Excalibur-Neptune report 2047356-TN-04-2

Task 2.1 1D fluid model tests

Ben Dudson, Peter Hill, Ed Higgins, David Dickinson, and Steven Wright

University of York

David Moxey

University of Exeter

June 16, 2021

Contents

1 Executive summary

The deliverable for this task is to define "Test cases for system 2-3", a 1D fluid solver with UQ and realistic boundary conditions. The aims of this is to provide a set of problems relevant to the Neptune use case, which run quickly enough to be used as part of a fast development cycle. These will be used to compare model implementations, and test approaches to Uncertainty Quantification (UQ) and preconditioning of complex nonlinear systems involving two-way coupling be-

tween plasma dynamics, atomic reactions, and plasma-wall interactions.

2 Single species

These tests exercise individual parts of the fluid solver, excluding interactions between fluids (e.g. plasma, neutrals) which complicate the full system of equations.

2.1 Nonlinear heat conduction

Heat diffusion is a standard problem in numerical methods, but heat conduction in plasmas has some features which can introduce challenges:

- The diffusion is strongly anisotropic, many (> 6) orders of magnitude difference between direction parallel and perpendicular to the magnetic field.
- The diffusion is nonlinear: In the collisional (Spitzer-Harm / Braginskii limit), the thermal diffusion coefficient depends on the temperature $T^{5/2}$.

Some useful scenarios which are relevant to the simulation of the tokamak edge and scrape-off layer are:

1. A domain with Dirichlet boundary conditions on both sides, fixing the temperature at both sides. The equation for the heat flux is

$$q = -\kappa_0 T^{5/2} \mathbf{b} \cdot \nabla T = \text{const}$$

   where κ_0 is a constant, calculated by integrating over a perturbed Maxwellian distribution function (see e.g Braginskii). This heat flux is along the magnetic field, which has a unit vector $\mathbf{b} = \mathbf{B}/|B|$. With a heat flux q constant along the magnetic field, the temperature $T(l)$ has an analytic solution

$$T_0^{7/2} - T^{7/2} = \frac{7}{2}\frac{ql}{\kappa_0}$$

where l is the distance along the domain, and T_0 is a constant.

2. A fixed (low) temperature at the "target" end of the domain, and a fixed input power q in the "upstream" boundary. This power input can be implemented either as a Robin boundary condition, setting the gradient of temperature T depending on the temperature at the boundary, or as a zero-gradient (Neumann) boundary condition with a source of power over an extended region.

3. A model which captures some of the effects of impurity and atomic interactions, and allows the heat flux $q\left(l\right)$ to be a function of distance along the domain l, is a variation on the Lengyel model [1], which is widely used in divertor modelling. The main feature of this approach is that by assuming that the radiated power is a function only of temperature, an analytic solution can be obtained. The divergence of the heat flux is related to the radiation loss:

$$\nabla \cdot \mathbf{b}q \simeq \frac{dq}{dl} = -n^2 Q\left(T\right)$$

where n is the plasma density, and the equality is exact in a constant magnetic field. By assuming that the pressure $nT = p$ is constant, density is also a function of temperature only.

$$q = -\kappa_0 T^{5/2} \underbrace{\frac{dT}{dl}}_{\frac{dq}{dl}\frac{dT}{dq}} = \kappa_0 T^{5/2} n^2 Q\left(T\right) \frac{dT}{dq}$$

and so

$$\int_{\text{target}}^{\text{upstream}} q dq = \kappa_0 p^2 \int_{\text{target}}^{\text{upstream}} T^{1/2} Q\left(T\right) dT$$

where the target and upstream are usually taken to be the wall and the tokamak outboard midplane or divertor entrance respectively, but can be any two points along the field line. Labelling "target" with subscript "t", and "upstream" with subscript "u":

$$q_u^2 - q_t^2 = 2\kappa_0 p^2 \int_{T_t}^{T_u} T^{1/2} Q\left(T\right) dT$$

By choosing an suitable cooling curve function $R\left(T\right)$, this equation can be compared to numerical solutions. One possible choice is the approximation

for nitrogen radiation used in Lipschultz 2016[2]:

$$Q = 5.9 \times 10^{-34} \frac{(T - 1\text{eV})^{1/2} (80\text{eV} - T)}{1 + 3.1 \times 10^{-3} (T - 1\text{eV})^2} \text{Wm}^3$$

for $1\text{eV} < T < 80\text{eV}$, and $Q = 0$ outside this range.

