

Towards kinetic axial-azimuthal Hall thruster simulations including ionization

IEPC-2025-0129

*Presented at the 39th International Electric Propulsion Conference, Imperial College London, London,
United Kingdom
14-19 September 2025*

Thomas A. Marks* and Alex A. Gorodetsky†
University of Michigan, Ann Arbor, MI, USA 48109

The open-source particle-in-cell code WarpX is used to perform 2-D axial-azimuthal Hall thruster simulations including a neutral population and self-consistent ionization using the direct simulation Monte Carlo (DSMC) method. The results of these simulations are assessed in comparison to similar simulations which use a static particle source. The time-averaged simulation outputs for the DSMC simulations differ from those without; in particular, the electron temperature and anomalous drag force are higher and the electric field is less strongly-peaked. The time-dependent outputs show strong ion density oscillations, consistent with the well-known Hall thruster breathing mode. The relative phasing of this oscillation with respect to other plasma properties is examined, with inconclusive results. Ion recombination at the anode is assessed and found to increase the amplitude of breathing mode oscillations. Numerical challenges encountered in these simulations are discussed, particularly the large fluctuations in particle count during thruster startup and over the course of the oscillations. The findings highlight the advantages of GPU computing and semi-implicit schemes in enabling kinetic Hall thruster simulations to reach new physical regimes, as well as avenues for future improvement.

I. Introduction

KINETIC methods are attractive for the modeling and simulation of low-temperature plasmas due to their ability to, in principle, capture all relevant physical phenomena without the need to invoke strong closure approximations. They are especially relevant for Hall-effect thrusters, where kinetic instabilities and non-Maxwellian effects prevent fluid methods from accurately capturing the performance and plasma properties of these widely-flown devices without extensive calibration against data.¹ Two of the most well-known Hall thruster phenomena related to instabilities are the breathing mode and anomalous electron transport. The former refers to low-frequency global ionization oscillations² which occur in most Hall thrusters and exhibit complex dependencies on operating conditions and testing conditions. The latter is the phenomenon of enhanced cross-field electron transport due to plasma instabilities driven by the azimuthal electron drift.³

One of the key challenges of Hall thruster modeling is the wide range of length- and timescales needed to resolve these phenomena at the same time. Breathing oscillations have frequencies measured in tens of kilo-Hertz, requiring simulation timescales between 10^{-4} and 10^{-3} seconds to resolve enough individual oscillations to provide useful time-averaged information. Simulating these long timescales typically requires fast, whole device models which treat the electrons as a fluid, neglecting small-scale physics^{4,5} In contrast, the kinetic simulations typically used to study anomalous electron transport need to resolve both the electron Debye length and plasma frequency ($\sim 10^{-5}$ m and $\sim 10^{-12}$ s, respectively). To reach steady state quickly and avoid the need to resolve the breathing mode, many 2-D axial-azimuthal and 3-D kinetic simulations use a static particle source instead of self-consistently modeling neutral dynamics and ionization.^{4,6,7} This choice allows the physics leading to anomalous mobility be analyzed at a reduced cost, but the insights gained may be limited due to the neglect of important physics. Other codes assume a stationary background neutral

*Postdoctoral research fellow, Department of Aerospace Engineering, marksta@umich.edu

†Associate professor, Department of Aerospace Engineering, goroda@umich.edu



density profile and apply a Monte Carlo Collision (MCC) algorithm to handle ionization.^{8–10} This is more physically realistic than a static particle source, but still neglects the time evolution of the neutral density, preventing ionization oscillations for developing. Still other kinetic Hall thruster simulations have included self-consistent ionization with an evolving neutral background population, but simulate 2-D axial-radial domains¹¹ on which the anomalous electron transport cannot be self-consistently resolved.

Experimental and computational studies suggest that the anomalous electron transport varies over the breathing oscillations and is sensitive to neutral dynamics.^{5,12,13} Additionally, the process by which the accelerating electric field establishes and localizes itself, which depends critically on the specifics of the anomalous electron transport^{14,15} cannot be addressed when the ionization profile is fixed in place. There is thus a need for axial-azimuthal kinetic Hall thruster simulations including self-consistent ionization.

In this paper, we apply the open-source particle-in-cell (PIC) code WarpX to this task. While designed for laser-plasma physics,¹⁶ WarpX has been gaining use in the low-temperature and electric propulsion plasma community. In a previous study,¹⁷ we applied WarpX to the benchmark problem found that a combination of GPU acceleration and judicious use of semi-implicit and energy-conserving algorithms¹⁸ led to large speedups, opening the door to the longer run-times needed to capture self-consistent ionization. Recently, other groups have used WarpX for low-temperature problems including ionization, including for a crossed-field electron beam-generated plasma¹⁹ as well as for a Hall thruster.⁹ In this latter case, as in many previous simulations, ionization was modeled using MCC assuming a static neutral density profile.

Leveraging the insights of our previous work, we use WarpX to simulate a benchmark-like Hall thruster incorporating both neutral dynamics and ionization with the direct simulation Monte Carlo (DSMC) method. We simulate the evolution of both neutrals and plasma species over 200 200 μ s of simulation time, ten times longer than the benchmark simulation. Our simulations exhibit the kHz-scale ionization oscillations characteristic of the breathing mode. We analyze the relative phasing of multiple plasma properties relative to the ion density oscillations, including that of the effective anomalous electron collision frequency, and additionally comment on some of the numerical and implementation challenges encountered, particularly around managing particle count. To our knowledge, these are the first kinetic axial-azimuthal simulations of the Hall thruster discharge plasma incorporating ionization using the DSMC method. By building on an existing open-source code, we hope that our work serves as a useful baseline on the path to the use kinetic simulations in the design and characterization of Hall thrusters.

This paper is organized as follows. In Sec. II, we describe WarpX as well as the extensions made to support the present work and the numerical and physical parameters of our simulations. Then, in Sec. III, we present the results of our simulations. We assess the impact of physical assumptions, including the role of ion recombination at the anode, and analyze how the magnitude of the electron transport varies over time in relation to the plasma density and temperature. We also discuss the computational cost of our simulations and compare our results qualitatively to experiments. Finally, we summarize our work in Sec. IV.

II. Methods

A. Particle-in-cell simulations

In this work, we use the open-source particle-in-cell (PIC) code WarpX²⁰ to run 2-D axial-azimuthal simulations of a Hall thruster. We base our simulation on the 2019 LANDMARK axial-azimuthal benchmark simulation,²¹ which we previously implemented in WarpX.¹⁷ Our setup is largely identical to that of the benchmark, with the only changes made being those necessary to achieve simulations with stable ionization oscillations. For this reason, we provide only a brief summary of the benchmark conditions here before discussing the details relevant to our study.

