

Bounce-averaged theory in arbitrary multi-well plasmas: solution domains and the graph structure of their connections

Ian E. Ochs 

Department of Astrophysical Sciences, Princeton University, Princeton, NJ, USA

Corresponding author: Ian Ochs, iochs@princeton.edu

(Received 29 May 2025; revision received 1 July 2025; accepted 2 July 2025)

Bounce-averaged theories provide a framework for simulating relatively slow processes, such as collisional transport and quasilinear diffusion, by averaging these processes over the fast periodic motions of a particle on a closed orbit. This procedure dramatically increases the characteristic time scale and reduces the dimensionality of the modelled system. The natural coordinates for such calculations are the constants of motion (COM) of the fast particle motion, which by definition do not change during an orbit. However, for sufficiently complicated fields – particularly in the presence of local maxima of the electric potential and magnetic field – the COM are not sufficient to specify the particle trajectory. In such cases, multiple domains in COM space must be used to solve the problem, with boundary conditions enforced between the domains to ensure continuity and particle conservation. Previously, these domains have been imposed by hand, or by recognising local maxima in the fields, limiting the flexibility of bounce-averaged simulations. Here, we present a general set of conditions for identifying consistent domains and the boundary condition connections between the domains, allowing the application of bounce-averaged theories in arbitrarily complicated and dynamically evolving electromagnetic field geometries. We also show how the connections between the domains can be represented by a directed graph, which can help to succinctly represent the trajectory bifurcation structure.

Keywords: plasma simulation, plasma confinement, fusion plasma

1. Introduction

In a magnetic confinement device, such as those used for fusion reactors, particles typically experience fast helical motion along magnetic field lines. These motions then form closed periodic orbits, sometimes with precession around an angular coordinate. For modelling many macroscopic phenomena that occur on relatively slow time scales compared with the orbit time scale, it is much more efficient to average the effect of the process over the orbit, rather than to solve the full equations of motion for the orbit, as this involves a much slower time scale calculation in fewer dimensions. This simplification is the basis for bounce-averaged Fokker–Planck

(BAFP) theory, which can be used to model collisional transport (Rosenbluth, MacDonald & Judd 1957; Hager *et al.* 2016), quasilinear wave-particle interactions (Stix 1975; Bernstein & Baxter 1981; Fisch 1987; Fisch & Rax 1992; Eriksson & Helander 1994; Herrmann & Fisch 1997; Herrmann 1998), and radiation emission (Bilbao & Silva 2023; Zhdankin, Kunz & Uzdensky 2023) and absorption (Ochs, Mlodik & Fisch 2024; Ochs 2024) in tokamak (Harvey & McCoy 1992), stellarator (Mynick & Hitchon 1986; Nemov *et al.* 1999; Velasco *et al.* 2020; d’Herbement *et al.* 2022) and mirror (BenDaniel & Allis 1962; Marx 1970; Matsuda & Stewart 1986; Egedal *et al.* 2022; Frank *et al.* 2024) plasmas.

Each closed orbit is associated with conserved constants of motion (COMs) in a lower dimensional space. For instance, the energy ϵ , magnetic moment μ and azimuthal momentum p_θ form the typical COMs for a steady-state axisymmetric magnetic field arrangement such as a tokamak or magnetic mirror. Since they are constant on each averaged trajectory, the COMs represent a natural coordinate system for solving BAFP problems. Thus, the first step is often to express the distribution function as a function only of the constants of motion $\mathbf{Z}$, averaging over trajectories that differ only by the phase of the motion.

This procedure often works well for simple field arrangements without local maxima in the fields. However, it runs into a problem if there are local maxima. Consider, for instance, particles in the one-dimensional (1-D) double well potential $\psi(x)$ shown in figure 1. This 1-D arrangement (with its associated two-dimensional (2-D) phase space) has the constant of motion $\epsilon = mv^2/2 + \psi(x)$. For particles with energy $\epsilon > \psi_0$, the value of ϵ defines a single trajectory shape that traverses both positive and negative x , and there is no problem defining $f(x, v) \rightarrow f(\epsilon)$ to average slow processes over the bounce motion. However, for $\epsilon < \psi_0$, the value of ϵ no longer contains complete information about the trajectory, which bifurcates into two trajectories: one trapped in the well at $x > 0$ and the other trapped in the well at $x < 0$. Since the occupancy of each well can, in general, be different, the function $f(\epsilon)$ is no longer well defined.

Multiple wells, however, do not mean that BAFP cannot be used to model multi-well configurations. The resolution lies in splitting the space into three domains, representing: (i) the passing trajectories at $\epsilon > \psi_0$; (ii) the trajectories trapped in the left well; and (iii) the trajectories trapped in the right well. Boundary conditions (Matsuda & Stewart 1986; d’Herbement *et al.* 2022) then must be enforced between the domains to ensure continuity of the distribution function and particle conservation.

These bifurcations of trajectories are common in mirror physics and the solution of their associated BAFP problems has appeared several times in the study of tandem mirrors (Cohen *et al.* 1980; Matsuda & Stewart 1986; Katanuma *et al.* 1986, 1987; Fowler, Moir & Simonen 2017), which are characterised by internal peaks in both the magnetic field and electric potential around the tandem end plugs. However, in solving these problems, the boundary conditions were typically imposed by hand for very specifically shaped fields, and were only ever attempted in two dimensions. In moving to three-dimensional (3-D) computational modelling of modern mirrors, with fields that can be significantly more complicated due to sloshing ion distributions (Egedal *et al.* 2022; Endrizzi *et al.* 2023; Frank *et al.* 2024), multi-mirror configurations (Be’ery *et al.* 2018; Miller, Be’ery & Barth 2021), ponderomotive plugs (Miller *et al.* 2023; Rubin, Rax & Fisch 2023; Ochs & Fisch 2023; Rubin, Ochs & Fisch 2024; Kolmes & Fisch 2024; Rubin & Fisch 2025; Ochs, Kolmes & Fisch 2025) or centrifugal forces (Bekhtenev *et al.* 1980; Cho *et al.* 2005; Schwartz *et al.* 2024; Kolmes, Ochs & Fisch 2025), determining the domains by hand at the beginning of




FIGURE 1. A one-dimension double well scalar potential $\psi(x)$. The energy ϵ uniquely defines a single trajectory for $\epsilon > \psi_0$. However, at $\epsilon < \psi_0$, there are two trajectories that have the same ϵ , corresponding to trapping in the two wells. Thus, the function $f(\epsilon)$ is not necessarily well defined below $\epsilon = \psi_0$.

the simulation may be difficult or impossible. Furthermore, while modern codes for neoclassical transport in stellarators (Nemov *et al.* 1999; Velasco *et al.* 2020) solve similar multi-well matching problems, they generally assume that the electrostatic potential is constant enough on a trajectory to not influence the trapping condition, which is distinctly untrue for mirror plasmas. Thus, it is important to formalise conditions and methods for establishing solution domains for arbitrarily complicated and arbitrarily high-dimensional multi-well BAFP problems, with wells that result from variations in both the magnetic field and electric potential. Establishing these conditions, particularly for mirror-like plasmas, is the goal of this paper.

To fulfil this goal, we begin in § 2, by reviewing the basics of BAFP theory. In § 3, we show how trajectories bifurcate for common mirror configurations, establishing notation and motivating the need for a general method to identify consistent domains. In § 4, we formalise the requirements for establishing consistent domains and show how to fully partition COM space $\mathbf{Z}$ so that the distribution $f(\mathbf{Z})$ is single-valued and well defined. In § 5, we establish the boundary conditions between these domains, allowing BAFP to be solved in the global space despite an arbitrary number of trajectory bifurcations. Because these boundary conditions can be complicated, with different domains matched at different parts of the boundary, in § 6, we show how the connections between the various domains can be simply visualised via a directed graph. Finally, in § 7, we discuss how these ideas enable the design of simulations for dynamically evolving multi-well plasma systems.

