

RESEARCH ARTICLE | OCTOBER 07 2025

Semi-implicit continuum kinetic modeling of weakly collisional parallel transport in a magnetic mirror

M. Dorf  ; M. Dorr  ; V. Geyko  ; D. Ghosh  ; M. Umansky  ; J. Angus 



Phys. Plasmas 32, 103903 (2025)
<https://doi.org/10.1063/5.0293154>

 CHORUS



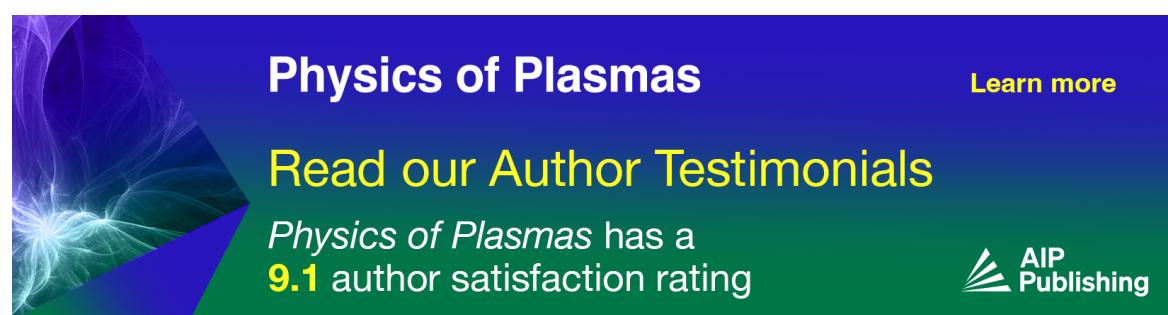
Articles You May Be Interested In

Implicit full-F simulations of neoclassical ion transport

Phys. Plasmas (August 2025)

A non-neutral generalized Ohm's law model for magnetohydrodynamics in the two-fluid regime

Phys. Plasmas (May 2023)



Physics of Plasmas [Learn more](#)

Read our Author Testimonials

Physics of Plasmas has a **9.1** author satisfaction rating



Semi-implicit continuum kinetic modeling of weakly collisional parallel transport in a magnetic mirror

Cite as: Phys. Plasmas **32**, 103903 (2025); doi: [10.1063/5.0293154](https://doi.org/10.1063/5.0293154)

Submitted: 25 July 2025 · Accepted: 25 September 2025 ·

Published Online: 7 October 2025



View Online



Export Citation



CrossMark

M. Dorf,^{a)} M. Dorr, V. Geyko, D. Ghosh, M. Umansky, and J. Angus

AFFILIATIONS

Lawrence Livermore National Laboratory, Livermore, California 94550, USA

^{a)}Author to whom correspondence should be addressed: dorfl@llnl.gov

ABSTRACT

We present implicit-explicit (IMEX) kinetic simulations of weakly collisional parallel plasma transport in magnetic mirror configurations using the continuum code COGEN. The numerical scheme employs a Jacobian-free Newton–Krylov method with algebraic multigrid preconditioning to overcome the severe time step limitations imposed by strong mirror forces in fully explicit schemes. Applied to parameters relevant to the Wisconsin HTS Axisymmetric Mirror experiment, the IMEX approach enables time steps up to 2.5×10^4 times larger than those permitted by explicit methods, resulting in a $2500 \times$ speedup in 1D–2V simulations of parallel transport with kinetic ions and Boltzmann electrons. Additionally, a reduced bounce-averaged model for a square mirror is implemented to support the computationally intensive fully kinetic simulations. The bounce-averaged formulation is used to evaluate the numerical convergence of the velocity-space discretization algorithms and to assess the role of the collision model by comparing simulations employing the nonlinear Fokker–Planck and the simplified Lenard–Bernstein–Dougherty collision operators.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC BY-NC-ND) license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>). <https://doi.org/10.1063/5.0293154>

I. INTRODUCTION

In the past several decades, magnetic fusion research has been heavily focused on large-scale toroidal confinement devices like tokamaks¹ and stellarators.² However, recent technological advances including the development of high-temperature superconducting (HTS) magnets appropriate for fusion applications,³ together with improvements in the theoretical understanding of magnetized plasma stability and transport, have enabled researchers to explore and experiment with reemerging alternative fusion concepts such as magnetic mirrors.^{4–7} Being more compact and thereby less expensive, magnetic mirrors attract particular attention from academia and private investors. A notable example of a modern mirror experiment is the Wisconsin HTS Axisymmetric Mirror⁴ (WHAM), which explores the benefits of strong magnetic fields with high mirror ratio for the classical weakly collisional confinement approach.

The success of alternative magnetized fusion concepts, such as magnetic mirrors, relies heavily on advanced modeling capabilities that can provide scientific insights and assess scalability to reactor parameters. This motivates the adaptation of advanced gyrokinetic computational tools, which have been extensively developed over the

past several decades for the mainstream tokamak approach, to mirror applications. Although modeling certain mirror-specific high-frequency processes, e.g., the drift-cyclotron loss-cone instability, requires resolving cyclotron ion motion,⁸ the analysis of lower-frequency processes with $\omega \ll \omega_{ci}$, such as collisional transport and gradient-driven drift modes, can be facilitated by employing the gyro-averaging formalism. Here, ω_{ci} denotes the ion cyclotron frequency. A recent study by Francisquez *et al.*⁹ outlines the challenges faced by continuum gyrokinetic codes in modeling HTS mirrors and demonstrates the application of the gyrokinetic code Gkeyll to simulate 1D–2V parallel transport in mirror geometry for parameters characteristic of the WHAM experiment.⁴ We also note 2D–2V gyrokinetic simulations,¹⁰ performed with the particle-in-cell GTC-X code, that investigate the penetration of divertor electrostatic biasing in the scrape-off layer of a field-reversed configuration, which shares linear-geometry features with magnetic mirrors.

Although the presence of an open-field line region in a tokamak edge implies a certain level of geometrical similarity with a magnetic mirror, the physics parameters can, nevertheless, be quite different. Classical mirror plasmas are significantly less collisional than the edge

of a tokamak. For instance, for the parameters characteristic of a mid-size tokamak (e.g., DIII-D) edge,¹¹ $n_i \sim 10^{19} \text{ m}^{-3}$, $T_i \sim 100 \text{ eV}$, $L_{\parallel} \sim 6 \text{ m}$, $m_i = 2m_p$, the Knudson number, measuring the ratio of the ion-ion mean free path, λ_{ii} , to a length scale in variation of background plasma profiles along the magnetic field, L_{\parallel} , is given by $K_{\text{DIII-D}} \sim 3$. Here, n_i is the ion density, T_i is the ion temperature, and m_i and m_p denote the ion and proton mass, respectively. In contrast, the WHAM mirror parameters, $n_i \sim 3 \times 10^{19} \text{ m}^{-3}$, $T_i \sim 8 \text{ keV}$, $L_{\parallel} \sim 2 \text{ m}$, $m_i = 2m_p$, correspond to $K_{\text{WHAM}} \sim 1.5 \times 10^4$. This means that the ion transit time is four orders of magnitude shorter than the ion-ion collision time, thereby making the problem of collisional mirror transport intrinsically multiscale in time integration. It is also important to note that the presence of loss-cone regions and high-energy beams in the velocity phase space requires the use of *full-F* computational methods that include background plasma evolution.

The temporally multiscale nature of a weakly collisional mirror transport can be effectively handled by making use of the bounce-average formulation that averages over the fast ion transit time.^{4,5,12} Although this approach can significantly facilitate numerical simulations of plasma dynamics inside the mirror, it lacks the ability to efficiently describe the region outside the trap and the coupling of the confined mirror region to plasma-facing components. Moreover, the standard bounce-averaged formulation assumes a single trapped region—an assumption that is violated in tandem magnetic mirrors^{13,14} and even in simple mirrors with sloshing ion distributions,¹⁵ where electrons can become trapped in local potential peaks near ion beam turning points. Therefore, first-principles gyrokinetic simulations are required for high-fidelity integrated modeling.

It is interesting to note that the ion transit time, $\tau_{\parallel,i} \sim 50$, and the electron transit time, $\tau_{\parallel,e} \sim 1$, at the tokamak edge are both much shorter than the characteristic time of the anomalous radial plasma transport, $\tau_{\text{transp}} \sim 1 \text{ ms}$. Furthermore, full-*F* simulations are required to describe substantial deviations of a plasma distribution function from a local Maxwellian. Nevertheless, modern HPC gyrokinetic codes are capable of performing millisecond full-*F* simulations of 2D-2V and even 3D-2V plasma transport using straightforward explicit time integration schemes.¹⁶ In contrast to tokamaks, mirrors are distinguished by the presence of a strong mirror force, $F_{\nabla B} = -\mu \nabla B$, which can impose a stringent CFL limit on the size of a stable time step for explicit time integration: $\Delta t_{\nabla B}^{\text{CFL}} < m_i \Delta v_{\parallel} / (\mu \nabla B)$. Here, Δv_{\parallel} is the cell size in the parallel velocity direction, $\mu = m_i v_{\perp}^2 / (2B)$ is the magnetic moment, and for a typical WHAM simulation, $\Delta t_{\nabla B} \sim 0.1 \text{ ns}$. As a result, explicit time integration of a weakly collisional mirror plasma on the ion-ion collisional timescale (tens of milliseconds) becomes prohibitively expensive. Recent explicit numerical simulations of 1D-2V collisional transport performed with the Gkeyll code required approximately 74 h of wall time on 288 cores to simulate 72 μs of plasma dynamics involving kinetic ions and Boltzmann electrons.⁹ Simulations with kinetic electrons, which are also considered in that work, are significantly more expensive due to stiff electron time scales. Although a substantial, 20 \times , speedup was demonstrated by introducing *ad hoc* softening of the ∇B force, explicit simulations spanning tens of milliseconds remain prohibitively costly for explicit methods.