Our PIC simulations contain three species: neutral xenon atoms, singly-charged xenon ions, and electrons. The problem domain, illustrated in Fig. 1, is two-dimensional, consisting of an axial (x) and an azimuthal (y) dimension, with axially-varying magnetic field $B(x)$ pointing in the radial (z) direction. The azimuthal dimension is assumed periodic for both particles and fields. The magnetic field has a maximum strength of B_{max} , which occurs at $x = x_{Bmax} = 0.75$ cm. We apply a voltage drop of V_0 between the anode at $x = 0$ and a virtual cathode at $x = x_c = 2.4$ cm. In our simulations we choose $B_{max} = 200$ G and $V_0 = 300$ V, compared to 100 G and 200 V, respectively, in the benchmark. During testing, we determined that these changes led to more consistent ionization behavior and more stable simulations. These operating conditions are also more typical of Hall thruster operation.^{22–24} We enumerate the simulations' physical and numerical parameters in Tab. 1.






Figure 1: (a) A schematic cutaway view of a Hall thruster, showing the coordinate system and orientation of our simulation domain. (b) An illustration of the domain, with important numerical parameters labeled.

Table 1: Numerical and physical parameters used in the simulations in this study.

Parameter	Symbol	Value
Propellant	-	Xenon
Mass of propellant atom	M	2.18×10^{-25} kg
Axial domain length	L_x	2.5 cm
Azimuthal domain length	L_y	1.28 cm
Axial location of max. magnetic field	x_{Bmax}	0.75 cm
Axial location of cathode	x_c	2.4 cm
Maximum magnetic field strength	B_{max}	200 G
Discharge voltage	V_0	300 V
Initial plasma density	$n_{e,0}$	10^{17} m^{-3}
Initial neutral density	$n_{n,0}$	10^{19} m^{-3}
Initial electron temperature	$T_{e,0}$	20 eV
Initial ion temperature	$T_{e,0}$	0.5 eV
Cathode electron temperature	$T_{e,c}$	10 eV
Anode temperature	T_a	2000 K
Injected neutral axial speed	$u_{n,0}$	460 m/s
Initial neutral particles per cell	$N_{ppc,n}$	75
Initial electron/ion particles per cell	$N_{ppc,e}$	75
Simulation time	t_{max}	200 μs
DSMC collision check interval	N_{dsmc}	1009 iters.
Diagnostic output interval	N_{output}	5000 iters.

B. Direct simulation Monte Carlo algorithm

To model electron-impact ionization of neutral atoms, we employ the Direct Simulation Monte Carlo (DSMC) capabilities of WarpX. In contrast to the Monte Carlo collision (MCC) approach, which simulates collisions between impacting particles and a target species represented as a static background density, DSMC directly models binary collisions between macroparticles.²⁵ Whenever collisions are to be performed, the algorithm first sorts particles based on which grid cell they lie in, then pairs particles with their collision partners. If there is a mismatch in the number of particles, the particles of the less numerous species are split to give each colliding particle exactly one partner. Then, for each pair, the algorithm computes the probability of a collision based on the particle energies and appropriate cross section data. Based on this, it then determines whether a collision occurs. Electron-impact ionization collisions create two electrons and one ion, with the weight of particle set to equal that of the colliding electron. The impacted neutral macroparticles' weight is then reduced by the same amount.

As the DSMC algorithm requires particles be sorted and paired every time it is run, it could potentially dominate the cost of our simulations if run every iteration. Fortunately, our simulation timesteps ($\sim 10^{-11}$ s) are much smaller than typical ionization times in Hall thrusters ($\sim 10^{-7}$ s). We therefore only perform DSMC ionization calculations every N_{dsmc} iterations. In the present work, we set N_{dsmc} to 1009.^a Additionally, we neglect electron-ion and electron-neutral binary collisions to better isolate the effects of ionization. We are interested in exploring the effect of these and other collisions in future work.

C. Initialization

Our simulations begin by randomly loading the domain with particles of each species. We distribute the particle positions uniformly and velocities according to Maxwellian distributions with specified temperatures. As in the benchmark simulation, the initial ion temperature is 0.5 eV, but we increase the initial electron temperature from 10 to 20 eV. In testing, this helped ensure that a sustainable burst of ionization happened before the charged particles left the domain. The neutral temperature is equal to the anode temperature, $T_{n,0} = T_a = 2000$ K, and the bulk axial velocity is 460 m/s. We discuss the reasons for choosing these values in the next section.

Our simulations begin with uniform particle densities. At initialization, we calculate the ion/electron particle weights $W_{e,0}$ assuming we have 75 particles of each species per cell and a density of 10^{17} m^{-3} ,

$$W_{e,0} = \frac{n_{e,0} L_x L_y}{N_{ppc,e} \Delta x \Delta y}.$$

We take the same approach to compute the initial neutral weight $W_{n,0}$, though we take the initial neutral density to be 10^{19} m^{-3} ,

$$W_{n,0} = \frac{n_n L_x L_y}{N_{ppc,n} \Delta x \Delta y}.$$

Other Hall thruster simulations in WarpX have accelerated convergence to steady state through carefully-chosen initial particle density profiles.²⁶ We do not take this approach in this work as we do not expect our simulations to reach a steady state due to the presence of the breathing mode, though such a technique may still be useful for smoothing startup transients.

D. Anode boundary conditions

We inject neutral atoms at every timestep assuming the neutrals have the same temperature as the anode and that flow through the anode gas distributor is choked. In this model, the axial speed of the neutrals is set by the sound speed at the anode temperature, or

$$u_{n,0} = \sqrt{\frac{5}{3} \frac{k_B T_a}{M}}, \quad (1)$$

where $u_{n,0}$ is the axial neutral speed, k_B is Boltzmann's constant, T_a is the anode temperature in Kelvins, and M is the mass of a xenon atom in kilograms. According to the basic predator-prey model of the breathing

^aWe choose a prime number to ensure that DSMC collision checks do not regularly coincide with other parts of the simulation occurring at fixed intervals, such as outputs to disk. Other potential effects that would trigger infrequently, such as resampling, could then be given intervals which are co-prime with that of the DSMC interval.



mode, the frequency of the breathing oscillations depends strongly on the axial speed of the neutral atoms.¹² This stems from the fact that the neutrals need time to refill the channel after having been depleted by ionization. Choosing a higher axial speed permits the resolution of more oscillations within a fixed time period, decreasing the time to convergence. We thus choose a somewhat high anode temperature of 2000 K to accelerate our simulations, which according to Eq. 1 gives the atoms an axial speed of 460 m/s. The axial neutral velocities then follow a truncated drifting Maxwellian momentum distribution function at the specified temperature and velocity. The flux of neutrals at the anode is computed by assuming a neutral density of 10^{19} m^{-3} , the same value we use to initialize the neutral population.