2. Bounce-averaged theory in COM coordinates

Here, we briefly review the basics of BAFP theory. We move quickly, with a primary focus on magnetic mirror geometries.

In a magnetic mirror, whether simple, tandem or centrifugal, particles in the mirror live on trapped trajectories. On the shortest time scale, the gyroperiod $1/\Omega_c$, the particle gyrates around the magnetic field line. On the next longest time scale, the bounce period $\tau_b \sim L_{\parallel}/v_{\parallel}$, the particle transits the mirror, completing a closed orbit.

Absent instabilities, particle deconfinement in such mirrors occurs due to collisional processes. Thus, on a collision time scale $\tau_c \gg \tau_b, 1/\Omega_c$, collisions push the particle onto new orbits. Eventually, the particle is pushed onto an orbit that leaves the device.

Obviously, it is inefficient (whether analytically or computationally) to resolve the gyroperiod, bounce time scales and complex trajectories when evaluating the slower diffusion processes that drive particles out of the device. Suitably averaging over these fast time scales to produce a lower-dimensional set of equations describing these slow processes is the goal of BAFP.

The first step in deriving the theory is to transform to coordinates that are constant along a trajectory. For an axisymmetric, non-relativistic plasma, a suitable set of COMs are:

$$\mu \equiv \frac{1}{2} \frac{mv_{\perp}^2}{B(\mathbf{x})}, \quad (2.1)$$

$$\epsilon \equiv \frac{1}{2}mv_{\parallel}^2 + \mu B(\mathbf{x}) + \psi(\mathbf{x}), \quad (2.2)$$

$$p_{\theta} \equiv m\Omega_r r + q\Phi(\mathbf{x}). \quad (2.3)$$

Here, μ is the magnetic moment of the particle, ϵ is the energy and p_{θ} is the azimuthal momentum. Additionally, $v_{\parallel}$ and $v_{\perp}$ are the velocity components parallel and perpendicular to the magnetic field, $\Omega_r(\Phi)$ is the rotation frequency of the magnetic surface, and q and m are the charge and mass of the particle. The potential energy term ψ is given by

$$\psi = q\phi_{\parallel} + \frac{1}{2}mr^2\Omega_r^2, \quad (2.4)$$

where $\phi_{\parallel}$ is a potential coordinate along the flux surface (Kolmes *et al.* 2024). Finally, Φ is the flux function, given by

$$\Phi = rA_{\theta}. \quad (2.5)$$

In looking at p_{θ} , note that the ratio of the first term to the second is given by

$$\frac{m\Omega_r}{qA_{\theta}} \sim \frac{\Omega_r}{\Omega_c}, \quad (2.6)$$

i.e. p_{θ} conservation implies that particles stay on their flux surface for particles far from the Brillouin limit (Rax *et al.* 2015). We assume in the following that we are in this regime, taking Φ as the third COM in place of p_{θ} .

To reach our six dimensions of phase space, we need three other variables. These we take to be the azimuthal angle θ , the gyro-angle α and the distance along a field line s . These variables (θ, α, s) are the variables we will integrate out to form the bounce-averaged theory. Integrating out α is the basis for gyro or drift kinetics; integrating out θ makes the theory axisymmetric and integrating out s results in the bounce average.

Now, in general, we will start with an equation of the form:

$$\sqrt{g_X} \frac{\partial f}{\partial t} = \frac{\partial}{\partial X^i} [\sqrt{g_X} \Gamma_X^i(f, \mathbf{X})] + \sqrt{g_X} S(\mathbf{X}), \quad (2.7)$$

where $\Gamma^i(f, \mathbf{X})$ is a flux operator, $S(\mathbf{X})$ is a source operator and $\sqrt{g_X}$ is the volume element in the space $\mathbf{X} \equiv (\mathbf{x}, \mathbf{p})$. This equation is in conservation form; integrating the left-hand side gives the total change in particle number in a given region, while

the first term on the right-hand side reduces to a surface integral of the flux out of the region.

To get our bounce-averaged Fokker–Planck equation, we first need to convert to the six-dimensional (6-D) space $\mathbf{Y}$ that contains the COMs, i.e. $\mathbf{Y} \equiv (\epsilon, \mu, \Phi, \theta, \alpha, s)$. Because the metric in phase space $\mathbf{X} \equiv (\mathbf{x}, \mathbf{v})$ is $g_{ij} = \delta_{ij}$, corresponding to a unit volume element, the volume element $d\mathbf{V} = \sqrt{g_Y} d\mathbf{Y}$ in the new coordinates is

$$\sqrt{g_Y} = 2 \sqrt{\left| \frac{\partial X^m}{\partial Y^i} \frac{\partial X^n}{\partial Y^j} \delta_{ij} \right|} = \frac{\sqrt{2}}{m^{3/2} \sqrt{\epsilon - \mu B - \psi}}. \quad (2.8)$$

Here, the factor of 2 comes from the fact that positive and negative $v_{\parallel}$ are condensed into the same portion of COM space. We also convert the flux Γ to Y space. For instance, for a Fokker–Planck flux operator:

$$\Gamma_X^i = A_X^i f + D_X^{ij} \frac{\partial f}{\partial X^j}, \quad (2.9)$$

we take

$$\Gamma_Y^i = \frac{\partial Y^i}{\partial X^m} A_X^m f + \frac{\partial Y^i}{\partial X^m} \frac{\partial Y^j}{\partial X^n} D_X^{mn} \frac{\partial f}{\partial Y^j}, \quad (2.10)$$

Now, we assume that f is independent of (θ, α, s) , i.e. take $f(\mathbf{Z})$, $\mathbf{Z} \equiv (\epsilon, \mu, \Phi)$. Then, we integrate the Fokker–Planck equation over the non-COM coordinates in $\mathbf{Y}$. The integrals over θ and α are trivial, leading to two factors of 2π . The integral over s is more non-trivial, but results in

$$\sqrt{g_Z} \frac{\partial f}{\partial t} = \frac{\partial}{\partial Z^i} (\sqrt{g_Z} \Gamma^i(f, \mathbf{Z})) + \sqrt{g_Z} S_Z(\mathbf{Z}), \quad (2.11)$$

where

$$\sqrt{g_Z} \equiv 4\pi^2 \int_{s_1}^{s_2} ds \sqrt{g_Y}, \quad (2.12)$$

$$\Gamma_Z^i \equiv \frac{1}{\sqrt{g_Z}} \int_{s_1}^{s_2} ds \sqrt{g_Y} \Gamma_Y^i, \quad (2.13)$$

$$S_Z \equiv \frac{1}{\sqrt{g_Z}} \int_{s_1}^{s_2} ds \sqrt{g_Y} S_Y, \quad (2.14)$$

and where s_1 and s_2 are the two outer limits of the orbit.

Noting that $v_{\parallel}(s) = \sqrt{2(\epsilon - \mu B(s) - \psi(s))/m}$, we see that the volume element $\sqrt{g_Z}$ is proportional to the total bounce time $\tau_b = \int_{s_1}^{s_2} ds/v_{\parallel}(s)$, while the operators Γ_Z^i and S_Z^i are averaged based on the time $ds/v_{\parallel}(s)$ they spend in each region. Thus, the phase-space volume average and temporal bounce average are equivalent.

