

UKAEA-CCFE-PR(23)145

S. Mijin, D. Power, R. Holden, W. Hornsby, D.
Moulton, F. Militello

ReMKiT1D - A framework for building reactive multi-fluid models of the tokamak Scrape-Off Layer with coupled electron kinetics in 1D

Enquiries about copyright and reproduction should in the first instance be addressed to the UKAEA Publications Officer, Culham Science Centre, Building K1/O/83 Abingdon, Oxfordshire, OX14 3DB, UK. The United Kingdom Atomic Energy Authority is the copyright holder.

The contents of this document and all other UKAEA Preprints, Reports and Conference Papers are available to view online free at scientific-publications.ukaea.uk/

ReMKiT1D - A framework for building reactive multi-fluid models of the tokamak Scrape-Off Layer with coupled electron kinetics in 1D

S. Mijin, D. Power, R. Holden, W. Hornsby, D. Moulton, F. Militello

ReMKiT1D - A framework for building reactive multi-fluid models of the tokamak Scrape-Off Layer with coupled electron kinetics in 1D

Stefan Mijin^{a,*}, Dominic Power^b, Ryan Holden^{a,c}, William Hornsby^a, David Moulton^a, Fulvio Militello^a

^a*United Kingdom Atomic Energy Authority, Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK*

^b*Blackett Lab., Plasma Physics Group, Imperial College London, London, SW7 2AZ, UK*

^c*School of Mathematics and Physics, University of Surrey, Guildford, GU2 7XH, UK*

Abstract

In this manuscript we present the recently developed flexible framework for building both fluid and electron kinetic models of the tokamak Scrape-Off Layer in 1D - ReMKiT1D (**R**eactive **M**ulti-fluid and **K**inetic **T**ransport in **1D**). The framework can handle systems of non-linear ODEs, various 1D PDEs arising in fluid modelling, as well as PDEs arising from the treatment of the electron kinetic equation. As such, the framework allows for flexibility in fluid models of the Scrape-Off Layer while allowing the easy addition of kinetic electron effects. We focus on presenting both the high-level design decisions that allow for model flexibility, as well as the most important implementation aspects. A significant number of verification and performance tests are presented, as well as a step-by-step walkthrough of a simple example for setting up models using the Python interface.

Keywords: framework; fluid; kinetic; electrons; tokamak; SOL; multi-fluid; collisional-radiative

PROGRAM SUMMARY

Program Title: ReMKiT1D

CPC Library link to program files: (to be added by Technical Editor)

*Corresponding author.

E-mail address: stefan.mijin@ukaea.uk

Developer's repository link: <https://github.com/ukaea/ReMKiT1D> and <https://github.com/ukaea/ReMKiT1D-Python>

Code Ocean capsule: (to be added by Technical Editor)

Licensing provisions: GPLv3

Programming language: Fortran, Python

Supplementary material: <https://doi.org/10.14468/fdq7-z869>

Nature of problem: The flexible generation and modification of 1D models pertaining to multi-fluid simulations of the tokamak Scrape-Off Layer (SOL) with electron kinetics and reaction support. This would then allow both for rapid iteration on reduced models as well as the evaluation of kinetic electron effects in equilibria and during transients, following the formalism previously developed for SOL-KiT[1]. The framework was not only envisioned as the successor to SOL-KiT, but a tool that would allow users to construct their own models coupled with electron kinetics capabilities.

Solution method: The framework is written heavily utilizing Object-Oriented design principles, in particular using a generalized version of the puppeteer pattern as presented by Rouson et al[2], as well as the heavy use of the strategy pattern/dependency injection. The Fortran code is MPI parallel and utilizes the PETSc library for implicit time-stepping. MPI parallelization is extended to distribution function Legendre harmonics, allowing for improved strong scaling. Initialization of the Fortran framework is done using JSON configuration files generated by an accompanying Python interface, and data analysis is standardized using widely used data formats such as HDF5.

Additional comments including restrictions and unusual features: The present manuscript focuses on the design and high-level implementation of the framework, as well as the demonstration of the workflow and various verification and performance benchmarking tests. Some details are avoided for the sake of brevity at various points, and these are meant to be available as part of the general code documentation or tutorials offered on the main repositories.

References

- [1] S. Mijin, A. Antony, F. Militello, SOL-KiT-Fully implicit code for kinetic simulation of parallel electron transport in the tokamak Scrape-Off Layer, *Comp. Phys. Comm.* 258 (2021)
- [2] D. Rouson, J. Xia, X. Xu, *Scientific Software Design: The Object- Oriented Way*, Cambridge University Press, 2011

1. Introduction

The Scrape-Off Layer (SOL) denotes the region of open magnetic field lines just outside of the core of magnetically confined fusion (MCF) devices, such as the tokamak. These open field lines impinge on material surfaces (walls, limiters, or divertor targets) in the device, leading to plasma-surface interactions, which can inject impurity species into the plasma. Furthermore, the region outside the core, often referred to as the edge, is cold enough for atomic, or even molecular, neutrals to persist and interact with the plasma in a non-trivial way. As such, this region of the device is particularly rich in multi-species plasma physics, with many reactions occurring between the species.

Furthermore, due to the scale separation of parallel (to the open magnetic field lines) and perpendicular transport, 1D analytical and numerical models of the SOL have been a mainstay since the early days of SOL simulations. The reduced geometrical/dimensional complexity allows for the treatment of more involved physics, while keeping run times short compared to 2D or 3D simulations. As such, a range of 1D codes geared towards exploring one or more aspects of the SOL have been in use over the last several decades, ranging from PIC codes[1, 2], to continuum fluid[3, 4, 5] and kinetic [6, 7, 8, 9] codes. These have been used for a variety of problems, and some of them will be reviewed here to provide context for the code to be presented in the bulk of this text.

One of the most common problems tackled by 1D codes is that of simulating equilibrium divertor regimes, with a particular focus on detachment[10, 11, 3] - the regime in which the plasma recombines significantly in front of the target, effectively producing a cloud of neutrals which shields it from exposure to the hot upstream conditions. Tying into this is the study of how target conditions are affected by transient phenomena upstream[12, 13], when particles and energy are transiently injected far away from the targets and allowed to propagate towards them and interact with any of the background plasma and neutral species present in the equilibrium conditions.

Another field of research is that of parallel transport in the SOL. While fluid models tend to use classical values of transport coefficients (e.g. heat conductivity or viscosity) that rely on local values of plasma parameters, due to the strong parallel gradients formed in the SOL, kinetic effects can come into play[14], modifying the classical values in both equilibrium and transient conditions. Examples of kinetic effects include heat flux suppression

and enhancement[15, 16, 17], as well as the modification of the target plasma sheath properties[18], which fluid codes struggle to include consistently, often resorting to measures such as heat flux limiting.

Finally, with the presence of many species, and atomic and molecular physics coming into play, the need arises to couple the various species through reaction rates, leading naturally towards collisional-radiative models (CRMs), which attempt to solve and reduce a set of coupled ODEs[19, 20, 21, 22, 23, 24]. Most often, however, these models are 0D, neglecting transport effects and taking reaction rates based on an underlying Maxwellian distribution for the plasma species. 1D models, especially kinetic models, when including collisional-radiative processes, can be used to study transport effects on the particle and energy balance governed by these complicated reactions[25].

The goal of the framework to be presented is to provide a relatively easy way to build models that can tackle most of the above issues, combining multi-fluid, collisional-radiative, and kinetic physics, particularly in the context of SOL plasmas. However, the presented framework could also be relevant to other fields where 1D fluid equations might need to be solved.

The ReMKiT1D (**R**eactive **M**ulti-fluid and **K**inetic **T**ransport in **1D**) framework consists of a core code written in Modern Fortran and controlled through JSON configuration files constructed using a Python package, allowing for rapid model iteration and flexibility. While the numerical approach is extremely customizable, with many features exposed to the user at a Python level, a significant number of numerical procedures and approaches are adapted from work previously done on the hydrogenic hybrid fluid-kinetic code SOL-KiT[26], which has been used to tackle some of the problems detailed above, namely the quantification of kinetic electron effects on parallel transport and reaction rates in equilibrium and transient conditions. However, while SOL-KiT is a purpose-built code, ReMKiT1D has been designed as a framework, such that models like the one implemented in SOL-KiT, can be easily built, used, and modified by a variety of users. Furthermore, it allows for improvements in performance by introducing a second parallelizable dimension in the electron distribution function harmonics (see Sections 4 and 6).

A brief overview of problem types and equations ReMKiT1D is designed to tackle will be given in Section 2, followed by the description of high-level features and concepts used in designing ReMKiT1D in Section 3. Those implementation details relevant to the 1D problems presented here, such as

grids, normalizations, and some operators, will be presented in Section 4, together with the description of the MPI communication in both spatial and harmonic directions. Details on Python-Fortran coupling and an example workflow are given in Section 5, before an extensive list of verification and benchmarking problems are presented in Section 6, including parallel performance scaling tests with the novel harmonic dimension parallelization. Planned and potential use cases and extensions of the framework are discussed in Section 7.

2. Target problem classes

With a focus on mathematical models arising in the research of transport along magnetic field lines in the SOL, ReMKiT1D has been designed to handle systems of nonlinear ODEs and PDEs arising from coupled fluid and kinetic models. These will be laid out in this section, without focusing on any individual model or implementation.

2.1. Systems of nonlinear ODEs

Nonlinear ODEs arise both from the spatial and velocity discretization of fluid and kinetic models of interest, as well as naturally in the context of collisional-radiative modelling. In general, given a vector of variables \vec{v} , the system of interest can be written as

$$\frac{d\vec{v}}{dt} = \mathbf{M}(\vec{v}) \cdot \vec{v} + \vec{\Gamma}, \quad (1)$$

where $\mathbf{M}(\vec{v})$ is the matrix of (nonlinear) coupling coefficients and $\vec{\Gamma}$ some constant vector. It is convenient to write the system in this way both for the purposes of defining the default implicit integration scheme used, and to draw parallels with the form of general collisional-radiative models [24]. In general, the systems of ODEs that occur in problems of interest are stiff, and implicit integration schemes together with potentially flexible operator splitting might be required (see Section 3).

2.2. 1D PDEs

Systems of hyperbolic conservation laws of the form

$$\frac{\partial X(\vec{x})}{\partial t} + \nabla \cdot \vec{\Gamma}_X(\vec{x}) = S_X, \quad (2)$$

where X is a conserved quantity, $\vec{\Gamma}_X$ the flux of that quantity, and S_X the source/sink of X , arise naturally in the modelling of multi-species fluids. Many well-known methods of solving such equations exist, reducing the system through discretization to a system of, in general, nonlinear ODEs. The complexity of the conservation laws comes from the forms of the fluxes and sources, which can be complex nonlinear functions of conserved quantities. Furthermore, cases where evolving the primitive (instead of conserved) quantities is simpler also arise. This can lead to parabolic equations, such as those that arise in heat conduction problems, as well as combined advection-diffusion problems.

Finally, equations without an explicit time derivative term can arise, as is the case with the elliptic problem of Poisson's equation

$$\Delta\varphi = f. \tag{3}$$

Further complications arise with the addition of complicated boundary conditions due to the interaction of the SOL plasma with the wall, such as the sheath boundary condition and recycling. These are mentioned in Appendix A, and for more details the reader is encouraged to consult previous work [26, 6, 2], particularly in the context of kinetic boundary conditions.

As such, it is of interest to at least attempt to cover as many of these cases as possible, which is simplified somewhat due to the 1D nature of the problems considered in this manuscript. Implementation details behind the 1D differential operators available in ReMKiT1D (as well as notes on user level customization) will be presented in Section 4.

2.3. 1D electron kinetic equation

Electron kinetic effects such as heat flux suppression, sheath boundary effects, as well as potential modification of reaction rates due to non-Maxwellian distributions, are of interest both in equilibrium and transient conditions. Following the approach in SOL-KiT[26], these effects are included as emergent physics through the option to solve the 1D electron kinetic equation,

$$\frac{\partial f}{\partial t} + v_x \frac{\partial f}{\partial x} - \frac{eE}{m_e} \frac{\partial f}{\partial v_x} = \left(\frac{\delta f}{\delta t} \right)_c. \tag{4}$$

Using an expansion of the distribution function into Legendre harmonics f_l , this results in an equation of the form

$$\frac{\partial f_l}{\partial t} = A_l + E_l + C_l \quad (5)$$

where A_l , E_l , and C_l represent spatial advection, velocity space advection due to the electric field, as well as collision operators, respectively. This enables fine control over the fidelity of the kinetic effects included, through both choosing the number of resolved distribution function harmonics, as well as controlling the individual operators evolving each harmonic. Fluid moments are obtained directly from moments of the individual harmonics, such that scalar moments are given by

$$\langle \phi \rangle = \int \phi(v) f(\vec{v}) d\vec{v} = 4\pi \int_0^\infty \phi(v) f_0(v) v^2 dv,$$

and vector moments by

$$\langle \vec{a} \rangle = \int \vec{a}(v) f(\vec{v}) d\vec{v} = \frac{4\pi}{3} \int_0^\infty a(v) f_1(v) v^2 dv$$

For more details on the expansion, the reader is directed towards the SOL-KiT model and other similar models in the literature[27, 28, 29, 30]. Here it will only be mentioned that the capability for including the Coulomb collision operator outside of the Lorentz approximation and the proper treatment of inelastic electron-neutral collisions is desirable and has been implemented in ReMKiT1D following the techniques established in SOL-KiT[26].

Given the broad range of problems and the aim for flexibility, the need for separating high-level concepts and implementation arises. ReMKiT1D endeavours to exploit low-cost abstraction and the concept of scalable design to produce a framework that is both fit for purpose in terms of adequately addressing the target problem classes above, as well as providing the user with low level control, high level quality-of-life features, flexibility in the models and numerics, and rapid iteration. This requires at least a conceptual separation of high-level concepts and those directly tied to the implementation of operators and solvers in 1D, and an attempt has been made to separate these concepts with the design of the source code as well. These high-level concepts and the 1D implementation will be presented in the following two sections.

3. High-level concepts and patterns

In this section the basic high-level pattern behind ReMKiT1D will be presented, providing both a bird’s eye view of the code, as well as some general and essential implementation details. However, exact implementations and interfaces are omitted for the sake of brevity and flow, and the choice is made to focus on a more descriptive presentation of the components and their relationships. To this end, both classes and objects are written in **PascalCaseBold** text to differentiate them from concepts with similar (and in some cases the same) names.

Where appropriate, Section 5 and the example workflow there will be referenced to provide additional concrete examples of concepts discussed in this section.

3.1. The Modeller-Model-Manipulator pattern

ReMKiT1D’s high-level algorithm is built on the generalization of the puppeteer pattern as presented by Rouson et al [31]. Figure 1 shows a simplified UML diagram of this pattern, as implemented in ReMKiT1D.

The fundamental object in the pattern is a central **Modeller** object, containing the variables that should be accessible to multiple components, as well as any supporting library wrapper routines (such as MPI or PETSc[32]). At this point, the exact implementation/data structure behind the variables is irrelevant, as long as the interface allows for them to be safely passed through to other components. In ReMKiT1D, encapsulation of the variable values in the **VariableContainer** has been avoided, in order to minimize the number of potential copies/memory allocation¹. More details on variables, communication, and external library interfacing is given in following sections.