In our baseline simulations, we remove neutral atoms and ions that cross the anode plane from the simulation. However, the dynamics in the near-anode region, including ion recombination at the anode, may have strong effects on the magnitude of the breathing mode oscillations.^{4,27,28} As such, we have also implemented ion recombination as an optional feature, which we leave disabled in our baseline simulations but examine later in Sec. C. Practically speaking, at each timestep the total weight of ions and neutrals that cross the anode plane is recorded. We then generate $N_{n,a}$ new neutrals at the anode with weights of each new particle, $W_{n,a}$ chosen to be as close to the initial neutral weight as possible.

$$N_{n,a} = \left\lceil \frac{\sum_i W_{p,i}}{W_{n,0}} \right\rceil$$

$$W_{n,a} = \frac{\sum_i W_{p,i}}{N_{n,a}}$$

Here, $\sum_i W_{p,i}$ is the sum of ion and neutral weights that cross the anode at any given timestep. The new neutrals are treated as if they had reflected off of the anode, and given velocities according to a one-sided Maxwellian with temperature equal to the anode temperature. In Sec. C, we perform some limited simulations to assess the impact of anode ion recombination on our results.

E. Cathode boundary conditions

Our simulations employ the same electrical and particle boundary conditions at the virtual cathode as the benchmark case.²¹ At each simulation timestep, we average the potential at the the virtual cathode line ($x = x_c$) and then correct the potential in the domain according to

$$\phi(x, y) = U(x, y) - \frac{x}{x_c} \bar{U}_c \quad (2)$$

where $U(x, y)$ is the potential from the electrostatic solver and \bar{U}_c is the average uncorrected potential along the cathode line. The benchmark simulation takes this approach to limit the impact of the cathode sheath on the rest of the simulation and to ensure a stable voltage drop of V_0 between anode and cathode.

To ensure current continuity, we log the net charge of particles that cross the anode boundary, injecting the result as electrons distributed uniformly along the cathode line. The injected electrons follow a three-dimensional Maxwellian distribution with an injection temperature, T_{ec} , equal to 10 eV. We have modified this functionality to account for varying particle weights due to the DSMC algorithm, which was not possible in the benchmark simulations. Our algorithm for determining the injection weights is the same is used for injecting neutrals, except we use per-particle weights close to the initial electron/ion weights, $W_{e,0}$, and use the net weight difference between ions and electrons as the total weight to inject at the cathode.

F. Implementation

We use the open source particle-in-cell code WarpX for our simulations.²⁹ While it was designed primarily the simulation of laser-plasma interactions in the context of particle accelerators, WarpX's capabilities are sufficiently general to model a diverse range of plasmas. WarpX has been fruitfully applied to problems in low temperature plasmas by ourselves¹⁷ and others,^{9,19,26} including Hall thruster simulation.

To implement the subroutines needed for 2-D Hall thruster simulations into WarpX, we leverage its user-defined callback functionality. This permits users to write Python functions that hook into the code's functionality at predefined points in the solver loop. For this work, we use two callbacks. The first applies the correction to the electrostatic potential in Eq. 2 and runs directly after WarpX solves for the electric



field. The second runs after the particles have been advanced but before they are deposited, and handles both the injection of electrons at the cathode and, if enabled, ion recombination at the anode.

Besides its active open-source development community and large suite of well-tested functionality, one of the main advantages of WarpX is its ability to accelerate computation using GPUs.²⁰ In our previous study,¹⁷ we demonstrated that leveraging WarpX in combined with new GPU hardware led to large speedups over the state of the art in 2-D Hall thruster simulation. For this reason, we run each of our simulations on one NVIDIA H100 GPU.³⁰ We employ mixed-precision in our simulations, with particle data stored as 32-bit floating point numbers and the electrostatic field solve using 64 bits throughout.

G. Grid resolution and time-stepping

To further accelerate simulations and to enable the testing of a range of numerical and physical parameters, we use the same semi-implicit and energy-conserving numerical scheme as in our prior work.^{17,18,31} This algorithm works similarly to the direct-implicit³² particle-in-cell algorithm and raises the effective plasma frequency by the square of a user-defined constant C_{SI} . Our previous study showed that this scheme allows the use of coarser grids and longer timesteps than are typically permitted in the PIC method without excessive numerical heating. We use the semi-implicit scheme throughout this work and refer to four different grid scale and timestep combinations, hereafter 1x, 2x, 4x, and 8x. Tab. 2 summarizes the numerical parameters of each case. The 1x simulation uses the same grid cell size and timestep as the benchmark,²¹ with the 2x, 4x, and 8x cases using grid cells and timestep sizes 2, 4, and 8 times larger, respectively.

Table 2: Case-dependent numerical and physical parameters used in the simulations in this study.

Parameter	Symbol	Value (1x)	Value (2x)	Value (4x)	Value (8x)
Timestep [s]	Δt	5×10^{-12}	1×10^{-11}	2×10^{-11}	4×10^{-11}
Semi-implicit factor	C_{SI}	1	8	16	8
Axial cells	N_x	512	256	128	64
Azimuthal cells	N_y	256	128	64	32
Axial cell size [m]	Δx	4.88×10^{-5}	9.77×10^{-5}	1.95×10^{-4}	3.90×10^{-4}
Azimuthal cell size [m]	Δy	5×10^{-5}	1×10^{-4}	2×10^{-4}	4×10^{-4}
Initial neutral weight	$W_{n,0}$	3.25×10^{12}	1.3×10^{13}	5.21×10^{13}	1.04×10^{14}
Initial electron/ion weight	$W_{e,0}$	3.25×10^{10}	1.3×10^{11}	5.21×10^{11}	1.04×10^{12}

H. Simulations

We list in Tab. 3 the simulation cases analyzed in this study. We include both simulations using a static particle source identical to that used in the benchmark as well as those which model a full neutral population and DSMC ionization. Simulations are additionally categorized according to their grid resolution and timestep using the classification scheme of Tab. 2. When presenting time-averaged results, we average over a time interval t_{avg} extending backward from the final simulation time t_{max} . We run three identical replicates of each of the DSMC simulations to quantify the amount of run-to-run variation stemming from the stochastic nature of the PIC algorithm.