3. Multiple wells and the need for multiple domains

Bounce-averaged theory provides a massive simplification to the collisional Vlasov equations. However, it relies on one major assumption: that f can be written as a function of the COM coordinates alone: $f = f(\mathbf{Z})$. As we will see, this condition is often violated in scenarios relevant for modern mirrors. When it is violated, a single




FIGURE 2. (a) Discretised magnetic field B and potential energy ψ as a function of axial segment n and (b) COM-space accessibility plot for electrons in a magnetic mirror with a (typical) outward-pointing electric field. In the COM-space accessibility plot, each line $n \in \{0, 1, 2\}$ represents the boundary below which particles do not have enough kinetic energy to enter that axial segment.

function $f(\mathbf{Z})$ on COM space is no longer adequate to describe the complete system. Then, the problem must be solved on a set of bounded domains, with appropriate boundary conditions enforced on the shared boundaries between domains.

To explore these issues, we will make two simplifications to the above mentioned theory. First, we will look at equations for a single value of Φ , looking only at the 2-D COM space of (ϵ, μ) . Despite this reduction in dimension, we note that the methods presented here generalise straightforwardly to 3-D COM space.

Second, we will consider a piecewise-constant series of potentials $\psi(n(s))$ and magnetic fields $B(n(s))$, with the n th piece of width Δs_n , so that the integrals in the bounce-averaged theory become sums. Of course, this is equivalent in the limit of $\Delta s \rightarrow 0$ and the number of partitions going to infinity, but working with a small discrete number of areas will allow us to develop the discussion much more clearly.

Consider, then, particle dynamics along a field line. As is often conventional, we will take the mirror dynamics to be symmetric about the mirror midplane, and so consider a sequence of $\psi(n)$, $B(n)$, with $n \in \{0, \dots, N-1\}$, and with $n=0$ corresponding to the midplane and $n=N-1$ corresponding to the mirror throat. We normalise the potentials and fields to their midplane values, with $\psi(0)=0$ and $B(0)=1$, and arbitrary units for ψ and ϵ .

Given sequences $\psi(n)$ and $B(n)$, we can determine the particles allowed in any given field line segment n . For a particle to be allowed in a region, it must have positive parallel kinetic energy, $v_{\parallel}^2 > 0$. From (2.2), we see that this requires

$$\epsilon \geq \mu B(n) + \psi(n). \quad (3.1)$$

Each pair $(\psi(n), B(n))$ thus determines an allowed region of (ϵ, μ) space, i.e. the region satisfying (3.1).

The simplest case is provided by electrons in a simple mirror (which also applies to both electrons and ions in a centrifugal mirror). An ambipolar potential generally forms (Pastukhov 1974; Najmabadi, Conn & Cohen 1984; Ochs, Munirov & Fisch 2023), confining electrons and repelling ions. Thus, for electrons, both ψ and B increase as a function of line segment index n , and so particles are best trapped at the midplane (figure 2). This is the ideal situation for BAFP theory: trajectories that




FIGURE 3. (a) Discretised magnetic field B and potential energy ψ as a function of axial segment n and (b) COM-space accessibility plot for ions in a magnetic mirror with a (typical) outward-pointing electric field. Because of the decreasing electric potential towards the edge of the device, some ions get trapped between the mirror throat and the midplane, i.e. at $n = 1$. These ions are referred to as ‘Yushmanov-trapped’.

can access segment $n = 1$ also access segment $n = 0$, so the distribution function $f(\epsilon, \mu)$ is clearly single-valued. To apply a BAFP theory, operator averages are performed over segment $n = 0$ for the particles above line 0 and below line 1, and over segments $n \in \{0, 1\}$ for the particles that are above line 1 and below line 2. A similar situation will occur whenever $\partial\psi/\partial s \geq 0$ and $\partial B/\partial s \geq 0$ for the entire field line from the midplane to throat. Such a plasma can even be modelled using midplane momentum coordinates, since every allowed trapped trajectory reaches the midplane ($n = 0$).

A somewhat more complex case is provided by ions in a simple mirror. Because the potential is repulsive, ψ now decreases as a function of n (figure 3). Thus, a class of ions can develop which are trapped off the midplane in the potential well, despite the higher magnetic field there. These are known as ‘Yushmanov-trapped’ particles (Yushmanov 1966; Post 1987). In a plasma with Yushmanov-trapped particles, a midplane-momentum coordinate theory becomes insufficient, because these particles exist in the plasma but do not reach the midplane. However, the distribution function $f(\mathbf{Z})$ is still clearly single-valued, so BAFP theory can still be used, solved on a single domain in COM space. The operator averages are performed over $n = 1$ for the Yushmanov particles, over $n = 0$ for the particles above line 0 and below line 1, and over $n \in \{0, 1\}$ for the passing particles that traverse both segments.

As the magnetic geometry and potentials become more complex, the single-valuedness can quickly break down. For instance, consider a case where the magnetic field is constant, but the potential has an intermediate peak (figure 4). Now, as in the 1-D double well in figure 1, ϵ and μ are not sufficient to define a particle trajectory, since a particle with insufficient ϵ to traverse the peak can be trapped on one side or the other. For instance, in figure 4, in region r_1 where (ϵ, μ) is below line 1 and above line 2, the particle can either be on a trajectory that is trapped in segment $n = 0$ or a trajectory trapped in segment $n = 2$. There is no reason for particles on these different trajectories to share a value for the distribution function; thus, $f(\epsilon, \mu)$ fails to be single-valued in this region. (A similar bifurcation occurs for a local maximum of the magnetic field; the only difference is that changing B would change the slope of the lines in (ϵ, μ) space, rather than their ϵ -intercept.)




FIGURE 4. (a) Discretised magnetic field B and potential energy ψ as a function of axial segment n and (b) COM-space accessibility plot for a scenario with a constant magnetic field and an internal potential maximum. This scenario exhibits a bifurcation of trajectories around $n = 1$, and thus the bounce-averaged theory for this space requires three solution domains.




FIGURE 5. Directed graph structure of boundary conditions for the field configuration in figure 4.

When such a bifurcation of the trajectory happens, the domain on which the BAFP equations are solved must be split, to keep $f(\epsilon, \mu)$ single-valued on each domain. Each domain consists of (a) an area in (ϵ, μ) space and (b) a continuous subset of the field line segments n over which the bounce-average is taken. For the case in figure 4, three domains are required: the domain d_1 consisting of region r_2 , bounce-averaged over field line segments $n \in \{0, 1, 2\}$; the domain d_2 consisting of regions r_0 and r_1 , bounce-averaged over field line segments $n \in \{0\}$; and the domain d_3 consisting of region r_1 , bounce-averaged over field line segments $n \in \{2\}$.

To solve the problem, the domains must be stitched together. This requires two boundary conditions at the boundary (line 1). First, $f(\epsilon, \mu)$ must be continuous at the boundary. Second, the phase-space flux must be continuous, i.e.

$$n_i \Gamma_Z^i|_{d_1} = n_i \Gamma_Z^i|_{d_2} + n_i \Gamma_Z^i|_{d_3}, \quad (3.2)$$

where n_i is a shared normal vector to the surface in (ϵ, μ) space.

We can see from the flux conservation equation that the two domains d_2 and d_3 merge into d_1 . This suggests a representation of the relationship between the domains as a directed graph, with the directionality going from the domain representing a single passing trajectory to the domains representing bifurcated trajectories (figure 5).




FIGURE 6. (a) Discretised magnetic field B and potential energy ψ as a function of axial segment n and (b) COM-space accessibility plot for a scenario exhibiting trajectory bifurcation without a local maximum in either potential or magnetic field. This ‘Yushmanov trajectory bifurcation’ process is discussed in more detail in [Appendix A](#).