The main relationships in the **Modeller-Model-Manipulator** pattern are between the **Modeller** and the **Model** objects it contains, as well as between the **Modeller** and the **Manipulator** objects it contains. These will be discussed in detail in the next two subsections, but the following short summary should illustrate the responsibilities of the different pattern members:

¹Functional programming patterns and different compiler optimization might be another way to do this, but this has not been done for the version of the code presented here.

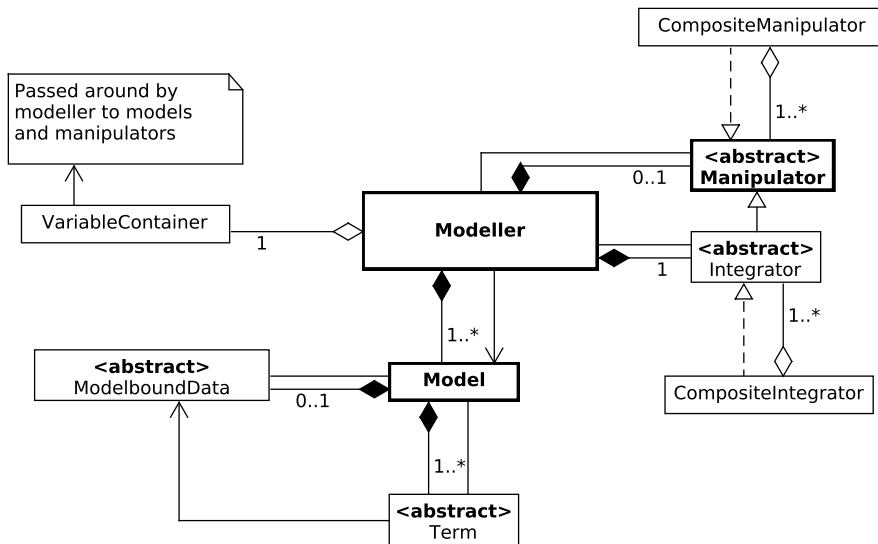


Figure 1: A simplified UML diagram of the high-level structure of ReMKiT1D, showcasing the **Modeller-Model-Manipulator** pattern, with those three components in bold edged boxes. A **Modeller** has one or more **Models** and is coupled with its **Manipulator** component so that the **Manipulator** can call **Modeller** functions and modify the variables. A special case of the **Manipulator** is the **Integrator**, which uses the **Modeller-Model** interface to evolve variables in time.

- **Models** contain, update, and evaluate **Terms** - objects representing additive terms in the various equations of Section 2. Any one term is associated with one variable, which it nominally evolves. A **Model** is also allowed to have data specifically tied to it, and by default only accessible to the **Terms** within it.
- The **Manipulator** modifies variables based on calling various **Modeller** routines. This can be anything from performing an integration step on a set of **Models/Terms** while obeying communication rules to evaluating and storing individual terms into diagnostic variables. The main concept behind the **Manipulator** class is the enabling of high-level dependency injection.
- The **Modeller** provides a simple interface to the external world and contains the **Manipulators** and **Models**. It can then be fitted into a time-stepping/output loop to complete the time integration procedure.

3.1.1. Models and Terms

As noted above, **Terms** represent additive terms S_i in equations of the form $dv/dt = \sum_i S_i$ or $\sum_i S_i = 0$. They are always associated with an evolved variable, and can be evaluated into a variable data structure conforming to the evolved variable (see variable subsection below and Section 4). At the moment, ReMKiT1D differentiates between general explicit **Terms** which can only be evaluated, and matrix **Terms**, which can both be evaluated and passed into PETSc. The present version of the code focuses solely on the matrix terms, and all examples in this paper will focus on them. However, some development has gone into building the architecture for the addition of matrix-free/explicit **Terms**, and that will be noted where appropriate.

The **Models** in Section 5, for example, each contain a single term, representing spatial derivative operators used to evolved coupled variables, resulting in a wave equation.

Let both v and u be implicit variables (those that have a mapping to a global implicit vector for use in PETSc, see below for more details). Then a matrix term describing the evolution of v represents the RHS of

$$\frac{dv_i}{dt} = M_{ij}u_j, \quad (6)$$

where summation on repeated indices is assumed. The matrix M_{ij} , whose functional dependence on various variables and time is dropped here for readability, is then composed of multiplicative components and a stencil S_{ij} so that

$$M_{ij} = cT(t)F_iR_iC_jS_{ij}, \quad (7)$$

where c and F_i are some constant² and a constant array that depends only on the index i , which in practice involves functions of just the spatial/velocity space coordinates and $T(t)$ is an explicit multiplicative time dependence. Finally, R and C are row and column functions of some variables in general. In practice, these are set to products of variables raised to some power like

$$R_i = \prod_n v_{n,i}^{p_n},$$

²This constant is usually associated with some normalization quantities - see Section 5

where n now indexes different variables. Together with stencils, these fully define a matrix term, noting that every component other than a stencil is optional. Stencils come in different forms, representing the discretization of different operators. Some will be presented in Section 4. The simplest stencil is a diagonal stencil, with an example of a Python code snippet used to generate the matrix **Term** representing

$$\frac{dv}{dt} = \frac{u}{2w^2}$$

given as:

```

1 import RMK_support.simple_containers as sc
2
3 eVar = "v"
4 iVar = "u"
5 nConst = 0.5
6
7 #Set R = 1/w^2
8 vData = sc.varData(reqRowVars=["w"], reqRowPowers=[-2.0])
9
10 newTerm = sc.GeneralMatrixTerm(evolvedVar=eVar, implicitVar=
    iVar, customNormConst=nConst, varData=vData, stencilData=sc.
    diagonalStencil())

```

Note that the above variables (v, u and w) do need to be registered in the **VariableContainer** object, which will be covered below. The above will then generate the matrix $M_{ij} = \delta_{ij}/(2w_i^2)$ that can then be used to evolve v using an implicit method built with the PETSc library. More complicated stencils will be presented in Section 4, but an example, relevant to Section 5, is a spatial gradient stencil. Thus, stencils represent discretizations of various differential or integral operators one might encounter while dealing with fluid or electron kinetic problems.

Terms like the one above are then grouped into implicit and general groups, with the only difference between the groups that implicit groups can only contain matrix **Terms**, and can be evolved using the fixed-point integrator (see next section and Appendix B), while general groups can only be evaluated for use in solvers that only require a RHS vector³. These groups are housed in **Model** objects and can be updated and evaluated from the

³In order to be able to evaluate a group, all **Terms** must have the same evolved variable, otherwise the group can only have its matrices passed to PETSc if it is an implicit group

parent **Modeller** object. In general, a **Term** can also use the evaluation of other **Terms** in the **Model** or even borrow their matrices for modification. An example of where this is useful is writing terms that depend on taking velocity space moments of other terms (see Appendix A).

The **Model** can also contain a **ModelboundData** object, which, in principle, calculates and stores data accessible to the **Model** and its **Terms**. Some examples of **ModelboundData** will be covered in the following sections.

3.1.2. Manipulators and Integrators

The fundamental idea behind **Manipulators** is formalizing high-level dependency injection through enabling callbacks to the **Modeller**⁴. **Manipulators** then allow for the direct manipulation of variable data in the **Modeller**.

Manipulators are stored in a **CompositeManipulator** object, which the **Modeller** calls directly, and each **Manipulator** is also associated with a priority (0 being the highest). This way, one can control when certain **Manipulators** are called. For example, one might want to call a **Manipulator** as often as possible as it modifies a variable that is used in some internal iteration of an integrator. On the other end of the spectrum, one might want to just call the **Manipulator** before outputting data, if the **Manipulator**'s task is extracting diagnostic variables such as term evaluation.

Several important data access **Manipulators** are implemented at the moment. They include term evaluation **Manipulators**, that store the evaluation value of a **Term** in one variable, useful for analysis, as well as debugging. Similarly, an extraction **Manipulator** is available for accessing **ModelboundData** values that can fit into regular variables (more on this in the next sections).

Finally, the most important **Manipulators** are the **Integrators**, which have their own specialized container in the **CompositeIntegrator** object. These objects are responsible for the time integration, and the **CompositeIntegrator** controls any single integration call in two ways:

1. By applying any global time step control
2. By calling individual **Integrator** components in accordance with precisely defined integration steps

⁴One can also think of this as an implementation of the well-known strategy pattern

The application of the global time step control is done by re-scaling the initial time step in accordance with some rule (see Appendix B for an example).

Integration steps are defined by the following:

- The associated **Integrator** object
- The fraction of the global time step (the total time step requested in the integration call) associated with the step
- Evolution and update rules (e.g. which **Models** and **Term** groups should be evaluated or how often to update non-linear terms)

In combination with the grouping of **Terms**, this gives the user full control over any potential operator splitting through simply defining each step in sequence, and associating the **Models** (and optionally even individual **Term** groups) to be used in the evolution. Furthermore, the control over update frequency of both individual non-linear **Terms** as well as **ModelboundData** opens up performance optimization opportunities at the expense of accuracy, all accessible at the highest level of the interface.

At the moment, two integrator types are implemented in ReMKiT1D, an explicit Runge-Kutta integrator up to 4th order⁵, as well as a first order Backwards Euler method with fixed-point iterations, borrowing heavily from SOL-KiT's integrator. It is briefly described in Appendix B.

3.2. Variables, communication, and Derivations

Variable data is stored in **VariableContainer** objects, which are responsible for providing data indexing as well as features pertaining to the two general types of variables in ReMKiT1D:

- Implicit variables - these are allowed to be evolved by matrix terms and can be declared as their RHS implicit variables,
- Derived variables - while these cannot be implicitly evolved, they can have derivation rules associated with them.

⁵Arbitrary Butcher tableau support has been implemented, but not available in the Python interface as of v1.0.x.

A **VariableContainer** is equipped with routines to generate local (in the MPI sense) flattened vectors of variables for use in PETSc routines. The only variables that are included in these local vectors are implicit variables.

In many use cases, there are variables which aren't suitable to be included in an implicit solve, but must be calculated using existing variables. The way ReMKiT1D deals with this is through **Derivations**, which act as generalized function wrappers that take in multiple variables and calculate one⁶. A derivation rule is then a combination of a derivation and a list of required variables, and we refer to variables calculated using this approach as derived variables.

In general, **Derivations** wrap impure functions, allowing for changes to their internal state. However, most derivations available in ReMKiT1D are written avoiding side-effects, with some tree-based calculation derivations, covered below, written explicitly with pure functions to enable compiler optimization.

While the list of available **Derivation** classes is too long to cover here, it is worth noting that they can be combined both additively and multiplicatively through corresponding composite **Derivation** objects. More involved examples include derivations that take moments of distribution function variables, or specialized derivations for polynomial functions of multiple variables (see Appendix A for an example of where this is useful).

Furthermore, variables are by default assumed to live on just the spatial grid (see Section 4.1 for details), but can be specified to be distributions that also have velocity and Legendre harmonic indices, or simply scalar variables.

The implementation of communication across the spatial domain and in distribution function harmonics will be presented in detail in the next section, though it is worth noting here that some **Derivations** might require knowledge of data calculated or evolved by MPI processes other than the local one. An example would be a central difference operator derivation, which will require knowledge of spatial halo values before being able to correctly calculate the difference. Thus, there is danger of out-of-order communication and derivations. This is true in particular with scalar variables, which are always associated with a primary host process, and are broadcast to all others. If a **Derivation** on one process requires a scalar variable living on another

⁶The implementation of the interface is a little different, with **Derivations** taking in data and indices associated with the variables that are required




Figure 2: An example of how ReMKiT1D handles the calculation and communication of derived variables. **Left:** A graph representation of variable dependencies in this example. Here a and b are implicit variables and can be safely communicated first. c is derived using only implicit variables a and b and is thus given depth 0; d is not an implicit variable, but does not have a derivation rule associated with it (perhaps being evolved explicitly), and is also given depth 0. e and f depend on depth 0 and 1 derived variables, and thus have depths 1 and 2, respectively. **Right:** The derivation-safe communication algorithm - derivations are called on each depth only once the depth before it has been communicated, ensuring no out-of-order operations. **Example:** If c requires taking the central difference derivative of a , the halo of variable a must first be communicated.

process, the broadcast **MUST** happen before the derivation call in order for the correct value to be used. In ReMKiT1D, this is ensured by associating a derivation depth with every variable, and always using a communication-safe derivation call, as follows:

- Implicit variables are given a derivation depth of -1; they are always safe to communicate and are the first to be communicated.
- Derived variables that only require implicit variables (or don't have any required variables) are given a depth of 0; they can be calculated once implicit variables have been communicated.
- All other derived variables are given depth equal to $d + 1$, where d is the highest depth of all variables required by a given derived variable.

Thus, variables of depth d are calculated only after variables of depth $d - 1$ have been communicated.

The above algorithm ensures safe calls to all **Derivation** routines, under the condition that there are no cyclical dependencies. This can be represented by a directed acyclic graph, as shown in Figure 2.

Finally, it is worth noting that a variable-like **ModelboundData** object is available, which stores derived variables only. Those variables, unlike in the **VariableContainer** can also represent single harmonics, making them useful for some kinetic algorithms (see Appendix A). However, the above communication-safe derivation call is not available, so care should be taken that any derived variables in such a **ModelboundData** object are added in the correct order. In practice this is rarely an issue, since most derived variables in **ModelboundData** tend to be calculated using variables in the **VariableContainer**, while only a minority require variables that only live in the **ModelboundData**.