Note that the above model assumes a constant magnetic field, so that $\nabla \cdot \mathbf{b} = 0$. For conventional large aspect-ratio tokamaks this can be a reasonable approximation, but for spherical tokamaks (such as MAST-U, STEP) the variation in B has a significant impact. The Lengyel model is only a very rough guide to real experiments, but is useful here as an analytic solution to test against.

Moving beyond 1D into 2D and 3D tokamak geometry test cases: Analytic tokamak equilibria with X-points can be created based on work by Cerfon and Freidberg (a useful tool was created by John Omotani[3]. These equilibria were used in [4] to develop benchmark cases, and test a range of numerical schemes.

2.2 Uniform source, outflow boundary

Now adding fluid equations, for a single species particle density, energy (or pressure or temperature), and momentum along the magnetic field.

A 1D domain, with two boundaries:

- No-flow upstream. This can be implemented as a symmetry boundary: Zero-gradient density, pressure and temperature, and zero-value flow velocity and heat flux boundary conditions.
- Free outflow. This can be implemented in finite difference/finite volume methods by extrapolating all quantities with e.g. constant gradient. To preserve positive definite density and temperatures, it can be useful to extrapolate the logarithms of these thermodynamic quantities.

Inside the domain a uniform source of particles and (internal) energy, which then flow towards the free outflow boundary.

This system should evolve to a steady state in which the particle flux increases linearly with distance from the no-flow boundary.

2.3 Half source, outflow boundary

This is a variation on the above test, which tests the ability of the numerical scheme to resolve abrupt changes in sources, and the resulting abrupt changes in gradients.

The boundary conditions are the same as for the previous test, but the sources only fill half of the domain closest to the no-flow boundary. In this case the particle source is uniform in the first half of the domain, so the particle flux starts from zero at the no-flow boundary, and increases linearly with distance until reaching the end of the source region, and is then constant between the end of the source region and the free outflow boundary.

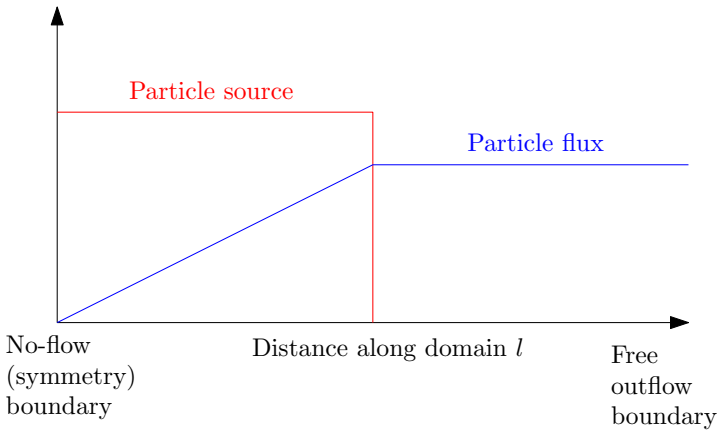

Figure 1: Sources and particle flux in "Half source, outflow boundary" case

2.4 Half source, sheath boundary

This has the same sources as the previous case, but now tests the implementation of a sheath boundary condition. Two important components of this boundary condition in a fluid model are:

1. Imposed out-flow velocity at the sound speed, generally called the Bohm condition, or Bohm-Chodura in a magnetised plasma. The flow speed into the sheath depends in general factors such as the currents and voltages, whether species should be considered isothermal or adiabatic, and can be quite complex in situations with multiple species (see e.g. [5]). A simple and widely used approximation is the isothermal sound speed (see e.g. Stangeby):

$$c_s = \sqrt{\frac{e\left(T_e + T_i\right)}{m_i}}$$

   where the electron and ion temperatures T_e and T_i are in units of eV, m_i is the ion mass (in kg), and the sound speed c_s is in units of m/s.

2. Heat conduction to the target. In addition to the energy loss expected because the fluid is flowing to the target, there is also typically an additional loss of energy. Physically what is happening at the target is that slow electrons are reflected, while energetic electrons reach the target, so that the electrons are rapidly cooled. Some of this energy is transferred to accelerating ions into the target, via the sheath voltage. The energy flow through the sheath is often characterised by the energy flux:

$$q_{\text{sheath}} = \gamma_{sh} enT c_s$$

   with a different γ_{sh} for each species. The actual power flow to the target is complex, depending like the flow speed on the electric fields and currents, and the sheath heat loss coefficientγ_{sh} can vary considerably (by an order of magnitude), particularly during transients.

