Development of the Plasma Thruster Particle-in-Cell Simulator to Complement Empirical Studies of a Low-Power Cusped-Field Thruster

by

Stephen Robert Gildea

B.S., Aerospace Engineering, University of Florida (2007)

S.M., Aeronautics and Astronautics, Massachusetts Institute of Technology (2009)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
February 2013

Author .....
Department of Aeronautics and Astronautics
October 12, 2012

Certified by .....
Manuel Martínez-Sánchez
Professor of Aeronautics and Astronautics
Committee Chair/Thesis Supervisor

Certified by .....
Nicolas G. Hadjiconstantinou
Professor of Mechanical Engineering
Committee Member

Certified by .....
Dr. William A. Hargus, Jr.
Research Scientist, Air Force Research Laboratory
Committee Member

Certified by .....
Paulo Lozano
Associate Professor of Aeronautics and Astronautics
Committee Member

Accepted by .....
Eytan H. Modiano
Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Development of the Plasma Thruster Particle-in-Cell Simulator to Complement Empirical Studies of a Low-Power Cusped-Field Thruster

by

Stephen Robert Gildea

Submitted to the Department of Aeronautics and Astronautics
on October 12, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Cusped-field plasma thrusters are an electric propulsion concept being investigated by several laboratories in the United States and Europe. This technology was implemented as a low-power prototype in 2007 to ascertain if durability and performance improvements over comparable Hall thruster designs could be provided by the distinct magnetic topologies inherent to these devices. The first device tested at low-powers was eventually designated the "diverging cusped-field thruster" (DCFT) and demonstrated performance capabilities similar to state-of-the-art Hall thrusters. The research presented herein is a continuation of these initial studies, geared toward identifying significant operational characteristics of the thruster using experiments and numerical simulations.

After a review of hybrid, fluid, and particle-in-cell Hall thruster models, experimental contributions from this work are presented. Anode current waveform measurements provide the first evidence of the distinct time-dependent characteristics of the two main modes of DCFT operation. The previously named "high-current" mode exhibits oscillation amplitudes several factors larger than mean current values, while magnitudes in "low-current" mode are at least a full order smaller. Results from a long-duration test, exceeding 200 hours of high-current mode operation, demonstrate lifetime-limiting erosion rates about 50% lower than those observed in comparable Hall thrusters.

Concurrently, the plasma thruster particle-in-cell (PTpic) simulator was developed by upgrading numerous aspects of a preexisting Hall thruster model. Improvements in performance and accuracy have been achieved through modifications of the particle moving and electrostatic potential solving algorithms. Data from simulations representing both modes of operation are presented. In both cases, despite being unable to predict the correct location of the main potential drop in the thruster chamber, the model successfully reproduces the hollow conical jet of fast ions in the near plume region. The influences guiding the formation of the simulated beam in low-current mode are described in detail.

A module for predicting erosion rates on dielectric surfaces has also been incorporated into PTpic and applied to simulations of both DCFT operational modes. Two data sets from high-current mode simulations successfully reproduce elevated erosion profiles in each of the three magnetic ring-cusps present in the DCFT. Discrepancies between the simulated and experimental data do exist, however, and are once again attributable to the misplacement of the primary acceleration region of the thruster. Having successfully captured the most significant erosion

profile features observed in high-current mode, a simulation of erosion in low-current mode indicates substantially reduced erosion in comparison to the more oscillatory mode. These findings further motivate the completion of low-current mode erosion measurements, and continued numerical studies of the DCFT. Additionally, PTpic has proven to be a useful simulation tool for this project, and has been developed with adaptability in mind to facilitate its application to a variety of thruster designs — including Hall thrusters.

Thesis Supervisor: Manuel Martínez-Sánchez
Title: Professor of Aeronautics and Astronautics

Acknowledgments

I dedicate this thesis to my wife Erin, son Thomas, and to both of our families. There were many different instances of family members stepping forward to help us during the tougher moments of our stay in the Boston area. Through each obstacle, there was no doubt that we had their love and support. I would especially like to thank my Mom, my Dad, my step-father Glen, brother Mark, and Erin's parents Tom and Dori for helping out as much as you did. I could not physically have done this without all of your help. Many other family members and friends have been supportive in a number of ways, and I just want to say here that I am deeply appreciative. You have all made our time here in Boston an overwhelmingly positive and memorable one.

To my adviser, Manuel, I am thankful for the opportunity to have learned so much from him. Aside from his broad and thorough technical knowledge of so many things, his demeanor and approach to research are traits that I hope to see develop in myself as I begin my postgraduate career.

Contents

Chapter 1

Introduction

Interest in small satellites continues to grow as newly demonstrated capabilities lead to more ambitious mission objectives. In addition to cost savings associated with lighter payloads, more launch opportunities are available for systems that can be accommodated by multiple launch vehicles. With this flexibility, the use of small satellites allows more independent objectives to be met for a given cost, increasing accessibility to space. Advantages are offset somewhat by more stringent constraints on propulsion systems, diminishing or eliminating the maneuverability of small satellites in orbit. Satellites piggybacking with larger payloads are then forced to operate exclusively in the insertion orbit, and concerns about the accumulation of orbital debris may limit the number of launch opportunities for small satellites in the future. These difficulties could be mitigated by extending the benefits of electric propulsion to this class of satellites, increasing their utility to research, industrial, and governmental communities.

The study of more efficient and durable low-power Hall thruster concepts is an active area of research in many countries[1–5]. This is motivated by the operational heritage of Hall thrusters in space, with proven efficiencies and lifetimes greater than 50% and 7000 h[6,7]. Figure 1-1(a) illustrates the standard configuration of a Hall thruster, characterized by an annular discharge chamber, anode, and neutral injector. An electromagnetic coil typically supplies the magnetic field, which is dominantly radial throughout the chamber. Magnetic fields obtain maximum values of several hundred gauss near the exit plane. Typically, a noble gas such as xenon is injected into the thruster near the anode while a cathode located downstream emits low energy electrons. The anode is biased positive relative to cathode potential, initiating ionization

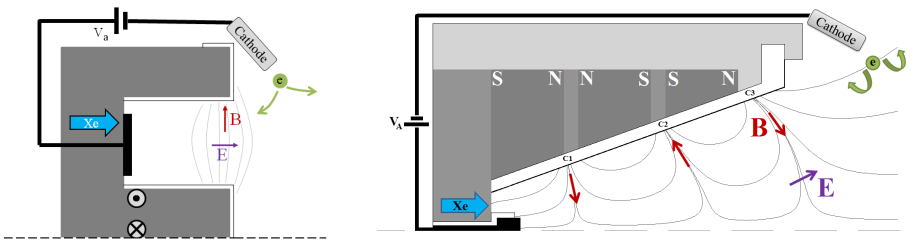

(a) Schematic representation of a Hall thruster.

(b) DCFT schematic with representative magnetic field lines overlayed.

Figure 1-1: Schematics representing Hall and cusped-field thrusters.

of the gas via collisions with electrons accelerating upstream. Electromagnets are used to establish a mostly radial magnetic field with a maximum of several hundred gauss near the exit plane, providing partial confinement of electrons in a closed Hall drift perpendicular to the applied fields. The much heavier ions are not confined by the magnetic field, and are instead accelerated away by the electric field, providing thrust. The electric field in the region of electron confinement is needed to balance the magnetic force on the electron current. Without the magnetic field to hinder the axial mobility of electrons, they would have too short a residence time in the thruster and would be lost to the anode without the necessary ionization taking place.

Scaling Hall thruster designs to low powers while maintaining high efficiencies has proven challenging, especially below 200 W, as shown in Table 1.1[a]. Alternatively, plasma thrusters with markedly different magnetic topologies and field strengths have also been examined. The cylindrical hall thruster (CHT) concept has been studied at low to moderate power levels, but recent efforts have focused on low-power applications [9,10]. The magnetic field strength is larger than in traditional Hall thruster designs, with a maximum value of $\approx 1\ kG$. Near the exit, the magnetic field lines are mostly radial, converging axially further upstream along the centerline. This topology is similar to the end-Hall thruster [11], though the field is much stronger in the CHT. Also, the anode remains annular in CHT designs, to give electrons a longer cross-field path to the anode than if it were placed on the axis of symmetry. More detailed aspects of the design depend on whether the magnetic field is provided by electric currents or permanent

[a] In Table 1.1 it is necessary to make a distinction between "lifetime" and "soft-failure" times because some authors provide useful thruster lifetimes in excess of soft-failure times. A "soft-failure" [8] is said to occur when the insulator is compromised but the thruster and satellite remain capable of operating within specifications.

Table 1.1: Survey of low-power Hall thrusters for which lifetime and performance data were found in the literature. The primary reference is listed with the thruster designation, while additional sources are indicated in corresponding columns. The listed anode powers correspond to operating conditions used during the long duration tests, while representative anode efficiencies are listed in cases where efficiency measurements at the tested anode power were not available.

Thruster Desig.	Anode Power [W]	Anode Eff.	Anode Voltage [V]	Soft-Fail. Time [h]$\times 10^3$	Predicted Lifetime [h]$\times 10^3$	Test Duration [h]$\times 10^3$
SPT-50 [13]	320	47%	170-180	$>2.5^{\times}$	$\geq 2.5^{\times}$	.83
KM-45 [14]	310	40-50% [1]	310	3.5-4	3.5-4	1
KM-32 [2]	200	30-40%	250	2-3	3^{*}	.5
BHT-200 [15]	200	43.5%	250	1.3-1.5	>1.7	>1.7
HT-100 [16]	175,100 [17]	25% [4]	300	.3 [18],>.45 [17]	1.5^{**} [18]	.45
SPT-30 [19]	150	26% [20]	283	$.6^{\times}$	$.6^{\times}$	-
SPT-20M [21]	$\leq$100	$\leq$38% [5]	-	.59-.91	4	1

$\times$ Details of test or method leading to the prediction not found in literature.
$*$ Center pole is measurably eroded before hour 500 of testing.
$**$ Predicted lifetime pertains to 100 W anode power, rather than the tested 175 W condition.

magnets, or how the magnetic poles are arranged [12]. The anode efficiency of the thruster varies between 13-38% over a power range of 60-220 W, depending on the thruster design, cathode keeper current, and which parameters are included in the efficiency calculation [10]. Although no results quantifying erosion rates or lifetime for the CHT have been reported, a smaller surface-area-to-volume ratio in the discharge chamber has been suggested as one reason a CHT would be more durable than a Hall thruster operating at an equivalent anode power.

1.1 The Diverging Cusped-Field Thruster

The diverging cusped-field thruster (DCFT) is another low-power alternative presently being investigated. As shown in Figure 1-1(b), permanent ring-magnets of alternating polarity downstream of a magnetic pole piece give rise to several regions containing convergent magnetic field lines, referred to as cusps. In addition to the three ring-cusps, point-cusps are established on the thruster axis to either side of each ring-cusp — at the anode and downstream of C3, for instance. The ring-cusps are referred to in their order proceeding downstream from the anode: the first cusp (C1), second cusp (C2) and third cusp (C3). Unlike CHTs and Hall thrusters, the

anode is cylindrical, located on the thruster axis with magnetic field lines mostly normal to its surface. Rather than cross-field electron diffusion, electron mobility near the anode is impeded by a magnetic mirroring effect. Field strengths along the axis exceed 4 kG near the anode and 1 kG between C1 and C2, but decrease to zero where each ring-cusp separatrix intersects the centerline. A separatrix is a special surface which can be thought of as the divider between different magnetic cells. More rigourously, a separatrix is defined as a surface on which the magnetic vector potential is zero valued. A more detailed map of magnetic field strength in the thruster chamber is provided in Figure 1-2.

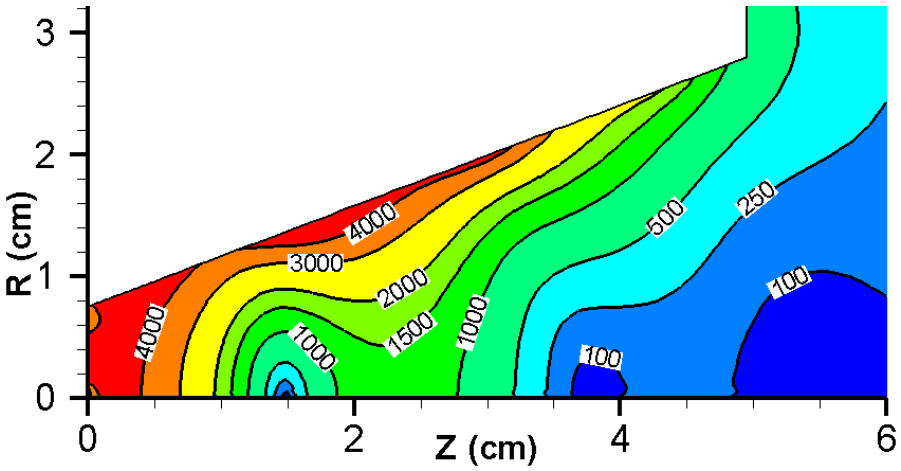

Figure 1-2: The magnetic field magnitude, plotted in units of gauss, within the interior of the DCF thruster. These are numerical calculations made using the MAXWELL SV software package. The downstream face of the anode is located at the z=0.5 cm mark, but is not shown here.

The use of cusped magnetic fields in plasma thrusters did not begin with the DCFT concept. DC discharge ionization chambers in ion thrusters have long utilized magnetic cusps for plasma confinement[22]. An electrodeless traveling wave cusped-field accelerator was also investigated as early as 1962 by G. S. James et al. in conjunction with early Hall thruster concepts[23]. However, a more direct precursor to the DCFT was the development of high efficiency multistage plasma (HEMP) thrusters by Kornfeld et al. in Germany[24]. HEMP thrusters also use permanent magnets to create magnetic cusps in the thruster chamber, and testing thus far has focused on power levels near 1.5 kW for applications on telecommunications satellites[25]. Engineering models of the HEMP 3050 have reported efficiencies of 45% with thrust and specific impulse values of 50 mN and 3000 s[24]. More provocative is the lifetime estimate, stated to

be in excess of 18,000 h based on reported erosion rates that are smaller than those measured in comparable thrusters by a factor larger than 200 [26].

The DCFT was developed to determine if the use of cusped magnetic field topologies could provide longer lifetimes to low-power thrusters with performance capabilities comparable to Hall thrusters. Performance similar to the BHT-200, a commercially developed and flight-proven thruster, was demonstrated at Busek in 2007. At a flow rate of 8.5 *sccm*, anode efficiencies were measured between 40-45% at anode powers between 185-275 W [27]. More recent measurements are likely in agreement with the initial results when the uncertainties in each data set are considered [28]. For a brief introduction to basic concepts and terminology related to propulsion, refer to Appendix A. Two operational modes were also identified, termed the high-current (HC) and low-current (LC) modes. These names describe how, for a given anode flow rate, the anode current decreases when the anode voltage is increased above a threshold value. The change of modes is also accompanied by a change in the visual appearance of the plume, as shown in Figure 1-3. Later, the HC-mode was found by the author to be accompanied by large-amplitude, low-frequency anode current oscillations, while those same oscillations were absent or greatly reduced in the LC-mode [29,30]. Additional studies have included spectroscopy of the DCFT plasma [31], the effects of alternate magnetic topologies near the exit plane [32,33], and particle-in-cell simulations [34]. Emissive probe measurements [32,35,36] and ion velocimetry [37] provide strong evidence that the location of the acceleration region in cusped-field thrusters is strongly correlated with the location of the separatrix nearest the exit plane. A similar design with a cylindrical chamber has also been investigated by the Stanford Plasma Physics Laboratory [36]. A more detailed summary of prior investigations of the DCFT is provided in Section 1.1.2 following a brief description of estimated plasma parameters in this device.