3.1. Ambipolar fields and the need for an automatic method

Historically, especially in the study of tandem mirrors with thermal barriers, these domain relations have been imposed at the beginning of the problem and left fixed for the computation (Matsuda & Stewart 1986). However, this is likely to be insufficient for accurate simulations of modern mirrors, which often have potential profiles determined by sloshing kinetic ion distributions (Egedal *et al.* 2022; Endrizzi *et al.* 2023; Frank *et al.* 2024). Correctly determining the resulting potentials requires solving for the electric potential in each region self-consistently over the course of the simulation so as to enforce quasineutrality, i.e. to enforce

$$\sum_s Z_s \int_s f_s d\mathbf{Z} = 0 \quad (3.3)$$

for each region (Frank *et al.* 2024; Kolmes *et al.* 2025). As a result, it might not be known at the beginning of the simulation where the potential peak might occur, and thus over which field line segments and regions of (ϵ, μ) to define the domains. Thus, dynamic detection of such domains is a necessity.

Making matters worse, the bifurcation of trajectories does not generally require such an easy-to-spot feature as a local peak in the electric or magnetic field (which is the basis for such calculations in stellarator neoclassical transport codes (Nemov *et al.* 1999; Velasco *et al.* 2020)). Consider the field arrangement in [figure 6](#). In segment $n = 1$, this has neither a local maximum in the magnetic or electric field, and thus does not seem like it would represent a barrier between local wells. Nevertheless, as can be seen in [figure 6](#), there is a region of (ϵ, μ) space that can be accessed by particles with $n \in \{0, 2\}$, but not for $n = 1$. Thus, this field set-up also has a trajectory bifurcation and must be solved using multiple grids.

The conditions required for the particular situation in [figure 6](#), which we term a ‘Yushmanov trajectory bifurcation’, can actually be solved for ([Appendix A](#)); it turns out that such a bifurcation will occur when

$$\frac{\partial \psi}{\partial B} < 0 \quad \text{and} \quad \frac{\partial^2 \psi}{\partial B^2} > 0. \quad (3.4)$$

Nevertheless, in anticipating arbitrary field arrangements involving multiple wells and, particularly, in looking forward to non-relativistic plasmas in three COM dimensions, it is clearly useful to have a generalisable method that can take the allowed regions (in COM space) for each part (in $\mathbf{x}$) of the plasma, and return the requisite domains and boundary conditions. The development of such a method is the subject of the next sections.

4. Formalising the requirements for cohesive domains

As a starting point, we assume that we have a set of segments $\mathcal{N} = \{0, \dots, N-1\}$, with allowed areas of (ϵ, μ) space defined for each segment $n \in \mathcal{N}$, of the form $\epsilon \geq b_n(\mu)$. This will in general chop the (ϵ, μ) space up into a number of different regions r_i , with edges defined by the boundary functions $b_n(\mu)$.

We note that in formulating the problem in this way, the generalisation to 3-D (ϵ, μ, Φ) space is straightforward, since each allowed region will now be given by $\epsilon \geq b_n(\mu, \Phi)$, with the boundaries between regions given by 2-D surfaces rather than 1-D curves. The following discussion thus applies immediately to more complicated and higher-dimensional problems, though to keep things clear, we will restrict the discussion to non-relativistic (ϵ, μ) space.

Now, consider a specific region r_i . This region will have a set $\mathcal{N}_i \subseteq \mathcal{N}$ of the segments for which r_i is an allowed region. Ordering the set $\mathcal{N}_i$ in terms of increasing n , we can in general then split $\mathcal{N}_i$ into maximal continuous subsequences $\mathcal{C}_i^m$, defined such that each $\mathcal{C}_i^m$ has no breaks in the integer sequence of n values. The set of all $\mathcal{C}_i^m$ for a given region r_i we denote $\mathcal{C}_i$. For a given r_i , the size (cardinality) of the set $\mathcal{C}_i$ represents the number of distinct trajectories in that region. For instance, returning to figure 6, the indicated region in (ϵ, μ) , which we can label as region r_0 , has $\mathcal{N}_0 = \{0, 2\}$, $\mathcal{C}_0^0 = 0$, $\mathcal{C}_0^1 = 2$ and $\mathcal{C}_0 = \{\mathcal{C}_0^0, \mathcal{C}_0^1\} = \{\{0\}, \{2\}\}$. The fact that $|\mathcal{C}_0| = 2$ means that there are two distinct trajectories; namely, those trapped in axial segment $n = 0$ and those trapped in axial segment $n = 2$.

At this point, it is helpful to define a population, denoted p_i^m , as the combination of the region r_i (a volume in COM space) and a maximal continuous subsequence of allowed axial segments $\mathcal{C}_i^m$ in that region. The population defines everything other than the boundary conditions about the bounce-averaged Fokker–Planck problem in COM space: namely, it defines the region over which the problem is to be solved and the axial segments over which the bounce average is to be taken. Boundary conditions then come from stitching together connecting populations.

However, solving the Fokker–Planck problem for each population so defined would be extremely inefficient, since there can be many regions even for problems (such as those in figures 2 and 3) that can be solved on a single domain. Thus, our goal will be to build up collections of compatible populations into maximal simulation domains.

Thus, define a domain d as a collection of connected and pairwise-compatible populations $\{p_i^m\}$. To define ‘connected’ and ‘pairwise-compatible’, it is helpful to consider what happens as we try to extend the domain to a new population. Thus, consider a population p_i^m in domain d , with region r_i and continuous sequence $\mathcal{C}_i^m$. Then, try to extend the domain d by adding a population with region r_j bordering r_i . What are the conditions that allow this to happen?

First, it only makes sense to extend the domain if there is a population on r_j that contacts trajectories from p_i^m ; i.e. if

$$\exists \mathcal{C}_j^n \in \mathcal{C}_j : (r_i \text{ borders } r_j) \quad \text{and} \quad (\mathcal{C}_j^n \cap \mathcal{C}_i^m \neq \{\}). \quad (4.1)$$

If this condition is not satisfied, then trajectories in region r_j are located in a spatially distinct part of the plasma from those trajectories described by p_i^m and thus do not interact in a Fokker–Planck manner with the population p_i^m . Thus, populations p_i^m and p_j^n are said to be connected if

$$(r_i \text{ borders } r_j) \quad \text{and} \quad (\mathcal{C}_j^n \cap \mathcal{C}_i^m \neq \{\}). \quad (4.2)$$

However, connectedness is not sufficient to ensure that the domain can be extended. Consider, for instance, the regions r_1 and r_2 on either side of line 1 in [figure 4](#). Then, consider the population $p_2^0 \equiv (r_2, \{0, 1, 2\})$. Going to region r_1 , we see that there are two populations that are connected to p_2^0 : $p_1^0 \equiv (r_1, \{0\})$ and $p_1^1 \equiv (r_1, \{2\})$. Precisely because there are multiple connected populations in region r_1 , the domain cannot be extended, because this multiplicity means that trajectories bifurcate at this point and a boundary condition (3.2) must be enforced at line 1. Thus, for p_i^m in region r_i to be compatible with a population p_j^n in region r_j , we must have

$$\nexists l : (l \neq n) \quad \text{and} \quad (\mathcal{C}_j^l \cap \mathcal{C}_i^m \neq \{\}). \quad (4.3)$$

However, this is still not sufficient, since this compatibility relation must be reciprocal; condition (4.3) shows that p_2^0 is incompatible with the p_1^0 and p_1^1 in [figure 5](#), but not that p_1^0 and p_1^1 are incompatible with p_2^0 . Thus, for p_i^m in region r_i to be compatible with a population p_j^n in region r_j , we must also have

$$\nexists l : (l \neq m) \quad \text{and} \quad (\mathcal{C}_i^l \cap \mathcal{C}_j^n \neq \{\}). \quad (4.4)$$

If, for populations p_i^m and p_j^n , conditions (4.3)–(4.4) hold, then the populations are compatible and can exist in the same domain.