3.2.1. Tree-based calculation Derivations

A particularly flexible type of **Derivation** is worth covering in slightly more detail at this point. This is based on an expression tree calculation, where each node (represented in the Fortran code by the **CalculationNode** class) can have a particular set of properties:

- Whether the node is additive or multiplicative with respect to the results of its children
- A single constant to add to/multiply the results of the children, depending on whether the node is additive or multiplicative
- An associated variable name from the **VariableContainer** - only relevant for leaf nodes, where it is treated as the result of the node's non-existent children
- A unary transformation, to be applied to the result of the node

By default, nodes are multiplicative, with a constant of unity, and no unary transformation. Most basic functions, such as *exp* and *log*, are implemented as unary transformations. Unary transformations can also have associated parameters, allowing for added flexibility. An example is raising variables to

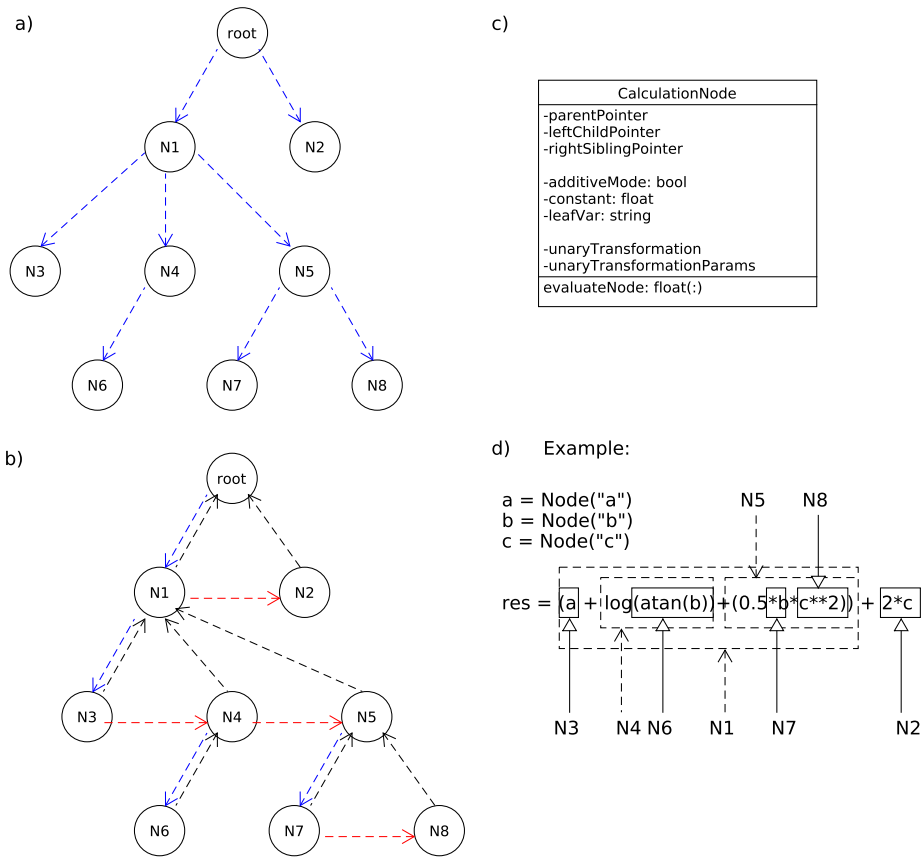


Figure 3: Example of a calculation tree. a) General expression tree diagram we are looking to represent. Each node should add/multiply the results of its children and potentially apply some unary transformation to that result. b) Left-child right-sibling representation of the calculation tree. Blue arrows connect the parents with their leftmost child, red connect nodes with their sibling to the right, and black arrows point back to the parent. c) UML representation of the **CalculationNode** class. See text for more details. d) Example of a Python expression that would generate the example tree for ReMKiT1D to use. Leaf nodes are pointed out using solid line arrows and boxes, and the composite nodes with dashed line arrows/boxes. The root node and nodes N1, N4, and N5 are two additive nodes, a node with an applied function, and a multiplicative node with a constant, respectively. Note that, in this case, all three variables participating in leaf nodes must conform, i.e. be of the same size (e.g. all must be fluid variables or distribution variables, but not a mix).

integer- or real-valued powers, which are transformations where the power is a transformation parameter. Another useful parameterized transform is

the *shift* transform, which cyclically shifts the flattened array representation of a variable/node evaluation result some number of entries. This allows for representation of finite difference/volume operators through appropriate combinations of *shift* transforms. Transforms are also supplied that can be used for the contraction of distribution variables into variables that are only defined on the spatial grid, and vice-versa. In this way, preparation work has been done for general non-matrix **Terms** to be implemented with a simple Python level interface in the future.

The result of a node evaluation is given by

$$R_i = f_i(c_i + \sum_{j \in C_i} R_j), \quad (8)$$

for an additive node and

$$R_i = f_i(c_i \prod_{j \in C_i} R_j), \quad (9)$$

for a multiplicative node, where f_i is the unary transformation, c_i is the constant, and C_i is the list of the children of node i . The expression tree is represented in the Fortran code using a left-child right-sibling tree structure, where each child node also has a reference to its parent. The two above expressions then become

$$R_i = f_i(c_i \cdot_i R_{L_i} \cdot_{P_i} R_{S_i}), \quad (10)$$

with \cdot_i now representing the binary operation (addition or multiplication) associated with node i , L_i the left child of i , S_i the right sibling of i , and P_i the parent node of i . Both representations of the tree structure, as well as the UML element for the **CalculationNode** class, and an example Python expression generating an expression tree are given in Figure 3.

A few notes on the implementation details of this particular type of **Derivation** are in order, as there are a number of (Fortran) peculiarities that need to be addressed. Firstly, copying derived types with pointer components is error prone, as the copy's pointers will not point to the same object as the original's pointers, but to the pointers themselves. This can lead to surprises when objects go out of scope, causing all copies of them to have their pointer components dangling. This can be avoided using Fortran's allocatable components. However, since the unary transformations are procedure pointer components, it is impossible to completely avoid this issue of

dangling pointers. The way ReMKiT1D’s implementation gets around these behaviours is through flattening the expression tree, and unpacking it only once evaluations are required. One can argue that the procedure pointer issue can be avoided by checking association at each *evaluateNode* call, and associating the pointer with the correct procedure if it became un-associated. Unfortunately, this is technically a side-effect, and would disallow the use of pure *evaluateNode* functions, which was one of the aims when designing this particular expression tree representation.

4. Implementation in 1D

A number of implementation details are unique to the 1D case dealt with by ReMKiT1D. These include the default grids in the code, variable communication, as well as some of the finite volume operators, and will be presented in this section. Several implementation details that are not necessarily tied to the 1D approach will also be presented here, such as the handling of collisional-radiative processes.

4.1. Grids

ReMKiT1D deals with three types of variables⁷:

- Scalar variables
- Fluid variables - the default variable type, lives only on the spatial grid
- Distribution variables - live on both the spatial and velocity grids (including Legendre harmonics)

ReMKiT1D contains default implementations of the spatial and velocity grids used to both determine the numbers of degrees of freedom for each variable, as well as construct some default operators. In principle, the user does not have to use the default operators and can construct custom stencils (see below).

⁷Four if single harmonic variables in variable-like model-bound data objects are considered, but these cannot be explicitly included in the main variable container object in v1.0.x.

4.1.1. Spatial and velocity grids

The default spatial grid in ReMKiT1D consists of a 1D array of cells i , which can be thought of as stacked truncated cones, such that their base areas are given by cell face Jacobians J_i , allowing for the representation of flux tubes of varying cross-section in the SOL. Together with the cell widths dx_i , these Jacobians determine cell volumes $V_i = dx_i(J_i + J_{i-1})/2$. Alongside the grid of cells i , a dual/staggered grid of cells $i + 1/2$ is defined. The volumes and face Jacobians of these cells are calculated so that they fit into the regular grid. Thus the right face Jacobian of cell $i + 1/2$ is the cell centre Jacobian/area of cell $i + 1$, $J_{i+1/2} = (J_i + J_{i+1})/2$, and the volume is given by $V_{i+1/2} = dx_i(J_{i-1/2} + J_i)/4 + dx_{i+1}(J_{i+1/2} + J_i)/4$. An exception to this are dual grid cells near boundaries, which can be extended so that the dual grid cell volumes cover the same volume as the regular cells, which is useful when a boundary condition on an outer cell face should be applied to variables defined on the dual grid. A schematic of the default ReMKiT1D spatial cells is shown in Figure 4.




Figure 4: The default ReMKiT1D spatial grids. Black cells are cells on the regular grid, with their cell centres labelled by whole numbers i and marked by black stars, and their right cell face Jacobians given by J_i . The dual/staggered grid points are marked with red stars and the corresponding cells with dashed red lines. Dot-dashed blue lines mark an example of an extended boundary dual cell, in this case assuming that the left face of cell i in the schematic is an outer boundary face.

Fluid variables thus live either on the regular or the dual grid, with the possibility of linear interpolation between them. Distribution variables, on the other hand, either live entirely on the regular grid, or their even l harmonics live on the regular and their odd harmonics on the dual grid, or vice-versa. This is because some of the more notable moments of even harmonics include densities and energies, which naturally live in cell centres/on the regular grid, while the odd harmonics produce moments that determine various fluxes.

When the spatial grid is periodic, the regular and dual grids have the same number of cells. Otherwise the dual grid has one fewer cell, as the outer boundaries of the domain are always assumed to be on regular cell boundaries. In terms of the actual array indexing, a variable that lives on the staggered grid with index i corresponds to the right cell face of regular cell i , or $i + 1/2$ in Figure 4.

The velocity grid is a 1D array of cells with defined cell centres representing discrete velocity magnitudes, and is accompanied by the harmonic index, similar to the implementation in SOL-KiT[26]. The distribution function is assumed to be azimuthally symmetric, so that the harmonic number l is the only one that needs to be tracked⁸. As such, distribution variables are fully determined by a spatial, a harmonic, and a velocity index, making them effectively 1D2V.

The layout of implicit fluid and distribution variables in a flattened array passed to PETSc is covered in Appendix B.

4.1.2. MPI communication

ReMKiT1D utilizes MPI parallelism, with each rank responsible for evolving/calculating its own local variable data in the following way. The spatial domain is simply decomposed, with halo exchange coupling the different spatial partitions. However, in order to obtain speedup when multiple distribution harmonics are included, ReMKiT1D also allows partitioning in the harmonic domain, though in a less efficient manner due to some fundamental constraints. Thus, the domain decomposition when evolving distribution variables is inherently 2D. Each column represent a spatial partition and each row represents a harmonic partition. Hence, spatial(harmonic) information is exchanged between columns(rows). This is shown in Figure 5.

⁸Although support is built in for an eventual inclusion of the m numbers as well

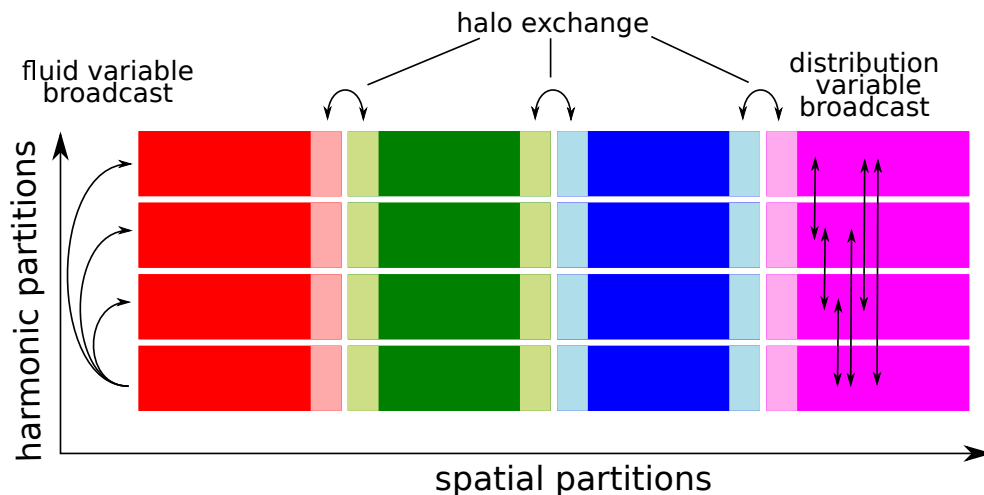


Figure 5: Schematic of a spatial-harmonic domain decomposition. Each cell represents one MPI process, and each column (shown in different colours) corresponds to a single set of spatial coordinates. Within a column fluid variables live in the first process, while harmonics are distributed across the column processes. Three of the four types of communication in ReMKiT1D are shown: halo exchange - exchanging spatial information along rows; fluid variable broadcasts - broadcasting from each column’s root process to other processes in the column; distribution variable broadcast - broadcasting data from all column processes to all other processes in the same column.

The following rules for local variable data are observed:

- Fluid variables live in the first row of each column, where they are evolved/calculated and broadcast to other processes in their respective column. However, derived variables need only be communicated if necessary (as deemed by the user and set by the problem), and are calculated on each process independently following the communication-safe algorithm in Figure 2.
- Distribution variables are spread across all processes, with their harmonics partitioned within processor columns. Due to how some operators might require the entirety of the distribution function at a single spatial location, harmonics within a column are broadcast to all processors in that column, which makes column communication more expensive than the row halo exchange.
- Scalar variables are assigned a host processes, such that they are broadcast from that process to all others. This is only necessary if the scalar

variable’s value is expected to depend on some spatial information. An example is extracting the value of a variable in the last spatial cell and passing it to all other processors.

Based on the above rules for variables, four types of variable communication exist in ReMKiT1D, with three shown in Figure 5. These are

- Halo exchange - where both fluid and distribution variables are exchanged within their respective processor rows
- Fluid variable broadcast - where fluid variables are broadcast from each processor column’s root to the other column processes
- Distribution variable broadcast - where distribution variable harmonics are broadcast from the processes that evolves them to all other processes in the same column
- Scalar variable broadcast - broadcast from single host process to all others (not shown in Figure 5)

Other minor communication routines exist, primarily to check integrator convergence criteria by performing a logical reduction operation over all processes.

The two domain decomposition dimensions are not equivalent, as might be expected from a standard 2D domain decomposition. This is solely due to the fact that we are dealing with variables of varying dimensionality. In particular, harmonic decomposition is more communication-heavy. However, due to the large number of practically dense matrix operators, such as collision operators, speedup can be gained even with this increased cost of communication, as will be shown in Section 6.

4.2. Operators in 1D

In order to illustrate how matrix term stencils work, a number of operators will be reviewed in this section. However, this will not be an exhaustive list, in particular when it comes to various collision operator stencils and kinetic boundary conditions, the implementation of which is adopted from SOL-KiT[26].

As shown in the Python code snippet in Section 3, the simplest (and default stencil) is the diagonal stencil $S_{ij} = \delta_{ij}$. It does, however, allow for the

specification of evolved spatial cells/harmonics/velocity cells, effectively including only the relevant diagonal terms in the matrix. The diagonal stencil also automatically determines whether interpolation/extrapolation is necessary. This is important when staggered grids are used and the evolved and implicit variable are not defined on the same grid. In this case the diagonal stencil will automatically linearly interpolate from the implicit variable's grid onto the evolved variable's grid.

The second most important stencil group are the various spatial difference stencils, starting with the central difference stencils, where both the implicit and the evolved variable must be defined on the same grid (regular or dual), and the implicit variable is interpolated onto the corresponding cell boundaries. An example is the central difference divergence stencil, where i and j here are spatial cell indices

$$S_{ij}u_j = \frac{J_i u_{i+1/2} - J_{i-1} u_{i-1/2}}{V_i}, \quad (11)$$

where $u_{i+1/2}$ is the variable u interpolated on the boundary of cell i . For variables defined on the dual/staggered grid, the interpolation is performed accordingly, and the above expression can be used with $i \rightarrow i + 1/2$. For a gradient operator, the Jacobians are ignored, and the stencil becomes simply

$$S_{ij}u_j = \frac{u_{i+1/2} - u_{i-1/2}}{dx_i}. \quad (12)$$

In the case of differences on a staggered grid, where the implicit and evolved variable live on different grids, the expressions look the same, but now $u_{i+1/2}$ is not interpolated, since it is the actual value of the implicit variable, considering it lives on the boundaries of cell i . In practice, this amounts to forward/backward difference, depending on whether the implicit variable lives on the dual/regular grid, respectively.