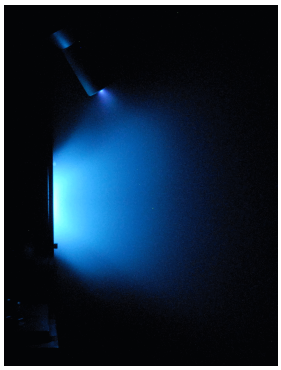

(a) Thruster operating in high-current mode.

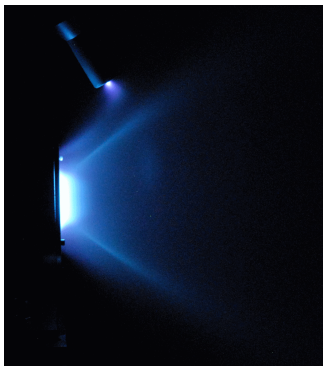

(b) Thruster operating in low-current mode.

Figure 1-3: DCF Thruster operating in Chamber 1 at the Edwards AFB AFRL.

1.1.1 Representative Plasma Parameters

Representative plasma densities and temperatures in Hall or cusped-field thrusters are $n_e = 10^{11}$–10^{13} $[cm^{-3}]$ and $T_e = 20$ $[eV]$[b]. Only transitory confinement is desirable: the goal of a plasma thruster is to produce and direct ions downstream of the thruster with maximum effectiveness. Electrons gain energy as they are emitted from the cathode and travel upstream toward the anode. Inhibited from streaming freely to the anode by the magnetic field, electrons collide primarily with neutral particles. However, the collisions that do occur are not sufficient to equilibrate temperatures between the energetic electrons and the heavier species while traversing through the thruster. The disparity in mass between electrons and the xenon nuclei accounts for the poor energy transfer between the two populations. Measured neutral temperature values in Hall thrusters are of the same order as ceramic wall temperatures[38]. Minimum ion temperatures in the plume and near-plume regions are typically on the order of 1 $[eV]$[39,40], though values closer to 0.1 $[eV]$ have also been measured[41]. In regions where ion energies are mixed, ion temperatures can be an order of magnitude larger — though, in these cases a temperature provides a weaker description of ion distributions. Localized electron temperature maxima are measured between $15 - 40$ $[eV]$[42] meaning that the relative magnitudes of temperatures are summarized concisely as $T_e \gg T_i \gg T_n$. Due to the preferential depletion of high-energy electrons by wall losses (necessary to overcome the sheath potential) and inelastic processes, the

[b] As a reference, the number density of ambient air is 2.5×10^{19} $[cm^{-3}]$, 1 $[\mu Torr] = 1.3 \times 10^{-9}$ [atm], and 1 $[eV] = 11,604.5$ $[K]$

ionization fraction is finite and the remaining neutral density is important to consider. Some sense of the non-equilibrium state of electrons is garnered from the fact that the Saha equation predicts nearly complete single-ionization of xenon for T_e on the order of 1 $[eV]$. Since the neutral population is important to consider in Hall thrusters, even in the presence of electron temperatures exceeding 10 $[eV]$, high-energy portions of electron distributions are not restored quickly enough by electron-electron collisions — or any other processes.

From the simple estimates given above, some important length and time scales can be estimated. Definitions of the terms used to characterize plasmas in this section are provided in Appendix B. The Debye length is of the order of tens of micrometers, and the plasma frequency is approximately 60×10^9 $[^{rad}/_s]$ — corresponding to a period of about a tenth of a nanosecond. The maximum electron cyclotron frequency in Hall thrusters is about 4×10^9 $[^{rad}/_s]$, while the corresponding value in the original DCFT design is 90×10^9 $[^{rad}/_s]$. Electrons have Larmor radii ranging from $.8 - .03$ $[mm]$ for field strengths between $200 - 5000$ $[G]$, while none except the lowest energy ions have any chance of being magnetized at the higher field strengths. The end of Appendix B includes a more detailed discussion of ion magnetization. The Bohm speed for xenon ions with the estimated parameters is approximately $2 - 4$ $[^{km}/_s]$.

1.1.2 Review of DCFT Research

In order to be a viable option to small satellites as a source of propulsion, and even an alternative to Hall thrusters in the low-power regime, the previously discussed performance levels of the DCFT should be maintained or improved upon while concurrently demonstrating advantages in lifetime and lessened beam divergence. A detailed understanding of the physics involved in cusped-field thruster operation is obviously vital to identify which combinations of geometry and magnetic topology will provide these attributes, if any. This section summarizes how past investigations have contributed to our present understanding of the DCFT, and identifies areas of interest to be studied in more detail.

Questions as simple as "How many cusps are really necessary?" and "What should the magnetic field strength at the anode and near the walls be?" are both of fundamental importance, though no satisfactory answers can be given as yet. In most cases, only qualitative arguments, or comparisons between experiments run using a baseline design and a slightly perturbed vari-

ant, can be used. For instance, we suspect that the field strength at the anode should be large to discourage excessive electron collection by the anode. However, it is difficult to assign an upper bound on how strong the field should be because the thruster cannot operate without collecting at least as much current at the anode as there is ion current in the beam. A simple test designed to gauge the sensitivity of performance on the number of cusps between the anode and exit consisted of moving the anode upstream to the point between C1 and C2 in Figure 1-1(b), where the field strength reaches its maximum value along the central axis. Measurements of anode current at identical flow rates and voltages, compared to baseline conditions, suggest a significant decrease in efficiency[c], as the anode power nearly doubled for all flow rates tested. Therefore, it appears that each cusp present in the original design is significant in determining the net impedance of electrons to the anode. The effect of additional cusps has not been examined experimentally, but may offer diminishing returns when cost, complexity, and mass are considered. Also of interest, the HC-mode oscillation frequency increased by roughly the same factor as the length between the last cusps and the anode was decreased. These observations support arguments linking the HC-mode oscillation frequency to the time scale associated with the neutral time-of-flight from the anode to the ionization region [29,30].

A conceptually similar approach was taken to understand why the hollow structure of the plume is present in *both* modes, despite their distinct visual appearances, and to gain insight into which aspects of the design contribute to the presence of the observed bifurcation. Initially, the conical shape of the thruster chamber was suspected. However, the first tests of a cylindrical variant of the MIT design at Stanford University [36] showed the same distinct plume shape as seen in Figure 1-3(b). This surprising result required a different explanation, which shifted focus toward considering how electrons travel from the cathode to the thruster interior. It was realized that nearly all electrons in the plume would end up in the last downstream cusp, C3, because it is the locus for all magnetic field lines going into the plume.

Conceptually, and not to be taken too literally, each field-line in the plume provides a path for electrons to reach C3 — where their magnetization suggests an analogy relating each field-line to a current carrying wire with some small but finite resistance. These tubes of electron

[c] In order for the anode efficiency not to decrease, the beam current would have needed to increase by the same margin, which would likely require a large increase in the utilization efficiency. This scenario would be unexpected, but has not been explicitly ruled out by existing data from DCFT tests.

current all converge toward the separatrix in C3. No field line upstream of the separatrix reaches the plume, just as no line downstream of the separatrix penetrates upstream of it. As field-line density increases along the separatrix approaching the wall — continuing with the wire analogy — the number of resistors in parallel between two nearby surfaces, and along the direction of the $\underline{B}$-field, increases. Based on this heuristic argument, the regions of maximum parallel conductivity should be the separatrices, meaning that the electric field at the exit should be roughly perpendicular to the separatrix in C3.

This picture is consistent with the observed plume structure in the conical and cylindrical cusped-field thrusters, where the last separatrices share similar contours. The easiest way to test this hypothesis, short of building a new thruster, was to manipulate the magnetic topology using an electromagnet positioned just past the exit plane. However, the shape of the separatrix is less sensitive to the electromagnet closer to the thruster wall — and therefore nearer to the permanent magnets — so that a complete modification of the separatrix was not possible using this approach. Nonetheless, ion current density measurements in the plume clearly show that some fraction of the beam is focused by steering the inner portion of the separatrix toward being perpendicular to the center axis[33]. Furthermore, these modifications of the separatrix do not have a significant effect on the measured ion energies in the plume. Concurrent measurements of ion velocities in the original design — recorded using the non-perturbing laser-induced fluorescence (LIF) diagnostic — provide further experimental support of the significance of the magnetic separatrix[43]. A sample of these data are provided in Figure 1-4(a). Kinetic simulations of the DCFT, shown in Figure 1-4(b), also predict that the separatrix can be regarded as a line of nearly constant potential.

The initial goal of the DCFT design was to investigate alternative magnetic topologies as a way to provide longer lifetimes to low-power plasma thrusters — specifically as an alternative to low-power Hall thrusters. Direct measurements of ion current density j_i in the chamber were obtained by mounting probes in a modified insulating cone. Qualitatively, these probe data indicate that the vast majority of ions reaching the wall do so in the vicinity of the cusps[44]. Quantitative interpretations of the data are hindered by disturbances of the bulk plasma by the probes, and by the presence of a strong and often oblique magnetic field along the wall. The end result of ion losses to the wall is erosion of the insulator, which was measured by operating the

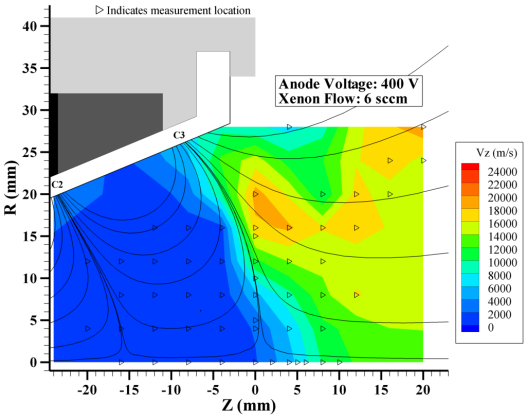

(a) Ion axial velocity component measurements for the DCFT operating in the LC-mode. These data were originally published by MacDonald, et al.[43] The ion velocity increases sharply across the separatrix. Measurement locations including the radial component cannot be obtained upstream of the exit plane due to line-of-sigh constraints, but measurements downstream of the exit indicate that the bulk ion velocities are roughly perpendicular with the separatrix contour.

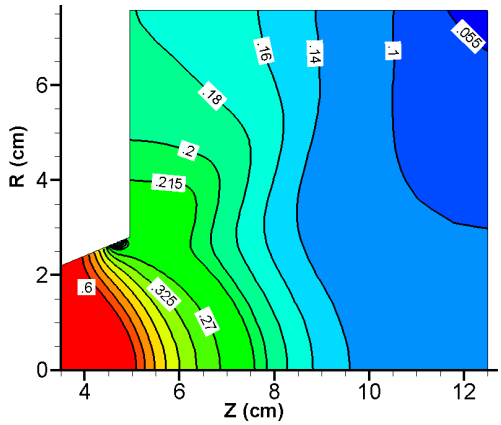

(b) Normalized potential distributions in the near plume of the DCFT. Potential is normalized by 550 V and unlabeled lines between 0.6 and 0.325 step down in increments of 0.055. A comparisons of constant potential contours with the field lines in C3 — shown in Figure 1-1(b) — confirms that the kinetic model predicts an electric field mostly perpendicular to the last separatrix[34].

Figure 1-4: Experimental and simulated data strongly suggest the last separatrix has a strong role in shaping the electric field near the exit plane of the thruster.

thruster continuously for over 200 [h] at an Air Force Research Laboratory (AFRL) facility and measuring surface profiles beforehand and afterwards. The erosion rate in the original DCFT design was less than half of rates measured in comparable Hall thrusters and localized within cusps[45]. Surprisingly, the highest erosion rate was measured to occur in C2, rather than at the exit as observed in Hall thrusters. Measurements in the LC-mode would have been preferred, but the thruster would revert to HC-mode after several minutes of operation. This effect has only been consistently observed in one vacuum chamber at this particular AFRL facility — and prevented detailed measurements of j_i in the plume for LC-mode as well[29]. The reason for this preference is that less erosion is expected to occur in LC-mode, especially in the vicinity of C2. This hypothesis is based on the fact that the HC-mode is suspected to be a manifestation of an ionization instability, consuming a significant fraction of neutrals present in the chamber. As the plasma evacuates from the chamber, the depressed potential structure in each cusp could conceivably steer low energy ions preferentially toward the walls where they are accelerated through the sheath potential. Field lines connected to C2 occupy the largest volume in the interior of the thruster — see Figure 1-1(b) — suggesting that the greatest net erosion occurs in the second cusp in HC-mode because a significant volume of ions drains into this region during discharge oscillations. Furthermore, electron temperatures in the steadier LC-mode are thought to be lower than in the HC-mode, which is plausible based on a comparison of xenon spectra measured in each mode[31]. A lower electron temperature allows for a smaller potential difference between the wall and plasma bulk, reducing the expected ion energy at the walls — and the maximum erosion rate as well.

The existence of two modes of operation began as a peculiarity, and remains one of the least understood aspects of DCFT operation. For a fixed flow rate, as the anode voltage is increased, the thruster typically reaches a threshold at which a sharp transition from an appearance depicted in Figure 1-3(a) gives way to the less diffuse plume in Figure 1-3(b). The name of this second mode is motivated by the fact that the DC anode current I_a decreases by as much as one-half the current in the previous mode, and at the same time as the visual change. The opposite transition, from LC to HC-mode, is brought about by increasing flow-rate at a constant anode voltage. The transition of the waveform for $I_a(t)$ is shown clearly in Figure 1-5, where the gradual disappearance of the current spike is observed as the flow rate decreases[30].

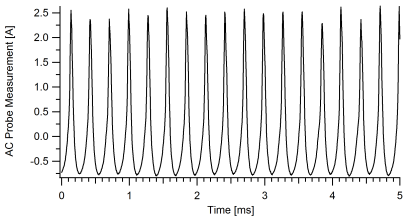

(a) Example of a common HC-mode AC waveform, marked by regular large amplitude current peaks.

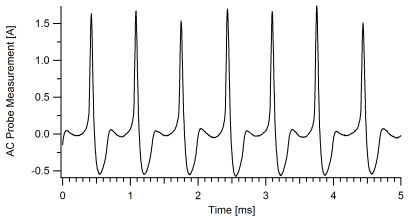

(b) Another commonly encountered HC-mode AC waveform. Here, every other peak is suppressed.

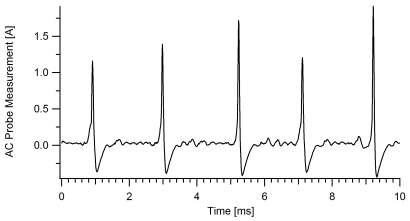

(c) This waveform represents the thruster very near transition to the LC-mode. When the flow rate is lowered further, the large peak no longer occurs and the trace takes on the appearance seen in Figure 1-5(d)

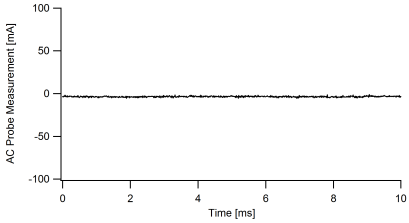

(d) Example a LC-mode AC current trace. Peak-to-peak amplitudes are typically on the order of several milliamperes.

Figure 1-5: AC coupled waveforms measured during DCFT transition from LC-mode to HC-mode. The transition was induced by increasing the xenon flow rate incrementally at constant anode voltage. The thruster operated in LC-mode in Figure 1-5(d), and stepped through the intermediate states until reaching the full blow HC-mode shown in Figure 1-5(a).