Unfortunately, it can be shown ([Appendix B](#)) that the compatibility property is not transitive. Thus, to add a population to the domain, connectedness must only be checked for one region, but compatibility must be checked with all populations already in the domain. This non-transitivity is why the domain was defined to be a connected and pairwise-compatible set of populations.

To completely partition the space for solution of the Fokker–Planck problem, we must assign every population to a domain. These domains are then connected by boundary conditions along shared boundaries. As shown in [Appendix B](#), this decomposition is not, in general, unique. Nevertheless, an example algorithm for performing a valid decomposition is described in [Appendix C](#).

At the end, each domain d_a (which is a set of populations) will be associated with a set of connected regions $\mathbb{R}_a$ and a single continuous set of axial segments $\mathbb{C}_a$. Thus, for each domain d_a , the BAFP problem is to be performed over the area in $\mathbf{Z}$ that is spanned by $\mathbb{R}_a$, with the bounce average taken over the axial segments contained in $\mathbb{C}_a$.

5. Enforcing boundary conditions

Now that the space is partitioned into domains, we must define boundary conditions for each domain. As in § 3, we need to enforce continuity and particle conservation at each boundary; the key aspect will be to identify which populations must be matched.

There are two types of boundaries. External boundaries are defined by the edge of accessible $\mathbf{Z}$ for the total modelled space. These boundaries are either reflecting, if they represent the boundary of 0 kinetic energy that defines accessibility,

or approximately absorbing (Dirichlet with a value of 0), if they represent the loss cone. (Though it may be noted that in more collisional mirrors (Mirnov & Riutov 1979; Rognlien & Cutler 1980) or mirrors with significant secondary electron emission from the walls (Konkashbaev, Landman & Ulinich 1978; Skovorodin 2019), a non-zero loss cone boundary may be appropriate.)

The remaining boundaries are internal, representing matching conditions between domains. In general, if we look at a domain d_a , each internal boundary can be defined by two regions; the region r_i (with population p_i^m in domain d_a) and the region r_j . If r_j does not represent a region within the domain d_a , then this represents an outer boundary of the domain d_a , where the boundary conditions must be applied. We can denote this boundary by b_{ij}^a .

Now, note that of the two regions r_i and r_j , one of these regions will always have a set of accessible axial segments that is a strict superset of the other region; i.e.:

$$\mathcal{N}_i \subset \mathcal{N}_j \quad \text{or} \quad \mathcal{N}_j \subset \mathcal{N}_i. \quad (5.1)$$

Thus, we can define region r_i to be ‘higher access’ than r_j if $\mathcal{N}_i \supset \mathcal{N}_j$, and ‘lower access’ if the reverse holds. In general, trajectories can bifurcate in going from the higher-access region to the lower-access region, but not *vice versa*.

Now let us establish the boundary conditions for boundary b_{ij}^a between region r_i (in domain d_a) and r_j (not in domain d_a). We start by identifying the higher-access region. Then, taking the population for region r_i in domain d_a to be p_i^m , we start by finding the connected populations in region r_j . This set of populations will define a set $\mathcal{D}_{ij}^a$ of domains connected to domain d_a at that boundary, i.e. the domains which contain those connected populations. If r_i is the higher-access region, then the boundary conditions now take the form:

$$f(\mathbf{Z})|_{b_{ij}^a, d_a} = \underbrace{f(\mathbf{Z})|_{b_{ji}^b, d_b} = \dots}_{d_b \in \mathcal{D}_{ij}^a} \quad (5.2)$$

$$n_i \Gamma_Z^i|_{b_{ij}^a, d_a} = \sum_{d_b \in \mathcal{D}_{ij}^a} n_i \Gamma_Z^i|_{b_{ji}^b, d_b}. \quad (5.3)$$

Note that if r_j has no populations that are connected to p_i^m , then $\mathcal{D}_{ij}^a = \{\}$ and the boundary conditions reduce to the reflecting boundary conditions of an external boundary.

If, instead, r_i is the lower-access region, then there will be only one population p_j^n that is connected to p_i^m in domain d_a . Denote the domain associated with p_j^n as d_b . Then, we can establish the boundary conditions as in the previous paragraph, but starting with p_j^n in d_b . Explicitly, we first find all the populations p_i^l in region r_i connected to p_j^n and thus construct the set of domains $\mathcal{D}_{ji}^b$ containing those connected populations. Note that $\mathcal{D}_{ji}^b$ will, by definition, include d_a . Then, the boundary conditions are

$$f(\mathbf{Z})|_{b_{ji}^b, d_b} = \underbrace{f(\mathbf{Z})|_{b_{ij}^c, d_c} = \dots}_{d_c \in \mathcal{D}_{ji}^b} \quad (5.4)$$

$$n_i \Gamma_Z^i|_{b_{ji}^b, d_b} = \sum_{d_c \in \mathcal{D}_{ji}^b} n_i \Gamma_Z^i|_{b_{ij}^c, d_c}. \quad (5.5)$$

The above mentioned procedure defines all internal boundary conditions for the BAFP problem.

6. Graph structure of domain connections

The rules outlined in § 4 and § 5 are sufficient to properly set up a BAFP problem, to solve it either analytically or (more likely if application is necessary) computationally. However, it is also useful when thinking about a problem to be able to quickly visualise how the various domains fit together. The abovementioned rules, which are focused on conditions at each boundary (of which each domain can have quite a number), do not make developing this intuition particularly straightforward. The goal, then, is to quickly encode the boundary condition connections into the connections between domains, which can then be visualised.

The condition for two domains to be connected is simple: domains d_a and d_b will be connected if

$$\exists p_i^m \in d_a, p_j^n \in d_b : (p_i^m \text{ and } p_j^n \text{ are connected}). \quad (6.1)$$

Furthermore, we can establish a ‘direction of connection’, based on whether p_i^m or p_j^n is in a higher-access region. We draw the direction from the higher-access region to the lower-access region, representing the direction in which trajectories can bifurcate. In other words, if $\mathcal{N}_i \supset \mathcal{N}_j$, the direction of connection is from d_a to d_b . Note that if two domains connect at multiple boundaries, the direction of connection can, in principle, be bi-directional.

With such connections established, the boundary condition structure of a BAFP problem can then be represented as a graph. This graph can help to clarify the multi-well structure of the BAFP problem, in particular, encoding the bifurcation structure of trajectories.

For instance, [figure 7](#) shows a ‘tiered well’, where as ϵ decreases, trajectories first bifurcate at the local potential and magnetic maximum at $n = 3$, and then bifurcate again at the local potential and magnetic maxima at $n = 1$ and $n = 5$. Although there are only three resulting regions in (ϵ, μ) space, the separate wells mean that there are seven separate domains ([figure 8](#)). The boundaries between these domains are connected as shown in [figure 9](#). Domain d_6 represents the domain of trajectories which transit the entire device. Then, this bifurcates into domains d_4 and d_5 , which represent the trajectories trapped on either side of $n = 3$, but which pass over the potential maxima at $n = 1$ and $n = 5$. Finally, the trajectories bifurcate again at these final potential maxima. In this way, the graph structure succinctly encodes the domain connections that result from the well structure.