If solving a single fluid, i.e a single density, energy and momentum equation, a common assumption is that the ions and electrons have equal densities ($n = n_e = n_i$), temperatures ($T = T_e = T_i$) and velocities ($v = v_e = v_i$). This means that the total pressure $p = enT_e + enT_i = 2enT$. The sheath velocity is then

$$c_s = \sqrt{\frac{2eT}{m_i}}$$

In a fluid model with ratio of specific heats $\gamma = 5/3$ (typical for plasma simula-

tions), the energy flow is

$$q_{fluid} = \left(\frac{5}{2}p + \frac{1}{2}m_i v^2\right) v + q_{cond}$$

where q_{cond} is the heat conduction. At the sheath where $v = c_s$ we have:

$$\begin{aligned} q_{fluid} &= \left(\frac{5}{2}2enT + \frac{1}{2}m_i\frac{2eT}{m_i}\right) c_s + q_{cond} \\ &= 6enTc_s + q_{cond} \end{aligned}$$

so by comparing q_{fluid} with q_{sheath}, if $\gamma_{sh} = 6$ then there should be no heat conduction through the sheath. More typical values are γ_{sh} between 6.5 and 9. The additional heat flux can be used to calculate the temperature gradient at the sheath in a Robin-type boundary, or removed from the last cell through the boundary with the target.

This more complex system still has analytic solutions, depending only on the given input power, and the upstream density.

3 Recycling

The half source, sheath boundary test case is a reasonable model of a low density plasma (what Stangeby describes as a "simple SOL"), but is missing a crucial ingredient for tokamak edge and divertor simulation: the interaction with neutral gas.

When plasma meets a material surface a sheath is formed, a small region with strong electric fields, which accelerates ions towards the surface. When ions hit the surface they tend to pick up electrons and become neutral atoms. Some of those atoms will reflect from the surface, others will stick to the surface and perhaps combine into molecules, or become embedded inside the surface before diffusing out again. The details depends on what the wall is made of (e.g. carbon or tungsten), but the key feature is that the majority of the impinging ions ($> 99\%$) will typically come back into the plasma as neutral atoms or molecules.

The atoms and molecules which come off the material surface back into the

plasma typically encounter high electron temperatures, and are quickly ionised and converted into ions again. These ions then return to the wall, where they again become neutral atoms. The majority of the ions flowing to the wall are typically undergoing this *recycling* process, so that the flow of ions from upstream (regions away from the wall and divertor) are typically small relative to the flow of ions to the surface.

This recycling process of continually ionising atoms removes energy from the electrons, both in overcoming the ionisation potential, and also in radiation from the relaxation of intermediate excited states of the atoms. This radiation can also be critical to reducing heat flows to the surfaces, and is enhanced by a feedback mechanism: As the plasma (electrons) cool, the radiated energy per ionisation rapidly increases at low temperatures (< 5eV). This can lead to the plasma "detaching" from the wall, reducing the heat and particle fluxes, in a process similar to the condensation instability seen in space plasmas.

3.1 Recycling sources

A straightforward modification to the 1D single fluid model used in section 2, is to add sources to represent the recycled flux of particles.

- The power source is kept away from the "target", to represent the flow of heat from the main plasma.
- Rather than the particle source being at the same location as the power source, the particle source is now put close to the target.

A reasonable choice of particle source is an exponential decay, with highest source at the target. This is to represent the mean free path of neutral atoms entering the plasma. It's important to note that the decay length for this source should be significantly ($10 - 100$ times) longer in the direction along the magnetic field than the ionisation mean free path. This is because neutral atoms are travelling away from the wall, but the magnetic field is at a shallow angle (typically a few degrees) to the wall. Neutral atoms may only travel a short distance from the wall, but this can correspond to a relatively long distance along the magnetic field in the 1D domain simulated here.

To represent the loss of energy from ionisation and the associated excitation radiation, a sink of energy proportional to the particle source can be added. The ionisation potential for hydrogen atoms is 13.6eV, and dissociation potential for molecules about 4.5eV per molecule, but the energy lost per ionisation varies significantly: Neutral species can enter the plasma with some energy (e.g Franck-Condon energy, around 3.5eV but with wide variation), lowering the effective energy cost; as discussed above, excitation radiation increases the effective energy loss per ionisation. Typical values used in the literature are around 30eV, which is probably a reasonable value to use for this test case.

With this fixed particle source (and so fixed power sink), some care should be taken that the temperature is not driven to zero by removing more power than is put in. This is quite straightforward if a fixed energy cost per ionisation is used, as then the total power sink can be balanced against the power input. If a temperature-dependent ionisation is used, however, then this can be more difficult.