Additionally, the peak-to-peak magnitudes of $I_a(t)$ in each mode are plotted in Figure 1-6 for comparison. As suggested earlier in the discussion of erosion, there is strong evidence to support that the neutral refill timescale determines the HC-mode oscillation frequency. Furthermore, the width of the current pulse is consistent with estimates of the residence time of ions leaving the chamber at the Bohm speed. However, the mechanism responsible for triggering the instability is not known. One possibly critical piece of information is that the cylindrical variant tested at Stanford has only been observed to operate with a steady anode current, likely equivalent to the LC-mode of the MIT DCFT. This suggests that, even if the instability mechanism cannot be identified, it may be avoidable at the operating conditions of interest. This latter situation is much preferred, as there may actually be a multitude of contributing factors capable of initiating the instability leading to HC-mode.

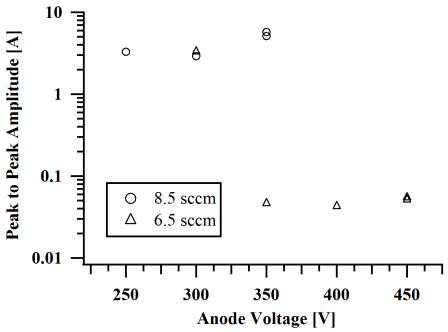

Figure 1-6: The peak-to-peak magnitude of anode current oscillations is given over a range of anode voltages. The transition from HC to LC-mode occurs when the anode voltage is increased from 300 V to 350 V for the 6.5 sccm case.

1.2 Numerical Hall Thruster Models

At some level, the Boltzmann equation is the starting point for any model attempting to describe the behavior of plasmas in electric propulsion devices[d]. For Hall thrusters[e] the plasma consists of three constituent species: electrons, ions, and neutrals of a high-purity noble gas — typically xenon. Within these subdivisions, neutrals and ions exist in various excited electronic states [31], and multiply charged ions may also be present in significant quantities [46]. Macroscopic electric $\underline{E}$ and magnetic $\underline{B}$ fields attempt to confine and accelerate the plasma as well. Current densities are not large enough, nor do electric fields vary fast enough, to necessitate accounting for induced magnetic fields in most cases[f]. However, the high mobility of electrons in the plasma has a dominant effect on the electric field, necessitating solutions dependent upon the distribution of charges, as well as on boundary conditions.

Plasma interactions with confining surfaces are also significant. Secondary electron emission resulting from electron impacts on dielectric surfaces has been observed to have a non-negligible effect on discharge parameters [48], and enhanced conductivity across the $\underline{B}$-field due to wall collisions is frequently cited as a contributing factor to explain higher-than-classical levels of electron mobility perpendicular to $\underline{B}$ [49]. Erosion of internal surfaces by ions lost to the walls

[d] The validity of describing Coulomb collisions by the binary Boltzmann collision operator is discussed briefly in Section 1.2.3.

[e] Many of the descriptions given of Hall thrusters provided here are equally valid for cusped-field thrusters. The main differences enter in discussions of the magnetic field strength, topology, geometric design, and the variation in neutral density as the channel area increases toward the exit.

[f] Non-Intrusive measurements of magnetic field during thruster operation suggest that induced-field effects may have an influence on the precise location of the acceleration region, though it is not a dominant effect [47].

reduces efficiency and ultimately determines the lifetime of Hall thrusters. Predicting erosion profiles during the thousands of hours required of flight qualified thrusters is especially challenging. Accurate numerical studies must attempt to extrapolate future profiles from predicted ion fluxes based on an incomplete understanding of erosion processes [15,50,51]. Even then, measurements needed to benchmark results are limited in availability due to the high costs of operating a thruster and vacuum chamber, or by proprietary or otherwise restricted data. The operation of the cathode and its coupling to the thruster and plume is also important to consider [52,53]. Finally, as the vast majority of testing is done here on Earth, the effects of the testing environment, such as the vacuum chamber size and the persistence of background gases despite the substantial pumping capabilities of many facilities, should be kept in mind when interpreting and comparing data [54,55].

From the present discussion, it is clear that a complete model of Hall thrusters based entirely on first principles would include the simultaneous solution of many multi-dimensional Boltzmann equations, coupled to each other through particle interactions and/or Maxwell's equations — surface interaction effects would also need to be included, and the cathode would need to be in the domain of the problem. Clearly, a completely accurate solution to such a problem is intractable analytically and impossible numerically — considering the availability of computational resources and a limited knowledge of particle-particle and particle-surface interactions. Like nearly all physical situations of interest, researchers are forced to make approximations to the rigorous equations in order to make predictions and understand important trends — while at the same time relying on experiments to validate or repudiate arguments and help identify further areas of interest. It is important, however, to remember and justify the approximations made, in order to maintain a clear picture of the limitations and validity of resulting predictions.

With this in mind, the goal of this section is to describe three broad classes of models used to study Hall thrusters. The validity of the fluid description of plasmas, as often applied to Hall thrusters, is discussed through a detailed examination of a model recently developed at the Jet Propulsion Laboratory. This is significant because this model may represent a trend away from more complicated kinetic and hybrid models within the electric propulsion community. The basics of kinetic particle-in-cell (PIC) models are then discussed, with emphasis placed on the

model developed at the MIT Space Propulsion Laboratory (SPL). This model has and will be used to study the cusped-field thruster concept, which is the main topic of this thesis. First, to bridge the gap between fluid and kinetic models, the hybrid approach will be discussed. A hybrid model often applied to Hall thrusters is known as HP-Hall, and was also developed at the MIT SPL[56]. Hybrid Hall thruster models have been widely adopted as a research tool in the electric propulsion research community[15,57–61]. These descriptions are by no means exhaustive, but do provide a sense of the different approaches taken to understand the detailed physics of plasma thrusters and improve their designs.

1.2.1 Hybrid-PIC models: HP-Hall

A compromise between the complexity of PIC models and the oversimplifications made by fluid models is realized in hybrid thruster models. Typically, electrons are considered as a continuum in local equilibrium, while one or more heavy species is modeled as a set of discrete particles. The hybrid methodology developed by Fife has become widespread in the field of Hall thruster research and development[15,58–61]. In hybrid models, the electron equations are subject to the same limitations as other fluid models when it comes to describing rarefied or transitional flows. These limitations are discussed in Sec. 1.2.2. The following discussion draws from the work of Parra, et al.[62] and Fife's PhD thesis[56].

The governing equations for electrons are solved along and perpendicular to magnetic field lines. Thus, a coordinate whose value is constant along magnetic field lines is desirable, and is provided by a scalar function λ whose gradient is normal to magnetic field lines at all points. The resulting relationship between the partial derivatives of $\lambda(r, z)$ is provided in Equation 1.1. Azimuthal components of the magnetic field are assumed to be zero, and the only nonzero component of the vector potential $\underline{A}$ is A_θ. Contributions to the curl of the magnetic field from plasma current densities and unsteady electric fields are neglected, requiring that $\nabla \times \underline{B} = 0$. As a result, $B_r = -\frac{\partial A_\theta}{\partial z}$ and $B_z = \frac{1}{r}\frac{\partial}{\partial r}(rA_\theta)$.

$$\nabla\lambda \cdot \underline{B} \equiv 0 = -\frac{\partial A_\theta}{\partial z}\frac{\partial \lambda}{\partial r} + \frac{1}{r}\frac{\partial}{\partial r}(rA_\theta)\frac{\partial \lambda}{\partial z} \tag{1.1}$$

$$\frac{\partial}{\partial \hat{n}} f(\lambda) = \frac{\partial \lambda}{\partial \hat{n}}\frac{\partial}{\partial \lambda} f(\lambda) = rB\frac{\partial}{\partial \lambda} f(\lambda) \tag{1.2}$$

As a result, any λ satisfying Equation 1.1 can be a stream function, and Fife chooses $\lambda = -rA_\theta$. This definition is useful for calculating derivatives perpendicular to field lines. For some function of λ, the derivative normal to the field lines is given in Equation 1.2.

In cylindrical coordinates with azimuthal symmetry, electrons are strongly magnetized with a conductivity along magnetic field lines that is much greater than their perpendicular conductivity. Electron inertia terms are ignored, the electron temperature T_e is assumed constant along field lines, and the equilibrium electron momentum equation is projected along the magnetic field. The result of the approximation is a balance between pressure and electrostatic forces, leading to the conclusion that a thermalized potential ϕ^* is constant along field lines[63]. The thermalized potential can only vary across field lines, so it is a function of the magnetic stream function λ only, shown in Equation 1.3. In Equation 1.3, ϕ is the electrostatic potential and n_e is electron density normalized by an arbitrary factor.

$$\phi^*(\lambda) \equiv \phi - \frac{T_e}{e}\ln(n_e) \tag{1.3}$$

A generalized Ohm's law is used to describe electron flow perpendicular to the magnetic field using an effective electron collision frequency, given in Equation 1.4. This equation ignores any azimuthal variations in the scalar pressure, as well as the azimuthal electric field. It can be recast as an equation for the perpendicular electron current density $\underline{j}_{e,\perp}$ in terms of functions along field lines by using Equation 1.2 and 1.3[62]. Here, ν_e is the total electron collision frequency.

$$\begin{aligned} \underline{u}_{e,\perp} &= \frac{\nu_e}{m_e n_e \omega_c^2}\left[-\frac{\partial P_e}{\partial \hat{n}} - e n_e E_{\hat{n}}\right] \\ \underline{j}_{e,\perp} &= -n_e \frac{r\nu_e}{\omega_c}\left[\frac{\partial T_e}{\partial \lambda}\left[\ln(n_e) - 1\right] + e\frac{\partial \phi^*}{\partial \lambda}\right] \end{aligned} \tag{1.4}$$

The cross-field mobility of electrons in Hall thrusters is known to vary widely along the length of the thruster, from near the lower classical limit to above the Bohm value[64–66]. The total electron frequency includes contributions form (e-n) collisions, wall collisions, a factor for anomalous diffusion, and other factors deemed important for a particular application. Extensive tuning of the anomalous contribution is necessary to match experimental results[15,57,60]. The collision rates, or any other expressions with dependencies on ion and neutral parameters are

calculated using outputs from the PIC submodel, such as the neutral density, ion density and ion flux. The electron density is determined by enforcement of charge neutrality throughout the simulation domain. Therefore, the model cannot be applied past the entrance of sheaths, where space charge effects become important.

In order to satisfy the boundary conditions arising from the applied electric field, the total current from each species is calculated by integrating current densities over magnetic field lines intersecting the walls of the thruster. For a magnetic line of force, which is really a surface with azimuthal symmetry, the total electron I_e and ion currents I_i intercepted by a given field line are related through the anode current I_a based on charge conservation, as shown in Equation 1.5. Here, λ_a, λ_c, V_a, and A are the values of the magnetic stream function at the anode and cathode, the applied anode voltage, and the anode area — respectively. For dielectric walls, no net electric current can be lost to the walls in the steady state, so there should be no net wall current included in Equation 1.5. By integrating over a field line intersecting an insulating boundary, the need to solve for currents parallel to $\underline{B}$ within the bulk is bypassed. However, this does limit the application of the model to magnetic topologies where both ends of all field lines intersect insulating walls. The applied potential difference is enforced by calculating the thermalized potential in Equation 1.3 along field lines passing through the anode and cathode. The potentials at these two location are inputs of the model, allowing ϕ^* to be calculated directly once T_e and n_e are known. At the cathode, n_e is supplied from the PIC model while at the anode, I_e determines n_e, as shown in Equation 1.6. However, I_a remains unknown up to this point.

$$I_a = I_i(\lambda) - I_e(\lambda) \tag{1.5}$$

$$\phi^*\left(\lambda_a\right) = \phi\left(\lambda_c\right) + V_a + \frac{T_e}{e}\ln\left(\frac{I_a - I_i}{eA\sqrt{\frac{T_e}{2\pi m_e}}}\right) \tag{1.6}$$

$$\frac{\partial \phi^*}{\partial \lambda} = f\left(\lambda, \frac{\partial T_e}{\partial \lambda}\right) - I_a \tag{1.7}$$

$$\phi^*\left(\lambda\right) = \phi^*\left(\lambda_o\right) + \int_{\lambda_o}^{\lambda} f \; d\lambda - I_a\left(\lambda - \lambda_o\right) \tag{1.8}$$

In order to close the system of equations, the electron temperature along each field line must be determined. This is actually the most involved process of the entire method, so the details are not presented here. The full details are presented and described in the PhD Thesis of Fife[56], but it amounts to integrating the electron energy equation over magnetic field lines and solving the resulting non-linear partial differential equation numerically. For stability reasons, the electron temperature must be solved with a time-step 100-1000 times shorter than the PIC interval for heavy particles[62]. Knowing $T_e(\lambda)$, the discharge current can now be determined. Solving for $\frac{\partial \phi^*}{\partial \lambda}$ after integrating the perpendicular Ohm's Law over a field line, and enforcing charge conservation leads to a relationship in the form of Equation 1.7. This equation can be further integrated over λ to give Equation 1.8, which provides $\phi^*(\lambda)$ for a given value of I_a. Solving for the value of the anode current that satisfies Equation 1.8 for $(\lambda, \lambda_o) = (\lambda_a, \lambda_c)$ ensures that the potential boundary conditions are satisfied. These equations are solved at each discrete iteration in time over a desired period, providing predictions of the bulk properties of the plasma throughout the interior of the thruster and into the near plume.

1.2.2 Fluid Models

Generally, the utility of a fluid model derived from the Boltzmann equation relies on the accuracy that can be obtained in closing the set of moment equations through assumed expressions for each necessary higher-order moment, or a more rigorous mathematical approach. A detailed description of the link between the Boltzmann equation of a species and its corresponding fluid equations is given in Appendix C, where the validity of certain methods of closure are also discussed. The applicability of a fluid model to a given problem therefore rests entirely on the strength of approximations needed to equalize the number of equations and unknowns, which often varies from one situation to the next. The use of an ansatz can only be justified if its form has been validated to within a desired accuracy through comparisons to pertinent experimental data or credible numerical solutions of more general models. Even then, the same accuracy can not always be expected when applying the model based on this ansatz to a different problem, and the resulting errors will be unquantifiable without comparisons to other data.

Rather than guessing or stipulating the best method of closure, a systematic approach is to expand the single particle distribution function f_s governed by the Boltzmann equation in

terms of some small parameter ϵ about a zeroth order distribution f_0. This method is due to Chapman and Enskog, and is thoroughly described in the classic text by Chapman and Cowling (1970)[67]. Truncation of the expansion allows each surviving correction term to be found, in principle, so that f_s can be approximated to a known order in ϵ. The approximate forms for any moment can then be determined by evaluating the appropriate integrals, leaving a set of partial differential equations governing the physically meaningful lower-order moments that are valid in situations where ϵ remains sufficiently small.

A more general expansion is given by Harold Grad in terms of Hermite polynomials[68]. Depending on the desired accuracy, the resulting approximate distribution is postulated to depend on a chosen number of lower-order scalar moments of f_s, but not on their gradients. A detailed discussion of the relationship between the Grad method and Chapman-Enskog expansion is provided by H. Struchtrup (2005)[69], among others. Yet another method, also more general than the Chapman-Enskog expansion, consists of either an expansion of the distribution function in terms of spherical harmonics or cartesian-tensors — the two methods being mathematically equivalent. Two sets of authors applying the cartesian-tensor expansion to plasmas are Shkarofsky et al. (1966)[70] and Mitchner & Kruger (1973)[71].

A similar approach was taken by Chew, Goldberger, and Low[72], expanding distribution functions in terms of Larmor radii. They find that a closed set of hydrodynamic equations may be obtained if the magnitude of the third moment of the velocity distribution is neglected — this is referred to as the pressure-transport tensor by the authors and is also related to the heat flux vector. In Ref. 72, they also assume that electric and magnetic fields are perpendicular, which should be a good approximation away from sheaths in devices similar to Hall thrusters. Presumably, closure can also be attained by specifying the form of the third moment in terms of known quantities. The main point here is that a useful fluid description of a plasma in a magnetic field can be obtained under certain circumstances, but the validity of the closed set of equations hinges on the magnitude or assumed form of the highest order moment that is retained.