Of course, the ‘tiered well’ is a very straightforward example with a very clear logic to the domain connections. The established rules for establishing cohesive domains and determining the connections between domains really prove their value when considering more complicated scenarios. For instance, consider the field arrangement in [figure 10](#). One might scratch their head for quite a while coming up with a set of consistent domains to describe this problem, but the rules in § 4, as implemented in the algorithm in [Appendix C](#), can quickly come up with a set of consistent domains ([figure 11](#)). Furthermore, the methods described here can also quickly reveal the connections between these domains ([figure 12](#)). Thus, it can be seen that formulating the problem in this way opens up the possibility of solving arbitrarily complicated multi-well BAFP problems.

7. Conclusion

In this paper, we have shown how to construct consistent domains for the solution of BAFP problems, so that in each domain, a single point in COM space




FIGURE 7. (a) A 'tiered well' field arrangement and (b) COM accessibility plot. In this scenario, in order of decreasing ϵ , trajectories first bifurcate around axial segment $n = 3$, then bifurcate again around segments $n = 1$ and $n = 5$.




FIGURE 8. Domain decomposition for the 'tiered well' in figure 7. Each domain d_a consists of a list of populations, and as a result has an associated set of regions $\mathbb{R}_a$ (the shaded area in each plot) and continuous set of axial segments C_a , given in the title of the plot. Here, several domains (e.g. d_0 , d_1 , d_2 and d_3) share the same region in (ϵ, μ) , but represent different populations because they occur at different positions along the field line.




FIGURE 9. Graph structure of domain connections for the 'tiered well' in figure 7, with the domains defined in figure 8. The forks in the graph are associated with the trajectory bifurcations, first at $n = 3$ (splitting d_6 into d_4 and d_5), then again around segments $n = 1$ and $n = 5$.




FIGURE 10. An arbitrary complicated field arrangement, with many crossings in the accessibility boundary lines.




FIGURE 11. An algorithmically solved consistent domain decomposition of the scenario in figure 10, showing the shaded set of regions $\mathbb{R}_a$ and continuous set of axial segments $\mathbb{C}_a$ for each domain d_a .




FIGURE 12. Graph structure of domain connections for the scenario in figure 10, with the domains defined in figure 11. The graph structure is much more complicated than for the tiered well scenario in figures 7, 8 and 9, but if one chooses any set of connected nodes, it is possible to identify the boundary where the conditions are enforced.

corresponds to a single trajectory. We also showed which boundary conditions were enforced between the resulting domains and how these boundary connections (associated with trajectory bifurcations) could be succinctly summarised via directed graphs.

Though the focus was primarily on magnetic mirrors, the generality of the derived conditions should allow bounce-averaged theories to be applied to a variety of plasma systems with well-defined COM, including tokamaks and quasisymmetric stellarators. Crucially, the clear rules for identifying consistent domains should enable the gridding and boundary-matching process to be automated, allowing bounce-averaged theories to be applied even in systems where wells (and their associated trajectory bifurcations) arise dynamically and unpredictably over the course of the simulation. This is the case, for instance, in the simulation of magnetic mirrors with highly kinetic sloshing ion distributions or arrangements in which Yushmanov trapping becomes significant. Of course, there are other issues to be solved in making bounce-averaged simulations work efficiently for such plasmas, including the development of stable and efficient methods for the ambipolar solve discussed in § 3.1, as well as efficient methods for choosing between the various possible domain decompositions so as to retain maximal efficiency and continuity. Nevertheless, the present work provides a crucial stepping stone towards enabling BAFP codes to simulate these more complicated and relevant modern systems, providing a self-consistent alternative to much more expensive codes that resolve the parallel dynamics.

Acknowledgements

The author thanks Nat Fisch, Alex Glasser, Thomas Foster and Elijah Kolmes for valuable conversations.

Editor Cary Forest thanks the referees for their advice in evaluating this article.

Funding

This work was supported by ARPA-E Grant No. DE-AR0001554.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Differential condition for Yushmanov trajectory bifurcation

Here, we derive the non-relativistic condition for a trajectory bifurcation to occur when the potential is increasing and the magnetic field is decreasing as a function of axial coordinate s (or *vice versa*); i.e. when

$$\text{sgn} \left(\frac{\partial B}{\partial s} \right) = -\text{sgn} \left(\frac{\partial \psi}{\partial s} \right). \quad (\text{A.1})$$

We call this a ‘Yushmanov trajectory bifurcation’ because it is closely related to the off-midplane ‘Yushmanov-trapped’ ions that are well known from mirror theory (Yushmanov 1966; Post 1987).

Thus, recall figure 6. The reason a trajectory bifurcation occurs in this case is that lines 0 and 2 intersect below line 1. We can easily write this condition in terms of

B_n and ψ_n . The intersection point μ^* of lines 0 and 2 is given by

$$\mu^* B_0 + \psi_0 = \mu^* B_2 + \psi_2 \Rightarrow \mu^* = \frac{\psi_2 - \psi_0}{B_0 - B_2}. \quad (\text{A.2})$$

Then, the value of the accessibility boundary at this μ^* must be lower for line 0 than for line 1:

$$\mu^* B_0 + \psi_0 < \mu^* B_1 + \psi_1. \quad (\text{A.3})$$

We can turn the condition into a differential condition by taking the axial segments to be equally spaced with spacing Δs , and taking $\Delta s \rightarrow 0$. This requires working to second order in Δs :

$$\psi_0 = \psi_0; \quad (\text{A.4})$$

$$\psi_1 = \psi_0 + \frac{\partial \psi}{\partial s} \Delta s + \frac{1}{2} \frac{\partial^2 \psi}{\partial s^2} (\Delta s)^2; \quad (\text{A.5})$$

$$\psi_2 = \psi_0 + \frac{\partial \psi}{\partial s} (2\Delta s) + \frac{\partial^2 \psi}{\partial s^2} (2\Delta s)^2. \quad (\text{A.6})$$

The same expansion is enforced for B . With this prescription, (A.3) becomes

$$\frac{\partial^2 B / \partial s^2}{\partial B / \partial s} > \frac{\partial^2 \psi / \partial s^2}{\partial \psi / \partial s}. \quad (\text{A.7})$$

Thus, if (A.1) and (A.7) hold at a point, there will be a trajectory bifurcation there.

Alternatively, conditions (A.1) and (A.7) can also be put into s -independent form by treating ψ as a function of B and expanding $\psi(B)$, yielding the equivalent but simpler conditions for a trajectory bifurcation:

$$\frac{\partial \psi}{\partial B} < 0 \quad \text{and} \quad \frac{\partial^2 \psi}{\partial B^2} > 0. \quad (\text{A.8})$$

To see these conditions in action, we can take fields of the form:

$$B = c_B e^{a_B s}; \quad \psi = c_\psi (1 - e^{a_\psi s}), \quad (\text{A.9})$$

with $a_B > 0$ and $a_\psi > 0$ to satisfy (A.1). Then, (A.7) and (A.8) reduce to the condition that $a_B > a_\psi$.

The accessibility boundaries for fields of the form of (A.9) are shown in figure 13 for $s \in \{0, 1, 2\}$. It can be seen that the condition $a_B > a_\psi$, corresponding to (A.7), correctly describes the onset of the bifurcation.

Appendix B. Non-transitivity of compatibility and non-uniqueness of domain decomposition

In this appendix, we show that the compatibility conditions are not necessarily transitive and that the domain decomposition is not necessarily unique. To show this, an example field configuration suffices.