Note that the above stencils do not include the contribution from the domain boundaries, which can be dealt with a corresponding boundary condition stencil for both divergence and gradient operators. For a divergence boundary condition, the following form is assumed

$$S_{ij}u_j = \pm \delta_{i i_b} \frac{J_{i_b} F_b u_b}{V_{i_b}}, \quad (13)$$

where i_b is the spatial index of the boundary cell (with J_{i_b} denoting either the left or right face Jacobian in this case), and the sign depends on whether the

boundary is the left(+) or right(-) domain boundary. F_b is the flux Jacobian variable, and both it and u are linearly extrapolated onto the boundary. This form assumes that the flux through the boundary is given as $F_b u_b$, with the flux Jacobian variable living on the regular grid. The boundary condition operator written in this form allows for the application of a lower bound to the flux Jacobian, which is useful, for example, in setting the Bohm condition at the divertor target (see Appendix A). For a gradient operator both the flux and face Jacobians are ignored, and the stencil becomes simply

$$S_{ij}u_j = \pm\delta_{i_b}\frac{u_b}{dx_{i_b}}. \quad (14)$$

Note that the above divergence boundary term is non-linear, due to both F and u being variables. More involved stencils might have derivation rules associated with them, or might require access to model-bound variables. An example is the spatial diffusion stencil, which requires both the evolved and the implicit variable to live on the regular grid, and takes in a derivation rule in order to calculate the diffusion coefficient D

$$S_{ij}u_j = \frac{1}{V_i} \left(J_i D_{i+1/2} \frac{u_{i+1} - u_i}{dx_{i+1/2}} - J_{i-1} D_{i-1/2} \frac{u_i - u_{i-1}}{dx_{i-1/2}} \right). \quad (15)$$

Kinetic/distribution variable stencils tend to be more involved, and will not be covered in detail in this manuscript. An example of this complexity is the velocity space derivative operator, representing the term

$$\frac{\partial f_l}{\partial t} = \frac{\partial}{\partial v} (C(v) f_l), \quad (16)$$

where the function $C(v)$ can be specified in the code as either a fixed velocity space vector, or a model-bound single harmonic variable. Similarly, f_l can be interpolated onto the velocity space cell boundaries, and the linear interpolation coefficients can be specified in a similar way to $C(v)$, defaulting to interpolating directly onto the cell boundaries. An example where both custom interpolation and $C(v)$ are useful is for the Chang-Cooper-Langdon scheme[33] for the Coulomb collision operator (used in the model in Appendix A).

Other kinetic stencils include

- A moment stencil - represents taking the n -th moment of a harmonic $4\pi \int_0^\infty f_l v^{n+2} dv$.

- Velocity space diffusion - with the ability to specify single harmonic model-bound data diffusion coefficients.
- Logical boundary condition - a stencil representing the logical boundary condition, as derived for SOL-KiT[26].
- Spatial difference stencil - $\partial/\partial x$ operator for harmonics, behaving either as a central difference or a staggered (forward/backward) difference stencil, depending on where the individual harmonics are defined.
- Boltzmann collision operator stencil - based on the SOL-KiT Boltzmann collision integral implementation and using the collisional-radiative model-bound data (see next section)
- Other niche stencils - such as a stencil taking the moment of a kinetic term, or stencils designed for particular Coulomb collision operator terms.

A particularly convenient stencil from a user’s perspective is the custom 1D fluid stencil, which gives the high-level user effectively low-level access with a very flexible interface. This stencil is defined as

$$S_{ij} = \sum_k \delta_{i,j+s_k} X_{k,i} v_{k,i} w_{k,i}, \quad (17)$$

where s_k defines the k -th relative stencil entry position. $X_{k,i}$ is the i -th entry of the fixed stencil component for the k -th relative stencil entry, while v_k and w_k represent **VariableContainer** and **ModelboundData** fluid variables that can be included as individual stencil columns. In this way, combined with various **Derivations**, the user can represent most fluid stencils. An example would be a three point stencil, where $s = [-1 \ 0 \ 1]$, meaning that each spatial location requires information from its left neighbour, itself, and its right neighbour. This three point stencil can then be used, for example, to represent the diffusion stencil in equation (15) by setting

$$X_{1,i} = \frac{J_{i-1/2}}{V_i dx_{i-1/2}}, \quad X_{3,i} = \frac{J_{i+1/2}}{V_i dx_{i+1/2}}, \quad X_{2,i} = -X_{1,i} - X_{3,i},$$

where the diffusion coefficient was assumed to be 1, for simplicity⁹. One could

⁹Otherwise there would be a need to define three derived variables to use as v_k or w_k in equation (17), for which one could use calculation trees


then add complexity by accounting for boundary conditions, for example by setting $X_{1,i}$ to 0 at the left boundary and $X_{3,i}$ to 0 at the right boundary.

More details on the available stencils and their options can be found in the User Manual and within the code documentation.

4.3. Collisional-radiative model-bound data

While one could write a collisional-radiative model (in the ODE sense from Section 2) purely using derived variables and diagonal stencils, these models tend to contain many transitions and terms, so the cognitive load on a user implementing them term-by-term would be high. When one factors in the special treatment needed for Boltzmann collisions, it becomes natural to group collisional-radiative data to allow for efficient **Term** generation. The **ModelboundCRMDData** class serves this purpose, unifying **Transition** objects and inelastic collision mapping, and enabling a simple interface for **Term** generation.

It is also worth noting here that species data in ReMKiT1D can be grouped into species objects, each being associated with a species name and integer ID, and containing data on the species charge and mass. Most importantly, it also allows for the association of certain variable names to a species, which further facilitates the automatic generation of **Terms**. The abstract **Transition** class assumes that each species is associated with an integer ID when defining the ingoing and outgoing states. For example, the integer ID 0 is always associated with electrons, and negative IDs are generally used for ion species, while positive IDs tend to denote neutrals species. As such, one could represent the ionization reaction



as a **Transition** from ingoing to outgoing states/species¹⁰ $[0 \ 1] \rightarrow [0 \ 0 \ -1]$, assuming that the hydrogen atom is given the ID 1 and H^+ the ID -1 . Then, depending on the concrete **Transition** class, the following quantities are accessible through the abstract interface:

- The reaction rate as a function of position

¹⁰In the context of the CRM model-bound data these two terms are used interchangeably. It is assumed that each tracked state is represented by a species in the model.

- The momentum loss rate as a function of position - usually available only for a small subset of electron induced transitions
- The energy loss rate as a function of position - generally associated with the energy loss of electrons
- Cross-section - electron impact cross section associated with the transition and used when constructing Boltzmann collision operators (can be a function of position in the general case, see below for detailed balance transition)
- Transition energy - either a fixed value or a value for each spatial cell based on the ratio of energy loss and reaction rates

Different **Transition** classes are available to the user, allowing for control over how the above quantities are calculated and used. The following are available and widely used as of v1.0.x:

- **SimpleTransition** - the simplest possible transition with a fixed transition energy and rate
- **DerivedTransition** - a transition with the reaction rate calculated using a **Derivation**. Uses a fixed transition energy by default, but can also associate **Derivations** with the momentum and energy loss rates
- **FixedECSTransition** - a transition with a fixed transition energy and cross-section (which can be specified for any number of electron distribution harmonics). The rates are calculated using the cross-section, together with inelastic grid mappings (see below)
- **DBTransition** - a transition obtained using detailed balance (see SOL-KiT paper[26]) with another transition which has a cross-section associated with it. The resulting cross-section held by this **Transition** object is thus also a function of position.

In order to use Boltzmann collision operators or any of the transitions with associated cross-sections, inelastic transition grid data must be generated, in the same way as in SOL-KiT, by specifying a set of transition energies. Then transitions such as the **FixedECSTransition** and stencils such as the Boltzmann collision operator stencil can calculate rates and operators that

obey particle and energy conservation. Term generator objects can be created in ReMKiT1D, that can scan the **ModelboundCRMDData** object and generate all terms corresponding to the particle and/or energy loss/gain rates due to collisions, as well as any Boltzmann collision operators. For example, particle sources can be generated using the reaction rate data of each transition, multiplying them with the ingoing state densities corresponding to the transition¹¹ and by the population change due to the transition. In the above electron impact ionization example the population change for electrons and ions is +1, while it is -1 for the atoms. In general, these produce terms of the form

$$\left(\frac{\partial n_b}{\partial t}\right)_T = P_b^T K^T \prod_{b' \in I_T} n_{b'}, \quad (18)$$

where b is the species index (and can be associated with any species in either the ingoing or outgoing state lists), I_T is the ingoing state list for transition T , K^T is the reaction rate of the transition, and P_b^T is the population change of species b in transition T . These sorts of generated terms are diagonal stencil terms using model-bound rate data and are implicit in the final ingoing state density (given as the first associated variable for that species). Other term generators can be found in the code documentation and examples.

5. Python-JSON-Fortran interface and example workflow

ReMKiT1D's core is written in Fortran, and the code is initialized through a JSON configuration file using the json-fortran library[34]. The motivation behind this approach is the combination of human readability and established IO libraries for JSON in both Fortran and Python. This section will cover IO both with JSON and HDF5 files before presenting an example of how a simple advection simulation with two equations can be generated in the Python interface.

5.1. IO with JSON and HDF5 through Python interface

As noted above, ReMKiT1D is fully initializable using just a JSON configuration file. The JSON keys are defined in the Fortran code, and a Python

¹¹Some transition rates that use cross-sections already include one electron density factor, so these are automatically dropped by generators

interface that generates the corresponding JSON entries is provided. The main object in this Python interface is the **RKWrapper**, which is responsible for generating the configuration file and provides a convenient interface for creating ReMKiT1D runs without knowing the appropriate JSON keys. To illustrate the config.json file format, a snippet of the file produced for the advection example to follow is provided below. The snippet contains settings for HDF5 output as well as MPI communication.

```

1 "HDF5": {
2     "filepath": "./RMKOutput/RMK_advection_test
   /",
3     "outputVars": ["n", "n_dual", "T", "T_dual",
   "u", "u_dual", "time", "W"]
4 }
5 "MPI": {
6     "commData": {
7         "haloExchangeVars": ["n", "n_dual", "
   u_dual", "u"],
8         "scalarBroadcastRoots": [],
9         "scalarVarsToBroadcast": [],
10        "varsToBroadcast": []
11    },
12    "numProcsH": 1,
13    "numProcsX": 4,
14    "xHaloWidth": 1
15 }

```

Data output in ReMKiT1D is performed using the HDF5 library, with HDF5 datasets generated for each variable selected for output (see above snippet for example). Each .h5 file produced this way corresponds to one time step, and the Python interface provides routines for reading these files into xarray datasets, allowing for easy data processing.

HDF5 files can also be used to initialize variable values, instead of using the initial conditions in the config.json file. Furthermore, the code can be instructed to have each processor dump a restart HDF5 file, which can then be used to restart runs from defined checkpoints.

5.2. Simple advection workflow example

In order to illustrate the Python-level workflow involved with producing a ReMKiT1D config.json file, a simulation setup solving the following equations will be explored

$$\frac{\partial n}{\partial t} = -\frac{\partial u}{\partial x}, \quad (19)$$

$$m_i \frac{\partial u}{\partial t} = -\frac{\partial(nkT)}{\partial x}, \quad (20)$$

where m_i is here taken to be the hydrogen mass, and the flux is somewhat unconventionally written as u . Before moving on to the details of the workflow example, it is useful to note the default normalization scheme used in ReMKiT1D. While the users are free to set their own normalization constants for each term, the code comes with the following default normalization, used in all pre-built models, borrowing heavily from SOL-KiT[26].

The three independent normalization quantities are the reference density n_0 (in m^{-3}), temperature T_0 (in eV), and ion charge Z_{ref} . These are usually set to 10^{19}m^{-3} , 10eV, and 1, respectively. From these all derived normalization quantities are obtained:

- Velocity (both for the velocity grid and flow speeds) is normalized to $v_{th} = (2eT_0/m_e)^{1/2}$
- Time is normalized to the reference ion-electron collision time $t_0 = v_{th}^3/(\Gamma_{ei}^0 n_0 \ln \Lambda_{ei}(T_0, n_0)/Z_{ref})$, where $\Gamma_{ei}^0 = Z_{ref}^2 e^4/(4\pi m_e^2 \epsilon_0^2)$ and the Coulomb logarithm is taken from the NRL Plasma Formulary[35].
- Length is normalized to $x_0 = v_{th} t_0$
- The distribution function units are in n_0/v_{th}^3
- The electric field is normalized to $m_e v_{th}/(et_0)$
- Transition energies are normalized to T_0
- Heat flux is normalized to $m_e n_0 v_{th}^3/2$
- Cross-sections are normalized to $1/(x_0 n_0)$

In normalized units, the above equations become

$$\frac{\partial n}{\partial t} = -\frac{\partial u}{\partial x}, \quad (21)$$

$$\frac{\partial u}{\partial t} = -\frac{m_e}{2m_i} \frac{\partial(nT)}{\partial x}, \quad (22)$$

which is what will be implemented in the example that follows. The following code snippet initializes the wrapper and sets up IO and MPI settings:

```

1 from RMK_support import RKWrapper, Grid, Node, treeDerivation
2 import RMK_support.simple_containers as sc
3 import RMK_support.IO_support as io
4
5 #initialize wrapper object
6 rk = RKWrapper()
7
8 #set IO paths
9 rk.jsonFilepath = "./config.json" # Default value
10 hdf5Filepath = "./RMKOutput/RMK_advection_test/"
11 rk.setHDF5Path(hdf5Filepath) # The input and output location
    of any HDF5 files used/generated by the code
12
13 #set MPI properties
14 numProcsX = 4 # Number of processes in x direction
15 numProcsH = 1 # Number of processes in harmonic direction
16 haloWidth = 1 # Halo width in cells
17 numProcs = numProcsH * numProcsX
18 rk.setMPIData(numProcsX, numProcsH, haloWidth)

```

The run being generated will be run on 4 MPI processes, all of which are in the spatial direction, as there are no distribution variable to be parallelized in the harmonic direction. The halo width used is the default one, as no operator requires a halo wider than one cell.

The grid is initialized using the following code

```

1 xGridWidths = 0.025*np.ones(512) # widths of each spatial
    cell
2
3 #no velocity grid necessary - default values for the grids
4 vGrid = np.ones(1)
5 lMax = 0
6 gridObj = Grid(xGridWidths, vGrid, lMax,
    interpretXGridAsWidths=True)
7
8 rk.grid = gridObj

```

which initializes a grid of length $L = 12.8x_0$.