In fluid models of plasmas, a modified formulation of the Navier-Stokes equations that includes forces due to electric and magnetic fields are commonly utilized. The forms of the pressure tensor and heat flux vector are derived rigorously in the theory of Chapman and

Enskog[67], based on first order deviations of the local distribution function from equilibrium. As a result, the Chapman-Enskog distribution function includes corrections that are of first order in terms of the local Knudsen numbers[73] (pp. 2-3). The local Knudsen number (Kn) for a given measurable ρ is defined as the ratio of the mean-free-path (MFP) to the local scale length of gradients in ρ. The local scale length of a given quantity L_ρ is defined in Equation 1.9. Based on these arguments, the accuracy of fluid equations relying on this method of closure deteriorates as the local values of Kn become sufficiently large. Bird states that fluid models should be abandoned when Kn exceeds 0.2, but points out that significant errors are also present at a lower threshold of 0.1. Still, other authors suggest a lower cutoff of Kn $\gtrsim .05$[69] (pp. 11-12). The same author defines a cutoff for the definition of a rarefied gas as Kn $\gtrsim .01$. Note that an appropriate scale length is generally not the same as the characteristic size of the system. Based on these considerations, applying fluid models to Hall thruster plasmas is not likely to provide quantitatively accurate information unless the higher order moments are well represented by their assumed mathematical forms. However, it is often the only option by which simple, qualitative, relationships can be established — and can provide intuitive insights into complex or unfamiliar situations. It is better, then, to apply fluid models to Hall thrusters with an understanding and appreciation for their limitations.

The advantages of properly applying a controlled closure method are now apparent. By keeping track of assumptions and understanding how they affect the form of the distribution function, the breakdown of approximations can be detected and even quantified to a certain degree. The simplest method is to estimate Knudsen numbers beforehand, and then recalculate them from the solution for comparison. More complex methods are based on the effects certain combinations of moments have on characteristics of the approximate distribution function. For example, the distribution resulting from the Chapman-Enskog expansion is the product of a Maxwellian distribution f_M and polynomials of the molecular velocity components[69] (p. 96). Therefore, unlike a true distribution function, approximate functions may attain substantial negative values for certain combinations of moments — a clearly unphysical result. This concept of "moment realizability" has been used by Levermore et al.[74] to develop a method for assessing the validity of the Navier-Stokes equations as they are solved numerically. This is pertinent to plasma thruster modeling because it is common practice in the relevant literature to conduct

an analysis in terms of simplified Navier-Stokes equations ($Kn \ll 1$), or even Euler equations ($Kn \to 0$), despite the fact that the situation being studied is the expansion of a low-density plasma into the vacuum of space. Specific examples of this sort are presented in the next section.

$$L_\rho \equiv \frac{|\rho|}{|\nabla \rho|} \tag{1.9}$$

Hall 2De

Introduced toward the end of 2009[g], Hall 2De[75] is a plasma fluid model of Hall thrusters put forward by researches at the Jet Propulsion Laboratory, and is an offshoot of OrCa2D — a code developed to simulate the behavior of hollow cathodes. This 2-D (R-Z) model is meant to improve upon the capabilities of HP-Hall, described in moderate detail in Section 1.2.1 of this report. The main conceptual differences between HP-Hall and Hall 2De are the decision not to model ions and neutrals as particles, and the inclusion of finite resistivity effects in the electron momentum equation parallel to magnetic field lines. This modifies how the conservation of current is enforced, and the numerical solution is expedited by the use of a mesh whose boundaries are parallel and orthogonal to magnetic lines of force. By generalizing the parallel electron momentum equation, the authors claim that they are able to expand the computational domain further than with HP-Hall: far into the plume, in areas with decreased parallel conductivity, and in regions with more general magnetic topologies.

As mentioned at the onset, ions and neutrals in Hall 2De are not modeled as particles. Neutrals are assumed to be completely collisionless, and their flow is solved geometrically[h] based on the assumption of free molecular flow[76,77]. Ions are governed by continuity and momentum equations, and closure for them is provided by assuming a constant ion temperature and a scalar pressure tensor. The ion momentum equation ignores magnetic forces and includes terms for friction with the neutrals based on assuming a "quasi-Maxwellian" distribution for species when evaluating the Boltzmann collision operator. Ion-neutral (i-n) charge-exchange

[g] This is not the first plasma fluid model applied to a Hall thruster, but rather a recently developed example of one.

[h] Neutral densities are permitted to vary in time, but the neutral velocity field is not. It is unclear to the author what consequences this has on the simulated dynamics of the discharge, or if the model is more appropriately interpreted in the steady state.

and ionization due to electron-neutral (e-n) impacts contribute to the friction term. Multiple reaction paths allowing for the presence of up to triply-charged ions are included, each class of ions being governed by its own momentum equation.

Thus far, the treatment of the heavy species represents a simplification, rather than a generalization, of the HP-Hall model, aside form the inclusion of triply charged ions. However, accounting for parallel electron resistivity replaces the use of a thermalized potential along field lines with a more generalized Ohm's law shown in Equation 1.10. The thermalized potential condition utilized in HP-Hall is recovered by forcing the resistivities η and η_{ei} to zero, relying on a balance of electric ($E_\parallel$) and pressure forces ($\nabla_\parallel P_e$) along $\underline{B}$ instead. Here, $\eta_{ei} \equiv \frac{m_e \overline{\nu}_{ei}}{e^2 n_e}$ is the resistivity due to an average electron-ion collision frequency $\overline{\nu}_{ei}$, and η is the total resistivity, including contributions from $\overline{\nu}_{ei}$, electron-neutral collisions ν_{en}, wall collisions ν_{ew}, and an anomalous Bohm collision frequency ν_B. The expression for ν_B is shown in Equation 1.11, where α is specified as a function of position along the thruster axis (z) to match empirical data. A term for the parallel ion current density ($j_{i,\parallel}$) is also included in the equation governing the perpendicular electron current density. The electron energy equation is presented without derivation, but it appears that electron drift energy ($\frac{1}{2} m_e u_e^2$) terms are neglected. Thermal electron conductivity is included parallel and perpendicular to the magnetic field.

$$j_{e,\parallel} = \frac{1}{\eta}\left(E_\parallel + \frac{\nabla_\parallel P_e}{e n_e} - \eta_{ei} j_{i,\parallel}\right) \tag{1.10}$$

$$\nu_B \equiv \alpha(z)\omega_{c,e} \tag{1.11}$$

Unlike HP-Hall, where current conservation is enforced by taking advantage of the fact that $\underline{B}$-field lines intersect the inner and outer walls of Hall thrusters, Hall 2De enforces current conservation within each cell in the domain by making use of the quasi-neutrality assumption ($n_e \equiv n_i$). This assumption requires that $\nabla \cdot \underline{J} = 0$ for local charge conservation, where $\underline{J} \equiv \underline{j}_e + \underline{j}_i$. This equation is integrated over each cell volume with the aid of the divergence theorem. In a manner analogous to the procedure of HP-Hall, this allows the potential distribution to be determined subject to the boundary conditions provided by the anode and cathode. At this step, the alignment of cell boundaries with lines parallel and perpendicular to magnetic field

lines is exploited, simplifying the algebra by eliminating half of all dot-products and improving the numerical behavior of the model. Therefore, the main benefit of Hall 2De is a direct result of its more refined treatment of the electron moment equation. By solving for the parallel component of $\underline{j}_e$, it is no longer necessary for each field line to extend from the inner wall to the outer wall of the Hall thruster, or for the value of the magnetic stream function λ — see Section 1.2.1 — to be taken into account (it can assume non-unique values across multiple separatrices present in more complex topologies, like the magnetic field of the DCFT).

In 2010, this additional capability was used to model the BPT-4000 thruster to help investigate why no measurable erosion occurs past $\approx 5,600$ h of ground testing[78]. Although specific details of the BPT-4000 are restricted, the authors do state that the "quasi-one-dimensional approximation for electrons. . .does not permit the numerical simulation of the thruster plasma in the specific magnetic field topology exposed by the erosion of the BPT-4000 channel." It would be reasonable to speculate that this means magnetic field lines exposed during erosion intersect the same wall (likely the outer wall) that they originate from, or no other wall in the domain, thus invalidating the relationship between the ion, electron and anode currents assumed by HP-Hall, in the form of Equation 1.5. This is the case for all magnetic field lines in the DCF thruster, as well as field lines in cylindrical Hall thrusters that are downstream of the last magnetic surface to intersect the stunted center-pole piece. Thus, Hall 2De extends the methods for solving the electron energy equation utilized in HP-Hall to include a more general class of magnetic topologies. However, it is not necessary that the ions and neutrals be governed by fluid models. This decision was motivated by the desire to avoid the statistical noise inherent in PIC models, however this can be mitigated by appropriate averaging techniques and the use of smaller super-particle sizes.

Hall 2De simulations of the BPT-4000 and a 6 kW laboratory thruster are compared to results obtained with HP-Hall as a preliminary benchmark of the model. The results along the centerline show good agreement for some parameters, though the agreement between the two models is worse in 2D plots of n_e. Like all fluid models, and many PIC models too, simulation results have a strong dependence on the specified profile of $\alpha(z)$. However, by expanding the domain up to several multiples of the thruster channel width past the exit plane, they are forced to provide values of α much further downstream of the exit than is typically done. They

find that the profile and peak magnitude of α in the plume have a strong effect on plasma parameters and the electron-ion collision frequency as well. As a result, they conclude that the issue of electron transport in Hall thrusters cannot be resolved without considering "the two-dimensional character of the electron flow field in the thruster plume".

After some refinement of their method, they focused on benchmarking Hall 2De against measurements[78] and to examine the apparent lack of erosion in the BPT-4000. They continue to provide profiles of the Bohm frequency factor α within the thruster chamber, but then find it necessary to independently, and therefore inconsistently, vary the electron Hall parameter Ω_e in the plume region by several orders of magnitude. By varying both $\alpha(z)$ and specifying an inconsistent value of Ω_e in the plume, reasonable agreement with measurements were achieved. However, different profiles for α were able to produce similar results, consistent with results obtained using HP-Hall[60]. This leads them to state that "the far plume remains a region of elusive physics that are currently not captured by either Hall 2De or other state-of-the-art simulation codes."

It is also plausible, though, that the application of fluid equations in the far plume region contributes significantly to the observed discrepancies. As a consequence of proceeding with a system of fluid equations that is closed in an uncontrolled manner, it is very difficult to discern which portions of the errors arise from a breakdown of the fluid equations, and which are caused by deviations from classical electron diffusion. The links between simulated results and physics in the real device must also be distorted to some degree when the electron Hall parameter is specified independent of the total electron collision frequency.

Unfortunately, this scenario is often unavoidable in contemporary plasma physics. Probe measurements, when they can be obtained in regions of interest, are typically much harder to interpret than they are to obtain — especially in a moving plasma, near surfaces, and in magnetic fields with significant gradients. Optical measurements are extremely powerful because they do not significantly perturb plasma thruster plasmas, but can also be limited in the resolution, location, and variety of measurements that can be made. Kinetic models, at least in principle, are capable of providing the desired data directly, but they typically cannot be implemented without invoking assumptions to reduce the scale of the problem. Consequently, even existing PIC models require restraint in interpretation, and have not been overwhelmingly

successful at reproducing experimental observations[i]. All of these are reasons why there remains much to be learned about plasma physics, and why it is reasonable to think that considerable improvements over existing technologies can me made.

The following discussion is meant to provide a summary of limitations common to Hall thruster fluid models, and to balance arguments provided as justifications for the theoretical basis of Hall 2De. These are listed to support the notion that fluid models can often contribute to the identification and study of the dominant physical properties of a discharge, but that they should not be counted on to accurately predict operational details. However, it is not always clear which aspects of a given fluid model are the most believable, absent reliable measurements for comparison. For example, Nicoletopoulos and Robson demonstrate how a fluid model may or may not reproduce the main features of the so-called Franck-Hertz experiment[j], depending on the form of the heat flux ansatz and other decisions related to fluid model closure [80]. It would be interesting to use the heat flux ansatz suggested by these authors in a simple Hall thruster fluid model, because their closure method is intended to better represent the effects of inelastic collisions.

First, the validity of using a fluid description of ions in Hall 2De is justified by Knudsen number estimates based on thruster dimensions, and by the fact that some measurements indicate that the ion velocity distribution function f_i has an approximately Maxwellian form [75]. However, their Kn estimates were only favorable at ion temperatures 1-2 orders of magnitude lower than values associated with the measured ion distribution functions they cite. Second, Kn values were underestimated by comparing MFP values to thruster dimensions, rather than to scale lengths of gradients in plasma parameters [73] (p. 2). Additionally, estimates of ion velocity and energy distributions extracted from LIF measurements are also known to be non-Maxwellian — with low or high energy tails or a bimodal form [29,38,39,41,81–84]. These points are meant to emphasize that Hall thruster fluid models, even ones implemented as carefully as Hall 2De, should be interpreted with caution — especially in simulations extending many thruster-lengths into the plume. Fluid equations derived on the premise of $Kn \ll 1$ become invalid long before the continuum assumption breaks down, as it is the assumed form of the distribution function or,

[i] The limitations of PIC models discussed further in Section 1.2.6

[j] J. Franck and G. L. Hertz were awarded the 1925 Nobel Prize in Physics "for their discovery of the laws governing the impact of an electron upon an atom" [79].

equivalently the higher order moments, that determines when a particular set of fluid equations is invalidated. Assuming the Fourier ansatz for the heat flux vector, the exact form of which is derivable from a Chapman-Enskog expansion, is equivalent to choosing a system of fluid equations valid only when $Kn \ll 1$. Corrections to the heat flux vector arise when higher order effects are accounted for [67] (pp. 280–293) [69] (pp. 62–70). Comparisons of the Fourier heat flux expression with PIC-DSMC simulations of 1D RF discharges demonstrate the qualitative and quantitative errors associated with applying it to rarefied plasmas [85]. Additionally, the evaluation of collision terms based on an assumed form of the distribution function represents an unquantifiable error when the distribution function used to evaluate the collisional terms is not the same as the function used to evaluate all the other moments, including the heat flux vector. An illuminating discussion of many of these topics from a mathematical standpoint is available in Ref. 86.

1.2.3 Particle-in-Cell Models

In situations where the fluid equations do not provide an adequate plasma model, the more general kinetic problem must be addressed. This amounts to solving the Boltzmann equation, which requires numerical solutions in all but the simplest cases. Even with advanced numerical techniques, however, solving for the distribution function f directly from the integro-partial differential Boltzmann equation applied to engineering problems with complex boundary conditions is usually not practical. A more intuitive approach is to simulate the trajectories of a finite set of discrete particles over the domain of interest, and to construct the distribution function by simply counting particles. Often times, the distribution function is bypassed entirely, and higher order moments are calculated by weighted averaging. This section serves to introduce the general concept of a particle-in-cell (PIC) model by describing the procedures that make such an undertaking feasible, in preparation for a more detailed description of the SPL-PIC plasma thruster model in Section 1.2.4.