Here, the feasibility of multiscale collisional mirror transport simulations is explored by making use of the implicit-explicit (IMEX) time integration framework implemented in the continuum

gyrokinetic code COGENT, originally developed for tokamak edge modeling.¹⁷ The IMEX time integration algorithm used in COGENT¹⁸ is based on semi-implicit additive Runge–Kutta (ARK) methods¹⁹ and can provide consistent high-order time integration, including implicit treatment of selected stiff terms. It employs the Jacobian-free Newton–Krylov (JFNK) approach²⁰ to handle nonlinearities and utilizes preconditioning to improve convergence properties. Declaring the gyrokinetic Vlasov term as implicit enables stable time integration with stepping over the stiff time scales including advection, i.e., acceleration, in the parallel velocity direction due to strong mirror force and the parallel streaming. On the other hand, efficient preconditioning of the Vlasov operator (even in a simple case with a fixed advection velocity) requires solving a linear system in the full phase space dimension with a nonsymmetric and indefinite coefficient matrix. COGENT solves this linear system using the Approximate Ideal Restriction (AIR) option in the BoomerAMG algebraic multigrid solver contained in the Hypre linear solver library.²¹ The use of multigrid methods to solve nonsymmetric indefinite systems has historically been highly problematic, but the recent development of the AIR approach,^{22,23} including several variants, provides a way to extend the benefits of multigrid algorithms beyond the symmetric, positive-definite systems for which they are more commonly used. To further enhance the efficiency of the implicit Vlasov solver, COGENT allows the preconditioner to be defined using a lower-order discretization (first-order upwind, UW1) than that used for the Vlasov operator itself (fifth-order upwind, UW5). The low-order preconditioner yields a sparser matrix, for which robust AMG solver performance is observed, while also providing good efficiency even for relatively large time steps.

The IMEX COGENT time integration scheme is applied to 1D-2V simulations of parallel collisional transport using parameters characteristic of the WHAM experiment. The simulations employ the kinetic Vlasov equation for the ion species coupled with the nonlinear ion-ion Fokker–Planck collisions and a high-energy beam source. A Boltzmann electron model is used to describe self-consistent electrostatic potential variations. The implicit simulations enable time steps up to 2.5×10^4 larger than the CFL-limited time step and demonstrate speedups of about 2500 \times and 1800 \times compared to the corresponding explicit COGENT simulations and previously reported explicit Gkeyll simulations with Boltzmann electrons,⁹ respectively. Additionally, a fully implicit, simplified bounce-averaged (0D-2V) model for a square mirror (the BASM model) is implemented in COGENT to provide insights into the more computationally expensive 1D-2V simulations. The BASM model is verified by reproducing analytic results for the Pastukhov problem^{24,25} and is used to examine the numerical convergence of the velocity-space discretization algorithms, which are challenged by the presence of steep gradients at the loss-cone boundary for a weakly collisional mirror. The BASM model is also used to assess the role of the ion-ion collision model by comparing simulations using the Fokker–Planck operator²⁶ with those using the Lenard–Bernstein–Dougherty (LBD) model,^{27–29} which is often employed in full-*F* gyrokinetic modeling due to its simplicity.

This paper is organized as follows: the results of numerical simulations obtained using the simplified bounce-averaged model and the fully kinetic model are presented in Secs. II and III, respectively, and Sec. IV provides the conclusions.

II. BOUNCE-AVERAGED (0D-2V) SIMULATIONS

A. Simulation model

The bounce-averaged square-mirror (BASM) model describes a weakly collisional plasma confined in a square magnetic mirror. Under these assumptions, the kinetic distribution function for a plasma species can be considered uniform along the magnetic field, $f = f(v_{\parallel}, \mu)$, and the governing kinetic equation takes the following form:

$$\frac{\partial f}{\partial t} = C[f] + \Gamma_{src}[f] - \Gamma_{sink}[f]. \quad (1)$$

Here, $C[f]$ is the collision operator, Γ_{src} represents a source term, e.g., high-energy beams, and

$$\Gamma_{sink}[f] = (v_{\parallel}/L_{\parallel})f, \quad (2)$$

describes prompt parallel losses from the loss-cone region, e.g., due to collisional scattering. Here, m , v_{\parallel} , and $\mu = mv_{\perp}^2/(2B)$ denote the particle mass, parallel velocity, and magnetic moment, respectively; L_{\parallel} is the trap length, and the weakly collisional assumption implies that $\lambda \gg L_{\parallel}$, where λ is the collisional mean free path. The sink term in Eq. (2) is applied to the velocity-space region outside the loss-cone boundary given by

$$mv_{\parallel}^2/2 \geq \mu B_0(R_m - 1) + q\Phi_m. \quad (3)$$

Here, q is the particle charge, B_0 is the magnetic field at the mirror center, R_m is the mirror ratio, defined as the ratio of the maximum magnetic field (at the mirror throat) to the minimum field (at the mirror center), and Φ_m is the confining electrostatic potential. The boundary in Eq. (3) cuts through the cells of the (v_{\parallel}, μ) velocity grid (see Fig. 1), and the following numerical implementation of the sink term is adopted:

$$\Gamma_{i,j}^{sink} = \frac{v_{\parallel,i}}{L_{\parallel}} \times f_{i,j} \times \frac{dS_{i,j}^{loss}}{dv_{\parallel}d\mu}, \quad (4)$$

where $dS_{i,j}^{loss}$ denotes the area of a phase-space cell in the loss-cone region. For the case where the loss-cone boundary intersects a cell

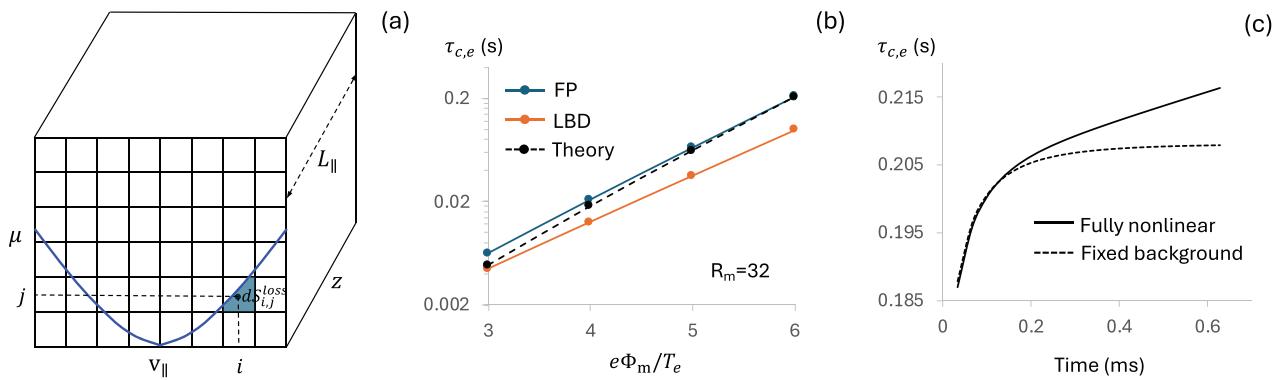


FIG. 1. Bounce-averaged (0D-2V) simulations. (a) Schematic of the bounce-averaged square-mirror (BASM) model. The blue curve illustrates the loss-cone boundary for $\Phi_m = 0$. For the case where the loss-cone boundary intersects a cell (i, j) , the shaded area, $dS_{i,j}^{loss}$, is used in Eq. (4). (b) Collisional losses of the electrostatically confined electron species for $R_m = 32$, $n = 10^{19} \text{ m}^{-3}$, and $T_e = 940 \text{ eV}$. The simulation results for the electron confinement time, $\tau_{c,e} = n_e |dn_e/dt|^{-1}$, obtained using the FP (blue) and LBD (orange) e-e collision models are compared with the analytical prediction in Eq. (11). (c) Time history of the electron confinement time obtained using the full nonlinear FP model and the fixed-background FP model for $e\Phi_m/T_e = 6$.

(i, j) , a second-order accurate approximation is used to compute $dS_{i,j}^{loss}$ [see Fig. 1(a)]. For all other cells, $dS_{i,j}^{loss} = dv_{\parallel}d\mu$.

The collision models considered in this work for like-species collisions include the nonlinear Fokker-Planck (FP) model described elsewhere²⁶ and the reduced Lenard-Bernstein-Dougherty (LBD) model,²⁹ which is often employed in full-F gyrokinetic modeling due to its simplicity,

$$C_{LBD}[f] = \nu \frac{\partial}{\partial v_{\parallel}} \left[(v_{\parallel} - U_{\parallel})f + \frac{T}{m} \frac{\partial f}{\partial v_{\parallel}} \right] + \nu \frac{\partial}{\partial \mu} \left[2\mu f + 2 \frac{T}{B} \mu \frac{\partial f}{\partial \mu} \right]. \quad (5)$$

Here, the particle number density (n), parallel flow velocity (U_{\parallel}), and the temperature (T) moments are defined as

$$n = \frac{2\pi}{m} \int f B dv_{\parallel} d\mu, \quad (6)$$

$$U_{\parallel} = \frac{2\pi}{nm} \int v_{\parallel} f B dv_{\parallel} d\mu, \quad (7)$$

$$T = \frac{2\pi}{3n} \int \left[(v_{\parallel} - U_{\parallel})^2 + \frac{2\mu B}{m} \right] f B dv_{\parallel} d\mu. \quad (8)$$

The collision frequency in Eq. (5) is given by

$$\nu = \frac{4\pi^{1/2} n Z^4 e^4}{3m^{1/2} T^{3/2}} \ln \Lambda, \quad (9)$$

such that the LBD collision model matches the temperature relaxation rate of the Boltzmann operator.²⁸ Here, Z is the species charge state and e is the elementary charge. The Coulomb logarithm used for the LBD and FP models is defined in the present work as

$$\ln \Lambda = 23 - \ln \left(2^{1/2} Z^3 n_0^{1/2} T_0^{-3/2} \right), \quad (10)$$

where $n_0(\text{cm}^{-3})$ and $T_0(\text{eV})$ are the density and temperature normalization factors used in the simulations, respectively.

