difference is merely in the boundary conditions, which are zero derivative with respect to this potential as opposed to zero value as they would be in the case of an electric dipole. It can be shown that variation in this potential is purely sinusoidal in the azimuthal direction, corresponding to an $n=1$ mode. Given this, we can then write down the differential equation for this potential:

$$\frac{\partial^2}{\partial r^2}\psi+\frac{1}{r}\frac{\partial}{\partial r}\psi-\frac{1}{r^2}\psi+\frac{\partial^2}{\partial z^2}\psi=S. \tag{3.15}$$

This problem can be solved in a semi-analytic manner by separation of variables. This results in a series expansion in terms of cosines or hyperbolic cosines and Bessel functions. This solution is described in more detail in the appendix.

Solution of this problem in ArbiTER is relatively straightforward. The formula in Eq. (3.15) is written in equation language, with the differential operators $\partial^2/\partial r^2$, $\partial/\partial r$, and $\partial^2/\partial z^2$ defined in such a way as to implement zero-derivative boundary conditions on all outer boundaries and zero-value boundary conditions on the central axis. Radius is input as a 1D profile function, which is then distributed across the entire domain. The source term is a product of two 1D profile functions (one for r and one for z) which are distributed across the domain before being multiplied.

The analytic and numerical solutions are compared in Fig. 6. In this case, the analytic solution was a superposition of 30 terms in the series expansion, and the numerical solution was calculated at a grid resolution of 100 radial points and 101 axial points. The radius and length of the cylinder are equal to 1, and the dipole is at a vertical position of .3. A 1% discrepancy between the solutions exists for the data shown, however, this discrepancy was found to scale with the convergence conditions specified as command-line options to the matrix solver. It is therefore presumed that this discrepancy can be made asymptotically small.

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

3.5 Extended domain method in x-point geometry

In the 2DX emulation topology (i.e. for a single-null x-point geometry plasma), there is a branch cut next to the x-point across which a phase-shift periodic boundary condition is applied, i.e. the eigenfunction on one side of the branch cut is matched to that on the other side times a phase factor which accounts for the toroidal periodicity of the modes [23]. In this section we present an alternative to this technique. The motivations for this and the situations where such a technique may have an advantage are a topic for future discussion. For purposes of this paper, it is sufficient to discuss the methods by which such a technique are implemented in the ArbiTER code.

In the extended domain method, each field line in a tokamak plasma is extended for multiple poloidal transits. With each successive poloidal transit, the phase shift between adjacent flux surfaces accumulates, and eventually has a suppressive effect on the instability of interest. The effect is related to the integrated magnetic shear. This results in an eigenmode that is roughly centered in a domain constituting multiple poloidal transits, with amplitude decaying exponentially to either side. If the amplitude decays to a sufficient degree before reaching the boundaries of the extended domain, the boundary conditions at the ends of the domain have little effect on growth rates or mode structure, so the introduction of such an artificial boundary does not affect the validity of the computational result. The extended domain method is commonly applied in codes which operate on closed flux surfaces, and is closely relate to the well-known ballooning formalism [32] However, the present application to complex separatrix-spanning topologies is, to the best of our knowledge, original.

This type of method can be implemented easily in ArbiTER by using a third dimension to represent the number of transits in the extended domain. By introducing an offset in the linkage across the branch cut, one end of the edge in each layer connects it to the next layer. A similar domain lacking the extra dimension is constructed to load in profile functions. An operator linking the reduced-dimensionality domain to the extended domain allows these profile functions to be projected onto the extended domain in a convenient manner. The layout of topological subdomains used in this method is shown in Fig. 7.

Results are shown in Figs. 8 and 9. Fig. 8 shows the sum of the amplitudes for all three periods of the extended domain, superimposed on a single period of the domain. The resulting eigenmode is localized near the separatrix. In Fig. 9, the amplitude in the full extended domain is shown. In this plot, we can see the decay of the amplitude as one moves away from the center period. We can also see that the amplitude outside the separatrix is largely localized to the center period.

4 Summary

A new eigenvalue solver, the ArbiTER code, has been developed. It is capable of solving linear partial differential equations (including certain classes of integral operators) in

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

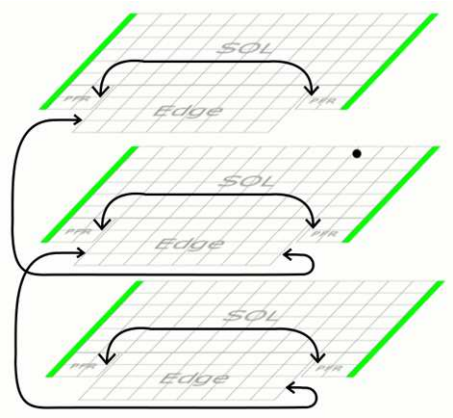


Figure 7: Layout of regions used in the extended domain method.