Table 3: Settings for the main simulation cases performed for this work.

Category	Grid	V_0 [V]	B_{max} [G]	t_{max}	t_{avg}	# Sims	Notes
Benchmark	1x	200	100	20 μ s	8 μ s	1	Benchmark of Charoy et al. ²¹
Static	1x	300	200	20 μ s	8 μ s	1	-
Static	2x	300	200	20 μ s	8 μ s	1	-
DSMC	2x	300	200	200 μ s	50 μ s	3	-
DSMC	4x	300	200	200 μ s	50 μ s	3	-

We note that the H100 GPUs we use did not have sufficient memory to complete 1x simulations including ionization. On the other hand, 8x simulations completed rapidly and were highly enabling for code



development, debugging, and tuning numerical and physical parameters such as the particle temperatures, the magnetic field strength, the discharge voltage, and the DSMC resampling interval. However, in our previous work we found that the 8x simulations greatly modify the spectrum of the azimuthal electric field fluctuations and produce time-averaged results that diverge from the 1x results by more than 25% in some cases. For our DSMC simulations, we therefore present results only for the 2x and 4x resolutions. To better separate the effects of ionization, from those due to differences from the benchmark case, we perform two simulations with static particle sources, at 1x and 2x resolutions, which otherwise use identical parameters to the DSMC cases. Finally, we include the original benchmark simulation²¹ at 1x resolution for completeness.

III. Results and discussion

In this section, we first analyze the time-averaged results of our baseline simulation with ionization. Next, we compare the 2x and 4x simulations to one another and analyze the time-dependent evolution of multiple plasma properties. We also assess the impact of increasing the azimuthal domain length and of enabling ion recombination at the anode, and report the computational cost of our simulations.

A. Time-averaged properties

We plot in Fig. 2 the profiles of the ion density, electron temperature, axial electric field, and mean axial ion velocity for the 2x simulations with ionization, compared to the 1x simulation with a static ion source. These have been averaged over the azimuthal dimension as well as the averaging time specified in Tab. 3. The DSMC simulations exhibit plasma densities up to two times higher than those using a static ion source, without a strongly-localized peak. The electron temperature curve has a similar shape but is 50% higher when including ionization. The electric field has a maximum value twice as high and a more strongly-peaked shape when using a static particle source, which leads to a steeper ion acceleration profile. Next, in Fig. 3, we show the time-averaged axial profiles of the neutral particle density and velocity for the DSMC 2x simulation. The neutral density decays by a factor of ten from the left to the right side of the domain, while the velocity increases 50%, from a value of 525 m/s at the anode to 800 m/s at the cathode.

In both Fig. 2 and Fig. 3, we quantify the degree of run-to-run stochastic variation by plotting the mean and range of outputs from three replicate simulations. The run-to-run variation is only discernible for the ion density, for which the range of predictions spans less than 5% of the mean density. For the other quantities, the output variance less than 1% and is indistinguishable from the mean value.

One of the main objectives of kinetic PIC simulations is to characterize the degree of instability-induced electron transport. To this end, we calculate anomalous azimuthal force density, $R_{an,y}$. This term represents the amount of azimuthal drag force created by correlated fluctuations in the electron density, n_e , and azimuthal electric field, E_y , and is proportional to the covariance these terms:³³

$$R_{an,y}(t) = q_e \langle (n_e(t) - \langle n_e \rangle_t) (E_y(t) - \langle E_y \rangle_t) \rangle_t. \quad (3)$$

In the above, the notation $\langle \cdot \rangle_t$ denotes averaging over a time period t . We take $t = 1 \mu\text{s}$ and present results averaged over azimuthal dimension and the final quarter of the simulation time. Fig. 4a shows the anomalous force density for the simulations without DSMC ionization compared to that computed by Petronio et al. for the benchmark simulation.³³ We observe a qualitatively similar shape to that of the benchmark simulation but up to 50% more anomalous drag across the domain. Additionally, the location of maximum drag has shifted upstream by 3 mm. These differences result solely from the increased magnetic field strength and discharge voltage in our simulations. We then compare in Fig. 4b these results to those from our simulations including ionization at both the 2x and 4x grid resolutions. The simulations with DSMC ionization have anomalous drag forces ten times greater than those without. This may potentially lead to increased cross-field electron transport, which could account for the reduction in peak electric field observed in Fig. 2. The large electron temperatures may also be explained by an increase in instability-induced heating stemming ultimately from the larger azimuthal drag force. The 2x and 4x simulations are qualitatively similar when compared to the case with a static particle source, though the force density reaches peak values 15% higher in the 2x case compared to the 4x case.






Figure 2: Time-averaged profiles of the (a) ion number density, (b) electron temperature, (c) axial electric field, and (d) mean axial ion velocity for 2x simulations with DSMC ionization (blue) and with the static particle source (black). We plot the range of values from three replicate simulations with ionization as a shaded blue band. The dashed vertical gray line indicates the location of the cathode boundary at $x = x_c$, and the shaded vertical gray band shows the ionization region for the “Static” simulations.




Figure 3: Time-averaged profiles of the (a) neutral number density and (b) axial neutral velocity for the 2x DSMC simulations. Run-to-run variation from three replicates is plotted as a pale band but is not distinguishable from the mean values.




Figure 4: Comparison of the time-averaged anomalous azimuthal force density, $R_{an,y}$, between (a) the simulations using a static particle source (“Static”) and the benchmark^{21,33} (b) and the simulations with ionization (“DSMC”) at 2x and 4x grid resolutions. Note the change in axis scale between (a) and (b).

B. Time-dependent properties

We plot in Fig. 5a and 5b the spatially-averaged neutral and ion densities for three replicate simulations with ionization for 2x and 4x cases, respectively. In each case, our simulations succeed at capturing breathing oscillations. Each simulation shows a strong initial transient which transitions into a coherent oscillation. This damps further into stationary quasi-periodic oscillation which maintains a consistent amplitude until the end of the simulation. The oscillations are less strongly-damped in the 2x simulations, taking 150 μs to reach at a stationary oscillation compared to just 75 μs in the 4x cases. The stochastic run-to-run variation primarily affects the phasing of the oscillations, but not their amplitude or envelope.




Figure 5: Traces of the average neutral and ion density over time for three identical baseline simulations in the (a) 2x and (b) 4x cases.