COGENT employs an implicit-explicit (IMEX) time integration algorithm¹⁸ based on semi-implicit additive Runge–Kutta (ARK) methods,¹⁹ which enables consistent high-order time integration and allows for the implicit treatment of selected stiff terms. Nonlinearities are addressed using the Jacobian-free Newton–Krylov (JFNK) approach,²⁰ with preconditioning applied to enhance convergence properties. The preconditioning of the stiff collision terms is developed, assuming that the Rosenbluth potentials for the FP operator and the distribution function moments (i.e., n , U_{\parallel} , T) in the LBD model [see Eq. (5)] remain fixed during a time step update. Under these assumptions, the collisional operator reduces to a linear advection-diffusion operator, for which a sparse matrix is assembled, and the Gauss–Seidel method is used to solve the resulting linear system. To improve consistency between the preconditioner and the full collision operator, thereby enhancing overall simulation performance, the same “fixed-moment” assumption is applied when evaluating the actual collision term in Eq. (1). Namely, although a second-order time integration scheme is employed, the moments of the distribution function are updated only once at the end of each time step. While such an approximation reduces the formal time integration accuracy, these effects might be minimal given that the simulations target long transport timescales. Finally, a second-order discretization in velocity space is employed for the FP and LBD collision operators.

B. Verification: Collisional mirror losses

The BASM model is verified by analyzing collisional losses of electrostatically confined species in a square magnetic mirror for a weakly collisional case (see Fig. 1). Here, an electron species is considered due to the relative simplicity of electron-ion collisions, which can be described using the Lorentz collision model. Electron–electron collisions are modeled with the full nonlinear Fokker–Planck operator. Good agreement [see Fig. 1(b)] is obtained with the approximate analytic solution in a weakly collisional regime with $R_m \gg 1$ and $\Phi_m/eT_e \gg 1$, originally derived by Pastukhov²⁴ and later corrected by Cohen,²⁵

$$\frac{dn_e}{dt} = -\frac{4}{\sqrt{\pi}} \frac{n_e \nu_e}{G(R_m)} \frac{T_e}{e\Phi_m} \exp\left(-\frac{e\Phi_m}{T_e}\right) I\left(\frac{T_e}{e\Phi_m}\right), \quad (11)$$

where $I(x)$ is related to the error function

$$I(x) = 1 + \frac{1}{2}(\pi x)^{1/2} e^{1/x} [1 - \text{erf}(x^{-1/2})], \quad (12)$$

and where

$$G(R_m) \approx \frac{2R_m + 1}{2R_m} \ln(4R_m + 2), \quad (13)$$

for $R_m \gg 1$. The electron collision frequency in Eq. (11) is given by

$$\nu_e = \frac{\sqrt{2\pi} n_e e^4}{m_e^{1/2} T_e^{3/2}} \ln \Lambda. \quad (14)$$

For comparison, the electron confinement time, $\tau_{c,e} = n_e |dn_e/dt|^{-1}$, is also computed using the LBD model for electron–electron collisions, which is found to substantially increase particle losses. This result is consistent with the recent work in Ref. 30, where a method to improve the LBD model for more accurate loss predictions is proposed.

TABLE I. Reference parameters used in COGENT simulations.

Parameter	Description	Value
n_0	Initial plasma density	10^{19} m^{-3}
T_0	Initial ion temperature	8.361 keV
T_e	Electron temperature	940 eV
m_i	Ion mass	$2m_p$
R_m	Mirror ratio	32
L_{\parallel}	Trap length	2 m
B_0	Magnetic field at the mirror center	0.5 T
E_b	Ion beam energy	25 keV
θ_{inj}	Beam injection angle	45°
Γ_b	Ion beam intensity	$2.3 \times 10^{15} \text{ s}^{-1} \text{ cm}^{-3}$

The simulations illustrated in Fig. 1 use plasma and magnetic mirror parameters inspired by recent numerical studies of the WHAM experiment^{4,9} (see Table I). The electron species is initialized with a Maxwellian distribution characterized by $n_0 = 10^{19} \text{ m}^{-3}$ and $T_e = 940 \text{ eV}$. The magnetic mirror ratio is $R_m = 32$, and the length of the magnetic trap is $L_{\parallel} = 2 \text{ m}$, corresponding to a weakly collisional regime with $\nu_e V_{Te}/L_{\parallel} = 2.6 \times 10^{-3}$. Here, $V_{Te} = \sqrt{2T_e/m_e}$ is the electron thermal velocity, and m_e is the electron mass. Note that for the 0D-2V BASM model, the trap length parameter enters the simulation only through the sink term in Eq. (2), where it controls the rate of prompt particle losses from the loss-cone region. No sources are included in the simulations, $\Gamma_{src} = 0$, and the electron distribution function evolves according to Eq. (1). The velocity grid resolution is given by $N_{v_{\parallel}} = 1024$ and $N_{\mu} = 1024$, and the velocity-domain extent corresponds to $-5V_{Te} \leq v_{\parallel} \leq 5V_{Te}$ and $0 \leq \mu \leq 12T_0/B_0$, where B_0 is the magnetic field at the trap center. The time step is set to $\Delta t = 3.3 \mu\text{s}$, which corresponds to a collisional CFL number of 1.5×10^3 for the Fokker–Planck operator. Numerical convergence of the simulation results with respect to the grid resolution, domain extent, and time step is confirmed.

Figure 1(c) shows the time history of the electron confinement time. Simulation results obtained using the full nonlinear FP operator for e–e collisions are compared with those from counterpart simulations, where e–e collisions are modeled as scattering off a fixed Maxwellian background corresponding to the initial electron distribution. Following initial agreement over a shorter (collisional) timescale, the full FP model yields an increasing confinement time compared to the fixed-background model over a longer (transport) timescale. This behavior can be explained as follows: particles that are scattered in the loss-cone region and lost from the trap have significantly higher-than-average energy for $e\Phi_m/T_e \gg 1$. As the energy-conserving collisions modeled by the full FP operator thermalize the electron distribution, its temperature decreases. This leads to improved electrostatic confinement, quantified by $e\Phi_m/T_e$, and thus to an increase in the confinement time. We also note small changes in the characteristic electron collision frequency in Eq. (14) due to the decreasing electron density and temperature. However, these effects are subdominant compared to the exponential sensitivity of the confinement time to the electrostatic potential. In contrast, when e–e collisions are modeled as collisions with a fixed Maxwellian background, energy is not conserved, and the electron temperature remains close to that of the background

distribution, i.e., the initial Maxwellian. Additionally, the density in Eq. (14) corresponds to the background Maxwellian and, therefore, remains fixed. As a result, a steady-state solution for $\tau_{c,e}$ is observed. It is instructive to note that the approximate analytical solution in Eq. (11) assumes collisions with a fixed Maxwellian background. In this context, Fig. 1(b) presents numerical simulation results obtained using the fixed-background model for both the FP and LBD e-e collision operators. In the latter case, the fixed-background model assumes constant density and temperature in Eqs. (5) and (9) corresponding to the initial Maxwellian.

C. High-energy beam relaxation

The BASM model is used here to provide insight into numerical simulations of high-energy beam fueling of a magnetic mirror system. In particular, we use it to examine the numerical convergence of the velocity-space discretization algorithms, which are challenged by the presence of steep gradients at the loss-cone boundary for a weakly collisional mirror. For simplicity, only a deuterium background distribution and beam ions are included in the simulation, and we omit the effects of ion-electron (cooling) collisions and self-consistent electric fields. As an illustrative example, we consider physical parameters inspired by recent numerical studies of the WHAM experiment:^{4,9} $B_0 = 0.5$ T, $R_m = 32$, $L_{\parallel} = 2$ m, $n_0 = 10^{19}$ m⁻³, $T_0 = 8.361$ keV, $m_i = 2m_p$, $E_b = 25$ keV, and $\theta_{\text{inj}} = 45^\circ$. Here, n_0 and T_0 correspond to the ion density and temperature of the initial Maxwellian background distribution, respectively, E_b is the energy of beam ions, and θ_{inj} is the beam injection angle. Beam fueling is described by the source term on the right-hand side of Eq. (1), where we arbitrarily adopt the following shape function:

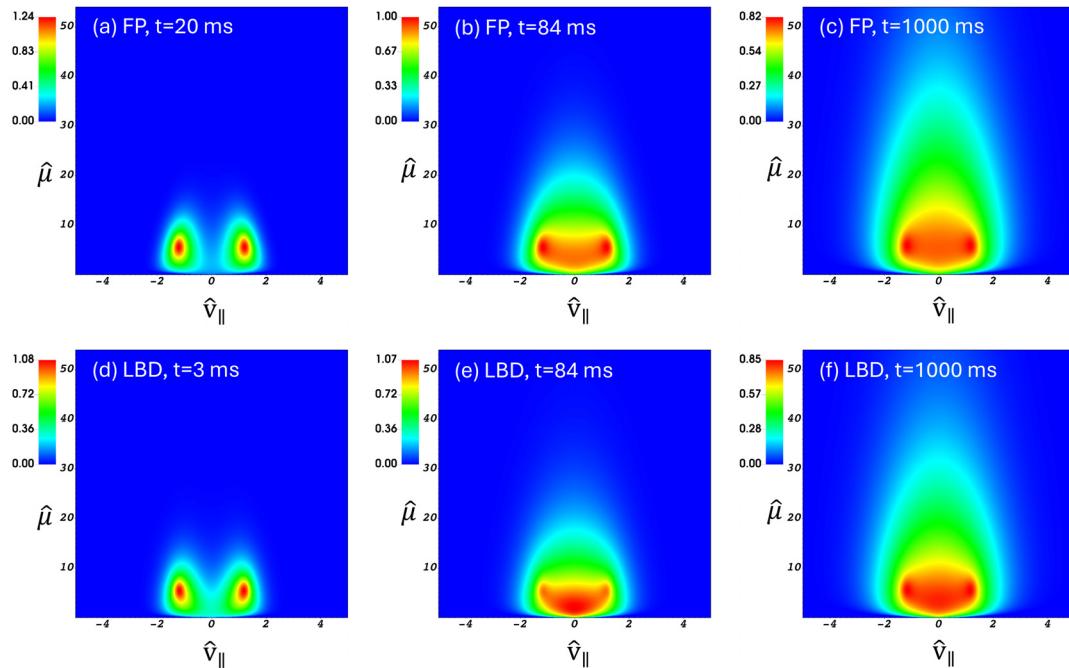


FIG. 2. High-energy beam relaxation obtained from 0D-2V BASM model simulations. The numerical results using the Fokker-Planck collision model [frames (a)–(c)] are compared with corresponding simulations using the LBD collision model [frames (d)–(f)]. Plotted is the normalized ion distribution function, $\hat{f}_i = (2T_0/m_i)^{3/2} \pi f_i/n_0$.