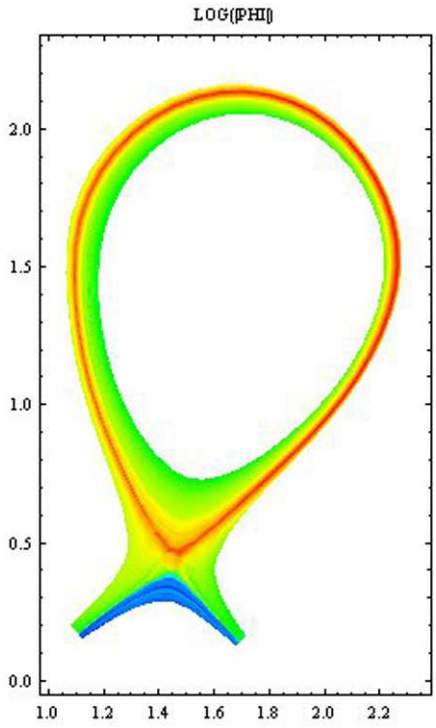


Figure 8: Eigenfunction from a resistive ballooning model, with extended domain solution superimposed onto 2D space. Horizontal axis is radius, vertical axis is elevation.

arbitrary topology and in an arbitrary number of dimensions. While its development is primarily motivated by problems in edge plasma physics, its capabilities are not limited to such applications.

The ArbiTER code has been tested in a variety of cases that demonstrate its capabili-

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

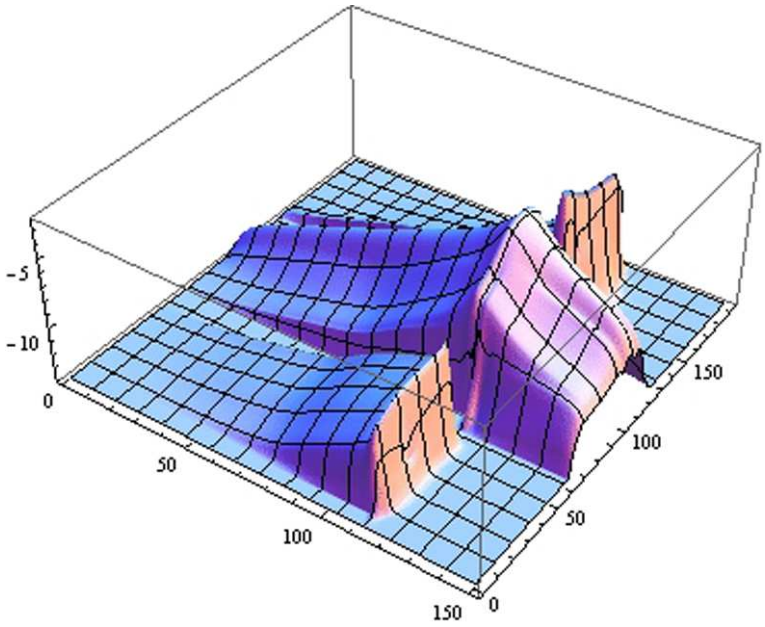


Figure 9: Eigenfunction from a resistive ballooning model in the extended domain method, in the original extended domain grid. Horizontal axes are grid coordinates in the poloidal (theta) and radial (psi) directions; the vertical axis is the logarithm of amplitude of electrostatic potential. Three periods of the extended domain are shown in the theta direction.

ties. These include kinetic plasma models, fluid models in snowflake or extended domain topology, and simple finite element analysis. For application to large scale, high dimensionality problems, the ArbiTER code implements the parallel version of SLEPc [29] with MPI. Ongoing and future work will explore advanced gyrokinetic models and the resulting computation requirements and parallel scaling.

The ArbiTER code shows tremendous potential both as a benchmarking aid and in primary physics studies. Its flexibility allows it to be applied in a wide variety of situations, and can thus be compared against nearly any code for which a relevant linear regime exists. Moreover, it can readily assimilate a variety of mathematical and numerical techniques, allowing essentially the same model equations to be solved by different methods, for instance comparing extended domain techniques to phase-shift periodic boundary conditions. Such capabilities make the ArbiTER code a valuable tool for exploratory research in plasma physics.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Fusion Energy Sciences under Award Number DE-SC0006562.