The PIC methodology was developed in the mid-1960's to study collisionless plasmas as an alternative to models based on calculating particle-pair forces from Coulomb interaction potentials. Comprehensive introductions to the method are provided in two widely utilized books by Birdsall & Langdon [87] and Hockney & Eastwood [88]. Rather than advancing trajectories

based on applied and resultant particle-pair forces, the charge of each particle is weighted to the nodes of a cell belonging to a spatial mesh. In the electrostatic case, Poisson's equation for the potential ϕ is solved on this grid, providing the electric field, which is then weighted back to the particle locations by the same method as before. The individual trajectories are integrated forward in time by an appropriately small amount, and the entire process is repeated. Hence, the computational complexity of PIC model scales linearly with the number of particles N in the domain, rather than as N^2 in the particle-pair case. The small characteristic lengths and times of plasmas set restrictions on the allowed time-steps Δt and cell sizes Δx required to maintain numerical stability and minimize systematic errors. For instance, the cell size should be smaller than the electron Debye length[k] λ_D, and the time-step must be smaller than the inverse of the maximum frequency — being either the electron plasma or electron cyclotron frequency in our case[l]. Even if Debye lengths are over-resolved by the mesh, the time-step should be small enough to prevent thermal electrons from traveling large distances relative to local mesh dimensions[88] (Sec. 9-2-3). These constraints have a strong influence on the approaches taken to model Hall and cusped-field thrusters.

Accounting for every atom or molecule in a PIC simulation is not even remotely feasible. In the DCFT, for instance, this would require tracking and moving many trillions of particles. Therefore, in order to replicate the physical densities of systems, computational particles are assigned statistical weights W, indicating the number of real particles represented by a virtual "superparticle". A typical superparticle represents 10^6 real particles, though this number can vary by orders of magnitude depending on the application. Statistical fluctuations are often prevalent features of PIC models, since the number of particles used is typically limited by computational resources and the amplitude of the noise varies as $N^{-1/2}$. In the simplest case, fluctuations about mean values in steady state simulations can be amplified by a factor of $W^{1/2}$ over physical values. In the case of plasma thruster models, quantities computed by counting particles in a volume or during a time step (such as number density, thrust, and current) will only register as multiples of the contribution from a single superparticle. Though this requires

[k] If λ_D exceeds the electron Larmor radius $r_{L,e}$, stability does not require the mesh to resolve $r_{L,e}$. However, to properly account for electron interactions within cells, Coulomb collisions would need to be included in the model. Modifications to Coulomb collision theory accounting for cases in which $r_{L,e} < \lambda_D$ are briefly discussed at the end of Sec. 1.2.7.

[l] The use of implicit, rather than explicit techniques can relax these constraints somewhat

more effort to interpret data, these annoyances can be overcome by time-averaging steady state results. The inherent graininess and noise associated with PIC results become more of a problem when used to study transient phenomenon or the effects of subtle perturbations.

Boundary conditions for particles are dealt with in a straightforward manner by specifying their treatment at the limits of the domain, such as absorption (deletion) or reflection. This allows the simulation domain to correspond closely or exactly with physical dimensions, and removes the need to specify a somewhat arbitrary wall-collisions frequency as in fluid models. Aside from kinematic effects, particles and boundaries can also interact in more complicated ways. Two important effects in Hall thrusters are secondary-electron emission and insulator erosion. Absent a more detailed subroutine, these phenomena may be included through the use of empirical models.

Inter-particle collisions were introduced to the PIC method in two distinct ways, though the methods are sometimes confused for one another due to the fact that authors will often times refer to one or the other as simply "the Monte Carlo method". The PIC method with Monte Carlo collisions, or PIC-MCC, was implemented when PIC was applied to more diverse problems, including weakly ionized plasmas, for which Coulomb collisions are unimportant and scattering collisions between electrons and a background neutral gas are dominant. A second collision model was combined with the PIC approach by Serikov, et al.[89] in 1999, called the direct simulation Monte Carlo, or the DSMC method. DSMC was developed by G. A. Bird in 1963[90] as a numerical representation of the Boltzmann collision operator for applications in rarefied gas and strong shock flows. More detailed information about the PIC-MCC and PIC-DSMC methods are provided in the following section. Over time, other methods have been proposed, but these two remain the most widely utilized in kinetic plasma models.

PIC-MCC & PIC-DSMC: Comparison and Limitations

The DSMC method was developed to simulate the effects of collisions in rarefied flows of neutral molecules. It is a particle method, meaning that a small sample of the actual atoms or molecules are represented by simulated particles carrying a statistical weight W. Over a time step shorter than a collision frequency, the particle trajectories are integrated forward in time. A computational grid is also used in DSMC, with cell dimensions shorter than the local MFP.

However, its purpose here is to divide the physical domain into discrete but small elements for the purpose of pairing particles for collisions. The assumption of molecular chaos inherent in the Boltzmann collision operator allows collisional processes to be modeled using random number generators, as follows. A comprehensive review of the DSMC method, including examples for many different collision types, is provided by K. Nanbu [91]. Further details are given in the book by G. A. Bird [73].

- A *maximum* number of collisions to occur in a given cell during a length of time Δt is calculated as $N_{max} = R_{max} V_c \Delta t$, where R_{max} is an upper-bound on the reaction rate of a processes, and V_c is the volume of the cell. This is done to avoid having to check every particle in the cell for the possibility of a collision, and is called the null-collision method [92] because the number of collisions carried out will be less than N_{max}.

- A particle pair is chosen randomly within the cell and a uniformly distributed random number $0 < R < 1$ is generated. The probability of a real collision between the pair is $P_c = \frac{g\ Q(g)}{g_{max}\ Q_{max}}$, so that if $R > P_c$ no real collision occurs, and the number of pairs remaining to be checked decreases by one. Otherwise, $R < P_c$ and an interaction does occur. The modifications of particle attributes is also determined randomly in most cases. In the expression for P_c, g is the magnitude of the relative velocity between the pair, and $Q(g)$ is the integral cross section corresponding to the relative speed. The value of the product $g_{max}\ Q_{max}$ is used to evaluate R_{max}.

The above process is repeated for each cell during each time step, where the details of the algorithm depend on the types of reactions accounted for, the particle species represented, and the particle weighting scheme — to name a few. However, if implemented properly the DSMC method, does more than resemble the action the of Boltzmann collision operator — it has been shown to be an appropriate numerical solution of the full Boltzmann equation [93]. As such, the DSMC method is applicable to dilute gases where particle interactions are binary.

Whereas DSMC was developed specifically with collisions in mind, the PIC method was initially developed to model collective effects in fully ionized plasmas, where the grid used to solve for the electric field smooths out short range interactions [94][m]. Collisional processes

[m] The force between particles as their positions coincide goes to zero, rather than diverging to infinity as with Coulomb's law

were added later to include effects of electrons interacting with background neutral gases, called PIC-MCC. Here, the velocities of the neutrals are assumed unchanged as a result of collisions with electrons, a good approximation considering the ratio of the electron and neutral masses. As in DSMC, a random number determines if an electron undergoes a collision based on a collision frequency or reaction rate. Then, the electron direction is reassigned randomly depending on the differential cross section.

In PIC-MCC, because there is no explicit particle pairing, average quantities of one species in a given cell are used in the calculation of the relative velocity. As such, PIC-MCC is really only appropriate for representing collisions between particles with disparate masses and velocities (i.e. $|\underline{v}_1 - \underline{v}_2| \approx |\underline{v}_1|$), $m_1/m_2 \ll 1$. Therefore, MCC is appropriate for (e-n) and, arguably, electron-ion collisions as well. For high energy ions in (i-n) collisions, MCC might be justifiable if a corresponding neutral particle were also scattered. However, the use of MCC for very-low-energy ions colliding with neutrals is not justified because of the sensitivity of the relative speed to the orientations of particle velocities. The same problem arises for interactions between particles of the same type, namely electron-electron (e-e), ion-ion (i-i), and neutral-neutral (n-n) collisions. In these cases, DSMC is needed for an accurate description. However, recent advances have blurred the distinction between these two methods even further.

Despite the limitations of the initial method, extensions of PIC-MCC have been proposed and used to represent Coulomb interactions. One method by Hockney separates the effects of short and long-range interactions. Long range forces are accounted for by the grid, similar to the standard PIC procedure, while the short range forces are accounted for by summing particle-pair forces [p. 267[88]]. The method by Takizuka and Abe[95] is often associated with PIC-MCC[94,96] but applies the DSMC approach and considers Coulomb collisions in a strictly binary fashion. DSMC (binary collisions) is not quite as accurate in this context as for non-Coulomb collisions due to the long range nature of the Coulomb force. However, in some cases the level of accuracy may be acceptable, as discussed in Appendix D. Nanbu also describes a model that adapts his treatment of the Boltzmann collision term to the Fokker-Planck operator for many small angle deflections[91]. J. Szabo, the author of the SPL-PIC plasma thruster model[97], implemented a model for (e-e) collisions using the Fokker-Planck expression represented as a Langevin type equation — and based on the model of Manheimer, et al.[98]. The approaches of Nanbu, Szabo,

and Manheimer, et al. are all based on approximate forms of the Boltzmann collision operator, and so do not represent a generalization to the inclusion of non-binary encounters[n]. A more distinct approach is suggested by Christlieb, et al. and described in a more recent PIC review by Verboncoeur[96]. Using clusters of localized particles in a *gridless* domain, the computational complexity needed to account for many-body interactions is reduced by only considering paired interactions within a cluster, and by pairing long-range interactions between clusters.

In PIC models of Hall thrusters, DSMC has usually only been used in simulations of the plume because of the importance of ion-neutral (elastic and charge-exchange) and (n-n) collisions[o]. These are the most important interactions to consider because it is necessary to predict the magnitude and energy composition of ion flux to spacecraft surfaces. Due to the complex geometries of spacecraft surfaces, these simulations were developed in three dimensions earlier on. For simplicity, electrons have been modeled as fluids. Possibly the first PIC-DSMC model applied to Hall thruster plumes was developed by D. Oh at MIT in 1997[101], and has been the foundation for the development of more refined plume models[102]. More recently, 2D PIC Hall thruster models have been described as applying DSMC collisions during plasma simulations[51]. However, insufficient details regarding the collision procedures are given in Ref. 51 to confirm that their use of the term DSMC is consistent with its meaning in this discussion. Although DSMC is distinct from MCC, the terms are sometimes used interchangeably. For instance, in Ref. 76 HP-Hall[p] is described as using DSMC collisions, when it actually uses the MCC procedure.

In most cases, measurements and detailed models of a thruster provide the input data for the plume models described above. However, the use and development of thruster models are more oriented toward understanding their physics and improving designs. Most Hall thruster models do not use DSMC to simulate collisions while charged particles are involved, likely due to the added complexity and computational cost of DSMC in comparison to MCC. This trend has some basis in physical arguments, such as the important role of (e-n) elastic and inelas-

[n] Two different equations are referred to as "Fokker-Planck" equations, or sometimes the "Landau-Fokker-Planck"[91] equation. One is quite general, and derivable from the Liouville equation [pp. 22-24[99]]. The other is contingent upon the assumption of numerous small deflections in velocity space, and is also derivable from the Boltzmann collision term [pp. 338-343[71]].

[o] Hirakawa does use DSMC to pre-simulate neutrals, but the technique is not used with the full plasma model[100]

[p] HP-Hall is discussed further in Sec. 1.2.1.

tic collisions. The ion mean-free-path is often long enough to consider ignoring their mutual collisions, and collisions between electrons are often neglected as well — though this may not always be justifiable. However, a true estimate of the electron *guiding center* MFP for various interactions is complicated by their strong magnetization within the thruster. Conversely, it is also reasonable to argue that (n-n), (i-n), and (e-e) collisions should be represented with DSMC due to the coupling that exists between species. However, the necessary cross-section, secondary electron yield, and erosion coefficient data over a complete range of energies is not available — limiting the quantitative accuracy that could be expected using less approximate algorithms. Furthermore, the anomalous enhancement of cross-field electron diffusion in many magnetized plasma discharges may offset gains in accuracy associated with more rigorous models of particle collisions[q]. Therefore, before a high-fidelity plasma thruster model can be made based on first principles, sub-models of surface interactions shown to influence the discharge must be greatly improved, and the mechanism of anomalous transport must be identified[r].

1.2.4 SPL-PIC Plasma Thruster Model: Initial Development

The SPL-PIC plasma thruster model is a PIC-MCC model developed by J. Szabo [97] and initially applied to modeling a miniature TAL-thruster (mini-TAL). Electrons, neutrals, and ions are simulated as statistically weighted particles. Their trajectories are numerically integrated forward in time according to classical mechanics using the leapfrog algorithm as modified by Boris [p. 59 [87]]. The model assumes symmetry about the central axis of the thruster, so only the radial R and axial Z coordinates are tracked. However, velocity components in the azimuthal θ direction are determined due to the importance of particle drifts perpendicular to both $\underline{E}$ and $\underline{B}$. The forces on particles are due to a static magnetic field and a macroscopic electric field calculated based on boundary conditions and the distribution of space charge in the domain. A schematic representation of the procedure is provided in Figure 1-7. The main features of the model relating to the representation of plasma physics are discussed here, with subsequent modifications by other authors given in a roughly sequential order. A kinetic approach to the

[q] This topic is addressed briefly in Section 3.5

[r] A fully three dimensional model may be able to address the common assertion that electron diffusion is enhanced by turbulent convection, driven by unsteady azimuthal electric fields. Two dimensional models that include some representation of azimuthal non-uniformities have been completed, though the results are not definitive and do not enable any quantitative models, as yet. [100,103–105]

problems is necessitated mainly by the large Knudsen values of each species in the thruster, as estimated in Section 1.2.7.

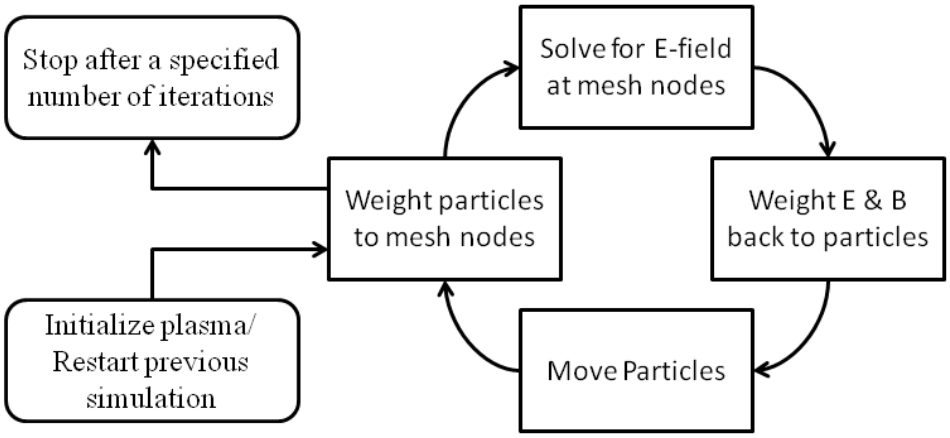

Figure 1-7: Procedural outline of the MIT plasma thruster PIC model.

The domain is specified by a structured non-orthogonal mesh, a simple example of which is shown in Figure 1-8. A grid generator[s] is used to create the mesh, and a transformation between physical (R, Z) and computational $[\eta, \xi]$ coordinates is pre-computed once for each cell. Individual cells must have dimensions comparable to the plasma Debye length in order to prevent plasma oscillations from becoming unstable numerically [p. 156 [87]]. The inverse plasma frequency ω_p sets the rate at which the potential solution on the grid must be updated, and the greater of ω_p and the electron gyro-frequency ω_c sets the timestep for electron trajectory integration. The heavier neutrals and ions allow for longer timesteps on account of their slower dynamics.

More details about the relationship between computational and physical coordinates are provided in Ref. 97, but several important definitions will be emphasized here to help clarify discussions elsewhere in this report. Each node's coordinates are represented in computational space by a pair of integers $[k, j]$. For example, the coordinates of the square point in Fig. 1-8 might be $(z, r) = (z_o, 0)$ and $[\xi, \eta] = [0, 0]$, while the coordinates of the node marked by the triangle might be $(z, r) = (z_{max}, r_{max})$ and $[\xi, \eta] = [N_Z - 1, N_R - 1]$. Again, N_Z indicates the number of nodes along the ξ-direction of the mesh, and N_R indicates the number of nodes in the η-direction. The continuous pair of computational coordinates $[\xi, \eta]$ are defined such

[s]SPL-PIC used a mesh generator originally supplied by M. Fife, though Tecplot also provides a robust mesh generator as an add-on.