FIGURE 13. Accessibility region boundaries in COM space for $s=0, 1, 2$ for fields of the form in (A.9), demonstrating the Yushmanov trajectory bifurcation condition ((A.1) and (A.7) or (A.8)). In panel (a), $a_B = a_\psi = 1$, and the lines all intersect at a single point, so there is (marginally) no bifurcation. In panel (b), $a_B = 1 < a_\psi = 1.2$, and so there is no bifurcation. In panel (c), $a_B = 1.5 > a_\psi = 1$, and so (A.7) (or (A.8)) is satisfied and there is a trajectory bifurcation. For all plots, $c_B = 1$ and $c_\psi = 2$.




FIGURE 14. Example field configuration and COM-space accessibility plot demonstrating non-transitivity of population compatibility and non-uniqueness of the domain decomposition.

Thus, consider the field arrangement in figure 14. This grid is characterised by six populations:

$$\begin{aligned}
 p_0^0 &: \{r_0, C_0^0 = \{0\}\}; \\
 p_1^0 &: \{r_1, C_1^0 = \{1\}\}; \\
 p_2^0 &: \{r_2, C_2^0 = \{0, 1\}\}; \\
 p_3^0 &: \{r_2, C_3^0 = \{0, 1, 2\}\}; \\
 p_4^0 &: \{r_4, C_4^0 = \{0\}\}; \\
 p_4^1 &: \{r_4, C_4^1 = \{2\}\}.
 \end{aligned}
 \tag{B.1}$$

There is a clear incompatibility between the populations p_3^0 in region r_3 , and the populations p_4^0 and p_4^1 in region r_4 . However, every other population is compatible with both population p_3^0 and p_4^0 . Thus, the compatibility condition is seen to be non-transitive.

This non-transitivity also leads to non-uniqueness of the domain set. Due to the bifurcation in trajectories between regions r_3 and r_4 , three domains are required. However, which populations go into these domains is a matter of choice. For instance, we can simply assign each of the troublesome bifurcated populations in region r_4 to their own domain:

$$d_0 = \{p_0^0, p_1^0, p_2^0, p_3^0\}; \quad d_1 = \{p_4^0\}; \quad d_2 = \{p_4^1\}. \quad (\text{B.2})$$

Alternatively, however, we can also isolate population p_3^0 instead of p_4^0 . Then, we get

$$d'_0 = \{p_0^0, p_1^0, p_2^0, p_4^0\}; \quad d'_1 = \{p_3^0\}; \quad d'_2 = \{p_4^1\}. \quad (\text{B.3})$$

Either of the domain decompositions are equally valid according to the conditions laid out in § 4 and thus the decomposition is not generally unique.

Appendix C. Algorithm to find a complete set of domains

A complete partitioning of the populations into domains can be performed as follows.

- (i) Define $\mathcal{U} = \{p_i^m\}$, the set of unsorted populations.
- (ii) Define $\mathcal{D} = \{\}$, the set of domains.
- (iii) If $|\mathcal{U}| > 0$, initialise a new domain $d = \{\}$; otherwise, terminate.
- (iv) Remove a population p_i^m from $\mathcal{U}$, add it to d .
- (v) Construct $\mathcal{R}_i$, the set of neighbouring regions of region r_i , where r_i is the region associated with populations p_i^m .
- (vi) Choose an $r_j \in \mathcal{R}_i$.
- (vii) Check if a population from r_j is already in d ; if yes, go to the next region in $\mathcal{R}_i$.
- (viii) Check if a population $p_j^n \in \mathcal{U}$ with region r_j is connected to population p_i^m (condition (4.2)) and compatible with all populations in d (conditions (4.3), (4.4)).
- (ix) If step (viii) is true, remove p_j^n from $\mathcal{U}$ and add p_j^n to d .
- (x) Repeat steps (vi)–(ix) until all regions in $\mathcal{R}_i$ have been checked.
- (xi) Repeat steps (v)–(x) for each population p_i^m in d .
- (xii) When all neighbouring regions of all populations have been checked, and no new populations can be added, add d to $\mathcal{D}$, and go to step (iii).

It is also useful to have a computer-friendly algorithm for step (viii) to determine compatibility and connectedness. To determine whether p_i^m has a compatible population in region r_j , one can do the following.

- (i) Count the number a of elements $\mathcal{C}_j^n$ of $\mathcal{C}_j$ such that $\mathcal{C}_j^n \cap \mathcal{C}_i^m \neq \{\}$.
- (ii) If $a = 1$, then choose the single $\mathcal{C}_j^n$ that intersects $\mathcal{C}_i^m$.
- (iii) Count the number b of elements $\mathcal{C}_i^k$ of $\mathcal{C}_i$ s.t. $\mathcal{C}_i^k \cap \mathcal{C}_j^m \neq \{\}$.

(iv) Then:

- (a) if $a = 0$, p_i^m is compatible with all populations in r_j , but not connected;
- (b) if $a > 1$, then p_i^m is connected to multiple populations in r_j , so there is no compatible population in r_j (there is a trajectory bifurcation going from r_i to r_j);
- (c) if $a = 1$ but $b > 1$, then p_j^n is connected to multiple populations in r_i , so p_j^n and p_i^m are not compatible (there is a trajectory bifurcation going from r_j to r_i);
- (d) if $a = 1$ and $b = 1$, then p_i^m is compatible with p_j^n in region r_j , with C_j^m given by step (ii). If r_i borders r_j , these populations are also connected.

If one wants to check whether a specific population p_j^l in region r_j is compatible with p_i^m , then one must additionally check that the compatible population p_j^n identified in step (ivd) is the same as the queried population p_j^l .

REFERENCES

- BE'ERY, I., GERTSMAN, A. & SEEMAN, O. 2018 Plasma confinement by moving multiple mirrors. *Plasma Phys. Control. Fusion* **60**, 115004.
- BEKHTENEV, A., VOLOSOV, V., PAL'CHIKOV, V., PEKKER, M. & YUDIN, YUN. 1980 Problems of a thermonuclear reactor with a rotating plasma. *Nucl. Fusion* **20**, 579–598.
- BENDANIEL, D.J. & ALLIS, W.P. 1962 Scattering loss from magnetic mirror systems – II. *J. Nucl. Energy. C* **4**, 79–88.
- BERNSTEIN, I.B. & BAXTER, D.C. 1981 Relativistic theory of electron cyclotron resonance heating. *Phys. Fluids* **24**, 108–126.
- BILBAO, P.J. & SILVA, L.O. 2023 Radiation reaction cooling as a source of anisotropic momentum distributions with inverted populations. *Phys. Rev. Lett.* **130**, 165101.
- CHO, T., *et al.* 2005 Observation of the effects of radially sheared electric fields on the suppression of turbulent vortex structures and the associated transverse loss in GAMMA 10. *Phys. Rev. Lett.* **94**, 085002.
- COHEN, R.H., BERNSTEIN, I.B., DORNING, J.J. & ROWLANDS, G. 1980 Particle and energy exchange between untrapped and electrostatically confined populations in magnetic mirrors. *Nucl. Fusion* **20**, 1421–1437.
- D'HERBEMONT, V., PARRA, F.I., CALVO, I. & VELASCO, J.L. 2022 Finite orbit width effects in large aspect ratio stellarators. *J. Plasma Phys.* **88**, 905880507.
- EGEDAL, J., ENDRIZZI, D., FOREST, C. & FOWLER, T. 2022 Fusion by beam ions in a low collisionality, high mirror ratio magnetic mirror. *Nucl. Fusion* **62**, 126053.
- ENDRIZZI, D., *et al.* 2023 Physics basis for the Wisconsin HTS axisymmetric mirror (WHAM). *J. Plasma Phys.* **89**, 975890501.
- ERIKSSON, L.-G. & HELANDER, P. 1994 Monte Carlo operators for orbit-averaged fokker–Planck equations. *Phys. Plasmas* **1**, 308–314.
- FISCH, N.J. 1987 Theory of current drive in plasmas. *Rev. Mod. Phys.* **59**, 175–234.
- FISCH, N.J. & RAX, J.-M. 1992 Interaction of energetic alpha particles with intense lower hybrid waves. *Phys. Rev. Lett.* **69**, 612–615.
- FOWLER, T., MOIR, R. & SIMONEN, T. 2017 A new simpler way to obtain high fusion power gain in tandem mirrors. *Nucl. Fusion* **57**, 056014.
- FRANK, S.J., *et al.* 2024, Integrated modelling of equilibrium and transport in axisymmetric magnetic mirror fusion devices. arXiv:2411.06644.
- HAGER, R., YOON, E., KU, S., D'AZEVEDO, E., WORLEY, P. & CHANG, C. 2016 A fully non-linear multi-species Fokker–Planck–Landau collision operator for simulation of fusion plasma. *J. Comput. Phys.* **315**, 644–660.