Next, in Fig. 6, we compare the time-dependent ion density curve to that of several other plasma properties. Here, we have averaged each quantity over the two spatial dimensions. For ease of comparison, we normalize these signals by shifting by their mean and scaling by the standard deviation according to

$$\bar{X} = \frac{X - E[X]}{\sqrt{\text{Var}[X]}}, \quad (4)$$

where X and \bar{X} are the un-normalized and normalized signals, respectively; $E[X]$ is the expected value or




Figure 6: Time-dependent traces of the spatially-averaged (a) neutral density, (b) electron temperature, (c) electric field, and (d) ion velocity compared to the ion density oscillations for one of the 2x simulations between $t = 100$ and $t = 160 \mu\text{s}$. The time points t_1 , t_2 , t_3 , t_4 correspond to 106, 112, 116, and 119 μs , respectively. The signal amplitudes have been normalized using Eq. 4. The raw signals are shown in a light color behind the darker curves, to which a low-pass filter with a cutoff frequency of 300 kHz have been applied.

mean of X , and $\text{Var}[X]$ is the variance of X . The neutral and ion densities have a clear phasing relationship, with a rise in neutral density preceding that in ion density by about 60 degrees.

We observe that the time-dependent electric field and electron temperature are noisy, making it difficult to discern a trend with respect to the ion density oscillations. As such, we apply a low-pass filter to these signals to reveal low-frequency behavior. Specifically, we use a Butterworth filter with a cutoff frequency of 300 kHz. Even with this filter, we do not recover a clear relationship between the ion density and the electric field and electron temperature. While less noisy, the relationship between the average ion velocity and ion density also ambiguous, though it appears vary in-phase with the density. Finally, we observe that the anomalous force density and ion density are strongly in-phase, though this is likely due to the direct multiplicative dependence of the force density on the plasma density.

Overall, with the exception of the neutral density, our results are ambiguous with respect to the relative phasing of oscillations in plasma properties, and we do not observe some of the relationships observed in experiments.^{12,34,35} This may be due in part to statistical noise. While we see low run-to-run variation for time-averaged quantities (indicative of low noise), noise plays a larger role in analyzing time-dependent behavior, as seen in the results for the electron temperature and electric field. Additionally, it may be that averaging spatially obscures some local variations which might have clearer phase relationships. To that end, we pick four times which correspond to different phases of the ion density oscillation— t_1 , t_2 , t_3 , and t_4 at 106, 112, 116, and 119 μs , respectively—and plot in Fig. 7 the azimuthally-averaged ion density, electron temperature, electric field and ion velocity at each of these four times.




Figure 7: Instantaneous azimuthally-averaged profiles of the (a) ion number density, (b) electron temperature, (c) electric field, and (d) axial ion velocity for the same 2x simulation and the same four times highlighted in Fig. 6

Interestingly, the instantaneous electric field profile (Fig. 7c) is more strongly-peaked than the time-averaged profile, and that the peak location in phase with the ion density oscillation in Fig. 6. At the low point of the oscillation (t_1), the peak is at its furthest-upstream. It then moves downstream until the high point of the oscillation (t_3) before moving back upstream in t_4 . In Fig. 7d, the steepest part of ion velocity curve (the “acceleration region”) can be seen moving slightly downstream from t_1 to t_3 , but instead of following the electric field back upstream in t_4 , it instead moves further downstream. The location and magnitude of the peak electron temperature also moves axially during an oscillation period, with the lowest and most-downstream value observed at the low point of the ion density oscillation and a movement upstream

toward a higher peak value as the ion density increases. Thus, the peak electron temperature varies in-phase with the ion density and the location of the peak varies out-of-phase.

It is interesting to compare these results qualitatively to experimental studies. Dale and Jorns used laser induced fluorescence (LIF) to infer the time variation of several plasma properties over a breathing mode oscillation in a magnetically-shielded Hall thruster.^{12,34} In contrast to our results, they observed a strong out-of-phase variation in the peak electron temperature and electric field with respect to the plasma density. For both of these quantities, the peak value moved upstream and increased as the plasma density decreased. Our results also differ from those recently obtained by Konopliv et al.,³⁵ which used optical emission spectroscopy (OES) to infer the time-dependent electron temperature in an SPT-70 Hall thruster. Like Dale and Jorns, they also found the electron temperature to oscillate out-of-phase with the discharge current and plasma density. They found, however, that changes in thruster operating conditions could suppress the electron temperature oscillations. It is therefore possible that our chosen plasma properties are not conducive to producing this behavior. Alternatively, given the noise present in the temperature traces, it may be that there were not enough particles in our simulations to accurately resolve any oscillations that may have been present. As we will show later, the particle count fluctuated strongly over the course of the breathing oscillations. This may have led to under-resolution of the electron velocity distribution function at the trough of the oscillation, the point at which we would expect the plasma temperature to peak. In turn, the under-resolution may have prevented the development of temperature oscillations.

Dale and Jorns also observed a strong variation in the anomalous electron collision frequency over the course of a breathing mode oscillation. While we follow Petronio et al. in computing an anomalous force density instead of an anomalous collision frequency, these quantities can be related via the ansatz commonly employed in fluid modeling^{33,36,37}

$$R_{an,y} \approx m_e n_e u_{ey} \nu_e \quad (5)$$




Figure 8: Instantaneous axial profiles of the anomalous collision frequency computed from 2x DSMC simulations. Times are as indicated in Fig. 6

We calculate and plot this quantity in Fig. 8. As observed by Dale and Jorns, the effective anomalous collision frequency varies strongly over the breathing mode cycle. In their results, however, the anomalous collision frequency reached a maximum value at the same time as the ion density, before collapsing to below the classical collision frequency at the trough of the oscillation. In contrast, our simulations show the anomalous collision frequency profile moving axially without changing in magnitude, with the exception of at t_1 , at the point of minimum ion density. Here, the sign of the azimuthal velocity reverses direction, causing the anomalous collision frequency to flip sign to negative.

Next, in Fig. 9, we plot the 2-D instantaneous azimuthal electric field and ion density for each of the highlighted time moments. Compared to the benchmark case,^{6,17,21} the azimuthal oscillations are much more incoherent and vary with time. The wavelength of the azimuthal oscillations also appears longer than in the benchmark. To quantify the effects of ionization on the azimuthally-propagating plasma waves, we plot in Fig. 10 and Fig 11 the wavelength and frequency spectra, respectively, of the azimuthal electric field and ion density oscillations. We compute these spectra at $x = 10$ mm, the location of which we indicate




Figure 9: Instantaneous 2-D electric field (left) and ion density (right) at the same four times highlighted in Fig. 6.