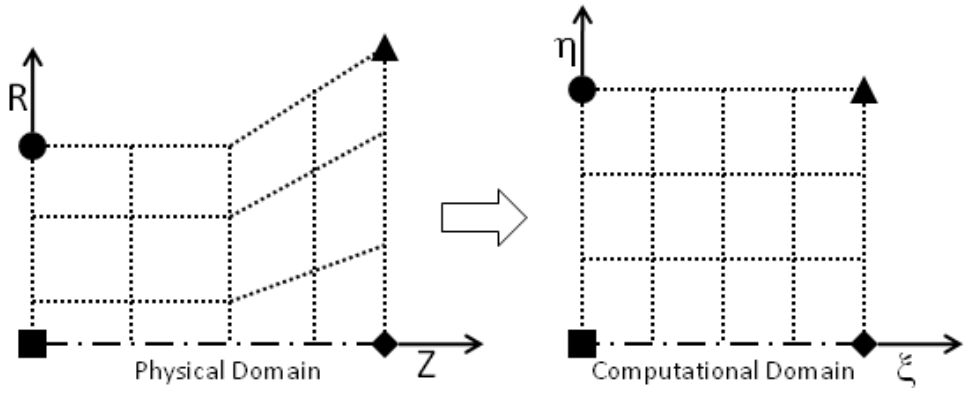

Figure 1-8: Simple example of how physical coordinates transform to computational coordinates using a structured, non-orthogonal mesh.

that $[\xi, \eta] = [k, j]$ at each node on the mesh — allowing coordinates within each cell to be conveniently defined as $[\alpha_k, \alpha_j] \equiv [\xi - k, \eta - j]$. On a regular, square mesh with spacing h, physical and computational coordinates could be determined easily using $\xi = \frac{z}{h}$ and $\eta = \frac{r}{h}$. However, PTpic makes no assumptions about the mesh geometry, aside from it being structured and mapping to a rectangle in computational space, which makes the determination of the computational coordinates one of the more intricate aspects of the model. This topic is discussed in Sec. 3.1.2. Each cell is labeled according by two integer coordinates $[k, j]$ as well, specifying the location of the node at the bottom-left corner of the cell.

Even with the use of particle weighting factors, the computational resources necessary to resolve the length and time scales for a device as small as the DCFT ($L \approx 5\ cm$) are not available to most researchers. For neutrals with thermal speeds on the order of 200 $^{m}/_{s}$ the simulation must span an interval of about 250 μs if the estimated time of convergence is taken to be the neutral residence time. The cyclotron and plasma frequencies necessitate timesteps on the order of 1 and 10 ps, respectively, requiring $2.5 - 25 \times 10^7$ iterations. If ω_p is the dominant frequency, then the potential solution must be updated every iteration[t]. For a representative Debye length on the order of $10^{-5}\ m$, this requires approximately 9×10^6 cells, and about four times as many nodes, to represent the DCFT domain. Typically, 20 or more simulated particles per cell, per species, is considered the minimum number to provide adequate statical

[t] A recent implementation of a 2D PIC Hall thruster model utilizing an implicit potential solver is described in Ref.[51]. The authors claim this permits using larger timesteps and avoids the need to resolve the Debye length.

representation of the distribution functions. This would require updating 540×10^6 particle coordinates and accounting for collisions when necessary. Assuming that all of this information could be stored in RAM, and that each particle and each node required only 1 flop per iteration, it would take 100 1 GHz processors 7 days to complete the simulation with perfect speed up through parallelization. Clearly, more than 1 operation will be necessary per particle — and especially per node as required to solve for electroquasistatic potential distributions — putting any simulation using the physical parameters of the plasma out of reach using existing hardware and software.

The primary novelty of Szabo's model is a careful manipulation of physical parameters allowing the completion of simulations in reasonable amounts of time while maintaining a connection to the real system. Identifying the two most inhibiting numbers in the paragraph above as the number of cells in the mesh and iterations required for convergence, the three factors contributing to these are the neutral residence time, the large characteristic frequencies, and the smallness of the Debye length compared to thruster dimensions. Szabo proposed artificially changing two physical constants to speed things up: *reducing* the simulated mass of heavy neutrals and ions m_H^*, and *increasing* the simulated free-space permittivity ϵ_o^*[u]. Asterisks will be used throughout to indicate the values of quantities in the simulation, which will often be compared to variables lacking an asterisk — meaning actual physical quantities. The use of an artificial mass ratio certainly precedes Szabo, going back to many early attempts at computational particle models and also some of the earliest Hall thruster PIC models [107]. The relationships between the simulated and physical constants are provided by the mass factor r and permittivity factor γ, as shown in Equation 1.12[v]. The direct consequences of these factors on the electron Debye length and plasma frequency are provided in Equation 1.13.

$$m_H^* \equiv \frac{m_H}{r} \qquad \epsilon_o^* \equiv \gamma^2 \epsilon_o \tag{1.12}$$

$$\lambda_D^* = \gamma \lambda_D \qquad \omega_p^* = \frac{\omega_p}{\gamma} \tag{1.13}$$

The effect of the permittivity factor is to reduce the electron plasma frequency and increase

[u] Independently, a plasma accelerator model created in the late 1960's also increased free-space permittivity [106].

[v] The mass factor I define here is the inverse of the mass factor used by Szabo.

the Debye length proportionately. The number of mesh cells required to resolve the Debye length decreases by γ^2, which also means that, for an equivalent average number of particles per cell, the total number of particles needed is also reduced by γ^2. The electron cyclotron frequency is unaffected by γ or r, but if the adjusted plasma frequency remains larger than ω_c, then the minimum particle timestep also increases by γ. In any event, the maximum allowable interval between successive calls to the potential solver increases by γ, reducing the number of times that this expensive task must be completed.

The mass factor is computationally beneficial in that heavy particle speeds are increased by $\sqrt{r}$ because their energies remain constant — the specified neutral temperature and applied voltage remain the same. Therefore, the neutral residence time is reduced by the same $\sqrt{r}$. The expected reduction in the number of iterations is a factor $\sqrt{r}$ if $\omega_c > \omega_p^*$ or $\gamma\ \sqrt{r}$ if $\omega_c < \omega_p^*$. In most Hall thrusters, with maximum field strengths typically no larger than 200 G, the second condition is met. However, in the DCFT, the reduced plasma frequency typically falls below the cyclotron frequency, though there is still a savings realized by not having to solve for the electric field every time particles are moved. Therefore, the total computational time should be reduced by an approximate factor $\gamma^2\sqrt{r}$ or $\gamma^3\sqrt{r}$ depending on the plasma parameters and $\underline{B}$-field strength. Although these changes make simulations numerically tractable, care must be taken in interpreting the results.

Increasing the permittivity and mass factors above unity has a number of important consequences to consider when implementing and interpreting PTpic. This is also discussed in Sec. 2.6 in Ref. 97. First, increasing ϵ_o^* reduces the ability of the plasma to maintain local charge neutrality and shield out applied fields. This is seen most clearly by considering the severe case shown in Equation 1.14. As γ becomes very large, keeping boundary conditions constant, the potential distribution relaxes to the vacuum field — as if no plasma were present. We can conclude that reasonable choices of γ will reduce the computational requirements to the maximum tolerable level, while fulfilling the condition $\lambda_D^* \ll L$, where L is a characteristic dimension of the thruster. Only then can the simulated plasma be expected to resemble the real one. There is no way to take into account all the effects of $\gamma > 1$, so the best choice is to pick as small a value as possible given the scope of the problem and availability of computational

resources.

$$\lim_{\gamma^2 \to \infty} \nabla^2 \phi = \frac{e(n_e - n_i)}{\gamma^2 \epsilon_o} = 0 \tag{1.14}$$

Decreasing the mass of the heavy particles creates some more issues regarding the resemblance of the simulated thruster to the real one. First, reducing the neutral particle mass while keeping the mass flow rate constant would increase the neutral flux by an entire factor r. Thus, to keep the neutral density at injection identical to the real case, the mass-flow rate $\dot{m}^*$ must be lower than the actual quantity by $\sqrt{r}$ because the heavy-particle speeds are already $\sqrt{r}$ faster than they should be. The combined effect of the reduced $\dot{m}^*$ and faster simulated velocities means that the simulated thrust value F^* should be interpreted as the predicted thrust of the actual device: $F^* = F$.

Enforcing the equivalence of the neutral density — and therefore the electron density if the ionization fraction remains the same — leaves an excess of heavy particle flux by a factor of $\sqrt{r}$. Therefore, for an equivalent fraction of neutrals to be ionized or even excited through collisions with electrons, the reaction rate must be increased by $\sqrt{r}$. Szabo accomplishes this by increasing the simulated total (e-n) cross section $Q^*_{en,T}$ — representing elastic scattering, ionization, and excitation — by the necessary $\sqrt{r}$. This has the side effect of increasing the (e-n) collision frequency ν^*_{en} by the same factor. In order to keep the effects of other collisions in the same proportion, namely (e-e) and anomalous scattering, these frequencies are also increased by $\sqrt{r}$. The net effect is to reduce the electron Hall parameter Ω^*_e by $\sqrt{r}$, which potentially has a very strong effect on the electron backstreaming current, and sets a sensible limit on how large r can be selected without completely demagnetizing electrons. Szabo acknowledges this in his thesis, stating that as long as Ω^*_e remains larger than 2 or 3, electrons should still behave as if magnetized. This is likely the strongest point of criticism of this model from the community[105,108,109].

Another important consequence of using $r > 1$ and $\gamma > 1$ is the weakening and stretching of the sheaths, the regions where the plasma and surface interact most directly. Hybrid and fluid models that assume quasi-neutrality ($n_i \approx n_e$) are unable to explicitly model sheaths because this assumption breaks down in these regions. They are forced to rely on analytical sub-models or specified boundary conditions to approximate the presence of the sheath. In the currently described PIC model, sheaths are able to be modeled, but their physics are strongly effected.

With the electron Debye length being the length scale over which a plasma can sustain charge separation based on its thermal energy, plasma sheaths have sizes of the order of the Debye length. By choosing a permittivity factor that is not too large, such that $\lambda_D^* \ll L$, exaggerated sheath effects can be mitigated, but not entirely avoided. Dealing with the effects of the mass factor, however, is again less tractable. Reducing the ion mass has the effect of increasing the ion current density to the wall, assuming that the bulk plasma density and temperature are unaltered by r and γ. As a result, the size of the potential drop from the plasma bulk to the wall — necessary to maintain equal time-averaged electron and ion current densities[w] — decreases in magnitude. Szabo's original model ignored this effect due to the smaller wall area in the min-TAL thruster. However, this is an important effect in thrusters with larger surface areas, like SPT-Hall thrusters and the DCFT. An attempt to correct this was made later by V. Blateau [110]. Changes made to Szabo's model are discussed in Section 1.2.5.

When particles encounter boundaries, they are either deleted or reflected depending on the kind of particle and type of boundary. For example, all particles are counted and deleted after crossing the free-space boundary in the plume, while particles that encounter the center axis are reflected specularly. Ions deposit their charge and are reflected diffusely with half their incident energy at all surface boundaries, while electrons are deleted after attributing their charge to the appropriate boundary for the field solver. After counting charges, the calculated and specified potential values at boundaries are used to solve for the potential distribution in the interior, using the successive over-relaxation (SOR) algorithm. In the plume, at free-space boundaries, the normal potential gradient is assumed to be zero.

Based on MFP estimates for each species in the mini-TAL, Szabo included a number of particle interactions in the model — all but one version of (e-e) collisions used the MCC method. Details about interactions for each species are provided in Appendix D and discussed further in Section 1.2.7. Three types of (e-n) interactions are modeled by MCC: elastic scattering, ionization, and excitation. Two ion interactions were also included: (i-n) scattering, (i-n) charge-exchange. Coulomb collisions between electrons and ions are *not* included in PTpic — though they were present in previous versions of SPL-PIC. The (i-n) collisions affect both

[w] Net, time-averaged current densities need only be zero at the surface of a perfect insulator. A body modeled as a perfect floating conductor requires only that the *total* current it collects have a zero value, on average.

species strongly because of their identical masses, so Szabo spread the accumulated effects evenly and randomly amongst neutrals in each cell, rather than pairing particles directly as in DSMC. Coulomb (e-e) collisions were represented using two different approaches: MCC and a Langevin model based on the Fokker-Planck equation. Collisions amongst ions were neglected, as were (n-n) and recombination interactions that do not involve boundary surfaces.

1.2.5 SPL-PIC Plasma Thruster Model: Further Development

After Szabo, several important changes and extensions to the code were made by V. Blateau [110] and J. Fox [111,112]. Fox also developed a different model using an unstructured mesh, variable particle weighting, and other features that are not present in the updated versions of Szabo's code [112,113].

Blateau extended the capability of Szabo's model to include SPT-type Hall thrusters. These generally have longer channels and ceramic walls that are electrically insulating rather than conducting. Aside from allowing the potential solver to handle more general boundary conditions, secondary electron emission (SEE) was added to account for the liberation of a "secondary" electron when a higher energy primary electron impacts the wall. The secondary is then accelerated into the plasma bulk through the electron-repelling sheath. The average number of secondaries emitted per incident primary is called the SEE yield, and is modeled as a monotonically increasing function of the primary energy. Based on the results of Szabo's study, Blateau did not include (e-e) his simulations, but did account for doubly-charged ions with reaction pathways from neutrals and singly-charged ions. To address the previously mentioned consequences of the mass factor on sheaths in the thruster, Blateau implemented a patch to try and make the simulation more similar to reality.

Sheath Patch

The reasoning behind the sheath patch, and the method used previously for addressing the issues caused by the reduced ion mass, are discussed in what follows. The lighter ion mass results in an increased ion flux $j_i^* = \sqrt{r}\, j_i$, which forces the electron flux to increase in order to maintain zero net current density to the walls. This is accomplished by reducing the electron-

repelling potential drop $\Delta\phi_s^*$ from the sheath entrance to the wall, given in Equation 1.15[x]. Blateau's patch focuses on the fact that more lower energy electrons are allowed to reach the wall than in the physical case. He estimates the effect of r by comparing an analytical equation for the real sheath drop ($\Delta\phi_s$) to a representation of $\Delta\phi_s^*$ extracted from the model. An approximate, analytical relationship between the physical and simulated sheath voltages is provided in Eq. 1.15. He then compares electron energies at the walls to the energy they would have needed to overcome the real sheath based on the previous estimate — and decides whether a particular electron will be reflected or allowed to impact the wall. In the case that a low energy electron is reflected, an electron with an appropriately high energy is located in the bulk and deleted — with its charge being attributed to the wall to maintain the magnitude of electron flux to the wall (leaving $\Delta\phi_s^*$ unaltered). This method rearranges the simulated energy composition of electrons lost to the wall in an improvised fashion. In cases where a low-energy electron is reflected, but a suitable high-energy electron cannot be located in the bulk, the analytically calculated difference between the real and computational sheath voltages is assumed too large, and reduced by 5% for subsequent iterations — reducing the number of negative charges deposited at the wall. Though the magnitude of the simulated sheath voltage is not directly addressed, reducing the number of negative charges reaching the wall for an unaltered number of positive charges will only serve to reduce the magnitude of $\Delta\phi_s^*$. Thus, this method is incapable of steering the simulated sheath strength toward the physical value. Including previous numerical studies of the DCFT[34,114], the Blateau sheath patch has been in place for all simulations of thrusters with ceramic walls using the SPL-PIC plasma thruster model. However, it was removed for all simulations presented in this thesis.