Basic variables are added in the following way

```
1 n = 1 + np.exp(-(gridObj.xGrid-np.mean(gridObj.xGrid)**2) #  
  A Gaussian perturbation  
2 T = np.ones(len(gridObj.xGrid)) # Constant temperature  
3  
4 # These will add both the variable 'v' and 'v_dual'  
5 rk.addVarAndDual('n',n,isCommunicated=True)  
6 rk.addVar('T',T,isDerived=True) # isDerived removes the  
  variable from the implicit vector  
7 rk.addVarAndDual('u',isCommunicated=True,primaryOnDualGrid=  
  True) # primaryOnDualGrid denotes that the main variable  
  is u_dual, and u is interpolated  
8 rk.addVar('time',isDerived=True,isScalar=True)
```

After the above code, the wrapper will have the following variables registered:

- 'n' - lives on the regular grid and is an implicit fluid variable (the default) - initialized as $n = 1 + \exp(-(x - L/2)^2)$
- 'n_dual' - lives on the dual/staggered grid and is derived by linearly interpolating 'n'
- 'T' - a derived fluid variable with no derivation rule associated with it, effectively making it constant - initialized to 1
- 'u_dual' - represents the flux, lives on the dual/staggered grid, and is an implicit fluid variable - initialized to 0
- 'u' - lives on the regular grid and is derived by interpolating 'u_dual'
- 'time' - an explicit scalar variable that the code will recognise as the time variable and use it in that way

To demonstrate how **Derivations** are added, the following snippet creates a calculation tree for the normalized total energy W and adds the corresponding **Derivation** to the wrapper

```
1 #add individual variables as nodes  
2 nNode = Node('n')  
3 uNode = Node('u')  
4 TNode = Node('T')  
5
```

```

6 massRatio = 1/1836 # approximate electron-proton mass ratio
7
8 #tree representation of normalized total energy calculation
9 wNode = 1.5*nNode*TNode + uNode**2/(nNode*massRatio) #
    assuming normalization to n_0*e*T_0
10
11 # Registering the derivation in the wrapper with the name "
    wDeriv"
12 rk.addCustomDerivation("wDeriv",treeDerivation(wNode))

```

Then one can add the variable to be calculated with this derivation with

```

1 rk.addVar("W",isDerived=True,derivationRule=sc.derivationRule
    ("wDeriv",['n','u','T']))

```

where the new variable 'W' is associated with the derivation rule "wDeriv" and requires the three variables that act as leaf nodes in the calculation tree. In general, the list of required variables depends on the type of derivation as well as the use case.

At this point we can start adding the **Models** and **Terms**, beginning with the continuity equation and the corresponding flux divergence term

```

1 # declare a new Model object
2 newModel = sc.CustomModel(modelTag="nAdvection")
3
4 #create a new general matrix term
5 divFluxTerm = sc.GeneralMatrixTerm(evolvedVar='n',implicitVar
    ='u_dual',customNormConst=-1.0,stencilData=sc.
    staggeredDivStencil())
6 newModel.addTerm("divFlux",divFluxTerm) # add a new term with
    tag "divFlux"
7 rk.addModel(newModel.dict())

```

where the implicit variable is 'u_dual' since it is the implicit flux variable, and the stencil is staggered as 'n' and 'u_dual' live on different grids. Similarly, the pressure gradient term is added as

```

1 newModel = sc.CustomModel(modelTag='pGrad')
2 #Required variable data for pressure
3 vData = sc.VarData(reqColVars=['T'])
4
5 gradTerm = sc.GeneralMatrixTerm(evolvedVar='u_dual',
    implicitVar='n',customNormConst=-massRatio/2,stencilData=
    sc.staggeredGradStencil(),varData=vData)
6 newModel.addTerm("gradTerm",gradTerm)
7 rk.addModel(newModel.dict())

```

where now the C array in equation (7) is set to the 'T' variable on line 3, making the staggered gradient stencil act on both 'n' and 'T', with the temperature value being lagged in time by the non-linear solver (has no effect in this example since it is a constant). Note that neither of the terms added has a boundary condition term. Since the grid is staggered, the main boundary conditions are on the divergence of the flux, and with no boundary condition term specified, these default to reflective boundary conditions (0 flux on boundaries). The only remaining setup concerns the time integration, starting with the definition of the implicit integrator and its addition to the **CompositeIntegrator** object

```

1 # the implicit BDE integrator that checks convergence based
   on the variables 'n' and 'u_dual'
2 integrator = sc.picardBDEIntegrator(nonlinTol=1e-12,absTol
   =10.0,convergenceVars=['n','u_dual'])
3
4 rk.addIntegrator("BE",integrator)

```

where the absolute solver tolerance is in units of machine precision¹², and the global integrator properties

```

1 # fixed timestep in this example
2 initialTimestep=0.1 # in normalized time units
3 rk.setIntegratorGlobalData(1, # number of allowed implicit
   term groups - grouping everything into one group per model
4                               1, # number of allowed general
   term groups
5                               initialTimestep)

```

where no time step control is specified, making the time step constant - $\Delta t = 0.1$. All **Models** are also set to allow only a single term group, as there are no diagnostic variables that would require evaluating individual terms or term groups, and there is no operator splitting in the integration of individual **Models**. In this simple example, a single integration step in the **CompositeIntegrator** is used, set in the following way

```

1 # a single integration step evolving all models
2 bdeStep = sc.IntegrationStep("BE")
3
4 for tag in rk.modelTags():
5     bdeStep.addModel(tag)

```

¹²More precisely, in the units of the Fortran intrinsic function epsilon() applied to the default ReMKiT1D real variable kind.

```
6
7 rk.addIntegrationStep("StepBDE", bdeStep.dict())
```

Finally the number of time steps is set¹³, as well as how often the code should output variable data

```
1 rk.setFixedNumTimesteps(10000)
2 rk.setFixedStepOutput(200)
```

The configuration file is then written by simply calling

```
1 rk.writeConfigFile()
```

The configuration file is then ready for use. Once the output files are generated they can be loaded into an xarray Dataset object using provided Python routine. The output of this advection test will be analyzed in the next section, and the Jupyter notebook with the test and analysis is available in the ReMKiT1D-Python repository.

6. Verification and benchmarking

In order to build confidence in the implementation of various operators both in the Fortran and the Python interfaces of ReMKiT1D a large number of tests have been performed. Some of those will be presented here, aiming to cover a wide range of use cases.

It should also be noted that the Fortran code for ReMKiT1D comes with its own unit test suite, built using the testing framework pFUnit[36], with those test integrated into the GitHub repository, while the Python package is tested using the pytest package. For more details on these tests see the corresponding repositories.

Beyond benchmarking, a number of parallel performance/scaling tests have been conducted, and those will also be covered in this section.

Table 1 lists all of the tests performed in upcoming sections, as well as the scripts associated with them, which are available either as parts of the Python package's examples or as supplemental materials for this manuscript.

6.1. Verification of fluid operators

A number of fluid operators are implemented in the code, as noted in Section 4. These are primarily divergence and gradient operators used to

¹³It is possible to set different time-stepping modes and output modes such as running until a certain elapsed normalized time is reached

Name and Section	Scripts
Fluid advection test 6.1.1.	ReMKiT1D_advection_test.ipynb
MMS test 6.1.2.	ReMKiT1D_MMS.ipynb
Kinetic advection test 6.2.1.	ReMKiT1D_kin_adv_test.ipynb
Coulomb collision operators 6.2.2.	ReMKiT1D_ee_coll_test.ipynb ReMKiT1D_ee_coll_test.ipynb ReMKiT1D_cold_ion_test.ipynb
Epperlein-Short test 6.2.3.	ReMKiT1D_ES_test.ipynb es_test.py es_verif.ipynb
Collisional-Radiative tests 6.3.	ReMKiT1D_crm_example.ipynb ReMKiT1D_kin_crm_test.ipynb
Epperlein-Short scaling tests 6.4.1.	es_test_weak_scaling.ipynb es_test_strong_scaling.ipynb
SOL-KiT-like scaling tests 6.4.2.	sk_comp_thesis.py sk_comp_thesis_kin.py ReMKiT1D_SK_comp_staggered_kin_thesis.ipynb ReMKiT1D_SK_comp_staggered_thesis.ipynb sk_comp_thesis_strong_fluid.ipynb sk_comp_thesis_strong_kinetic.ipynb

Table 1: Script names associated with each reported benchmarking and scaling test. All .py and those .ipynb files starting with ReMKiT1D are part of the RMK_support Python package’s examples. All other .ipynb files are part of the supplemental material for this manuscript.


represent terms such as advective and pressure gradient terms, as shown in the Python example in the previous section.

6.1.1. Simple advection

Figure 6 shows the result of running the advection test from the previous section for 1000 normalized times and comparing to the analytical result. Agreement is generally good, with the relative error during the simulation shown in Figure 6b. The relative error is defined here as

$$\delta n = \max(|n_{simulation} - n_{analytic}|/n_{analytic}),$$

where the maximum is taken along the spatial domain. Note that the two oscillatory features in the relative error in the figure come from the reflection of the wave at the boundaries.



(a) Comparison of analytic and numerical results at the end of the simulation $t = 1000t_0$

(b) Maximum relative error of numerical result during simulation

Figure 6: Advection test results compared with analytical solution

Note that the above implementation of the advection operators does not include any flux limiting or artificial viscosity. ReMKiT1D v1.0.x does allow for the inclusion of artificial viscosity using the calculation tree approach. Future implementations of non-matrix terms will address more complex flux-limiter schemes used for shock capturing.

6.1.2. MMS test of isothermal 2-fluid model

In order to test a slightly more involved problem, the following equations were implemented:

$$\frac{\partial n_s}{\partial t} + \frac{\partial \Gamma_s}{\partial x} = 0, \quad (23)$$

$$m_s \frac{\partial \Gamma_s}{\partial t} + \frac{\partial}{\partial x} (n_s k T_s + m_s \Gamma_s u_s) - Z_s e n_s E = 0, \quad (24)$$

$$\frac{\partial E}{\partial t} = -\frac{j}{\epsilon_0}, \quad (25)$$

where s is a species index, here either for electrons or deuterium ions, and $j = \sum_s Z_s e \Gamma_s$ is the total current, making the electric field equation, when solved implicitly, act as a current constraint (see SOL-KiT implementation). $\Gamma_s = n_s u_s$ is the particle flux, and Z_s is the species charge. The species temperatures T_s are left as constants. These equations can be suitably normalized to be

$$\frac{\partial n_s}{\partial t} + \frac{\partial \Gamma_s}{\partial x} = 0, \quad (26)$$

$$\frac{\partial \Gamma_s}{\partial t} + \frac{\partial}{\partial x} \left(\frac{m_e}{2m_s} n_s T_s + \Gamma_s u_s \right) - \frac{m_e}{m_s} Z_s n_s E = 0, \quad (27)$$

$$\frac{\partial E}{\partial t} = -t_0^2 \omega_p^2 j, \quad (28)$$

where t_0 is the normalization/electron-ion collision time and ω_p is the plasma oscillation frequency at the normalized density. The normalized temperature T_s is set to $0.5T_0$. In order to test the above equations, the following manufactured solution is used with reflective boundary conditions:

$$n_s = 1 + 0.1 \frac{x-L}{L}, \quad (29)$$

$$u_s = -0.01x \frac{x-L}{L^2}, \quad (30)$$

$$E = -\frac{1}{4} \frac{1}{n_e} \frac{\partial n_e}{\partial x} - \frac{1}{n_e} \frac{\partial}{\partial x} (n u^2), \quad (31)$$

$$(32)$$

where the fact that the temperature is equal to $0.5T_0$ is used explicitly and the electric field is calculated from the electron momentum equation. L is the length of the domain, and x the spatial coordinate, either on the regular or dual grid. Following the Method of Manufactured Solutions (MMS), the above solutions are inserted into the equations and the resulting source terms are added in order to push the solution towards the manufactured one. Note that the electric field equation is unaffected, as the manufactured solution assumes $j = 0$. Finally, the densities n_s are set to live on the regular grid, and the fluxes Γ_s and electric field E are set to live on the dual/staggered grid. The simulation is then run for several (≈ 3) sonic transition times L/c_s . L here is set to 10m and the normalised sound speed is $c_s = \sqrt{m_e T_e / m_i}$.

The errors are calculated as the maximum (within the domain) relative departure of the tracked quantities compared to the initial values at the end of the simulation, and are shown in Figure 7. When the manufactured solution is computed directly, in particular the $\partial(\Gamma_s u_s)/\partial x$ terms, the electric field converges poorly due to discrepancies at the system boundaries, see 7a. This is because the default operators used in this example assume that the boundary cells on the dual grid are extended, as shown in Section 4.1. Once this is taken into account and those gradient terms are modified in the manufactured solution, much better spatial convergence is obtained, see 7b.

6.2. Verification of kinetic operators

A number of kinetic operators are included in ReMKiT1D, with some used to compose more complex operators, such as the Coulomb collision operators (see Appendix A). A number of these operators will be subjected to verification tests in this section.

6.2.1. Kinetic advection

As noted in Section 4, on staggered grids the even harmonics live on the regular grid (cell centres) and the odd distribution harmonics live on the dual/staggered grid, so it is possible to write a simple advection test for the kinetic spatial advection operator that mimics the fluid advection setup by writing the equations for f_0 and f_1 without any fields or collisions

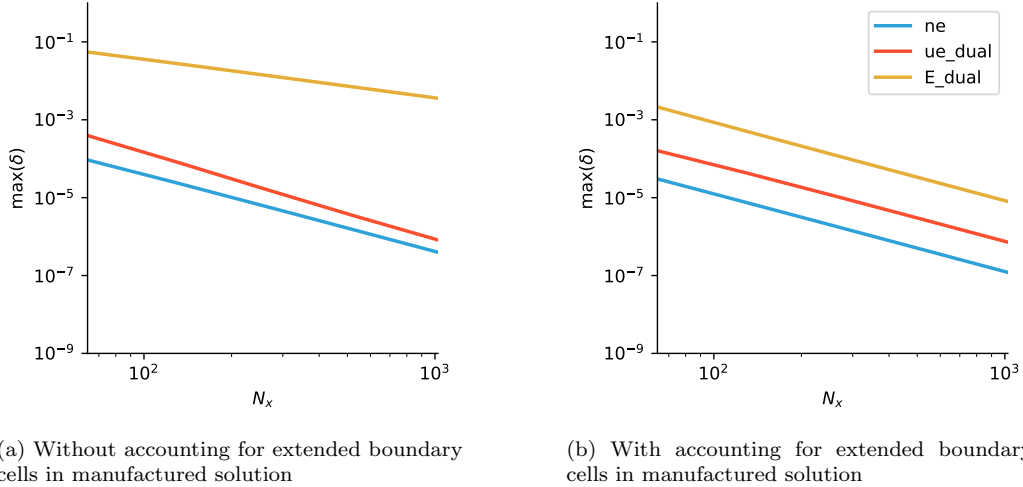


Figure 7: Convergence of MMS test on simple isothermal 2-fluid model with $N_x = 64 - 1024$. The plotted quantity is the maximum relative error of the respective variables compared to the manufactured solution after $t \approx 3L/c_s$

$$\frac{\partial f_0}{\partial t} + \frac{1}{3}v \frac{\partial f_1}{\partial x} = 0, \quad (33)$$

$$\frac{\partial f_1}{\partial t} + v \frac{\partial f_0}{\partial x} = 0, \quad (34)$$

which gives a wave equation for f_0 with wave speed $v/\sqrt{3}$. By initializing f_0 spatially as a Gaussian for all velocities v one can then test the numerical errors for each velocity grid. These are, as expected from keeping the same time and space discretization, worse for larger velocities, as shown in Figure 8. Similar to the fluid advection test the reflections from boundaries are seen as oscillations in the error. While Figure 8 shows a worryingly high error for high velocities, it should be noted that the Gaussian initial condition is the same for all velocities, so the distribution function is unphysically large at high velocities, where it would be orders of magnitude smaller than in the bulk, so in practice this error at high velocities contributes very little to the moments of the distribution function.