Figure 10: Azimuthal wavelength spectrum of the (a) azimuthal electric field and (b) ion number density, evaluated at $x = 10$ mm.




Figure 11: Frequency spectrum of the (a) azimuthal electric field and (b) ion number density, evaluated at $x = 10$ mm.

in Fig. 9. Compared to our simulations with a static particle source, including ionization increases the amplitude of the spatial azimuthal electric field oscillations by a factor of three to five. Additionally, we observe more spectral content at longer wavelengths, including at wavelengths as long as the domain size. The temporal oscillations display similar trends. In both ion density and electric field, the peak amplitude (not including the breathing mode) is ~ 10 times larger in the DSMC case than in the static case and occurs at a lower frequency. Notably, we observe in the temporal spectrum of the ion density oscillations the strong low-frequency kHz-range peak characteristic of the breathing mode. For our simulations, the dominant breathing mode frequency is between 55 and 60 kHz. This is two to three times greater than those typically observed in xenon thrusters at 300 V,³⁸ an intended consequence of our choice to increase the neutral speed by a similar factor from those observed experimentally.¹²

The higher oscillation amplitudes likely explain the large time-averaged anomalous force density in Fig. 4 compared to the static cases. Furthermore, the longer characteristic wavelengths match what we qualitatively observe in Fig. 9. Lastly, we do not observe a strong breathing mode signal in the spectrum of the azimuthal electric field. This supports our earlier supposition that the anomalous force density oscillations in Fig. 6 stem mainly from the density dependence of that term.

C. Effect of ion recombination at the anode

In Fig. 12, we show the effect of including ion recombination at the anode on the time dependence of the neutral and ion densities. We find that ion recombination induces strong, coherent oscillations, with a peak-to-peak ion density amplitude in 2x case of $1.5 \times 10^{18} \text{ m}^{-3}$ (580% mean). The mean ion density in the case with recombination is also 25% larger than without. Our simulations support the findings from studies using simplified fluid models^{4,27,28} and suggest that while ion recombination is not required for the onset of the breathing mode, it can enhance the development of strong, coherent oscillations and increase their amplitude.




Figure 12: Traces of the spatially-averaged average neutral and ion density over time with ion recombination at the anode enabled (dashed) and disabled (solid) in the 2x DSMC cases.

D. Computational cost

Lastly, we turn discuss the computational cost of our simulations, as well as some numerical issues encountered. In Tab. 4, we list the wall time for each of our simulations. As in our previous work,¹⁷ we find that GPU acceleration combined with grid coarsening are highly enabling for running kinetic Hall thruster simulations at reasonable computational costs, with the most expensive simulation (DSMC 2x) lasting just 4 days and 8 hours.

As mentioned in Sec. G, we were unable to complete a 1x simulation using DSMC due to memory constraints on the H100 GPU. Given the large density fluctuations we observed, particularly during the startup transient, it is likely that an explosion in the number of particles surpassed the memory capacity of the card. To investigate this, we plot in Fig. 13 the particle count and weight evolution of ions, electrons, and

Table 4: Computational cost in wall time for each of the simulations in Tab. 3. Each simulation ran on a single NVIDIA H100 GPU.³⁰

Simulation	Wall clock time
Benchmark, 1x	1 day, 14 hours
Static, 1x	1 day, 20 hours
Static, 2x	10 hours
DSMC, 2x	4 days, 8 hours
DSMC, 4x	1 day, 9 hours

neutrals over the course of one of the 2x DSMC simulations. We find that, indeed, the number of electron and ion particles per cell increases by a factor of 50 in the first few microseconds of the simulation, going from 75 to nearly 4000. In the depletion of particles that follows, the count drops to just eleven per cell. In contrast, by the end of the simulation the particle count fluctuates within the range of just 100 to 200 particles per cell, with a mean value of 157. This is about half of the benchmark simulation, suggesting that we may be operating in a somewhat under-resolved regime in terms of adequately capturing the particle velocity distribution functions (VDFs). The neutral particle count is more stable and eventually levels off at 261 particles per cell.

Fig. 13 shows that changes in neutral particle weight accounted for most of the observed neutral density oscillations, while the ion and electron densities remained constant. This stems from how the DSMC algorithm handles the weight difference between neutral atoms and electrons/ions. Each ionization event removes a small amount of weight from the much “larger” neutral particles while creating electron-ion pairs of constant weights. Neutrals thus steadily lose weight as they move through the domain, but are constantly refreshed by a flux of full-weight particles at the anode. With an ideal resampling algorithm, the electron and ion particle count would also be able to be kept constant while the weights fluctuated. This would produce both stable resolution of the particle VDFs and a reduced computational cost and memory burden for each simulation. While we experimented with particle merging to reduce the computational cost at the high point of the oscillation, we found that we were unable to produce stable simulations using it. Additionally, particle splitting is required to prevent severe under-resolution of the VDFs at the low point of the oscillations, particularly at early times. Future work should focus on examining multiple different resampling approaches in the context of reducing the amplitude of the particle count fluctuations while still capturing physical breathing oscillations.




Figure 13: (a) Evolution of the number of particles per cell for each species over the course of a 2x DSMC simulation. The line ‘i/e steady state avg’ indicates the average particle count of ions and electrons over the last 50 microseconds of the simulation. (b) The average weight per macro-particle over time for each species.

To conclude the discussion of computational cost, we note that the benchmark simulation conditions and those of our other simulation are idealized, and that simulating a more realistic system would be more expensive. Thrusters in the SPT family (SPT-50, SPT-70, and SPT-100) have discharge channel lengths the

same length as the entire benchmark simulation domain, requiring us to at least double the axial domain size, in order to capture the physics in the discharge channel and near-plume region. This would double the cost of the simulations. We also found in our spectral analysis in Sec. B that the amplitude of fluctuations on the scale of the azimuthal domain was high. This may suggest that the azimuthal length of our simulation domain is too short, and might need to be doubled or quadrupled to obtain satisfactory results. Other unrealistic parts of our simulation include the high neutral temperature and velocity, chosen to increase the breathing mode fluctuations frequency. Reducing these to realistic values would double the needed time to capture the same number of oscillations. Finally, the 2D nature of these simulations misses important physics—the amount of energy stored in plasma fluctuations decreases in 3D simulations due to additional loss mechanisms. This may help explain the large anomalous force densities and electron temperatures of our simulations, which are not reflective of experiment. Plasma-wall interactions and sheath physics are also neglected in our work but would be able to be resolved in a 3D simulation. The computational cost of 3D simulations is far larger than our 2D simulations, and would require scaling across tens or even hundreds of GPUs to be tractable. Thus, while kinetic Hall thruster simulations are more affordable than ever, there remain computational and algorithmic challenges to be overcome if they are to be useful in engineering design applications.