$$\Gamma_{\text{src}} = \Gamma_0 \exp \left(-\frac{(v_{\parallel} - V_{b,\parallel})^2 + \left(\sqrt{\frac{2\mu B_0}{m_i}} - V_{b,\perp} \right)^2}{2T_b/m_b} \right) + \Gamma_0 \exp \left(-\frac{(v_{\parallel} + V_{b,\parallel})^2 + \left(\sqrt{\frac{2\mu B_0}{m_i}} - V_{b,\perp} \right)^2}{2T_b/m_b} \right). \quad (15)$$

Here, $V_{b,\parallel} = V_{b,\perp} = \sqrt{E_b/m_i}$, $T_b = 200$ eV, and the normalization constant Γ_0 is determined such that the beam intensity, $\Gamma_b = (2\pi/m_i) \int \Gamma_{\text{src}} B dv_{\parallel} d\mu = 2.3 \times 10^{15}$ s⁻¹ cm⁻³, corresponds to a plasma density of approximately 3×10^{19} m⁻³ after 10 ms consistent with the analysis in Ref. 4.

Figures 2 and 3 show the results of numerical simulations obtained using the full nonlinear Fokker-Planck (FP) operator and the model LBD operator to describe ion-ion collisions. A nearly steady-state particle density is reached after approximately 1000 ms [see Fig. 3(a)], as beam fueling is balanced by collisional losses, i.e., scattering into the loss-cone region. A pronounced difference in the phase-space solution is evident between the collisional models, especially during the evolution stage [c.f., Figs. 2(b) and 2(e)], emphasizing the importance of using a detailed collision operator.

For the simulations illustrated in Figs. 2 and 3, the velocity grid resolution is given by $(N_{v_{\parallel}}, N_{\mu}) = (512, 768)$, and the velocity-domain extent corresponds to $-5 \leq \hat{v}_{\parallel} \leq 5$ and $0 \leq \hat{\mu} \leq 54$, where $\hat{v}_{\parallel} = v_{\parallel}/V_{T0}$, $\hat{\mu} = 4B_0\mu/T_0$, and $V_{T0} = \sqrt{2T_0/m_i}$. Numerical convergence studies with respect to velocity grid resolution are shown in Figure 4 for the case of the Fokker-Planck collisional model. Figure 4(a)

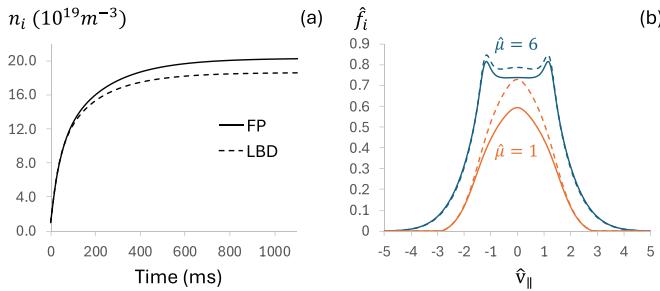


FIG. 3. Comparison of the Fokker-Planck (solid curves) and LBD (dashed curves) collisional models in the 0D-2V BASM-model simulations. Shown are (a) time history of the ion density and (b) the v_{\parallel} -lineouts extracted from the normalized ion distribution function, $\hat{f}_i = (2T_0/m_i)^{3/2} \pi \hat{f}_i / n_0$, for $\hat{\mu} = 1$ and $\hat{\mu} = 6$, at 1000 ms.

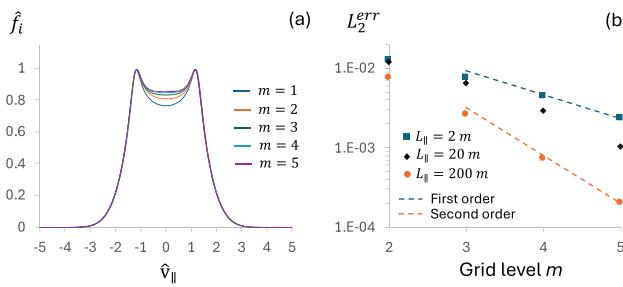


FIG. 4. Numerical convergence of the 0D-2V BASM model solution at 84 ms. Shown are (a) the v_{\parallel} -lineouts extracted from the normalized ion distribution function, $\hat{f}_i = (2T_0/m_i)^{3/2} \pi \hat{f}_i / n_0$, for $\hat{\mu} = 6$; and (b) Richardson extrapolation analysis of the numerical convergence for different values of the trap length, L_{\parallel} . Grid levels $m = (2, 3, 4, 5)$ correspond to the successful refinements of the coarse grid $(N_{v_{\parallel}}, N_{\mu}) = (128, 192)$ by factors of 2, 4, 8, and 16 in both velocity dimensions, respectively.

illustrates a lineout of the normalized distribution function, $\hat{f}_i = (2T_0/m_i)^{3/2} \pi \hat{f}_i / n_0$, along the v_{\parallel} -coordinate, obtained from simulations using a coarse velocity grid $N_{v_{\parallel}} = 128$ and $N_{\mu} = 192$ and its successful refinements by factors of 2, 4, 8, and 16 in both velocity dimensions. Although convergence of the numerical solution is apparent, the comparison of values at $v_{\parallel} = 0$ suggests that the convergence is only first-order, despite the use of second-order discretization in the numerical implementation of Eq. (1). First-order convergence is supported by a more rigorous analysis using the Richardson extrapolation method. Figure 4(b) plots the L^2 norm of the difference between the numerical solution \hat{f}_i obtained at a given grid spacing, h , and that at half the spacing, $h/2$, i.e., $L_2^{err} = \|\hat{f}_h - \hat{f}_{h/2}\|$.

The degraded convergence can be explained as follows. For a weakly collisional regime, $\nu_{ii}L_{\parallel}/V_T \ll 1$, the solution to Eq. (1) is distinguished by the presence of a narrow-layer region at the loss-cone boundary that mediates the transition between the confined and loss-cone regions. Here, V_T is the characteristic thermal velocity. The characteristic width of this region, δv , is determined from the balance between the sink term (prompt losses) and the collisional diffusion, $C_{ii}[f] \sim \Gamma_{sink}[f]$, and is given by

$$\delta v \sim (V_T \nu_{ii} L_{\parallel})^{1/2}. \quad (16)$$

If the narrow layer in Eq. (16) is not sufficiently resolved by the computational grid, degradation in numerical convergence properties can be expected. In the asymptotic limit, a weakly collisional solution can be obtained by imposing a zero Dirichlet boundary condition directly at the loss-cone boundary. However, not resolving the transition length effectively corresponds to imposing a zero boundary condition at the nearest cell center below the loss-cone boundary, rather than at its actual location. A cell-size error in the boundary condition specification can degrade convergence to first-order accuracy. It is instructive to note that embedded boundary methods^{31,32} can provide high-order accurate solutions for weakly collisional formulations that impose a zero Dirichlet boundary condition at a loss-cone boundary intersecting the grid cells.

For the numerical solution in Fig. 2(b), which is used for the convergence studies in Fig. 4, we have $T = 21$ keV, $n_i = 12 \times 10^{19} \text{ m}^{-3}$, and $\delta v/V_T \sim 6.3 \times 10^{-3}$, where Eq. (9) is used to estimate the ion-ion collision frequency in Eq. (16). At the same time, our most refined grid $(N_{v_{\parallel}}, N_{\mu}) = (2048, 3072)$ with $\Delta v_{\parallel}/V_T = 3.0 \times 10^{-3}$ and $\Delta v_{\perp}/V_T \approx 5.0 \times 10^{-3}$ does not sufficiently resolve the transition layer; therefore, degraded convergence may be expected. Here, $V_T = \sqrt{2T/m_i}$, and we used $\Delta\mu = mv_{\perp}\Delta v/B_0 \approx mV_T\Delta v/\sqrt{R_m}B_0$ at the loss-cone boundary. To further investigate this issue, we consider additional simulations in which the sink term is arbitrarily reduced by factors of 10 and 100 [see Fig. 4(b)]. For instance, this can be achieved by increasing the trap length, which leads to a higher collisionality parameter, $\nu_{ii}L_{\parallel}/V_T$, and a wider transition layer, δv . Second-order numerical convergence is readily observed in the case with the highest collisionality. For the medium-collisionality case, the convergence improves from first- to second-order as the grid resolution increases, and the transition layer becomes resolved.