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

A Analytic solution for dipole problem

Eq. (3.15) in Section 3.4 can be solved analytically by expanding in a series. This can be done in two different ways. The first is to expand as a series of Bessel functions in r, then solve the resulting differential equation in z for each component. This results in the equation:

$$\psi_1(r,z) = S\sum_{m=1}^{\infty} \frac{J_1(\mu_m r/a)}{(ha^2)[1-(1/\mu_m)^2]J_1^2(\mu_m)} \frac{\mu_m}{2a} S_m(z), \tag{A.1}$$

$$S_m = \begin{cases} \frac{\cosh(\mu_m z/a)\cosh(\mu_m(h-z_0)/a)}{(\mu_m/a)\sinh(\mu_m h/a)}, & z < z_0, \\ \frac{\cosh(\mu_m(h-z)/a)\cosh(\mu_m z_0/a)}{(\mu_m/a)\sinh(\mu_m h/a)}, & z \geq z_0, \end{cases} \tag{A.2}$$

where h is the size of the cylinder in z, a is the size in r, z_0 is the vertical position of the dipole, and μ_m is the mth root of $J_1'(x)=0$.

The second approach is to expand as a Fourier series in z, then solve the resulting differential equation in r for each component. This results in the equation:

$$\psi_2(r,z) = S\sum_{k=0}^{\infty} \cos\left(\frac{\pi k}{h}z\right)\cos\left(\frac{\pi k}{h}z_0\right)\frac{\pi k}{h^2}\left(K_1\left(\frac{\pi k}{h}r\right) - I_1\left(\frac{\pi k}{h}r\right)\frac{K_1'(\frac{\pi k}{h}a)}{I_1'(\frac{\pi k}{h}a)}\right) + S\frac{r/a^2+1/r}{2h}. \tag{A.3}$$

Each of these methods has a part of the domain in which convergence is slow with respect to the number of terms in the series. The first method works poorly for $z \approx z_0$. The second method works poorly for $r \approx 0$. These solutions can be spliced together by applying the condition:
If

$$\begin{cases} |z-z_0| \geq r, & \psi(r,z) = \psi_1(r,z), \\ |z-z_0| < r, & \psi(r,z) = \psi_2(r,z). \end{cases} \tag{A.4}$$

This results in a solution that converges rapidly at all points except the source, which is divergent to begin with. This yields a solution that can be used to benchmark a numerical model.

B Example of ArbiTER file setup

To provide a concrete example of how to set up a problem using ArbiTER, let us consider the problem of a 1D quantum harmonic oscillator equation:

$$-\lambda\Phi = \partial_x^2\Phi + V\Phi. \tag{B.1}$$

The computational domain for the field variable Φ, the profile function V, as well as the differential operator ∂_x^2 are defined by the topology language:

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

```
brick simple spacebase
dimensions nx       #Creates a 1D space of size nx
domain simple space
uses spacebase      #Makes the space available to define functions
link simple spacedxu #Creates mapping of space to itself offset by 1
links space space dir 3 postofs domin dir 1 one
link simple spacedxl      #As above
links space space dir 3 postofs domin dir 2 one
operator simple ddx2      #Combines these linkages for upper and
#lower derivatives into a second derivative operator
+space *(2+0j) *(-1+0j) +spacedxu +spacedxl *onodx2
```

Additional content in the topology file is then used to define the variables described in the above lines. The equation file then contains the model equation,

```
gg*(-1+0j)*PHI=(1+0j)*ddx2*PHI+(1+0j)*vv*PHI
```

the definition of the constant `onodx2`,

```
xx=dx
xx=xx*dx
xx=(1+0j)/xx
onodx2=xx
```

as well as the input parameters,

```
end end null
nx nx null
dx dx null
wci wci null
vv vv space
```

The grid file then contains a list of parameters and profile functions:

```
nx
99
dx
0.0816327
wci
1
vv
99
-16.
-15.3536
-14.7205
-14.1008
...
```

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

The profile function vv can be generated by a wide variety of means. In this case it is generated using a Mathematica script, but virtually any programming language can generate this information. Once these pieces of data are presented in the correct input files, it is then only necessary to run the program to solve the relevant model equation.

While this example only exercises a tiny fraction of the capabilities of this code, it does provide a concrete example of how to implement a simple physics model.