IV. Conclusion

In this work, we demonstrated kinetic 2-D Hall thruster simulations including a neutral population undergoing ionization. Leveraging a semi-implicit numerical method, we simulate a Xenon thruster on an axial-azimuthal simulation domain in conditions similar to the well-known benchmark of Charoy et al.²¹ at multiple grid resolutions. We compared these results to similar simulations in which ionization was modeled using a fixed, static particle source similar to the original benchmark case. Our simulations displayed stable breathing mode oscillations as well as time-averaged behavior very different from the benchmark simulation. In particular, the electron temperature and instability-induced drag forces were much larger, while the spectra of plasma fluctuations shifted toward higher amplitudes at lower frequencies and wavenumbers.

We additionally examined the relative phasing of multiple plasma properties with respect to the ion density oscillations. In contrast with some experimental studies, we observed no clear phasing relationship between the density and the spatially-averaged electric field or electron temperature, though both of these quantities displayed some axial displacement over the course of an oscillation period. We attribute these differences from experiment to both statistical noise stemming from fluctuating particle counts and the specifics of the benchmark case. Next, we found that, in line with previous numerical and theoretical studies, incorporating ion recombination at the anode increases the amplitude of the breathing oscillations. Finally, we assessed the computational cost of our simulations and examined the evolution of particle counts and weights over the course of the breathing mode. Despite the large fluctuations in particle count, particularly during the initial startup transient, we were able to complete each of our simulations in less than five days on a single GPU. Our results demonstrate that GPU acceleration and grid coarsening permit kinetic Hall thruster simulations to be run for longer time intervals than have previously been possible, laying the groundwork for the investigation of new physical and numerical phenomena.

Acknowledgments

This work was funded by Los Alamos National Laboratories under the project “Algorithm/Software/Hardware Co-design for High Energy Density applications” at the University of Michigan, and used computing resources provided by an AFOSR DURIP under Program Manager Dr. Fariba Fahroo and grant number FA9550-23-1-006. The authors acknowledge additional computational resources and support provided by Advanced Research Computing, a division of Information and Technology Services at the University of Michigan. This research used the open-source particle-in-cell code WarpX <https://github.com/BLAST-WarpX/warpx>. Primary WarpX contributors are with LBNL, LLNL, CEA-LIDYL, SLAC, DESY, CERN, Helion Energy, and TAE Technologies. We acknowledge all WarpX contributors. Our simulations were performed on WarpX version 25.07.



Data and code availability

The code used to to run Hall thrusters simulations in WarpX, as well as generate and analyze the results of this study, is available online at this Github URL: <https://github.com/archermarx/warpx-hall>. The data generated for this study are available from the authors on request.

References

- ¹Ioannis G. Mikellides and Alejandro Lopez Ortega. Challenges in the development and verification of first-principles models in Hall-effect thruster simulations that are based on anomalous resistivity and generalized Ohm's law. *Plasma Sources Science and Technology*, 28:48, 1 2019.
- ²Jean-Pierre Boeuf. Tutorial: Physics and modeling of Hall thrusters. *Journal of Applied Physics*, 121(1):011101, 01 2017.
- ³Igor D. Kaganovich, Andrei Smolyakov, Yevgeny Raitses, Eduardo Ahedo, Ioannis G. Mikellides, Benjamin Jorns, Francesco Taccogna, Renaud Guerout, Sedina Tsikata, Anne Bourdon, Jean-Pierre Boeuf, Michael Keidar, Andrew Tasman Powis, Mario Merino, Mark Cappelli, Kentaro Hara, Johan A. Carlsson, Nathaniel J. Fisch, Pascal Chabert, Irina Schweigert, Trevor Lafleur, Konstantin Matyash, Alexander V. Khrabrov, Rod W. Boswell, and Amnon Fruchtman. Physics of e×b discharges relevant to plasma propulsion and similar technologies. *Physics of Plasmas*, 27(12):120601, 12 2020.
- ⁴O. Chapurin, A. I. Smolyakov, G. Hagelaar, and Y. Raitses. On the mechanism of ionization oscillations in Hall thrusters. *Journal of Applied Physics*, 129(23):233307, 06 2021.
- ⁵L. Leporini, V. Giannetti, M. M. Saravia, F. Califano, S. Camarri, and T. Andreussi. On the onset of breathing mode in Hall thrusters and the role of electron mobility fluctuations. *Frontiers in Physics*, Volume 10 - 2022, 2022.
- ⁶J. C. Adam, A. Héron, and G. Laval. Study of stationary plasma thrusters using two-dimensional fully kinetic simulations. *Physics of Plasmas*, 11(1):295–305, 01 2004.
- ⁷W. Villafana, B. Cuenot, and O. Vermorel. 3-D particle-in-cell study of the electron drift instability in a Hall thruster using unstructured grids. *Physics of Plasmas*, 30:033503, 3 2023.
- ⁸Francesco Taccogna and Pierpaolo Minelli. Three-dimensional particle-in-cell model of hall thruster: The discharge channel. *Physics of Plasmas*, 25(6):061208, 06 2018.
- ⁹Lihuan Xie, Zhijun Zhou, Xin Luo, and Yinjian Zhao. The influence of adding Monte Carlo collisions on convergence in 3D PIC simulations of Hall thruster azimuthal instability. *Physica Scripta*, 100(7):075616, jul 2025.
- ¹⁰A. Marín-Cebrián, E. Ahedo, and A. Domínguez-Vázquez. Two-dimensional kinetic analysis of a hall thruster discharge with a null-magnetic point at the anode. *Journal of Physics D: Applied Physics*, 58(37):375201, sep 2025.
- ¹¹James Szabo. *Fully Kinetic Numerical Modeling of a Plasma Thruster*. PhD thesis, Massachusetts Institute of Technology, 2001.
- ¹²Ethan Dale. *Investigation of the Hall Thruster Breathing Mode*. PhD thesis, University of Michigan, 2019.
- ¹³Christine M Greve, Manoranjan Majji, and Kentaro Hara. Real-time state estimation of low-frequency plasma oscillations in Hall effect thrusters. *Physics of Plasmas*, 28:93509, 2021.
- ¹⁴Thomas A. Marks and Benjamin A. Jorns. Challenges with the self-consistent implementation of closure models for anomalous electron transport in fluid simulations of Hall thrusters. *Plasma Sources Sci. Technol.*, 32(4):045016, 4 2023.
- ¹⁵Thomas A. Marks and Benjamin A. Jorns. Evaluation of algebraic models of anomalous transport in a multi-fluid Hall thruster code. *Journal of Applied Physics*, 134(15):153301, 10 2023.
- ¹⁶J.-L. Vay, A. Huebl, A. Almgren, L. D. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D. P. Grote, M. Hogan, R. Jambunathan, R. Lehe, A. Myers, C. Ng, M. Rowan, O. Shapoval, M. Thévenet, H. Vincenti, E. Yang, N. Zaïm, W. Zhang, Y. Zhao, and E. Zoni. Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs. *Physics of Plasmas*, 28(2):023105, 02 2021.
- ¹⁷Thomas A. Marks and Alex A. Gorodetsky. GPU-accelerated kinetic hall thruster simulations in WarpX. *Journal of Electric Propulsion*, 4(1):34, May 2025.
- ¹⁸Dan C. Barnes. Improved C^1 shape functions for simplex meshes. *Journal of Computational Physics*, 424:109852, 2021.
- ¹⁹Nirbhav Singh Chopra, Mina Papahn Zadeh, Mikhail Tyushev, Andrei Smolyakov, Alexandre Likhanskii, and Yevgeny Raitses. Multimodal azimuthal oscillations in electron beam generated E×B plasma. *Physics of Plasmas*, 32(7):073511, 07 2025.
- ²⁰A. Myers, A. Almgren, L.D. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D.P. Grote, M. Hogan, A. Huebl, R. Jambunathan, R. Lehe, C. Ng, M. Rowan, O. Shapoval, M. Thévenet, J.-L. Vay, H. Vincenti, E. Yang, N. Zaïm, W. Zhang, Y. Zhao, and E. Zoni. Porting WarpX to GPU-accelerated platforms. *Parallel Computing*, 108:102833, 2021.
- ²¹T Charoy, J P Boeuf, A Bourdon, J A Carlsson, P Chabert, B Cuenot, D Eremin, L Garrigues, K Hara, I D Kaganovich, A T Powis, A Smolyakov, D Sydorenko, A Tavant, O Vermorel, and W Villafana. 2-D axial-azimuthal particle-in-cell benchmark for low-temperature partially magnetized plasmas. *Plasma Sources Science and Technology*, 28(10):105010, oct 2019.
- ²²John M Sankovic, John A Hamley, Thomas W Haaa, and Thomas W Haagt. Performance evaluation of the russian spt-100 thruster at nasa lerc. In *1993 International Electric Propulsion Conference*, page 094, 1993.
- ²³Richard R Hofer, Sarah E Cusson, Robert B Lobbia, and Alec D Gallimore. The H9 magnetically shielded Hall thruster. In *International Electric Propulsion Conference*, 2017.
- ²⁴Sarah E Cusson, Richard R Hofer, Robert B Lobbia, Benjamin A Jorns, and Alec D Gallimore. Performance of the h9 magnetically shielded Hall thrusters. In *2017 International Electric Propulsion Conference*, 2017.
- ²⁵G. A. Bird. Approach to Translational Equilibrium in a Rigid Sphere Gas. *The Physics of Fluids*, 6(10):1518–1519, 10 1963.