III. FULLY KINETIC (1D-2V) SIMULATIONS

In this section, we investigate the performance of our implicit numerical algorithms for the modeling of parallel plasma transport in a weakly collisional magnetic mirror. As an illustrative example, we consider the problem of high-energy beam fueling of a WHAM mirror system [see Sec. II C] that now includes a magnetic field profile, ion parallel streaming, and a self-consistent parallel electric field determined from the Boltzmann electron model.

A. Simulation model

Assuming the paraxial approximation with $B_z \gg B_r$ and $\nabla_{\parallel} \approx \partial/\partial z$, the time evolution of the ion distribution function, $f_i(z, v_{\parallel}, \mu)$, within a magnetic flux tube is governed by

$$\frac{\partial f_i}{\partial t} + v_{\parallel} \frac{\partial f_i}{\partial z} + \frac{\partial}{\partial v_{\parallel}} \left(\left[\frac{q_i}{m_i} E_z - \frac{\mu}{m_i} \frac{\partial B_z}{\partial z} \right] f_i \right) = C_{ii}[f_i] + \Gamma_{src}, \quad (17)$$

where E_z is the parallel electric field, and q_i and m_i denote the ion species charge and mass, respectively. In deriving Eq. (17), we used the fact that the Jacobian of the transformation to (v_{\parallel}, μ) velocity coordinates, $J_v = (2\pi/m_i)B$, cancels the spatially varying part of the Jacobian associated with a magnetic-flux-surface coordinate, $J_x \propto 1/B$. As a result, their product reduces to a constant factor, $J_v J_x \propto 2\pi/m_i$. The full nonlinear Fokker-Planck collision model is used for like-species ion-ion collisions, $C_{ii}[f]$. Collisions with electrons are not included in the present work, for simplicity. The electric field is obtained from the quasi-neutrality condition coupled to the Boltzmann electron model,

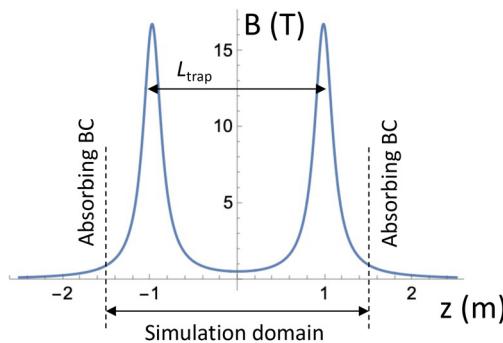


FIG. 5. Magnetic geometry used in 1D-2V fully kinetic simulations.

$$E_z = -\frac{T_e}{en_i} \frac{\partial n_i}{\partial z}, \quad (18)$$

where $n_i = (2\pi/m_i) \int dv_{\parallel} d\mu f_i B$ is the ion density and T_e is a uniform electron temperature.

An analytical approximation to the WHAM magnetic geometry (see Fig. 5) follows the approach in Ref. 9,

$$B_z(z) = \frac{\bar{B}}{\pi\gamma} \left\{ \left[1 + \left(\frac{Z - Z_m}{\gamma} \right)^2 \right]^{-1} + \left[1 + \left(\frac{Z + Z_m}{\gamma} \right)^2 \right]^{-1} \right\}, \quad (19)$$

where $\bar{B} = 6.5$ T, $\gamma = 0.124$, and $Z_m = 0.98$ m. The simulation domain extent is given by $-1.5 \text{ m} < z < 1.5 \text{ m}$, and a particle-absorbing boundary condition is imposed in the z -direction. The source function in Eq. (17) follows our choice in Eq. (15), where a spatially uniform normalization constant, Γ_0 , is now replaced with a spatially varying function $\Gamma_0 \exp(-z^2/L_b^2)$, where $L_b = 0.2$ m. The normalization constant Γ_0 is determined from the condition that it matches the particle production rate used in the BASM simulations in Sec. II C, i.e., $\Gamma_b = \int J_v J_x \Gamma_{src} dv_{\parallel} d\mu dz / \int J_x dz = 2.3 \times 10^{15} \text{ s}^{-1} \text{ cm}^{-3}$. Here, the integral along the z -coordinate is evaluated between the mirror throat points. The deuterium ion species, $m_i = 2m_p$, is initialized with a Maxwellian distribution, which has a uniform temperature $T_0 = 8.361$ keV and a nearly flat-top density profile with a steep decrease near the magnetic plugs, $n_i(t=0) = \bar{C} n_0 [\tanh((z_0 - z)/\bar{L}) + \tanh((z_0 + z)/\bar{L})]$, where $n_0 = 10^{19} \text{ m}^{-3}$, $z_0 = 0.75$ m, $\bar{L} = 0.1$ m, and the normalization coefficient $\bar{C} = 0.52$ is chosen to match the initial number of particles in the BASM simulations in Sec. II C. Following the simulations in Ref. 9, we set $T_e = 940$ eV.

Due to the substantial computational cost of the numerical simulations, we run them only for 20 ms, which roughly corresponds to the duration of the WHAM beam source. The results from the 0D-2V BASM model show that the phase-space solution is much more compact in velocity space at $t = 20$ ms compared to the final steady-state solution [cf. Figs. 2(a) and 2(c)]. Therefore, a smaller velocity-domain extent is used for the 1D-2V simulations: $-2.5 \leq \hat{v}_{\parallel} \leq 2.5$ and $0 \leq \hat{\mu} \leq 18$, where $\hat{v}_{\parallel} = v_{\parallel}/V_{T0}$, $\hat{\mu} = 4B_0\mu/T_0$, $B_0 = 0.5$ T, and $V_{T0} = \sqrt{2T_0/m_i}$. The baseline grid resolution is given by $(N_z, N_{v_{\parallel}}, N_{\mu}) = (256, 128, 192)$.

To illustrate the strong disparity in time scales, we consider the initial background plasma parameters as a reference, for which the

characteristic ion transit time is given by $\tau_{\parallel} = L_{\text{trap}}/V_{T0} = 2.2 \mu\text{s}$, while the characteristic ion-ion collision time is $\tau_{ii} = 1/\nu_{ii} = 100$ ms. Here, L_{trap} is the distance between the mirror throat points, and the collision frequency ν_{ii} is estimated using Eq. (9). A time step for stable explicit simulations is limited by the CFL constraint, $\Delta t_{\text{CFL}} \lesssim \min(\Delta t_{\parallel}^{\text{CFL}}, \Delta t_{\nabla B}^{\text{CFL}})$, where $\Delta t_{\parallel} \sim \Delta z/v_{\parallel}^{\text{max}} = 5.2$ ns and $\Delta t_{\nabla B}^{\text{CFL}} \sim m_i \Delta v_{\parallel} / (\mu_{\text{max}} |\nabla B|_{\text{max}}) = 0.1$ ns. It is clear that explicit time integration of a weakly collisional mirror plasma on the ion-ion collisional timescale becomes prohibitively expensive, and implicit time integration is required for first-principles kinetic modeling.

Here, the feasibility of multiscale collisional mirror transport simulations is explored using a second-order implicit-explicit (IMEX) time integration scheme, which enables implicit treatment of selected stiff terms. Specifically, the nonlinear Fokker-Planck operator, associated with the long ion-ion collisional timescale, is treated explicitly, whereas the Vlasov operator, involving fast advection in both the z and v_{\parallel} directions, is included as an implicit term. To handle nonlinearities in the Vlasov operator due to the presence of a self-consistent electric field, the Jacobian-free Newton-Krylov (JFNK) approach is employed²⁰ with the GMRES method used as a Krylov solver.³³ Although this approach alleviates the need to form or store the Jacobian matrix corresponding to the linearized system, it requires preconditioning for efficiency. In this work, we construct the preconditioner operator by neglecting the $\partial E_z / \partial f_i$ semi-dense contribution to the actual Jacobian operator. Physically, this means that our preconditioner operator does not capture collective modes involving the self-consistent electric field response, such as the plasma sound wave, $\omega = [(T_i + T_e)/m_i]^{1/2} k_{\parallel}$. Here, k_{\parallel} is a parallel wave vector and γ is an ion adiabatic index. We, however, note that for mirror systems, e.g., WHAM, where the ion temperature is much greater than the electron temperature, $T_i \gg T_e$, corrections to the linear-response Jacobian due to the self-consistent electric field model in Eq. (18) may be subdominant to the “fixed- E_z ” ion advection contribution.

Neglecting couplings between the ion distribution function and the electric field in the Jacobian preconditioner yields a 3D (1D-2V) nonsymmetric and indefinite sparse matrix for the discretized “fixed- E_z ” advection operator. COGENT solves this linear system using the Approximate Ideal Restriction (AIR) option in the BoomerAMG algebraic multigrid solver contained in the HYPRE linear solver library.²¹ The use of multigrid methods to solve nonsymmetric indefinite systems has historically been highly problematic, but the recent development of the AIR approach,^{22,23} including several variants, provides a way to extend the benefits of multigrid algorithms beyond the symmetric, positive-definite systems for which they are more commonly used. To further enhance the efficiency of the implicit Vlasov solver, COGENT allows the preconditioner to be defined using a lower-order discretization than that used for the Vlasov operator itself. In the weakly collisional regime, $\tau_{\parallel} \ll \tau_{ii}$, the collisional term in Eq. (17) is subdominant to the Vlasov term. Therefore, a numerical scheme with high spatial accuracy must be employed for the discretization of the Vlasov operator in order to mitigate truncation errors and minimize numerical pollution. In this work, we use a fifth-order upwind discretization (UW5) for the Vlasov operator in Eq. (17), whereas the preconditioner operator is constructed using a first-order (UW1) advection scheme. This lower-order preconditioner yields a sparser matrix, for which the most robust AMG solver performance is observed, while also maintains efficiency even for relatively large time steps.