3.2 Fluid neutral gas species

A more sophisticated neutral gas model represents the interaction of plasma and gas as two fluid species, one for the plasma and one for the neutral atoms. Further fluid species can also be added to represent molecules, short-lived species (eg H^-, H_3^+) or a number of vibrationally excited states. The network of reactions between these species can become quite complicated, even for something as "simple" as hydrogen.

Evolving neutral gas as a species modifies the reactions: The ionisation source of plasma particles is now evolving with the system state, rather than being imposed as an input parameter. The volumetric source S, in units of particles per second per cubic meter is

$$S = n_e n_a \left\langle \sigma v \right\rangle_{iz}$$

where n_e is the number density of electrons, n_a the number density of neutral atoms, and $\left\langle \sigma v \right\rangle_{iz}$ is the reaction rate in units of $\mathrm{m^3 s^{-1}}$. This rate is averaged over a Maxwellian distribution, and is typically derived from a 0-D collisional-radiative model, as an effective rate which averages over a number of

processes. These are tabulated in several databases; a commonly used one in plasma physics is Amjuel [6], used in the EIRENE monte-carlo code[1]. Reaction 2.1.5FJ in Amjuel describes ionisation of neutral atoms by electrons:

$$e + H(1s) \rightarrow e + H^+ + e$$

In addition to the ionisation, other reactions now need to be included. Most important is charge exchange (CX) between plasma ions and neutrals. In this reaction, the electron in a neutral atom swaps to the ion, represented as:

$$p + H(1s) \rightarrow p + H(1s)$$

(Amjuel reaction 0.1T). The left and right side of this equation are the same because the reaction swaps an ion and neutral atom but doesn't affect the number of particles of each species. This reaction often has a higher likelihood than ionisation, and provides a strong coupling between the plasma and neutral species. There is no net source or sink of particles for the plasma, but there is a source or sink of energy (equation 1) and momentum (equation 2), depending on the relative temperatures and flow velocities of the plasma ions (subscript 'i') and neutral atoms (subscript 'a'):

$$S_{\text{energy}} = \frac{3}{2} e \left(T_a - T_i\right) n_i n_a \left\langle \sigma v \right\rangle_{cx} \quad (1)$$

$$S_{\text{momentum}} = m_i \left(v_a - v_i\right) n_i n_a \left\langle \sigma v \right\rangle_{cx} \quad (2)$$

Once this reaction is included, the strong interaction between the plasma and neutral atom flows will tend to force the atoms into a very narrow layer close to the target, which can be challenging to resolve. Cross-field diffusion of neutral atoms (which can freely move across the magnetic field, unlike the plasma ions), together with refinement of the grid resolution near the target, are typically needed to widen and resolve this layer.

The handling of charge exchange reactions in a fluid model is problematic: Charge exchange produces a population of fast-moving neutral atoms with a quite different energy and momentum distribution to the atoms from the wall. Unless these atoms are strongly coupled to each other, which is rarely the case

[1]The EIRENE Fortran code reads the data tables from the Amjuel LaTeX source file.

in tokamaks except in specific high density locations, then a single fluid approximation is unlikely to be good. Occasionally in the literature another fluid species is used to represent the charge exchanged neutrals. In 1D there is also the challenge of how the transport of neutrals, particularly charge-exchanged fast neutrals, across the plasma should be represented: The plasma is not really a 1D tube, but a relatively thin sheet, from which neutrals can escape, carrying momentum and energy with them. Representing this, and the resulting interactions with in-vessel components like walls, baffles and pumps, requires at least a 2D representation.

3.3 Kinetic neutral gas species

In most of the tokamak the ratio of the mean free path to the local length-scales of density, temperature etc (the Knudsen number) is much larger than 1, putting the transport in the regime of rarefied gas dynamics, and implying that a kinetic (not fluid) treatment is needed. Unfortunately in other regions, typically near the target, the Knudsen number can be much less than 1, so that kinetic models become highly inefficient (because they're simulating a fluid situation). There are several kinetic plasma neutral models, EIRENE probably being the most well known, and several efforts ongoing to develop hybrid fluid-kinetic models.

Coupling a 1D plasma model to EIRENE would be a non-trivial task, not least because EIRENE is not publicly available or open source in any meaningful way[2]. It is likely that examining the SOLPS code, and benchmarking against SOLPS, would be the most direct way to achieve this.