For completeness, the grid parameters for this test are as follows. The spatial grid has normalized length $L = 12.8x_0$ with 128 cells, and the simulation is run with time steps of $\Delta t = 0.01t_0$, where we note again that

$v_{th} = x_0/t_0$. In this example $v_{max}dt/dx = 1$, resolving all of the wave speeds in the system, albeit poorly for higher v values, as evident from Figure 8. For more details the reader is directed to the relevant example script.




Figure 8: Absolute error of f_0 harmonic compared to analytical result due to spatial advection.

6.2.2. Coulomb collision operators


Coulomb collision operators are an important part in accurately modelling electron kinetic effects, and the implementation based on ReMKiT1D velocity space operators will be tested in this section. For more details on the functional form and numerical implementation, the reader is directed to previous work in SOL-KiT, as well as Appendix A and the relevant Python routines referenced in the scripts.

The first Coulomb collision operator to be tested is the isotropic electron-electron operator, a non-linear operator acting on f_0 , conserving particles and energy and pushing the distribution towards a Maxwellian. The conservative finite difference implementation of this operator should conserve particles exactly and energy up to non-linear tolerance. To test this, only the f_0 harmonic is evolved, with the following bump-on-tail initial condition (in normalized quantities)

$$f_0(t=0) = \frac{n}{(T\pi)^{3/2}} e^{-v^2/T} + \frac{0.1n}{(T\pi)^{3/2}} e^{-(v-v_{bump})^2/T},$$

where $n = 1n_0$, $T = 0.5T_0$, and $v_{bump} = 3v_{th}$, with standard normalization. This leads to the relaxation of the bump towards a Maxwellian with effective temperature $T \approx 6.07T_0$. This relaxation is plotted in Figure 9 on a log scale, with the x-axis being the energy grid (simply v^2 with the standard normalization). The time step used is $\Delta t = 0.05t_0$, and the simulations were ran up to $t = 60t_0$.

The velocity grid cell widths are given by $\Delta v_1 = 0.01535v_{th}$ and $\Delta v_i = c_v \Delta v_{i-1}$, with a total of 120 cells. The test was performed with two values of c_v , 1.025 and 1.03, which give total velocity grid lengths of approximately $11.27v_{th}$ and $17.25v_{th}$, respectively. As shown in Figure 9b, the shorter grid has a much worse energy conservation, given by the relative error of the effective temperature. This is because the temperature is considerably larger than the normalization temperature T_0 , leading to an under-resolved Maxwellian tail in the latter half of the relaxation on the shorter grid. While this resolution effect is important when there is a substantial evolution in the distribution tail during the simulation, if the solution is already close to equilibrium the error is less pronounced. However, care should always be taken that high energy electrons in kinetic runs are adequately resolved.



(a) Relaxation of bump-on-tail distribution to a Maxwellian

(b) Temperature error during bump-on-tail relaxation for two different grid lengths with the same number of cells

Figure 9: f_0 relaxation test under electron-electron collisions

The second collision operator to test is the electron-ion collision operator

for f_0 , which leads to temperature equilibration with fluid ions. To do this, both the collision operator and a term that takes its energy moment are added to the equations, to which an ion energy equation is now added, containing the moment term. For more information on the electron-ion operator, see Appendix A or references [17, 37]. To test the relaxation in the collisional limit, electron temperature is initialized at $T_e = 8\text{eV}$ and the ion temperature $T_i = 4\text{eV}$, with both densities set to the normalization density $n_0 = 10^{19}\text{m}^{-3}$. In this case, the equilibrium temperature is $T_A = 6\text{eV}$, and we expect both species to relax to that temperature. Following Shkarofsky[37], we define

$$\xi = \frac{n_i(T_e - T_i)}{n_e T_e + n_i T_i}, \quad (35)$$


$$t'_{ei} = \frac{8(n_e + n_i)}{3\sqrt{\pi}} \Gamma_{ei} \frac{m_e}{m_i} \left(\frac{m_e}{2kT_A} \right)^{3/2} t, \quad (36)$$

where Γ_{ei} is the standard Coulomb collision constant. Taking the small mass ratio assumption, the relaxation follows the simple differential equation

$$\frac{\partial \xi}{\partial t'_{ei}} = -\frac{\xi}{(1 + \xi)^{3/2}}, \quad (37)$$

to which an analytical solution is readily obtained. This solution is compared to the numerical result obtained with ReMKiT1D in Figure 10, showing good agreement. The simulation is ran on the short grid from the previous electron-electron collision example and with time steps $\Delta t = 0.1t_0$. As can be seen from the absolute error of the solution in Figure 10b, the final electron temperature is not exactly the same as the ion temperature, with the error well below 1%. This is likely due to finite velocity grid effects.

To test the electron-ion operator for $l = 1$, responsible for momentum exchange, the following setup was used. The electron distribution function was initialized as a Maxwellian at the standard normalization density and temperature ($T_0 = 10\text{eV}$, $n_0 = 10^{19}\text{m}^{-3}$), with the ions initialized at the same density, and flowing at the speed $u_i = 10^{-4}v_{th}$. The only terms solved were the cold ion electron-ion collision operator terms for $l = 1$ (see Appendix A or SOL-KiT), as well as terms in the ion momentum equations that represent the friction moments of the collision operator terms. In the slow ion limit, the distribution function has the following solution for its $l = 1$ harmonic



(a) Comparison of e-i temperature relaxation to analytical result

(b) Over-relaxation in e-i collision test

Figure 10: Temperature equilibration test under electron-ion collisions - due to finite grid effects

$$f_1 = -u_i \frac{\partial f_0}{\partial v}, \quad (38)$$

which is readily computed for a Maxwellian f_0 and is compared to the numerical simulation in Figure 11. The total momentum is conserved up to solver tolerance, and the error in the equilibrium electron flow speed (expecting it to equal u_i) is 0.36% on the same grid as the previous electron-ion $l = 0$ operator test.

Higher harmonic Coulomb collision operators are tested in the Epperlein-Short test to follow.

6.2.3. Epperlein-Short test

In order to fully test the electron-ion collision operators for higher harmonics the well-known Epperlein-Short (ES)[15, 38] test has been conducted. This entails a small electron temperature perturbation decay with electron-electron and electron-ion collisions enabled. Through the decay of the perturbation, the ratio of electron heat conductivity to the classical Braginskii[39] value can be inferred, either through a direct comparison or through examining the decay rate. The initial conditions are set to

$$T = T_0 + T_1 \sin(2\pi x/L),$$




Figure 11: Comparison of numerical result of equilibrium f_1 harmonics with slowly drifting cold ions to analytical result.

where the perturbation wavelength can be controlled by modifying the domain length L . The used grid was periodic and contained $N_x = 128$ spatial cells, $N_v = 120$ velocity cells with widths given by $\Delta v_1 = 0.0307v_{th}$ and $\Delta v_i = 1.025\Delta v_{i-1}$. In this example, the normalization temperature was set to $T_0 = 100\text{eV}$, while other normalization quantities are set to the default values. Density is set to the normalization value, T_1 was set to 0.1eV , and ion charge is left at 1. Following the approach used to benchmark SOL-KiT[26], the ES test was performed for 4 values of l_{max} , and the results are plotted against the same fit based on KIPP[7, 15] data in Figure 12. Here λ_{ei}^B can be

converted to normalized length just as with SOL-KiT - $\lambda_{ei}^B = 3\sqrt{\pi}/(4\sqrt{2})x_0$. Generally good agreement is found with previous SOL-KiT benchmarking, including the agreement with KIPP results, increasing confidence in the default collision operator implementation in the ReMKiT1D framework.




Figure 12: Convergence of ES test with number of harmonics. The fit is based on fit data from KIPP simulations published in [15] and previously used to test SOL-KiT.

6.3. Collisional-Radiative Model tests

In order to test the CRM capabilities a number of tests have been run for both fluid and kinetic electrons. The implementation of inelastic Boltzmann collision operators is adapted from SOL-KiT, and the reader is encouraged to refer to the original SOL-KiT paper for the explanation of the particle and energy conserving discretization. For the tests shown here, the atomic data is in-built Janev[40] hydrogen data¹⁴, together with NIST[41] spontaneous

¹⁴This data is hard-coded in the Fortran source code as an option, but the user is free to define their own data as well.

transmission rates.

Two fluid and one kinetic electron test have been performed. The fluid tests were focused on detailed balance and the Saha-Boltzmann equilibrium, as well as a qualitative comparison with existing kinetic electron simulations performed by Colonna et al[42]. All tests are done in 0D (one spatial cell). The velocity space used to calculate Maxwellian moments of the Janev cross-sections in the fluid case is composed of $N_v = 80$ cells with grid widths going from $\Delta v = 10^{-2}v_{th}$ to v_{th} on a logarithmic grid.

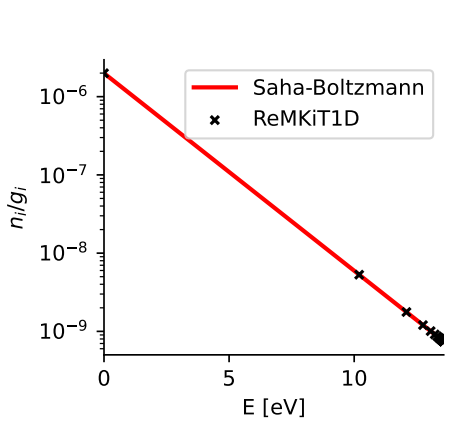
For all tests the following hydrogen reactions have been included:

- Electron impact excitation and de-excitation
- Electron impact ionization
- Three-body recombination


with the evolution test in Figure 13b also including radiative de-excitation and recombination.

For the two fluid tests the number of neutral states tracked was set to 25, including the ground state, while the electron temperature was set to $T_e = 1.72T_0 = 17.2\text{eV}$ (which corresponds to approximately 20000K). Default normalization was used in all tests. For the Saha-Boltzmann test the total density was set to $n = n_0 = 10^{19}\text{m}^{-3}$, with the initial atomic state distribution set to a Saha-Boltzmann distribution at half the electron temperature. The final atomic state distribution after a large number of time steps is shown in Figure 13a, showing equilibration at the expected Saha-Boltzmann distribution with $T = T_e$.

The second test was the qualitative replication of the $t = 10^{-8}\text{s}$ curve in Figure 8 of Colonna et al. For this purpose, the total density is initialized to $n = 733893.9n_0$, loosely corresponding to 1 atmosphere of pressure at 1000K and with an initial ionization degree of 10^{-3} . The electron temperature and initial atomic state distribution are set to the same as in the previous test. Finally, the electron distribution is clamped to its initial value. The result, shown in Figure 13b qualitatively agrees well with the corresponding curve in Figure 8 in the original reference. Finally, in order to test the conservation properties of the ReMKiT1D implementation of the SOL-KiT Boltzmann collision integrals for inelastic electron-neutral collisions, a long simulation was performed with 20 neutral states and with all non-radiative processes included. The velocity grid used had $N_v = 120$ cells with widths given by



(a) Saha-Boltzmann equilibrium with ReMKiT1D.



(b) Evolution of atomic state distribution for fixed initial electron Maxwellian distribution - to be compared with Figure 8 in [42]

Figure 13: Fluid tests of Collisional-Radiative Model with hydrogen. Here n_i/g_i is the population of state with principle quantum number i weighted by the state multiplicity ($g_i \propto i^2$ for hydrogen)

$\Delta v_1 = 0.01v_{th}$ and $\Delta v_i = 1.025\Delta v_{i-1}$. The electron temperature was normalized to $T_0 = 5\text{eV}$ and its initial value was set to T_0 . The initial electron density was set to 10^{19}m^{-3} and the initial neutral ground state density was set to 10^{18}m^{-3} , with no excited state populations. Only inelastic collision integrals were included, allowing us to test the conservation properties isolating only these quantities. In order to isolate discretization errors from time integration errors, a low non-linear iteration relative tolerance of 10^{-14} was used. As shown in Figure 14, both the energy and density relative errors are on the order of the iteration tolerance, even though the simulation was run for a macroscopically significant time and even though the total number of collision operators solved was above 400.

6.4. Parallel performance benchmarking

In order to test parallel performance scaling a number of tests have been performed, mostly focusing on scaling in runs with kinetic electrons, as those are both generally more expensive and conducive to scaling, as well as allowing us to test the scaling behaviour of parallelization in the harmonic direction. All performance scaling tests have been done on the ARCHER2 HPC machine.




Figure 14: Conservation of heavy particles and total energy in kinetic electron CRM simulation with 20 neutral states.

6.4.1. Epperlein-Short test - strong and weak scaling

A variant of the ES test used to verify collision operators has been used to test both strong and weak scaling, as well as investigate basic properties of harmonic parallelization available in ReMKiT1D.

For the strong scaling, the following parameters were used. The maximum harmonic number was set to $l_{max} = 7$, with $N_x = 128$ spatial cells with width $\Delta x = 0.1x_0$, and the simulations were run for N_t time steps with length

$\Delta t = 0.05t_0$. Other parameters were set to the same in the verification test. Due to different behaviour of spatial and harmonic parallelization, strong scaling was tested for 1,2, and 4 processes in the harmonic direction. The results are shown in Figure 15a, where it can be seen that the runs with more harmonic direction processes scale up to a higher number of cores. However, it should be noted here that adding processors in the harmonic direction by default produces a more difficult matrix solve problem due to higher harmonics having shorter timescales. This leads to sub-matrices belonging to processes responsible for high harmonic number naturally being stiffer, leading to the 4 harmonic direction processor run with 256 cores failing due to the solver. While this could potentially be solved by introducing more involved preconditioning, that is beyond the scope of the present manuscript.

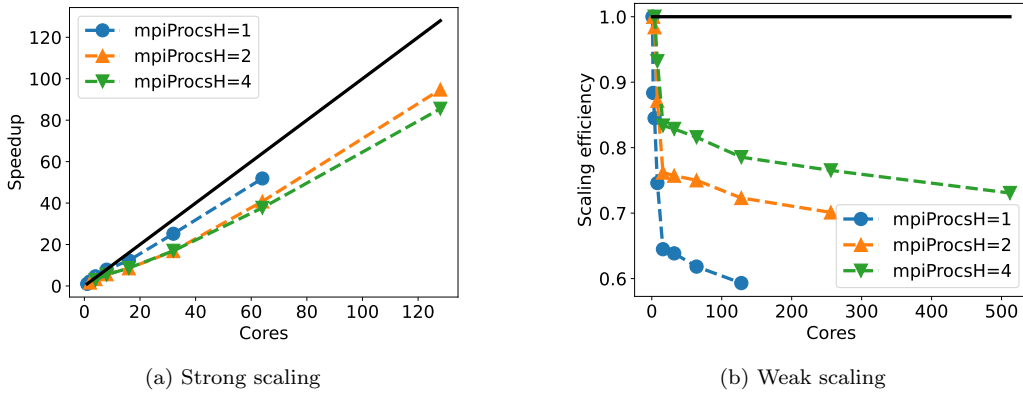


Figure 15: Strong and weak scaling for the ES test. The different colours/markers represent different numbers of processes in the harmonic direction.