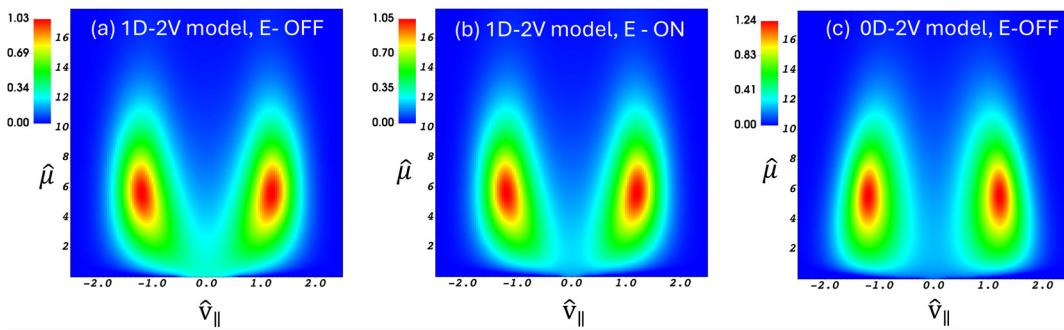


FIG. 6. Normalized ion distribution function, $\hat{f}_i = (2T_0/m_i)^{3/2} \pi f_i/n_0$, evaluated at the trap center at 20 ms. The results are shown for (a) fully kinetic 1D-2V simulations in the absence of E-fields; (b) fully kinetic 1D-2V simulations including a self-consistent E-field specified by the Boltzmann electron model; and (c) 0D-2V BASM-model simulations in the absence of E-fields. Frame (c) corresponds to a zoom-in of Fig. 2(a).

B. Zero E-field simulations

Here, simplified 1D-2V simulations are performed in the absence of electric-field effects. In this case, the preconditioner operator is consistent with the physical operator (modulo the difference in discretization order), allowing us to investigate the maximal benefits of our implicit algorithms. This simplified setup is also used to evaluate the convergence properties and numerical pollution associated with our spatial discretization.

The results of numerical simulations using the baseline grid resolution ($N_z, N_{v_{\parallel}}, N_{\mu} = (256, 128, 192)$) are shown in Fig. 6(a) for the normalized distribution function $\hat{f}_i = (2T_0/m_i)^{3/2} \pi f_i/n_0$ evaluated at the trap center. Qualitative agreement with the results of the BASM model is observed [cf. Figs. 6(a) and 6(c)]; however, the first-principles kinetic model captures additional features, e.g., the magnetic field profile. In particular, the peaks in the ion density profiles in Fig. 7(a) correspond to the beam turning points.

The numerical convergence studies of the final-state solution at 20 ms are illustrated in Fig. 8. Performing a mathematically rigorous analysis involving multiple levels of grid refinement in all phase-space dimensions, as done for the 0D-2V BASM model in Sec. II C, is too computationally expensive. Therefore, here, we consider the following set of computational grids: a coarse grid (128, 64, 96), a baseline grid (256, 128, 192), and three semi-fine grids with double the number of grid points in each individual direction, i.e., (512, 128, 192), (256, 256, 192), and (256, 128, 384). To gain insights into numerical

convergence, we examine two v_{\parallel} -lineouts in the (v_{\parallel}, μ) velocity space, taken at the center of the trap and corresponding to $\hat{\mu} = 6$ and $\hat{\mu} = 1$, which represent the high-energy beam and thermal background regions, respectively [see Fig. 8(c)]. Notably, in the beam region characterized by narrow peaks, numerical errors are more pronounced between the baseline and coarse grids compared to the background region, which exhibits a more diffuse profile. The results from simulations with the semi-fine grids show only a small difference relative to the baseline resolution in both regions.

As mentioned earlier, it is important to use a high-order (UW5) spatial discretization of the Vlasov operator to minimize numerical pollution. To further illustrate this point, we compare the results of our baseline and coarse grid simulations using the UW5 scheme with those obtained using a lower-order UW3 discretization (see Fig. 9). It is readily apparent that the excessive numerical diffusion produced by the UW3 scheme over long timescales, $t \gg \tau_{\parallel}$, leads to a completely spurious solution for the coarse grid case [see Fig. 9(a)]. While distinct beam features begin to emerge with the more resolved baseline grid [see Fig. 9(b)], the solution remains far from being quantitatively converged. In contrast, the UW5 scheme yields a converged velocity-space solution at the same (baseline) resolution.

We now discuss the computational efficiency of our IMEX time integration approach for the 1D-2V kinetic simulations (see Table II). The stable time step is limited by the CFL constraint associated with the explicitly treated Fokker-Planck operator. For the baseline grid resolution, the time step is initially limited by $\Delta t_{\text{init}}^{\text{CFL}} \approx 20 \mu\text{s}$ and decreases to $\Delta t_{\text{fin}}^{\text{CFL}} \approx 4.7 \mu\text{s}$ by the end of the simulation at $t_{\text{fin}} = 20 \text{ ms}$, as the collisional time scales shorten due to increasing ion density. The baseline grid simulation uses 512 cores of the NERSC system³⁴ and requires approximately 20 s and 7.5 s per $\Delta t_{\text{init}}^{\text{CFL}}$ and $\Delta t_{\text{fin}}^{\text{CFL}}$ steps, respectively. The JFNK solver tolerance is set to 10^{-4} , and it takes about 100 and 55 JFNK iterations for $\Delta t_{\text{init}}^{\text{CFL}}$ and $\Delta t_{\text{fin}}^{\text{CFL}}$, respectively. For comparison, the stable time step for a fully explicit time integration scheme is limited by $\Delta t_{\nabla B}^{\text{CFL}} \approx 0.19 \text{ ns}$, and each step requires approximately 0.9 s. Consequently, the IMEX simulations achieve a speedup of about 5×10^3 in the initial stage and 3×10^3 in the final stage. It is, however, instructive to note that the computational cost of an explicit simulation is dominated by the evaluation of the Fokker-Planck operator; disabling it reduces the wall time to about 0.15 s per step. Thus, explicit simulations can be accelerated by about a factor of 6 if the Rosenbluth potentials are evaluated infrequently.

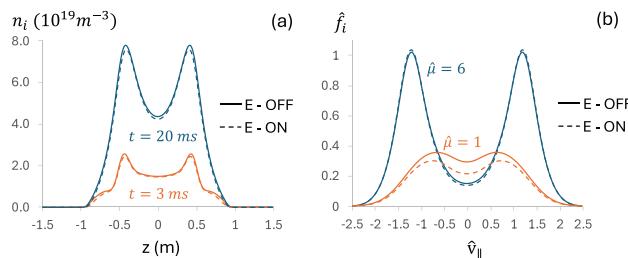


FIG. 7. Results of the fully kinetic 1D-2V simulations with (solid curves) and without (dashed curves) E-fields. Shown are (a) the ion density profile evaluated at 3 ms (orange) and 20 ms (blue); (b) the v_{\parallel} -lineouts extracted from the normalized ion distribution function, \hat{f}_i , at the center of the trap, shown for $\hat{\mu} = 1$ and $\hat{\mu} = 6$, at 20 ms.

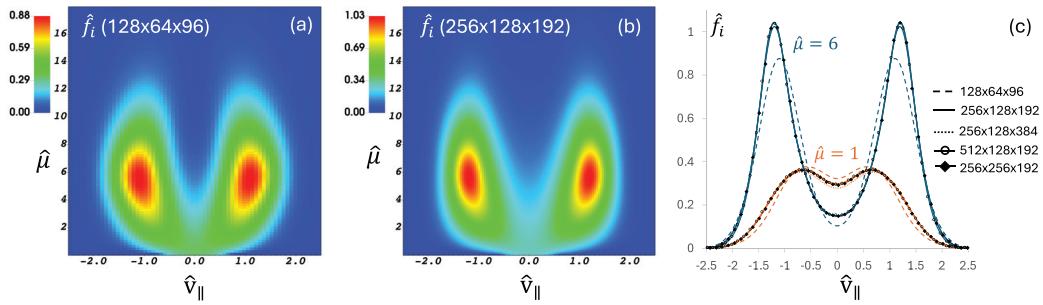


FIG. 8. Numerical convergence of the fully kinetic 1D-2V simulations without E-fields employing an UW5 discretization scheme for the Vlasov operator. The normalized ion distribution function is evaluated at the trap center at 20 ms.

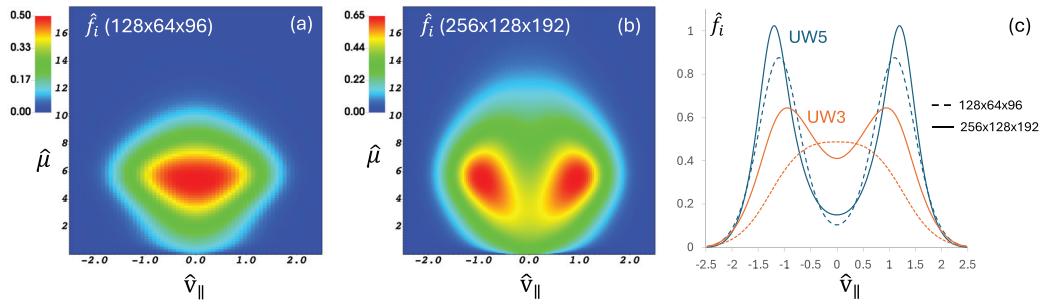


FIG. 9. Numerical convergence of the fully kinetic 1D-2V simulations without E-fields employing an UW3 discretization scheme for the Vlasov operator. The normalized ion distribution function is evaluated at the trap center at 20 ms. Frame (c) compares the v_{\parallel} -lineouts at $\hat{\mu} = 6$ for the UW3 (orange) and UW5 (blue) discretization schemes.

TABLE II. Performance of COGENT's IMEX scheme for simulations with the baseline grid resolution ($N_z, N_{v_{\parallel}}, N_{\mu} = (256, 128, 192)$ using 512 cores.