The main applications of such a 1D EIRENE coupling would be to compare against the fluid model; to use it as a platform for studying model order reduction of the neutral model to simpler and faster models; to test preconditioning strategies; uncertainty quantification involving coupled fluid-monte-carlo coupled algorithms; and to gain experience of carrying out the coupling, to apply to more complex 2D and 3D models. It is likely however that there are quicker, more direct ways of achieving these goals, such as using SOLPS directly.

[2] Access to EIRENE is in principle available from Juelich on request

4 Multiple species

Real plasmas of interest are not pure hydrogen, but contain a mixture of different species. A reactor will have deuterium and tritium isotopes, helium, wall materials such as tungsten and beryllium, seeded impurities such as argon, and trace impurities of e.g. oxygen. Not all of these need to be simulated in all cases, but the capability to model multiple species will become increasingly essential.

4.1 Separate electron and ion temperatures

The most important, but also probably simplest, variation on the models described in previous sections, is to separate the ion and electron temperatures. Close to the target, the low temperatures and high densities often lead to a strong coupling between electron and ion temperatures, but upstream this coupling becomes weaker: Typical measurements on present-day tokamaks find ion temperatures around a factor of two higher than electron temperatures around the outboard midplane.

Separating out the ion and electron temperatures involves:

- Evolving separate temperature or (internal) energy equations for the ions and electrons. The density and velocity equations are still the same, due to quasineutrality of the plasma and the absence of net currents in 1D simulations.
- Using different heat conduction coefficients for electrons and ions
- Separate sheath heat transmission coefficients (γ_{sh}, section 2.4)
- Carefully tracking the contribution of reactions to the energy balance of each species: Ionisation, for example, is an energy loss for electrons, but an energy gain for the ions, as energy is transferred from atoms into ions.
- Adding a coupling between electrons and ions (see Braginskii), due to collisions between them.

4.2 Neutral gas model

As mentioned in section 3.2, the plasma chemistry of even pure hydrogen plasmas can be quite complicated: There are a number of different pathways by which a ionisation (for example) can occur, some involving catalytic interactions between molecules, atoms and plasma ions. Atomic reaction rates are sensitive to the starting state of the molecule (e.g. high vibrational states are more likely to dissociate) or atom (highly excited states are more likely to ionise). In some cases it is likely that these states can persist for long enough (metastable states) that they can be transported around the domain and should be tracked as separate species. It is currently not well understood in the plasma community how complex the model needs to be, and what errors are made when simplified models are used.

4.3 Impurities

Finally, more complex models can be built to study the transport of multiple species, and their multiple charge states. For low-Z species each charge state is typically evolved as a separate species, with the reactions between them (typically ionisation, recombination, and charge exchange with hydrogen); higher Z materials like tungsten often need a charge-state "bundling" treatment, where a range of charge states is evolved as a single fluid, with effective reaction rates between bundled states. The model reduction needed to do this accurately and efficiently is an active area of research.

Once multiple species are included, particularly where their masses are comparable to each other, and where their concentrations are not "trace" level, the calculation of collisions between species becomes complicated. Hirshman and Sigmar published models; Zhdanov is a well known model; and there has been some work recently on improving these models and implementing them in simulation codes.

Unfortunately once models become this complex, analytic test cases can no longer be found. Instead the best option is probably a Method of Manufactured Solutions (MMS) test.

5 References

[1] L L Lengyel. Analysis of Radiating Plasma Boundary Layers (IPP 1/191). http://hdl.handle.net/11858/00-001M-0000-0027-6A04-1.

[2] Bruce Lipschultz, Felix I. Parra, and Ian H. Hutchinson. Sensitivity of detachment extent to magnetic configuration and external parameters. *Nucl. Fusion*, 56:056007, 2016, doi:10.1088/0029-5515/56/5/056007.

[3] John Omotani. Cerfon-freidberg geometry generator in python. https://github.com/johnomotani/CerfonFreidbergGeometry.

[4] M.Held, M.Wiesenberger, and A.Stegmeir. Three discontinuous Galerkin schemes for the anisotropic heat conduction equation on non-aligned grids. *Comp. Phys. Comm.*, 199:29–39, 2016, doi:10.1016/j.cpc.2015.10.009.

[5] D.Tskhakay and S.Kuhn. Boundary conditions for the multi-ion magnetized plasma-wall transition. *J. Nucl. Mat.*, 337–339:405–409, 2005, doi:10.1016/j.jnucmat.2004.10.073.

[6] D. Reiter. The data file AMJUEL:Additional Atomic and Molecular Data for EIRENE. http://www.eirene.de/html/amjuel.html, 2020.