For the weak scaling test, the problem was modified in order to avoid limitations due to the matrix solver. The total number of harmonics remains 8, but the length of the domain has been increased to $L = 80x_0$, the time step length reduced to $0.001t_0$ and the number of time steps increased to $N_t = 30000$. This way the main cost in the code was not the matrix solve, but the matrix construction. Scaling has been tested up to 4 ARCHER2 nodes, totalling 512 cores. The number of spatial cells was varied from 8 to 1024 in powers of 2, and the results are shown in Figure 15b for 1,2, and 4 processes in the harmonics direction. What can be seen is that in this example, where the code spends more resources on matrix construction instead of the solve, adding processors in the harmonic direction always improves performance.

The relative speedup from adding harmonics is shown in Figure 16, where it can be seen it is close to the ideal speedup, particularly for higher numbers of processes in the spatial direction.




Figure 16: Relative speedup in the weak scaling set of runs from the ES test when processors are added in the harmonic direction. The black dashed lines are ideal speedups.

6.4.2. SOL models - strong scaling

The final set of scaling tests is focused on testing more production-relevant models of the Scrape-Off Layer. Details of the models are given in Appendix A, and they will only briefly be summarized here. The models are loosely based on equations previously used in SOL-KiT, with the major difference being the use of the AMJUEL[43] database for effective hydrogen ionization and recombination. The electrons are treated either as a fluid or kinetically, and both options are tested here for scaling. The domain is reflective at the left boundary and has a sheath boundary condition at the right boundary, with a total domain length of $L = 10\text{m}$, with the spatial grid being finer closer

to the boundary. An effective heat flux of 3MW/m^2 is injected over $L_h = 3\text{m}$ upstream, while the ion temperature is assumed equal to the electron temperature. Standard normalization is used. The initial condition is based on a Two-Point Model solution[44], with the electron temperature upstream set to $T_u = 20\text{eV}$ and downstream to $T_d = 5\text{eV}$, and the upstream density set to $n_u = 8 \cdot 10^{18}\text{m}^{-3}$. Initial conditions assume no neutral particles, which are injected through the recycling flux at the target. Neutrals are diffusive, while the ion continuity and momentum equations are explicitly solved, including charge-exchange interactions with the neutrals.




Figure 17: Strong scaling with a fluid electron SOL model showing poor scaling due to low number of degrees of freedom per processor

The fluid runs are set up with $N_x = 512$ spatial cells and run until a time $t = 9000t_0$, with time steps adaptively set to approximately 10% of the shortest electron-ion collision time in the domain. Figure 17 shows the result of the scaling test up to 32 cores for this fluid problem. The scaling falls off very quickly due to limitations in the solver, with the simple explanation

being that there are not enough degrees of freedom per core for the matrix solver to scale properly.

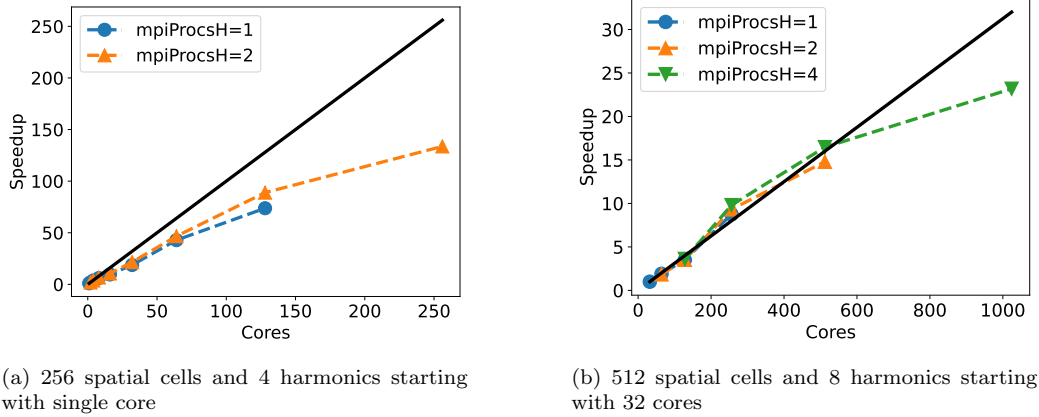


Figure 18: Strong scaling with a kinetic electron SOL model

For the kinetic electron test, two sets of runs were performed, a small scale set with $N_x = 256$ and $l_{max} = 3$ going from a single to 256 MPI processes, and a larger scale set with $N_x = 512$ and $l_{max} = 7$ going from 32 to 1024 MPI processes. Both sets were run up to $t = 50t_0$ with time steps adaptively set to 5% of the shortest electron-ion collision time. $N_v = 80$ velocity cells are used, with cell widths ranging from $0.05v_{th}$ to $0.4v_{th}$. The results of these scaling tests are shown in Figure 18. Unlike the Epperlein-Short test, adding processors in the harmonic direction for both sets of runs improves the scaling. This is likely due to this example having a more involved set of collision operators (due to the flowing ions), which shifts the cost of one solver iteration towards matrix building and away from the actual PETSc solver call. While Figure 18b seems to suggest a better-than-ideal speedup at first glance, this is simply due to the fact that that set of runs did not go down to serial due to the increased computational cost with 8 harmonics and 512 spatial cells, and the scaling seems to improve in the 64-512 core range compared to the 1-64 range.

The results presented above showcase the increased scalability of kinetic models compared to fluid electron models, especially with the novel harmonic dimension domain decomposition. In the example above, moving from fluid to kinetic, the number of degrees of freedom associated with implicit electron variables goes from 5 per spatial cell to 320-640 per spatial cell (4-8 harmonics with 80 velocity space cells).

7. Discussion

In this section the present limitations of the framework as well as potential use cases and future extensions will briefly be discussed. The focus will both be on the software design and numerical aspects, as well as on the achievable modelling with the current version of the software.

The main limitation in ReMKiT1D is its dimensionality, and the 1D aspect is baked into many parts of the framework. Another limitation is the hard-coded assumption that the distribution function is represented in a Legendre/Spherical Harmonic basis. However, basic conceptualization work is planned to explore the applicability of the Modeller-Model-Manipulator (3M) pattern in a way that is agnostic to both numerical methods and problem dimension. This would allow for solving the above two main limitations of the framework.

Even in 1D, the framework's main strength is its flexibility, allowing for rapid iteration on models in the SOL. Examples of planned or ongoing applications include:

- Equilibrium and transient simulations of the SOL, akin to those previously performed using SOL-KiT[18], but with flexible neutral and plasma models, as well as with multiple ion species, focusing on impurity transport and electron kinetic effects
- Exploration of different effective collisional operators that could be used in conjunction with both external and internal Collisional-Radiative Models to include impurity collisional effects on the electron distribution function in runs with kinetic electrons - this includes both simplified and high fidelity molecular deuterium effects
- The calibration of reduced models of SOL equilibria and transients against higher fidelity models both within ReMKiT1D and in other codes
- Training data production for machine-learning applications in the SOL

In order to properly handle some of the above applications, extensions to the framework might be required. Some of the extensions being planned or considered as options are:

- Full support for explicit **Term** objects, as well as improved options for explicit time-stepping
- Adaptive Boltzmann collision operator stencils that can handle reactions with varying energy costs, such as those arising from CRMs
- Full support of Python level custom stencils for kinetic operators, allowing users to generate their own stencils in velocity space (the architecture behind the custom fluid stencil can be adapted for this)
- Multi-linear interpolation support for atomic data for use in the in-built CRM model-bound data

While flexibility and design scalability are the main priorities in ReMKiT1D development, performance represents an increasingly important aspect, particularly since multi-node scalability has been demonstrated for kinetic electron simulations. In order to improve performance, the coupling with PETSc should be explored in more detail, particularly in terms of preconditioners, while PETSc's GPU support might also be an option for more demanding kinetic runs in the future. Other improvements in performance could be bundled with the generalization of the 3M pattern, with a future version of the code using more efficient data structures in terms of memory access.

Finally, the framework's Python interface is under active development with the aim to improve user-friendliness and introduce various quality-of-life features.

8. Conclusion

We have presented the new ReMKiT1D framework for the construction of 1D models of the tokamak Scrape-Off Layer. The framework was designed with the goal of combining flexibility and user-friendliness with the capability of handling electron kinetic effects coupled to multi-fluid models. In order to reach this goal, the main body of the framework is written in Object-Oriented Modern Fortran utilizing convenient abstractions, and coupling to a high level Python interface through the use of JSON and HDF5 files.

The design philosophy of ReMKiT1D is laid out, together with a workflow example demonstrating the use of the Python interface. Various verification tests are presented, together with parallel performance tests, which are focused on the novel parallelization in distribution function harmonics. It is

shown that this approach works and provides a significant improvement in the scalability of the framework when handling kinetic models. As for the verification tests, both individual operator tests as well as standard integrated tests such as the method of manufactured solutions or the Epperlein-Short test have been performed, showing expected agreement of the implemented models.

Finally, ongoing and potential uses as well as future extensions of the framework are discussed in detail, focusing both on the software engineering aspects as well as model development using the framework.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work used the ARCHER2 UK National Supercomputing Service (<https://www.archer2.ac.uk>). This work has been part-funded by the EPSRC Energy Programme [grant number EP/W006839/1]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ukaea.uk.

For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence (where permitted by UKRI, ‘Open Government Licence’ or ‘Creative Commons Attribution No-derivatives (CC BY-ND) licence’ may be stated instead) to any Author Accepted Manuscript version arising.

Appendix A. The SOL-KiT-like SOL models

Some modifications have been performed on the previously reported standard SOL-KiT model[26, 18] for the purpose of testing the implementation in ReMKiT1D, with one major modification being the use of AMJUEL[43] rates instead of the SOL-KiT-style embedded CRM in order to reduce computational costs significantly. The equations for both the fluid models and the kinetic electron model will be presented in this appendix for completeness, while the reader is directed to previous SOL-KiT publications for more details.

Appendix A.1. Fluid equations

A minor difference between the SOL-KiT equations in previous publications and the equations implemented in the ReMKiT1D SOL-KiT-like models is that the fluid equations in ReMKiT1D's implementation are in conservative form, utilizing the capability to implicitly calculate variables with no explicit time derivative in their equations to extract the temperatures and heat fluxes in a way that keeps implicit stability.

The electron fluid equations are given by:

$$\frac{\partial n_e}{\partial t} + \frac{\partial \Gamma_e}{\partial x} = S, \quad (\text{A.1})$$

$$m_e \frac{\partial \Gamma_e}{\partial t} + \frac{\partial}{\partial x} (n_e k T_e + m_e \Gamma_e u_e) + n_e e E = R_{ei}, \quad (\text{A.2})$$

$$\frac{\partial W_e}{\partial t} + \frac{\partial}{\partial x} [(W_e + n_e k T_e) u_e + q_e] + \Gamma_e e E = Q_e, \quad (\text{A.3})$$

where $\Gamma_e = n_e u_e$ and $W_e = 3n_e k T_e / 2 + n_e m_e u_e^2 / 2$. In order to facilitate future inclusions of multiple ion species the parallel transport coefficients for $q_e = \kappa_e \nabla k T_e$ and $R_{ei} = -0.71 n_e \nabla k T_e$ are calculated taking the Braginskii limit[39] ($m_e / m_i \rightarrow 0$ and $Z = 1$) using expressions from Makarov et al[45]. The particle source S results solely from ionization and recombination, and $Q_e = Q_h + Q_{en}$, where Q_h is the upstream heating term, and Q_{en} is the effective electron energy loss/source associated with ionization and recombination collisions. These are implemented using AMJUEL[43] rates H.4-2.1.5 and H.4-2.1.8 for the particle sources, and H.10-2.1.5 and H.10-2.1.8, together with the potential energy accounting for recombination (H.4-2.1.8. with 13.6eV), are used for Q_{en} . A specialized polynomial derivation is implemented in the framework to handle fits such as those in the AMJUEL database.

Note that T_e and q_e are actually treated as implicit variables using ReMKiT1D's capability to include temporally stationary variables in the implicit vector. This ensures stability due to the implicit nature of the scheme is kept, even if the plasma equations are solved in conservative form.

Ion equations are given as

$$\frac{\partial n_i}{\partial t} + \frac{\partial \Gamma_i}{\partial x} = S, \quad (\text{A.4})$$

$$m_i \frac{\partial \Gamma_i}{\partial t} + \frac{\partial}{\partial x} (n_i k T_e + m_e \Gamma_e u_e) - n_i e E = -R_{ei} + R_{CX}, \quad (\text{A.5})$$

where the assumption $T_i = T_e$ is taken as in the standard SOL-KiT model, and $R_{CX} = -\Gamma_i n_n K_{CX}$ is the charge-exchange friction, with K_{CX} using AMJUEL rate H.2-3.1.8, scaling the temperature dependence by 1/2 to account for tracking deuterium instead of hydrogen.

The neutrals are diffusive with their density evolved according to

$$\frac{\partial n_n}{\partial t} = \frac{\partial}{\partial x} \left(D_n \frac{\partial n_n}{\partial x} \right) - S, \quad (\text{A.6})$$

where the diffusion coefficient is set to $D_n = k\sqrt{T_n T_e}/(m_i K_{CX} n_i)$, with the neutrals assumed to have $T_n = 3\text{eV}$, corresponding to the Franck-Condon dissociation energy. The $\sqrt{T_e}$ factor comes from the ion thermal velocity, leading to the neutrals effectively having a diffusive temperature corresponding to a geometric mean between the Franck-Condon and ion temperatures.

Boundary conditions are set to reflective upstream, and sheath boundary conditions at the target, with $u_e = u_i = c_s = \sqrt{2kT_e/m_i}$ set by the Bohm condition and $q_e = \gamma_e \Gamma_e k T_e$, where the electron sheath heat transmission coefficient is approximately 4.86. Ions and electrons leaving the plasma are recombined on the surface and returned with 100% recycling. Finally, the electric field is solved for by implicitly solving

$$\frac{\partial E}{\partial t} = -\frac{j}{\epsilon_0}, \quad (\text{A.7})$$

where $j = e(\Gamma_i - \Gamma_e)$, which, together with the two momentum equations acts like a current constraint and ensures quasi-neutrality.

Appendix A.2. Electron kinetic equations

When electrons are treated kinetically the electron kinetic equations are solved for the evolution of a set number of distribution function harmonics instead of the three electron moment equations in the previous section. The equations have been discussed in detail in the context of SOL-KiT[26, 18], and will not be repeated here for the sake of brevity. A brief description of the terms that remain unchanged from the SOL-KiT implementation will

be given, with the effective cooling operators for use with AMJUEL rates introduced in more detail.