Parameter	$E - OFF$	$E - ON$
Implicit time step/wall-time per step	4.7 μ s/7.5 s	4.7 μ s/9.5 s
Explicit time step/wall-time per step	0.19 ns/0.9 s	0.19 ns/1.0 s
Number of JFNK iterations	≈ 55	≈ 70
(1.0×10^{-4} tol.)		

Finally, we note significantly enhanced performance of the IMEX simulations using the coarse grid resolution. At the end of the run, the stable IMEX time step is 19 μ s, the number of JFNK iterations is 65, and each time step requires approximately 1.8s on 512 cores. The entire 20 ms simulation completes in about 25 min. Although the numerical solution is not fully converged on the coarse grid, these fast, low-resolution runs are well suited for scoping studies and/or for providing informed initial conditions when searching for steady-state solutions.

C. Simulations including a self-consistent E-field

The results of the IMEX baseline-grid simulations including the effects of a self-consistent electric field [in Eq. (18)] are shown in Figs. 6(b) and 7. Compared to the corresponding results from simulations without an electric field, we observe a very similar solution in the

high-energy beam region and slightly degraded confinement in the thermal background region. The latter can be attributed to the presence of electrostatic forces that repel ions from the trap. At the same time, the presence of density peaks at the turning points of high-energy (sloshing) ions leads to the formation of a potential well that can help confine low-energy ions. To illustrate this point, we compare the results of simulations with and without beam injection. Figure 10 shows a zoom-in on the low-energy region, revealing enhanced confinement of low-energy ions.

For the case without beam injection, we observe the potential drop between the mirror center and the throat of $e\Delta\Phi/T_e \approx 7.6$, which is somewhat larger than the ambipolar Pastukhov potential inferred from the balance of electron and ion collisional losses from the trap, $\tau_{i,c} = \tau_{e,c}$. Specifically, estimating the ion confinement time as^{35,36} $\tau_{c,i} \sim 2.6 \log(R_m)\nu_i$, with the ion collision frequency ν_i given by Eq. (9), and using Eq. (11) for the electron confinement time, $\tau_{c,e} = n_e |dn_e/dt|^{-1}$, we obtain a Pastukhov potential of $e\Phi_m/T_e = 6.3$. Accurately capturing the ambipolar potential, however, requires modeling kinetic electron effects, including their reflection from sheath regions at the device ends and unlike-species collisions. These processes are beyond the scope of the present simulations using the Boltzmann electron model. We further mention substantial limitations³⁷ of the Boltzmann electron model in the expander region, i.e., outside the magnetic coil, where it predicts a potential drop $e\Delta\Phi \sim T_e \ln K$, which increases with the expansion ratio, K . Employing a kinetic electron response that captures cooling and

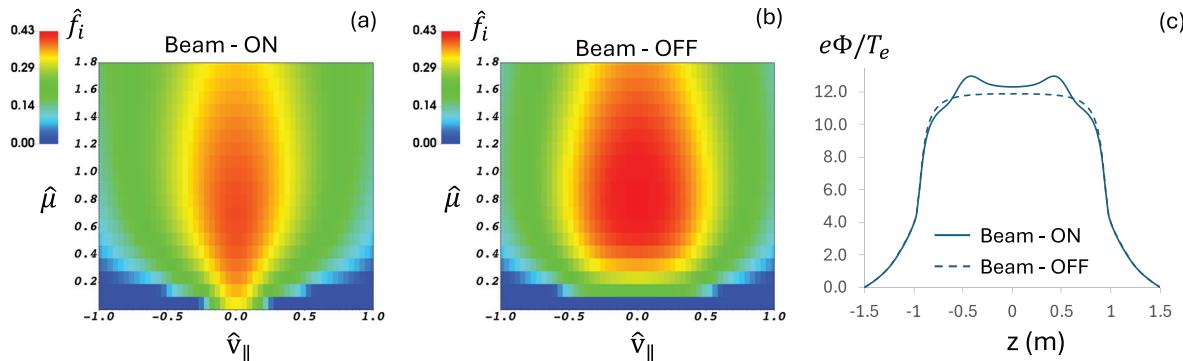


FIG. 10. Role of the sloshing-ion potential well in low-energy ion confinement. Shown are results from 1D-2V fully kinetic simulations with (a) and without (b) beam injection evaluated at 7 ms. Frame (c) illustrates the self-consistent electrostatic potential, $e\Phi/T_e$.

temperature anisotropy effects as well as electron reflection from the boundary sheath region would resolve this issue and yield a potential solution³⁷ that quickly saturates for large values of K . We also note that the baseline grid resolution in velocity space is insufficient to adequately resolve the distribution function outside the trap, and the numerical results in that region are not fully converged. As discussed in Ref. 9, this issue can be addressed, for instance, by introducing spatially varying normalization of the computational velocity space, which will be pursued in future work.

Including the self-consistent electric field in Eq. (18) introduces collective phenomena, such as plasma sound waves, which are not captured by the “fixed- E_z ” preconditioner model used in our simulations. These additional physical effects degrade the performance of the IMEX algorithms, with the most pronounced impact occurring at the beginning of the simulation when transient processes are strongest. To handle this initial stage, we employ an adaptive time-stepping algorithm from the PETSc library,³⁸ which adjusts the time step dynamically based on the solution behavior. As the simulation progresses and the system settles into a more quasi-steady state, the impact of the self-consistent electric field on solver performance diminishes. By the end of the run, we observe a modest slowdown of approximately 25% when comparing baseline-grid (256, 128, 192) simulations with and without electric fields (see Sec. III B), with a final stable time step of $\Delta t_{\text{fin}}^{\text{CFL}} \approx 4.7 \mu\text{s}$. Specifically, about 70 JFNK iterations are required to reach the same solver tolerance of 10^{-4} , and each step takes approximately 9.5s per step when run on 512 cores. This mild performance degradation is plausibly attributed to the condition $T_e \ll T_i$, as discussed in Sec. III A. The stable time step for fully explicit time integration remains limited by advection in the v_\parallel -direction, $\Delta t_{\text{V}\parallel}^{\text{CFL}} \approx 0.19 \text{ ns}$, with each step requiring approximately 1.0 s similar to the case without electric fields. Thus, the IMEX approach provides a speedup of over 2.5×10^3 by the end of the run. It is also informative to compare this performance with recent explicit simulations of the WHAM system using the GKeyll code.⁹ For the Boltzmann electron model and the original (non-smoothed) mirror force, GKeyll simulations employing the LBD collisional model require approximately 74 h on 288 cores to simulate 72 μs of physical time using the $(N_z, N_{v_\parallel}, N_\mu) = (288, 64, 192)$ grid. Normalized by the number of grid cells and processor cores, this is roughly 1.8×10^3 times slower than the

semi-implicit COGENT calculations employing the full nonlinear Fokker-Planck collision operator.

IV. CONCLUSIONS

In this work, we report on the development and application of implicit-explicit (IMEX) time integration methods for continuum kinetic modeling of weakly collisional parallel plasma transport in magnetic mirror configurations. The IMEX approach, based on high-order additive Runge-Kutta methods and Jacobian-free Newton-Krylov solvers with modern algebraic multigrid preconditioning, enables stable and efficient simulations that overcome the severe time step restrictions imposed by strong mirror forces in explicit schemes. As an illustrative example, we consider the problem of high-energy ion beam fueling of a weakly collisional mirror for parameters relevant to the WHAM facility. Both a simplified bounce-average model (BASM) for a square-shaped mirror and a fully kinetic model are employed to provide insights into ion species dynamics and to assess the performance of the numerical algorithms.

The fully implicit BASM model is verified by reproducing analytic results for the Pastukhov problem and is used to examine the numerical convergence of the velocity-space discretization scheme. In particular, degraded convergence, limited to first-order accuracy, is observed when the transition layer between the confined and loss cone regions is insufficiently resolved. The BASM model also highlights the importance of an accurate ion-ion collision model by revealing significant differences in the numerical solution obtained using the full Fokker-Planck collision operator vs the simplified Lenard-Bernstein-Dougherty (LBD) operator often used in gyrokinetic simulations.

The fully kinetic 1D-2V model demonstrates the practical benefits of the IMEX approach, which involves an implicit treatment of the Vlasov operator. Time steps approximately 2.5×10^4 times larger than the CFL-limited explicit step are achieved, yielding about $2500 \times$ speedup compared to explicit methods while maintaining numerical stability. The current “fixed- E_z ” preconditioner model does not account for collective oscillation processes such as plasma sound waves, which can limit its efficiency. Nevertheless, only modest performance degradation ($\sim 25\%$ slowdown) is observed when self-consistent electric field effects – modeled through quasi-neutrality and a Boltzmann electron response to the potential – are included. This

can be plausibly attributed to the relatively low electron temperature compared to the ion temperature in the parameter regime considered here, implying that corrections to the linear-response Jacobian from the self-consistent E-field model may be subdominant to the fixed-field ion advection terms. In addition, fifth-order upwind discretization of the Vlasov operator is found to be essential for minimizing numerical diffusion over long collisional transport timescales, whereas a third-order upwind scheme produces excessive numerical pollution.

The implicit technologies described in this work have the potential to enable comprehensive, integrated modeling of magnetic mirror systems using continuum gyrokinetic simulations. Our future work will focus on extending these methods to include radial nonuniformities, kinetic electrons, unlike-species collisions, and more complex magnetic geometries (e.g., tandem mirrors) relevant to next-generation mirror reactor designs.