- ²⁶Lihuan Xie, Zhijun Zhou, Xin Luo, and Yinjian Zhao. Effect of plasma initialization on 3D PIC simulation of hall thruster azimuthal instability. *Physica Scripta*, 99(7):095602, aug 2024.
- ²⁷Andrei Smolyakov, Oleksandr Chapurin, Ivan Romadanov, Yevgeny Raitses, and Igor Kaganovich. Theory and modelling of axial mode oscillations in Hall Thruster. In *AIAA Propulsion and Energy 2019 Forum*, 2019.
- ²⁸T. Lafleur, P. Chabert, and A. Bourdon. The origin of the breathing mode in Hall thrusters and its stabilization. *Journal of Applied Physics*, 130:053305, 8 2021.
- ²⁹Luca Fedeli, Axel Huebl, France Boillot-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhsa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, and Henri Vincenti. Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2022.
- ³⁰Nvidia Corporation. *NVIDIA H100 Tensor Core GPU Datasheet*, 2024. Accessed: 2024-06-18.
- ³¹Roelof Groenewald, Daniel C. Barnes, Mikhail Tyushev, Weiqun Zhang, Axel Huebl, Ales Necas, Andrei Smolyakov, Jean-Luc Vay, Toshiaki Tajima, and Sean Dettrick. New semi-implicit electrostatic particle-in-cell method to extend scope of the exascale WarpX code. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC24)*, 2024.
- ³²A. Bruce Langdon, Bruce I Cohen, and Alex Friedman. Direct implicit large time-step particle simulation of plasmas. *Journal of Computational Physics*, 51(1):107–138, 1983.
- ³³Federico Petronio, Alejandro Alvarez Laguna, Martin Jacques Guillon, Anne Bourdon, and Pascal Chabert. Refining the modeling strategy for anomalous electron transport in fluid simulations of Hall thrusters via insights from PIC simulations. *Physics of Plasmas*, 32(7):073513, 07 2025.
- ³⁴Ethan T. Dale and Benjamin A. Jorns. Non-invasive time-resolved measurements of anomalous collision frequency in a Hall thruster. *Physics of Plasmas*, 26, 2019.
- ³⁵Mary F. Konopliv, Lee K. Johnson, and Richard E. Wirz. Time-resolved electron temperature oscillations in hall thrusters. *Journal of Applied Physics*, 137(24):243303, 06 2025.
- ³⁶T. Lafleur, S. D. Baalrud, and P. Chabert. Theory for the anomalous electron transport in Hall effect thrusters. II. Kinetic model. *Physics of Plasmas*, 23:11101, 5 2016.
- ³⁷Thomas A Marks, Alejandro Lopez Ortega, Ioannis G Mikellides, and Benjamin Jorns. Self-consistent implementation of a zero-equation transport model into a predictive model for a Hall effect thruster. *AIAA Propulsion and Energy 2021 Forum*, 2021.
- ³⁸Vittorio Giannetti, Manuel Martín Saravia, Luca Leporini, Simone Camarri, and Tommaso Andreussi. Numerical and experimental investigation of longitudinal oscillations in Hall thrusters. *Aerospace*, 8, 2021.