The terms implemented from SOL-KiT are:

- Spatial and velocity space advection (Vlasov) terms
- Coulomb collision terms for electron-electron collisions for all harmonics¹⁵
- Coulomb collision terms for electron-ion collisions for cold flowing ions for harmonics with $l > 0$
- The logical boundary condition[26, 2] at the sheath using the previously developed harmonic formulation
- Diffusive heating terms upstream
- Secondary electron source/sink at low electron energy due to ionization/recombination
- Terms coupling the kinetic operators with the ion fluid equation through taking moments (e.g. R_{ei})

The only missing effect is the electron energy loss/gain terms due to ionization/recombination. These would normally be included as part of the Boltzmann collision operator for electron-neutral collisions, but the implementation used here has instead used effective rates from AMJUEL, so an effective cooling/heating operator needs to be implemented. The simplest implementation is a drag-like operator on f_0 , which, given an inelastic electron-neutral process with energy cost ϵ and rate K , can be written as

$$\left(\frac{\partial f_0}{\partial t}\right)_{\text{inel}} = -\frac{K\epsilon n_n}{m_e v^2} \frac{\partial}{\partial v} (v f_0), \quad (\text{A.8})$$

which can be shown to reproduce the energy loss rate using partial integration on taking the second moment. However, a simple implementation will not preserve this analytical property, so the following velocity space discretization (in standard normalization) is used

¹⁵Here the ability to define single harmonic variables in model-bound data comes in handy to implement Chang-Cooper-Langdon terms[33]

$$\left(\frac{\partial f_0}{\partial t}\right)_{\text{inel}} = -\frac{K\epsilon n_n}{v_k^2} \frac{C_k f_{0,k} - C_{k-1} f_{0,k-1}}{\Delta v_k}, \quad (\text{A.9})$$

where now $C_k = v_k^2 \Delta v_k / (v_{k+1}^2 - v_k^2)$ with boundary conditions $C_0 = 0$ and $C_{N_v} = \Delta v_{N_v}$. It is easy to show that this form gives the correct energy source. However, the number of particles is not necessarily conserved to machine precision due to the right boundary condition on C_{N_v} , which can be ensured by setting $C_{N_v} = 0$, instead moving the error to the energy source. In the tests presented here this was not done since the spurious density source is proportional to the value of the distribution function in the last cell, which is essentially 0. Either way, if the distribution function tail is well resolved this error will never play a role.

Appendix B. Implicit BDE integrator with fixed-point iterations

While the general approach for the Backward Difference Euler (BDE) integrator in ReMKiT1D borrows heavily from SOL-KiT[26], in the interest of clarity, it is useful to have a summary of the method here, as well as how time step length is controlled and how the variable data is stored in the implicit vector passed to PETSc matrix solvers.

Firstly, the variable placement in the implicit variable vector is done in the following way. Given a list of implicit variables v_n , which are either fluid (depend only on the spatial index) or distribution (depend on a spatial, harmonic, and velocity space indices) variables they are flattened and ordered in the implicit vector F as follows

$$\vec{F} = \begin{bmatrix} \vec{F}_1 \\ \vec{F}_2 \\ \vdots \\ \vec{F}_{N_x} \end{bmatrix}, \quad (\text{B.1})$$

where F_i is the sub-vector corresponding to spatial index i and is given by, for example,

$$\vec{F}_i = \begin{bmatrix} v_{1,i} \\ \vec{v}_{2,i}(h,v) \\ \vdots \\ v_{N,i} \end{bmatrix}, \quad (\text{B.2})$$

N is the total number of implicit variables, and where $\vec{v}_{2,i}$ is a distribution variable vector at spatial index i , which is further flattened first in the harmonic index h and the velocity space index v

$$\vec{v}_{2,i}(h,v) = \begin{bmatrix} v_{2,i,1,1} \\ v_{2,i,1,2} \\ \vdots \\ v_{2,i,1,N_v} \\ v_{2,i,2,1} \\ \vdots \\ v_{2,i,N_h,N_v} \end{bmatrix}, \quad (\text{B.3})$$

where N_h and N_v are the total number of distribution function harmonics and the number of velocity space cells.

Collating all individual term matrices, as presented in Section 3, one arrives at the global matrix equation

$$\frac{d\vec{F}}{dt} = M(\vec{F}) \cdot \vec{F}, \quad (\text{B.4})$$

where now $M(\vec{F})$ is in general a non-linear global PETSc matrix. The implicit BDE method with fixed-point iterations then discretizes the solution in time as

$$\frac{\vec{F}^{i+1} - \vec{F}^{i+1}}{\Delta t_i} = M(\vec{F}^{i*}) \cdot \vec{F}^{i+1}, \quad (\text{B.5})$$

where now i^* represents the value at the previous non-linear iteration of the solver. The integrator converges based on a set of convergence variables and a relative or absolute non-linear iteration error¹⁶. Finally, the global (at

¹⁶Usually based on L_2 norm of spatially-local values or largest L_1 norm

the **Composite Integrator** level) time step length Δt_i can be controlled through the application of a **TimestepController**, which scales the time step based on some spatially global criterion. An example is scaling an initial time step based on collisionality by multiplying it by the smallest value of (normalized) $T^{3/2}/n$, making sure that the shortest collisional times in the domain are always resolved. This is used in the SOL-KiT-like models in Appendix A.

Finally, the BDE integrator in ReMKiT1D is capable of recovering from failed matrix solves and cases where non-linear iterations fail by subdividing its time-step into smaller steps when a failure is detected. While crude, this method enables convergence when transient effects might momentarily make the matrix stiffer than expected. This integrator recovery is in addition to any time-step control defined at the global level.

References

- [1] D. Tskhakaya, On Recent Massively Parallelized PIC Simulations of the SOL, *Contributions to Plasma Physics* 52 (2012) 490–499.
- [2] R. J. Procassini, D. A. Knoll, Kinetically motivated boundary conditions for fluid models of scrape-off layer transport, *Journal of Nuclear Materials* 196-198 (1992) 363–368.
- [3] B. D. Dudson, J. Allen, T. Body, B. Chapman, C. Lau, L. Townley, D. Moulton, J. Harrison, B. Lipschultz, The role of particle, energy and momentum losses in 1D simulations of divertor detachment, *Plasma Physics and Controlled Fusion* 61 (2019).
- [4] G. L. Derks, J. P. K. W. Frankemölle, J. T. W. Koenders, M. van Berkel, H. Reimerdes, M. Wensing, E. Westerhof, Benchmark of a self-consistent dynamic 1D divertor model DIV1D using the 2D SOLPS-ITER code, *Plasma Physics and Controlled Fusion* 64 (2022) 125013.
- [5] E. Havlíčková, W. Fundamenski, F. Subba, D. Coster, M. Wischmeier, G. Fishpool, Benchmarking of a 1D scrape-off layer code SOLF1D with SOLPS and its use in modelling long-legged divertors, *Plasma Physics and Controlled Fusion* 55 (2013).
- [6] O. Batishchev, M. M. Shoucri, A. A. Batishcheva, I. P. Shkarofsky, Fully kinetic simulation of coupled plasma and neutral particles in scrape-off

- layer plasmas of fusion devices, *Journal of Plasma Physics* 61 (1999) 347–364.
- [7] A. V. Chankin, D. P. Coster, Benchmarks of KIPP: Vlasov-Fokker-Planck Code for Parallel Plasma Transport in the SOL and Divertor, *Contributions to Plasma Physics* 54 (2014) 493–497. doi:10.1002/ctpp.201410047.
- [8] M. Zhao, A. Chankin, D. Coster, An iterative algorithm of coupling the Kinetic Code for Plasma Periphery (KIPP) with SOLPS, *Computer Physics Communications* 235 (2019) 133–152. URL <https://doi.org/10.1016/j.cpc.2018.09.012>
- [9] K. Kupfer, R. W. Harvey, O. Sauter, M. J. Schaffer, G. M. Staebler, Kinetic modeling of scrape-off layer plasmas, *Physics of Plasmas* 3 (1996) 3644–3652. doi:10.1063/1.871957.
- [10] S. I. Krasheninnikov, A. S. Kukushkin, Physics of ultimate detachment of a tokamak divertor plasma, *Journal of Plasma Physics* 83 (2017) 155830501.
- [11] P. C. Stangeby, Basic physical processes and reduced models for plasma detachment, *Plasma Physics and Controlled Fusion* 60 (2018).
- [12] D. Tskhakaya, F. Subba, X. Bonnin, D. P. Coster, W. Fundamenski, R. A. Pitts, On kinetic effects during parallel transport in the sol, *Contributions to Plasma Physics* 48 (2008) 89–93.
- [13] I. Vasileska, X. Bonnin, L. Kos, Kinetic-fluid coupling simulations of ITER Type I ELM, *Fusion Engineering and Design* 168 (2021) 112407. URL <https://doi.org/10.1016/j.fusengdes.2021.112407>
- [14] W. Fundamenski, Parallel heat flux limits in the tokamak scrape-off layer, *Plasma Physics and Controlled Fusion* 47 (2005).
- [15] J. P. Brodrick, R. J. Kingham, M. M. Marinak, M. V. Patel, A. V. Chankin, J. T. Omotani, M. V. Umansky, D. D. Sorbo, B. Dudson, J. T. Parker, G. D. Kerbel, M. Sherlock, C. P. Ridgers, Testing non-local models of electron thermal conduction for magnetic and inertial confinement fusion applications, *Physics of Plasmas* 24 (2017) 092309.

- [16] A. V. Chankin, G. Corrigan, A. E. Jaervinen, Assessment of the strength of kinetic effects of parallel electron transport in the SOL and divertor of JET high radiative H-mode plasmas using EDGE2D-EIRENE and KIPP codes, *Plasma Physics and Controlled Fusion* 60 (2018).
- [17] D. Power, S. Mijin, F. Militello, R. J. Kingham, Ion–electron energy transfer in kinetic and fluid modelling of the tokamak scrape-off layer, *European Physical Journal Plus* 136 (11 2021).
- [18] S. Mijin, F. Militello, S. Newton, J. Omotani, R. J. Kingham, Kinetic and fluid simulations of parallel electron transport during equilibria and transients in the scrape-off layer, *Plasma Physics and Controlled Fusion* 62 (9) (2020) 095004.
- [19] D. R. Bates, A. E. Kingston, R. W. P. McWhirter, Recombination between electrons and atomic ions, I. Optically thin plasmas, *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 267 (1962) 297–312.
- [20] D. R. Bates, A. E. Kingston, Collisional-radiative recombination at low temperatures and densities, *PROC. PHYS. SOC* 83 (1964).
- [21] K. Sawada, T. Fujimoto, Effective ionization and dissociation rate coefficients of molecular hydrogen in plasma, *Journal of Applied Physics* 78 (1995) 2913–2924.
- [22] H. P. Summers, W. J. Dickson, M. G. O’Mullane, N. R. Badnell, A. D. Whiteford, D. H. Brooks, J. Lang, S. D. Loch, D. C. Griffin, Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional-radiative model for light elements, *Plasma Physics and Controlled Fusion* 48 (2006) 263–293.
- [23] D. Wunderlich, U. Fantz, Evaluation of state-resolved reaction probabilities and their application in population models for He, H, and H₂, *Atoms* 4 (12 2016).
- [24] P. T. Greenland, Collisional – radiative models with molecules, *Proc. R. Soc. Lond. A* 457 (2001) 1821–1839.

- [25] S. Mijin, F. Militello, S. Newton, J. Omotani, R. J. Kingham, Kinetic effects in parallel electron energy transport channels in the scrape-off layer, *Plasma Physics and Controlled Fusion* 62 (10 2020).
- [26] S. Mijin, A. Antony, F. Militello, SOL-KiT-Fully implicit code for kinetic simulation of parallel electron transport in the tokamak Scrape-Off Layer, *Computer Physics Communications* 258 (2021) 107600.
- [27] R. J. Kingham, A. R. Bell, An implicit Vlasov-Fokker-Planck code to model non-local electron transport in 2-D with magnetic fields, Vol. 194, 2004.
- [28] A. R. Bell, A. P. Robinson, M. Sherlock, R. J. Kingham, W. Rozmus, Fast electron transport in laser-produced plasmas and the kalos code for solution of the vlasov-fokker-planck equation, *Plasma Physics and Controlled Fusion* 48 (2006).
- [29] W. A. Hornsby, A. R. Bell, R. J. Kingham, R. O. Dendy, A code to solve the Vlasov-Fokker-Planck equation applied to particle transport in magnetic turbulence, *Plasma Physics and Controlled Fusion* 52 (7) (2010) 75011.
- [30] M. Tzoufras, A. R. Bell, P. A. Norreys, F. S. Tsung, A vlasov-fokker-planck code for high energy density physics, *Journal of Computational Physics* 230 (2011) 6475–6494.
- [31] D. Rouson, J. Xia, X. Xu, *Scientific Software Design: The Object-Oriented Way*, Cambridge University Press, 2011.
- [32] S. Balay, W. Gropp, L. C. McInnes, B. F. Smith, *Petsc, the portable, extensible toolkit for scientific computation*, Argonne National Laboratory 2 (17) (1998).
- [33] E. M. Epperlein, Implicit and conservative difference scheme for the fokker-planck equation (1994).
- [34] J. Williams, JSON-Fortran, <https://github.com/jacobwilliams/json-fortran> Last accessed 5 May 2023.
- [35] J. D. Huba, *Nrl plasma formulary* 19 (2013).

- [36] pFUnit, <https://github.com/Goddard-Fortran-Ecosystem/pFUnit>
Last accessed 5 May 2023.
- [37] I. Shkarofsky, T. Johnston, M. Bachynski, *The Particle Kinetics of the Plasmas*, Addison-Wesley Publishing Company, 1966.
- [38] E. M. Epperlein, Kinetic theory of laser filamentation in plasmas, *Physical Review Letters* 65 (1990) 2145–2148.
- [39] S. I. Braginskii, *Transport processes in a plasma* (1965).
- [40] R. K. Janev, U. Samm, D. Reiter, *Collision processes in low-temperature hydrogen plasmas*, 2003.
- [41] A. Kramida, Y. Ralchenko, J. Reader, N. A. Team, NIST Atomic Spectra Database (ver. 5.9), [Online], <https://physics.nist.gov/asd>
National Institute of Standards and Technology, Gaithersburg, MD.
DOI: <https://doi.org/10.18434/T4W30F>. Last accessed 13 May 2022 (2021).
- [42] G. Colonna, L. D. Pietanza, M. Capitelli, Coupled solution of a time-dependent collisional-radiative model and boltzmann equation for atomic hydrogen plasmas: Possible implications with libs plasmas, *Spectrochimica Acta - Part B Atomic Spectroscopy* 56 (2001) 587–598.
doi:10.1016/S0584-8547(01)00223-3.
- [43] D. Reiter, The data file AMJUEL: Additional Atomic and Molecular Data for EIRENE, <https://www.eirene.de/Documentation/amjuel.pdf> Last accessed 5 May 2023 (2020).
- [44] P. Stangeby, *The Plasma Boundary of Magnetic Fusion Devices*, Vol. 43, 2000.
- [45] S. O. Makarov, D. P. Coster, V. A. Rozhansky, A. A. Stepanenko, V. M. Zhdanov, E. G. Kaveeva, I. Y. Senichenkov, X. Bonnin, Equations and improved coefficients for parallel transport in multicomponent collisional plasmas: Method and application for tokamak modeling, *Physics of Plasmas* 28 (2021).