ACKNOWLEDGMENTS

The authors are grateful to C. Forest, S. Frank, Y. Petrov, L. Ricketson, D. Martin, and H. Johanses for fruitful discussions. This research was supported by the U.S. Department of Energy (Contract No. DE-AC52-07NA27344).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

M. Dorf: Conceptualization (lead); Investigation (lead); Software (equal); Writing – original draft (lead). **M. Dorr:** Conceptualization (equal); Investigation (equal); Software (equal); Writing – original draft (supporting). **V. Geyko:** Conceptualization (equal); Investigation (equal); Software (supporting); Writing – original draft (supporting). **D. Ghosh:** Conceptualization (equal); Investigation (supporting); Software (equal); Writing – original draft (supporting). **M. Umansky:** Conceptualization (equal); Investigation (equal); Writing – original draft (supporting). **J. Angus:** Conceptualization (equal); Investigation (equal); Writing – original draft (supporting).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹J. Wesson, *Tokamaks, International Series of Monographs on Physics* (Oxford University Press, Oxford, 2011).
- ²F. Warner, J. Alguacil, D. Biek, T. Bogaarts, G. Bongiovanni, V. Bykov, J. P. Catalán, R. K. Duligal, I. Fernández-Berceruelo, S. Giambrone, C. Hume, M. Hrecinuc, R. Kembleton, J. Lion, T. Lyytinen, J. A. Noguerón Valiente, I. Palermo, V. Queral, D. Rapisarda, W. J. Rutten, L. Sanchis, X. Sarasola, K. Sedlak, A. Snicker, D. Sosa, and F. R. Ugorri, “Overview of European efforts and advances in Stellarator power plant studies,” *Fusion Eng. Des.* **202**, 114386 (2024).
- ³D. Whyte, “Small, modular and economically attractive fusion enabled by high temperature superconductors,” *Philos. Trans. R. Soc. A* **377**, 20180354 (2019).
- ⁴D. Endrizzi, J. K. Anderson, M. Brown, J. Egedal, B. Geiger, R. W. Harvey, M. Ialovega, J. Kirch, E. Peterson, Y. V. Petrov, J. Pizzo, T. Qian, K. Sanwalka, O. Schmitz, J. Wallace, D. Yakovlev, M. Yu, and C. B. Forest, “Physics basis for the Wisconsin HTS axisymmetric mirror (WHAM),” *J. Plasma Phys.* **89**, 975890501 (2023).
- ⁵C. B. Forest, J. K. Anderson, D. Endrizzi, J. Egedal, S. Frank, K. Furlong, M. Ialovega, J. Kirch, R. W. Harvey, B. Lindley, Y. V. Petrov, J. Pizzo, T. Qian, K. Sanwalka, O. Schmitz, J. Wallace, D. Yakovlev, and M. Yu, “Prospects for a high-field, compact break-even axisymmetric mirror (beam) and applications,” *J. Plasma Phys.* **90**, 975900101 (2024).
- ⁶A. A. Ivanov and V. V. Prihodko, “Gas dynamic trap: Experimental results and future prospects,” *Phys. Usp.* **60**, 509 (2017).
- ⁷J. L. Ball, S. Mackie, J. G. van de Lindt, W. Morrissey, A. Perevalov, Z. Short, N. Schwartz, T. W. Koeth, B. L. Beaudoin, C. A. Romero-Talamas, J. Rice, and R. A. Tinguely, “Measurements of fusion yield on the centrifugal mirror fusion experiment,” *arXiv:2205.23047* (2025).
- ⁸A. Tran, S. J. Frank, A. Y. Le, A. J. Stanier, B. A. Wetherton, J. Egedal, D. A. Endrizzi, R. W. Harvey, Y. V. Petrov, T. M. Qian, K. Sanwalka, J. Viola, C. B. Forest, and E. G. Zweibel, “Drift-cyclotron loss-cone instability in 3-D simulations of a sloshing-ion simple mirror,” *J. Plasma Phys.* **91**, E85 (2025).
- ⁹M. Francisquez, M. H. Rosen, N. R. Mandell, A. Hakim, C. B. Forest, and G. W. Hammett, “Toward continuum gyrokinetic study of high-field mirrors,” *Phys. Plasmas* **30**, 102504 (2023).
- ¹⁰W. H. Wang, X. S. Wei, Z. Lin, C. Lau, S. Dettrick, and T. Tajima, “A gyrokinetic simulation model for 2D equilibrium potential in the scrape-off layer of a field-reversed configuration,” *Phys. Plasmas* **31**, 072507 (2024).
- ¹¹V. S. Chan, C. M. Greenfield, L. L. Lao, T. C. Luce, C. C. Petty, G. M. Staebler, and D.-D. Team, “DIII-D advanced tokamak research overview,” *Nucl. Fusion* **40**, 1137 (2000).
- ¹²Y. Petrov, R. Harvey, C. Forest, and J. Anderson, “Calculations of WHAM2 mirror neutron rates and FI transport using the GENRAY/CQL3D-M and MCGO-M codes,” in 49th EPS Conference on Contr. Fusion and Plasma Physics (European Physical Society, 2023).
- ¹³G. I. Dimov, “The ambipolar trap,” *Phys. Usp.* **48**, 1129 (2005).
- ¹⁴D. D. Ryutov, “Open-ended traps,” *Sov. Phys. Usp.* **31**, 300 (1988).
- ¹⁵D. D. Ryutov, H. L. Berk, B. I. Cohen, A. W. Molvik, and T. C. Simonen, “Magneto-hydrodynamically stable axisymmetric mirrors,” *Phys. Plasmas* **18**, 092301 (2011).
- ¹⁶D. Michels, P. Ulbl, W. Zhlobenko, T. Body, A. Stegmeir, T. Eich, M. Griener, G. D. Conway, F. Jenko, and A. U. Team, “Full-f electromagnetic gyrokinetic turbulence simulations of the edge and scrape-off layer of ASDEX Upgrade with GENE-X,” *Phys. Plasmas* **29**, 032307 (2022).
- ¹⁷M. Dorf and M. Dorr, “Continuum gyrokinetic simulations of edge plasmas in single-null geometries,” *Phys. Plasmas* **28**, 032508 (2021).
- ¹⁸D. Ghosh, M. A. Dorf, M. R. Dorr, and J. A. F. Hittinger, “Kinetic simulation of collisional magnetized plasmas with semi-implicit time integration,” *J. Sci. Comput.* **77**, 819–849 (2018).
- ¹⁹C. A. Kennedy and M. H. Carpenter, “Additive Runge–Kutta schemes for convection–diffusion–reaction equations,” *Appl. Numer. Math.* **44**, 139–181 (2003).
- ²⁰D. Knoll and D. Keyes, “Jacobian-free newton–Krylov methods: A survey of approaches and applications,” *J. Comput. Phys.* **193**, 357–397 (2004).
- ²¹See <http://www.llnl.gov/casc/hypre/> for “Hypre library is a software product of LLNL, CA, USA”.
- ²²T. A. Manteuffel, J. Ruge, and B. S. Southworth, “Nonsymmetric algebraic multigrid based on local approximate ideal restriction (LAIR),” *SIAM J. Sci. Comput.* **40**, A4105–A4130 (2018).
- ²³T. A. Manteuffel, S. Munzenmaier, J. Ruge, and B. Southworth, “Nonsymmetric reduction-based algebraic multigrid,” *SIAM J. Sci. Comput.* **41**, S242–S268 (2019).
- ²⁴V. P. Pastukhov, “Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential,” *Nucl. Fusion* **14**, 3 (1974).
- ²⁵R. H. Cohen, M. E. Rensink, T. A. Cutler, and A. A. Mirin, “Collisional loss of electrostatically confined species in a magnetic mirror,” *Nucl. Fusion* **18**, 1229 (1978).
- ²⁶M. A. Dorf, R. H. Cohen, M. Dorr, J. Hittinger, and T. D. Rognlien, “Progress with the COGENT edge kinetic code: Implementing the Fokker–Planck collision operator,” *Contrib. Plasma Phys.* **54**, 517–523 (2014).

²⁷M. Francisquez, J. Juno, A. Hakim, G. W. Hammett, and D. R. Ernst, “Improved multispecies Dougherty collisions,” *J. Plasma Phys.* **88**, 905880303 (2022).

²⁸P. Ulbl, D. Michels, and F. Jenko, “Implementation and verification of a conservative, multi-species, gyro-averaged, full-f, Lenard-Bernstein/Dougherty collision operator in the gyrokinetic code GENE-X,” *Contrib. Plasma Phys.* **62**, e202100180 (2022).

²⁹J. R. Angus, “On anomalous plasma transport in the edge of magnetic confinement devices,” Ph.D. thesis (University of California, San Diego, 2012).

³⁰M. H. Rosen, W. Sengupta, I. Ochs, F. I. Parra, and G. W. Hammett, “Enhanced collisional losses from a magnetic mirror using the Lenard-Bernstein collision operator,” *J. Plasma Phys.* 1–24 (published online 2025).

³¹H. Johansen and P. Colella, “A Cartesian grid embedded boundary method for Poisson’s equation on irregular domains,” *J. Comput. Phys.* **147**, 60–85 (1998).

³²W. Thacher, H. Johansen, and D. Martin, “A high order Cartesian grid, finite volume method for elliptic interface problems,” *J. Comput. Phys.* **491**, 112351 (2023).

³³Y. Saad, *Iterative Methods for Sparse Linear Systems: Second Edition* (Society for Industrial and Applied Mathematics, 2003).

³⁴NERSC, see <https://nersc.gov> for “National energy research scientific computing center” (2025).

³⁵G. F. Bing and J. E. Roberts, “End-losses from mirror machines,” *Phys. Fluids* **4**, 1039–1046 (1961).

³⁶J. Scheffel, J. Jäderberg, K. Bendtz, R. Holmberg, and K. Lindvall, “Axial confinement in the Novatron mirror machine,” *Nucl. Fusion* **65**, 066011 (2025).

³⁷M. Tyushev, A. Smolyakov, A. Sabo, R. Groenewald, A. Necas, and P. Yushmanov, “Drift-kinetic pic simulations of plasma flow and energy transport in the magnetic mirror configuration,” *Phys. Plasmas* **32**, 032514 (2025).

³⁸S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, “Efficient management of parallelism in object oriented numerical software libraries,” in *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen (Birkhäuser Press, 1997), pp. 163–202.