References

[1] A. Zeiler, J. F. Drake and D. Biskamp, Phys. Plasmas **4**, 991 (1997).
[2] B. I. Cohen, M. V. Umansky, W. M. Nevins, M. A. Makowski, J. A. Boedo, D. L. Rudakov, G. R. McKee, Z. Yan, and R. J. Groebner, Phys. Plasmas **20**, 055906 (2013).
[3] B. N. Rogers, J. F. Drake and A. Zeiler, Phys. Rev. Lett. **81** (1998) 4396.
[4] C. S. Chang et al., Journal of Physics: Conference Series **180** (2009) 012057.
[5] D. A. Russell, J. R. Myra and D. A. D?Ippolito, Phys. Plasmas **16**, 122304 (2009).
[6] F. D. Halpern, P. Ricci, B. Labit, I. Furno, S. Jolliet, J. Loizu, A. Mosetto, G. Arnoux, J. P. Gunn, J. Horacek, M. Kocan, B. LaBombard, C. Silva, and JET-EFDA, Nucl. Fusion **53**, 122001, (2013).
[7] J. Loizu, P. Ricci, F. D. Halpern, S. Jolliet and A. Mosetto, Nucl. Fusion **54**, 083033 (2014).
[8] K. Hallatschek, Plasma Phys. Control. Fusion **49**, B137 (2007).
[9] M.V. Umansky, X.Q. Xu, B. Dudson, L.L. LoDestro, J.R. Myra, Comp. Phys. Comm. **180**, 887 (2009).
[10] O. E. Garcia, V. Naulin, A. H. Nielsen, and J. J. Rasmussen, Phys. Plasmas **12**, 062309 (2005).
[11] P. Beyer, S. Benkadda, G. Fuhr-Chaudier, X. Garbet, Ph. Ghendrih, Y. Sarazin, Plasma Phys. Control. Fusion **49**, 507 (2007).
[12] P. Guzdar, R. Kleva, R. Groebner, and P. Gohil, Phys. Plasmas **11**, 1109 (2004).
[13] P. Ricci and B. N. Rogers, Phys. Plasmas **20**, 010702 (2013).
[14] R. H. Cohen, and X. Q. Xu, Contrib. Plasma Phys. **48**, 212 (2008).
[15] R. H. Cohen, B. LaBombard, D. D. Ryutov, J. L. Terry, M. V. Umansky, X. Q. Xu and S. Zweben, Nucl. Fusion **47**, 612 (2007).
[16] V. Naulin, A. Kendl, O. E. Garcia, A. H. Nielsen and J. Juul Rasmussen, Phys. Plasmas **12**, 052515 (2005).
[17] X. Q. Xu, R. H. Cohen, T. D. Rognlien, and J. R. Myra, Phys. Plasmas **7**, 1951 (2000).
[18] X. Q. Xu, W. M. Nevins, T. D. Rognlien, R. H. Bulmer, M. Greenwald, A. Mahdavi, L. D. Pearlstein and P. Snyder, Phys. Plasmas **10**, 1773 (2003).
[19] D. R. Hatch, P. W. Terry, F. Jenko, F. Merz, and W. M. Nevins, Phys. Rev. Lett. **106**, 115003 (2011).
[20] M. Kammerer, F. Merz, and F. Jenko, Phys. Plasmas **15**, 052102 (2008).
[21] P.J. Roache, Verification and Validation in Computational Science and Engineering, Hermosa Publishers, Albuquerque, NM (1998).
[22] W.L. Oberkampf and T. G. Trucano, Progress in Aerospace Sciences **38**, 209 (2002).
[23] D. A. Baver, J. R. Myra, M. V. Umansky, Comp. Phys. Comm. **182**, 1610 (2011).
[24] D.D. Ryutov, Phys. Plasmas **14**, 064502 (2007).
[25] H. Strauss, Phys. Fluids **24**, 2004 (1981).
[26] T. C. Hender, B. A. Carreras, W. A. Cooper, J. A. Holmes, P. H. Diamond, and P. L. Similon, Phys. Fluids **27**, 1439 (1984).

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press

[27] P. N. Guzdar and J. F. Drake, Phys. Fluids B **5**, 3712 (1993).
[28] E. A. Belli, J. Candy, Phys. Plasmas **17**, 112314 (2010).
[29] http://www.grycap.upv.es/slepc/
[30] http://www.mcs.anl.gov/petsc/
[31] J. R. Shewchuk, Delaunay Refinement Algorithms for Triangular Mesh Generation, Computational Geometry: Theory and Applications **22**(1-3):21-74, May 2002.
[32] J. W. Connor and J. B. Taylor, Phys. Fluids **30**, 3180 (1987)

https://doi.org/10.4208/cicp.191214.021015a Published online by Cambridge University Press