Note that none of the previous discussion, including Eq. 1.15, accounts for secondary electron emission (SEE). It has the equivalent effect of decreasing the net electron current density absorbed by the wall. Hence, the potential drop from the sheath to the wall will decrease in size when SEE is significant. This effect motivates an approximate upper-bound on the mass factor r_{max} in order to avoid overly strong effects from the mass factor. Expressing $\Delta\phi_s$ as fT_e, where f is some factor, the mass factor corresponding to $\Delta\phi_s^* = 0$ — in the absence of SEE — is $r_{max} \approx e^{2f}$. This bound for r is of little concern compared to the issue of maintaining

[x]The temperatures in this equation are expressed in electron-volts, where $\frac{kT}{e} \equiv T_{eV}$.

electron magnetization because $f > 5$ when there is no SEE. Hobbs and Wesson[115] show that the maximum ratio of secondary to primary electron fluxes under some simplifying assumptions is given by Eq. 1.16, so the sheath may become space-charge limited at artificially lower values of average secondary electron yield when a reduced ion mass is used. Note that $\Gamma_{max} < 0$ for a mass factor greater than about 3477, a reflection of the approximations used in deriving Eq. 1.16. In PTpic simulations, using large values of the mass factor r should cause $\Delta\phi_s^*$ to be less affected when SEE is space-charge limited. This has been observed in DCFT simulations, described further in Chapter 5. The empirical model used to determine the secondary electron yield in PTpic is described in Section 3.2.1, but it is worth mentioning here that PTpic does not impose any upper bounds on the secondary electron yield (SEY) Γ when an electron strikes a dielectric surface. The SEY in Eq. 1.16 refers to the average yield over all electrons at a particular location, whereas PTpic calculates and enforces a SEY value each time an super-electron strikes a dielectric surface.

$$\Delta\phi_s^* = \Delta\phi_s - \frac{T_e}{2}\ln r \tag{1.15}$$

$$\Gamma_{max} \approx 1 - 8.3\sqrt{r}\left(\frac{m_e}{m_i}\right)^{1/2} \tag{1.16}$$

J. Fox made numerous contributions toward improving the computational speed of the model and representing anomalous electron transport, discussed in Section 3.5. Without going into too many of the details, parallel versions of the most time consuming and repetitive algorithms in the code were modified or rewritten in order to take advantage of the speed-up offered by running simulations with many processors simultaneously. The potential solver was updated using a modified version of the red-black SOR algorithm with Chebychev acceleration. The solver was also re-benchmarked after its logical structure was modified to ensure convergent solutions were obtained at each instant in time — contrary to prior implementations. The particle initialization, neutral injection, electron injection from the cathode, particle moving and collision procedures were also updated to allow execution in parallel. The included collisions were: (e-n) scattering, excitation, and ionization; (i-n) scattering and charge-exchange; double ionization of neutrals and further ionization of ions to produce Xe^{2+}. Coulomb collisions were

not included in Fox's parallelization efforts, and so were not included in his PhD results, nor the author's previous work with the DCFT [34,114]. Coulomb collisions of any kind are also excluded from PTpic, and all results presented in this thesis.

1.2.6 Other Approaches Toward Tractability

A different approach aimed at reducing computational complexity while maintaining physical relevance has been proposed by researches in Europe [105]. Rather than altering the free-space permittivity and heavy particle mass, the size of the domain is reduced and the magnetic field is increased in order to preserve what the authors identify as the critical dimensionless parameters. Providing a concise description is made difficult by the fact that two versions of this approach have been put forward [108] which differ in their scaling of most parameters — simulated temperature, applied magnetic-field, thrust, and efficiency are some examples.

In addition to describing their own method for approaching PIC-MCC thruster simulations, the same authors provide an assessment of Szabo's approach — the main point of criticism centering around the fact that changing the mass ratio — combined with the accompanied cross section enhancement — alters the electron hall parameter "modifying completely the discharge characteristics". The authors also argue that changing the length and time scales of the plasma "make it impossible to recover the original system". However, previous results [34,97,110,112,116] and those presented in Chapter 5 indicate that this statement is somewhat of an exaggeration. Also, a similar neutral-dynamics acceleration procedure was recently applied in conjunction with the scaling method [117] without mentioning the same point of criticism. Completely accurate simulations of even small plasma thrusters are not presently feasible, using fluid or particle techniques. Both methods require accepting deficiencies in the representation of the diverse physical phenomena involved, and necessitate care and restraint in interpretation. However, both models can be extended to higher levels of accuracy as computational capabilities allow for less modifications of the physical system, though an improved understanding of plasma interactions with thruster surfaces will be required.

Recently, an objective comparison of the two methods was made by H. Liu, et al. [109], who decide Szabo's method provides more computational savings through the use of γ, but neglect using the mass factor. They instead choose to model neutrals as a fluid, though it unclear to

Table 1.2: Mass r and permittivity γ factors used in previous implementations of the SPL-PIC plasma thruster model. A primary reference may be listed with the thruster designation, and supplementary references are also provided as needed.

Thruster Designation	r	γ
mini-TAL [97,110]	250, 750 [122], 2500	5, 10
BHT-1000 [123]	2500, 1000 [124]	20
BHT-1500 [125]	2500 [121]	20 [121]
BHT-1500 [126]	625 [121]	20 [121]
MHT-9 [127]	625 [121]	20 [121]
P5	1000	50 [110], 40 [111,112,116]
DCFT	1000	1200 [114], 50 [34], 250

the author how this approach can address differences between electron and neutral timescales in simulations attempting to include neutral population dynamics. For their simulations of the SPT-100, Liu, et al. choose $\gamma = 40$, identical to Sullivan's choice for modeling the P5 thruster using the MIT plasma thruster model [116]. In 2005, Irishkov, et al. [118] presented a kinetic model they developed to study Hall thrusters which follows the logical development of Szabo's model very closely. They also make use of the permittivity factor, but do not describe any reduction of the simulated ion or neutral masses. To model the KM-37 thruster, Irishkov, et al. choose $\gamma \approx 22$. The mass and permittivity factors used with SPL-PIC to model several thrusters are given in Table 1.2. SPL-PIC results for the BHT-1000 are also presented in Ref. 119 with detailed comparisons to measurements. Also, a Hall thruster operating with bismuth as propellent [120] was simulated using $r = 2500$ and $\gamma = 20$ [121], and J. Fox describes modeling a high-power $(20 - 40\ [kW])$ Hall thruster in Ref. 111.

After summarizing the ways in which plasma physics and collisional processes have been altered, it is appropriate to question the value of applying a simplified particle model rather than a fluid model, especially considering the relative computational simplicity of the fluid models. Incomplete theoretical understanding of cross-field electron diffusion encourages this sentiment further. Neither approach can rightfully be expected to deliver quantitatively accurate predictions for all the quantities of interest. Tuning the models to match thrust measurements will invariably lead to a mismatch in anode current or skew predictions of ion flow to thruster

surfaces. In a simple sense, however, a particle model benefits from the fact that none of its simplifying assumptions are as arbitrary or uncontrolled as the closure choice required to formulate a solvable fluid model of Hall thrusters. This is equivalent to assuming the form of the distribution function beforehand, limiting all possible solutions to an overly restricted subset of phase space. On the other hand, a particle model attempts to solve the correct problem in a general, though statistically sampled, phase space. A particle model can never provide an answer that is inconsistent with the fact that the number of particles within a given volume of phase space is greater than or equal to zero. The same is not true of fluid models, and detecting the occurrence of these situations requires careful analysis in even the simplest problems [74]. Physical modifications to particle models only affect *how* the system evolves in time, as opposed to the permitted mathematical *form* of the solution. Additionally, as computational hardware capabilities improve, coinciding with the development of more efficient and scalable software models, the magnitudes of necessary changes to the physical systems being simulated will diminish. Fluid models are not as extendable in this sense, unless many higher order moments are incorporated in to existing models. These arguments are broad and open to criticism, but do attempt to differentiate between the shortcomings of each approach.

1.2.7 Collisional Processes & Mean-Free-Path Estimates

For binary collisions, the total reaction rate of a collisional process (P) between two groups of particles i and j with a relative speed g_{ij}, is given by Equation 1.17. The cross section $Q_{ij}^{(P)}$ for the process is generally a function of g_{ij} — n_i and n_j are the number densities for each species. From this, the frequency $\nu_{ij}^{(P)}$ with which a single particle of type i is subject to the interaction P is given by Equation 1.18. The total collisional frequency $\nu_{i,T}$ for a specified class of particles is the sum over the frequencies of all possible collisions with other species present in the environment. The mean free-path (MFP) of a particle between collisions is then given by Equation 1.19, taking into account the speed of the particle V_i. Representative MFP estimates can be calculated by using average values of g_{ij}, though the actual reaction rates, cross sections, collision frequencies and mean free-paths can be calculated — in principle — by integrating over the velocity space of each species involved in the collision. In the event that a collisional processes involves two particles of the same type, the result must be halved to compensate for

double counting.

The mean free-paths for several different collisional processes for ions, electrons and neutrals are estimated to aid in the discussion of why a thruster model utilizing a kinetic representation of the plasma will be more accurate. When Kn values do not satisfy $Kn \ll 1$, fluid models may not even provide qualitatively accurate predictions of certain quantities, such as the heat flux vector[85]. For simplicity, only singly charged ions are included in these arguments, though the method of analysis is easily extendable to doubly charged ions if cross section data are available.

$$R_{ij}^{(p)} = n_i n_j g_{ij} Q_{ij}^{(P)} \tag{1.17}$$

$$\nu_{ij}^{(P)} = \frac{R_{ij}^{(P)}}{n_i} = n_j g_{ij} Q_{ij}^{(P)} \tag{1.18}$$

$$\lambda_{i,T} = \frac{V_i}{\nu_{i,T}} \tag{1.19}$$

Based on data compiled by Szabo[97], the mean free-paths associated with a number of collisional processes can be estimated and compared to important macroscopic lengths in the simulation domain. These MFP estimates are made assuming $T_e = 5{-}25\ [eV]$, $n_e = 0.5 - 5.0 \times 10^{18} [m^{-3}]$, $n_n = 10^{18} - 10^{20} [m^{-3}]$, $T_i = 1\ [eV]$, and $T_n = 500[K] \approx 0.04\ [eV]$. Where two values are listed, the values correspond to the near-anode and exit plane. These numbers were selected based on preliminary simulation results of the DCFT[34]. Note that, for a strongly magnetized electron, the guiding-center MFP will be much less than the total distance traveled between collisions because the velocity component responsible for the gyro-motion of the particle does not contribute to a displacement of the guiding-center. As a result, the presence of a magnetic field in the chamber can cause particles that would otherwise have a MFP much larger than thruster dimensions to experience one or many collisions on their way through the device. The cross sections for each collision process considered here are listed and discussed further in Appendix D. A detailed discussion of Coulomb collisions is also provided in the same appendix. In the following paragraphs, cross-section estimates for specified particle interactions are presented, followed by a discussion of the average MFP values for each species. Recombination occurs at a negligible rate, and is not included in the model[97] (pp. 105-106).

Table 1.3: Cross section estimates for collisional processes in the DCFT based on the assumed values in Table D.1. The total cross section for (e-n) interactions includes excitation, ionization, and elastic scattering. The excitation and elastic scattering (es) cross sections are estimated to range from 0-0.5 $\times$ 10^{-15} cm^2 for the range of energies considered here.

Cross Section	$Q_{ij}^{(P)} \times 10^{15}$ cm^2
Q_{en}^{tot}	2.5-4.0
Q_{in}^{cex}	6.0-7.2
Q_{in}^{es}	1.5-3.4
Q_{nn}^{es}	5.5

Summary of MFP Estimates

Estimates of cross sections for various collisional processes are summarized in Table 1.3. These values are not adjusted to include effects from the mass factor. Cross section estimates for Coulomb collisions are provided in Table D.2, though all but one value there assumes thermal energies for electrons and ions. Higher energy particles will have much lower cross sections. In any event, the MFP is a more important metric of a collision's importance, discussed in the next section.

Discussion of Mean-Free-Path Analysis

Obtaining the values of Kn in the simulation domain without a solution requires estimates to be made. Measurements of ion velocity[43], ion current[44] and erosion rates[128] in the cusps of the DCFT indicate that the strongest gradients in these parameters are present in the magnetic cusps. Preliminary simulation results also confirm these observations[34,114]. The strong influence of the magnetic topology, especially the magnetic separatrices, suggests that the scale length of gradients in magnetic field strength L_B is an appropriate choice for estimating Kn values for each species in the flow. The scale lengths of the magnetic vector potential or stream function are not appropriate because their values are both zero on all separatrix surfaces. The field strength only goes to zero where the ring cusp separatrices intersect the thruster axis of symmetry.

Based on simulations of the magnetic field completed and analyzed using the Maxwell SV software package, a representative value of L_B in ring-cusps is 5 mm. This number increases

in between ring-cusps and past the exit of the thruster. Table 1.4 summarizes the estimated values of Kn calculated near the anode and at the exit of the DCFT, based on the previously discussed estimates of plasma parameters and cross sections. These results suggest that fluid models based on extensions of the Navier-Stokes equations are inappropriate for an accurate description of ions and electrons in a DCFT, especially at the exit where the acceleration and majority of ionization is thought to occur. The fluid description of neutrals is marginal at the exit, and will break down further as the plasma expands past the exit into vacuum. The plasma is sufficiently rarefied such that the closure of the system of equations for the mass, momentum and energy moments of the Boltzmann equation, based on the assumption of small deviations from equilibrium, is not valid.

Table 1.4: Knudsen numbers estimated based a scale length of 5 mm and cross sections described above. Ignoring Coulomb collisions, the values increase by factors of about 60 and 1.5 for ions and electrons, respectively. Based on these estimates, Coulomb collisions are more important for ions than electrons. Electrons are strongly magnetized, however, so the MFP of the guiding center will be lower than the total MFP shown here.

Species	Near Anode	Near Exit
Electrons	3	500
Ions	.03	5
Neutrals	.02	.14

In Chapter 7 of their book, Shkarofsky, et al.[70] discuss Coulomb collisions in great detail and give parameters for estimating when certain treatments of the Coulomb collisions are valid. For instance, the effects of many-particle interactions are well approximated by the binary collision assumption when macroscopic changes occur on a time scale longer than the inverse electron plasma frequency, and when $\omega_{c,e} < \omega_{p,e}$ — note that the condition $\omega_{c,e} > \omega_{p,e}$ is equivalent to $r_{L,e} < \lambda_D$ in an isotropic plasma. Even so, in many cases the necessary corrections are accounted for within insensitive logarithmic factors, and thus are unable to affect results beyond desired levels of accuracy. Even a basic summary of these effects is beyond the scope of this discussion, but several starting points on this topic are available in advanced textbooks, with references to more thorough arguments provided within [Shkarofsky, et al. (1966) pp. 17-19, also Ch. 7[70]; Montgomery & Tidman[99]; Reviews of Plasma Physics, Vo1. 1 (1965) pp. 105-111, 152-167[129]; Mitchner & Kruger (1973) pp. 59, 258, 328, 343[71]].

Again, for an isotropic plasma, the threshold magnetic field B_t for which $\omega_{c,e} = \omega_{p,e}$ is given by $B_t = \sqrt{\frac{m_e n_e}{\epsilon_o}}$. Based on simple estimates with $n_e = 10^{18}$ $[m^{-3}]$, $B_t \approx 0.1$ $[T]$. It is interesting to note that this field strength is exceeded by as much as a factor of 5 in the DCFT, represents an approximate maximum field-strength in cylindrical Hall thrusters, and exceeds maximum Hall thruster magnetic fields by about a factor of 5. In simulations using enlarged free-space permittivities $\epsilon_o^* = \gamma^2 \epsilon_o$, the threshold field strength is reduced by an entire factor γ. The significance of differences in field strengths relative to B_t in different thrusters was not examined in this thesis — nor were any implications of γ-induced reductions in B_t considered. However, it may be possible that this effect is important in the DCFT, which would require modeling electron-electron and electron-ion Coulomb interactions to asses any consequences arising from the electron Larmor radius supplanting the Debye length as the shortest scale-length in certain regions.