- HARVEY, R.W. & MCCOY, M.G. 1992 The CQL3D Fokker–Planck code. In *Proceedings of the IAEA Technical Committee Meeting on Simulation and Modeling of Thermonuclear Plasmas*, pp. 489–526.
- HERRMANN, M.C. 1998 Cooling alpha particles with waves *PhD thesis*. Princeton University.
- HERRMANN, M.C. & FISCH, N.J. 1997 Cooling energetic alpha particles in a tokamak with waves. *Phys. Rev. Lett.* **79**, 1495–1498.
- KATANUMA, I., KIWAMOTO, Y., ISHII, K. & MIYOSHI, S. 1986 Thermal barrier potential of a tandem mirror. *Phys. Fluids* **29**, 4138–4146.
- KATANUMA, I., KIWAMOTO, Y., SAWADA, K. & MIYOSHI, S. 1987 Fokker–Planck calculation of hot electron buildup in the thermal barrier region of a tandem mirror. *Phys. Fluids* **30**, 1142–1151.
- KOLMES, E.J. & FISCH, N.J. 2024 Coriolis forces modify magnetostatic ponderomotive potentials. *Phys. Plasmas* **31**, 112107.
- KOLMES, E.J., OCHS, I.E. & FISCH, N.J. 2025, Ion mix can invert centrifugal traps. arXiv:2504.18634.
- KOLMES, E.J., OCHS, I.E., RAX, J.-M. & FISCH, N.J. 2024 Massive, long-lived electrostatic potentials in a rotating mirror plasma. *Nat. Commun.* **15**, 4302.
- KONKASHBAEV, I.K., LANDMAN, I.S. & ULINICH, F.R. 1978 Possibility of decreasing the electron heat flux from open traps. *Soviet Phys. JETP* **47**, 501.
- MARX, K.D. 1970 Effects of spatial variations on collisional losses in a mirror-confined plasma. *Phys. Fluids* **13**, 1355–1371.
- MATSUDA, Y. & STEWART, J. 1986 A relativistic multiregion bounce-averaged Fokker–Planck code for mirror plasmas. *J. Comput. Phys.* **66**, 197–217.
- MILLER, T., BE'ERY, I. & BARTH, I. 2021 Rate equations model for multiple magnetic mirrors in various thermodynamic scenarios. *Phys. Plasmas* **28**, 112506.
- MILLER, T., BE'ERY, I., GUDINETSKY, E. & BARTH, I. 2023 RF plugging of multi-mirror machines. *Phys. Plasmas* **30**, 072510.
- MIRNOV, V.V. & RIUTOV, D.D. 1979 Linear gasdynamic system for plasma confinement. *Tech. Phys. Lett+*, **5**, 279.
- MYNICK, H. & HITCHON, W. 1986 A bounce-averaged fokker–Planck code for stellarator transport. *Nucl. Fusion* **26**, 425–438.
- NAJMABADI, F., CONN, R. & COHEN, R. 1984 Collisional end loss of electrostatically confined particles in a magnetic mirror field. *Nucl. Fusion* **24**, 75–84.
- NEMOV, V.V., KASILOV, S.V., KERNBICHLER, W. & HEYN, M.F. 1999 Evaluation of $1/\nu$ neoclassical transport in stellarators. *Phys. Plasmas* **6**, 4622–4632.
- OCHS, I.E. 2024 Synchrotron-driven instabilities in relativistic plasmas of arbitrary opacity. *Astrophys. J.* **975**, 30.
- OCHS, I.E. & FISCH, N.J. 2023 Critical role of isopotential surfaces for magnetostatic ponderomotive forces. *Phys. Rev. E* **108**, 065210.
- OCHS, I.E., KOLMES, E.J. & FISCH, N.J. 2025 Preventing ash from poisoning proton-boron 11 fusion plasmas.
- OCHS, I.E., MLODIK, M.E. & FISCH, N.J. 2024 Electron tail suppression and effective collisionality due to synchrotron emission and absorption in mildly relativistic plasmas. *Phys. Plasmas* **31**, 083303.
- OCHS, I.E., MUNIROV, V.R. & FISCH, N.J. 2023 Confinement time and ambipolar potential in a relativistic mirror-confined plasma. *Phys. Plasmas* **30**, 052508.
- PASTUKHOV, V. 1974 Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential. *Nucl. Fusion* **14**, 3–6.
- POST, R. 1987 The magnetic mirror approach to fusion. *Nucl. Fusion* **27**, 1579–1739.
- RAX, J.M., FRUCHTMAN, A., GUEROULT, R. & FISCH, N.J. 2015 Breakdown of the Brillouin limit and classical fluxes in rotating collisional plasmas. *Phys. Plasmas* **22**, 092101.
- ROGNLIEN, T. & CUTLER, T. 1980 Transition from Pastukhov to collisional confinement in a magnetic and electrostatic well. *Nucl. Fusion* **20**, 1003–1011.
- ROSENBLUTH, M.N., MACDONALD, W.M. & JUDD, D.L. 1957 Fokker–Planck equation for an inverse-square force. *Phys. Rev.* **107**, 1–6.
- RUBIN, T. & FISCH, N.J. 2025 Ponderomotive barriers in rotating mirror devices using static fields. *Phys. Plasmas*. arXiv:2502.02008.

- RUBIN, T., OCHS, I.E. & FISCH, N.J. 2024 Flowing plasma rearrangement in the presence of static perturbing fields. *Phys. Plasmas* **31**, 082109.
- RUBIN, T., RAX, J.M. & FISCH, N.J. 2023 Magnetostatic ponderomotive potential in rotating plasma. *Phys. Plasmas* **30**, 052501.
- SCHWARTZ, N.R., ABEL, I.G., HASSAM, A.B., KELLY, M. & ROMERO-TALAMÁS, C.A. 2024 MCTrans++: a 0-d model for centrifugal mirrors. *J. Plasma Phys.* **90**, 905900217.
- SKOVORODIN, D.I. 2019 Suppression of secondary emission of electrons from end plate in expander of open trap. *Phys. Plasmas* **26**, 012503.
- STIX, T. 1975 Fast-wave heating of a two-component plasma. *Nucl. Fusion* **15**, 737–754.
- VELASCO, J., CALVO, I., PARRA, F. & GARCÍA-REGAÑA, J. 2020 KNOSOS: a fast orbit-averaging neoclassical code for stellarator geometry. *J. Comput. Phys.* **418**, 109512.
- YUSHMANOV, E.E. 1966 Confinement of slow ions of a plasma with positive potential in a mirror trap. *Soviet Phys. JETP* **22**, 409.
- ZHDANKIN, V., KUNZ, M.W. & UZDENSKY, D.A. 2023 Synchrotron firehose instability. *Astrophys. J.* **944**, 24.