1.3 Summary of Contributions

In Chapter 2 results from three experiments that helped shape the basic understanding of the DCFT are presented. The first focuses on the identification of the oscillatory bifurcation that accompanies transitions from LC-mode to HC-mode. Before this was investigated, the visual appearance of the plume, and the time-averaged anode current value were the main indicators of mode operation. After this, a follow up study showing how trends in oscillation waveforms, amplitudes, and frequencies vary as operational inputs are changed is discussed. In Sec. 2.3, results from a 200 $[h]$ test enabling the measurement of the erosion rate distribution within and between all three cusps are given. These measurements established that erosion rates in the DCFT are 50% lower than rates observed in any comparable Hall thruster near the beginning of its operational lifetime.

Chapter 3 provides important descriptions of various aspects of PTpic, focusing on the integration particle trajectories and their subsequent interactions with different simulation boundaries and other particles via collisions. A number of significant developments from this work are described in Chapter 3, including a correction to how the leapfrog method is implemented, and a new approach to tracking particles as they move on the mesh, avoiding errors that were able to persist in the first version of the code the author received. In Sec. 3.2.2, the erosion

module that was created for PTpic based on the works of S. Cheng and J. Yim is described, as well as the available inputs that the user can provide as inputs. The last major topic from this chapter discusses how the ionization and charge-exchange algorithms had to be modified in order to conserve mass and charge during every iteration, and in every cell. The challenging aspect of that problem was working within the constraints of how the model was previously parallelized by J. Fox[111].

All of Chapter 4 is devoted to improvements made to the performance and accuracy of the potential solver. A more accurate set of finite difference equations for the integral 9-point method are derived and have been successfully implemented and benchmarked for certain situations. In favor of a previous iterative solver, whose issues with convergence are discussed, a more scalable direct solver based on the ScaLAPACK software package has been used productively in simulations. Finally, two correct approaches to properly representing boundary conditions on a charged floating body in the presence of a plasma are discussed, and specific proposals for solving the problem numerically are described.

The application of the newly developed model to the DCFT in both the low and high-current modes is presented and discussed in Chapter 5. Here, aside from various shortcomings whose causes have likely been identified, the reproduction of two important characteristics of the DCFT by PTpic are examined in detail: specifically the model's representation of hollow ion jets in the near plume and erosion profiles predicted by the sputtering module are described in Chapter 3. Finally, in Chapter 6 the adaptability of PTpic to new designs, as well as identifiable limitations, are described. The main conclusions of this thesis are presented in Chapter 7.

Chapter 2

Fundamental DCFT Experiments

Edwards AFB AFRL "Chamber 1"

This vacuum chamber is described here because it is used in multiple experiments discussed in this chapter. Chamber 1 is cylindrical, with a length of 4.1 m and a diameter of 2.4 m. The vacuum chamber is made of non-magnetic stainless steel, and the inner surface is covered by flexible sheets of graphite. Each end of the channel is protected by angled graphite panels to mitigate back sputtering, and the floor downstream of the thruster is covered by a graphite blanket for the same reason. The vacuum is established and maintained primarily by two 1.2 m gaseous helium two-stage cryogenic (15 K) pumps, with liquid nitrogen cooled (75 K) baffles. The total xenon pumping speed is 48,500 L/s. Pumping was supplemented by a turbomolecular pump, to affect gases such as hydrogen. Uncorrected chamber pressure was measured using a cold cathode gauge, calibrated for nitrogen, and reduced by a correction factor of 2.9 to account for xenon as the dominant gas constituent.

2.1 Identifying the Oscillation Bifurcation & HC-mode Plume Measurements

The experiments described in this section were first reported in a paper presented at the 2009 International Electric Propulsion Conference (IEPC) [29], and were performed in Chamber 1. An automated 2D (R-Θ) stage was used to move Faraday and retarding potential analyzer (RPA) probes within the plume of an operating DCF thruster to measure the characteristics of the

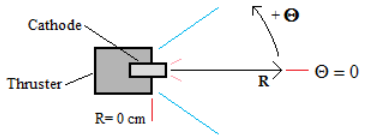

Figure 2-1: Top view of the 2D (R-Θ) coordinate system. The $\Theta = 0$ line is aligned with the DCFT axis of symmetry. The cathode is above the probing plane.

plasma being ejected. The arm mounted on the rotary stage is a stepper motor drive linear translational stage capable of traversing up to 80 $[cm]$ from the center of rotation of the rotary stage. This combination of rotary and linear motion allows for the continuous sweeping of plasma probes in the R-Θ plane. Further details of the experimental facilities are provided by Nakles et al. in Ref. 130.

Fig. 2-1 illustrates the coordinate system used to present the probe data. The thruster was fired for a total of 18 $[h]$ during this investigation. Thermocouples located on the thruster exterior near the anode and outer case enclosing the permanent magnets indicate that the temperature did not exceed 200 $[°C]$ at those locations during all tests performed for this investigation. Measured chamber pressure did not exceed 3.5 $[\mu Torr]$. A BHC-1500 hollow cathode was positioned with the plasma orifice center located 27.5 $[mm] \pm 1$ $[mm]$ downstream of the exit plane and 31.8 $[mm] \pm 1$ $[mm]$ outside the maximum radius of the thruster at the exit plane, as in previous investigations[27,131,132]. Furthermore, the thruster body and cathode potentials were allowed to float with respect to chamber ground during the tests described in this section. In several tables, for instance Table 2.1, anode voltage V_a, anode mass flow rate $\dot{m}_a$, anode current I_a, cathode mass flow rate $\dot{m}_c$, keeper current I_k, and keeper voltage V_k are provided to summarize the operating condition of the thruster while anode current waveform and plume measurements were acquired.

2.1.1 Current Oscillation Measurements

This section presents time resolved anode currents as recorded with an oscilloscope and processed with a signal analyzer. They demonstrate a conspicuous difference between the anode current behavior in HC and LC-mode. The anode current waveforms shown in Fig. 2-2, representing the DCFT operating in HC-mode, are characterized by high magnitude, low frequency oscillations. For all high-current operating conditions examined here, the lowest harmonic of

the frequency spectrum is between 3.5 $[kHz]$ and 4.0 $[kHz]$. The value of the lowest harmonic frequency appears to be unaffected by the operating condition of the thruster, in agreement with the theoretical frequency of highly nonlinear breathing modes in plasma discharges given by Barral and Ahedo[133]. They suggest that the frequency of current oscillations is approximately equal to the quotient of the neutral thermal speed and a characteristic thruster length. The residence time of a neutral xenon particle with a thermal speed corresponding to a temperature of 1000 $[K]$ over the 5 $[cm]$ axial length of the DCFT is approximately 0.2 $[ms]$. The frequency corresponding to the inverse of this time is 5 $[kHz]$, in close agreement with the lowest harmonic frequency of anode current oscillations for the DCFT operating in HC-mode.

The waveforms in Fig. 2-3, representing the DCFT operating in LC-mode, do not exhibit the same oscillatory behavior as seen in HC-mode. The large oscillations present in HC-mode are reduced by several orders of magnitude. The narrow dispersion of ion energies in Fig. 2-9 within regions of maximum current density, shown in Fig. 2-5(d), may result from the lack of strong oscillations in LC-mode, as laser-induced fluorescence studies of Hall thrusters have shown[134]. Previously[27], the DCFT was found to operate more efficiently in LC-mode than in HC-mode — though within uncertainty margins. At the time this study was completed, at the end of the summer in 2008, the lack of strong oscillations in LC-mode further suggested that it should be the preferred mode of operation.

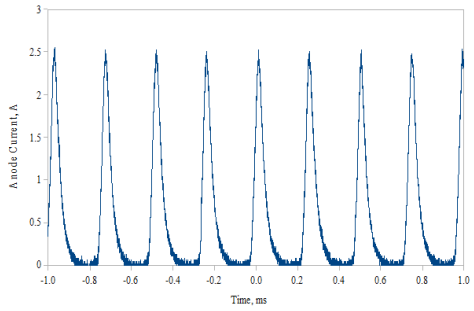

(a) Time resolved anode current. The thruster operating condition for these data is given in Table 2.1.

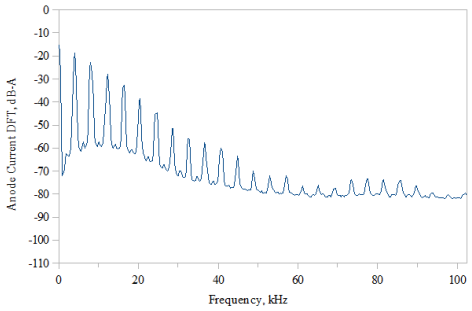

(b) Anode current spectrum for the waveform shown in Fig. 2-2(a).

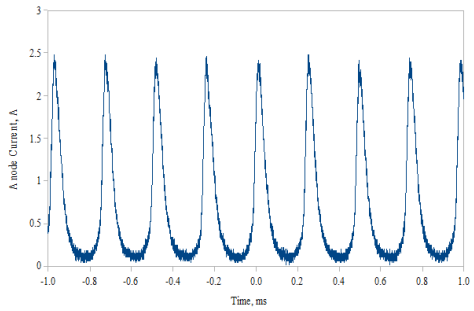

(c) Time resolved anode current. The thruster operating condition for these data is given in Table 2.2

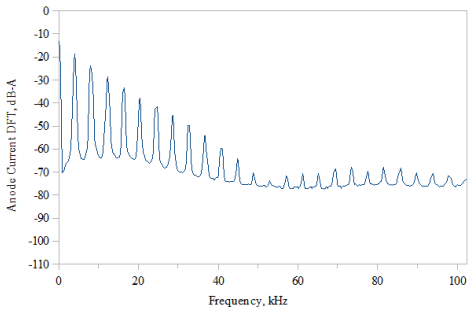

(d) Anode current spectrum for the waveform shown in Fig. 2-2(c)

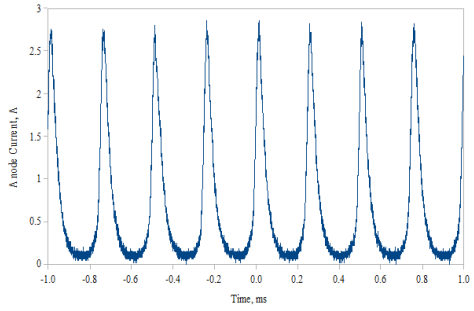

(e) Time resolved anode current. The thruster operating condition for these data is given in Table 2.3

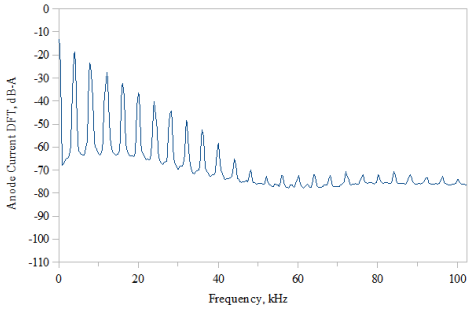

(f) Anode current spectrum for the waveform shown in Fig. 2-2(e)

Figure 2-2: Time resolved anode currents and spectra for HC-mode.

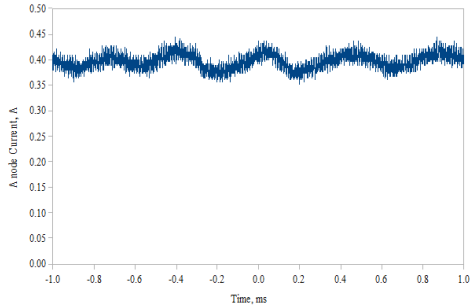

(a) Time resolved anode current. The anode and cathode settings, when these data were recorded, were identical to those provided in Table 2.3.

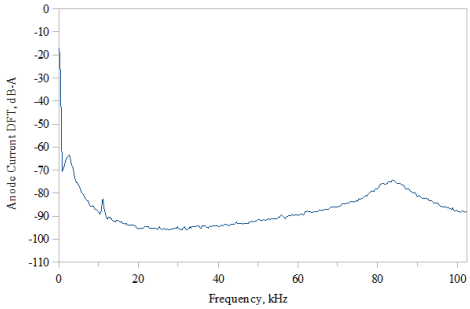

(b) Anode current spectrum for the waveform shown in Fig. 2-3(a). These data were recorded while the thruster warmed, after which it transitioned to the high-current mode, as in Fig. 2-2(e).

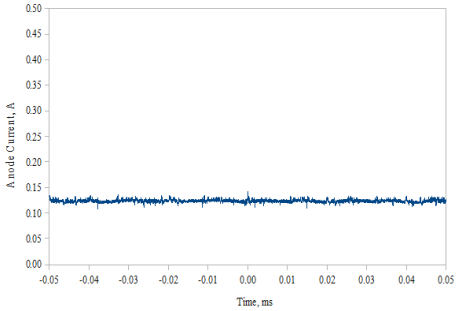

(c) Time resolved anode current. The thruster operating condition for these data is given in Table 2.7.

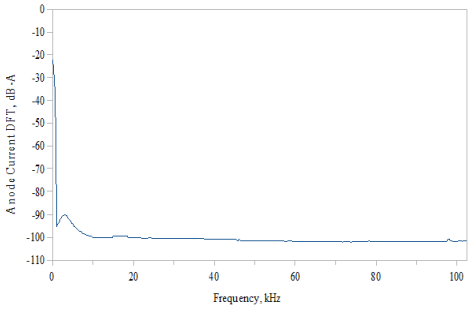

(d) Anode current spectrum for the waveform shown in Fig. 2-3(c).

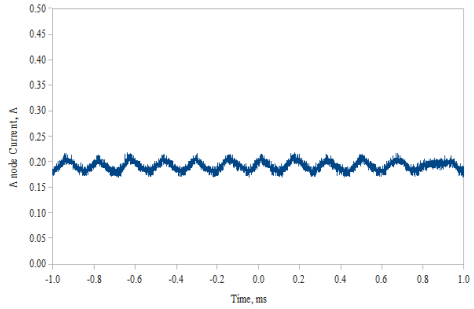

(e) Time resolved anode current. The anode and cathode settings, when these data were recorded, were identical to those provided in Table 2.4.

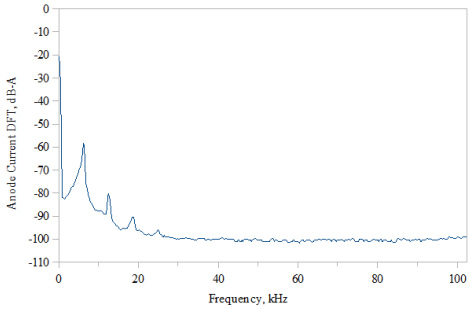

(f) Anode current spectrum for the waveform shown in Fig. 2-3(e).

Figure 2-3: Time resolved anode currents and spectra for LC-mode.

2.1.2 Ion Current Density & Energy Flux Distributions

This section contains Faraday cup (FC) results for selected operating conditions, plotted as current per steradian in Fig. 2-5. The values plotted in Fig. 2-5 are the product of j and R^2, where j is the local ion current density and R is the measurement radius. The FC data points are angularly resolved from $\pm 120^\circ$ off axis in 1° intervals in both directions, taken at two different radii from the thruster exit plane at each condition. The Faraday cup used to measure current density is pictured in Fig. 2-4(b). Ion current is collected and measured on the inner disk, measuring 8.3 $[mm]$ in diameter. A concentric guard piece, measuring 22.5 $[mm]$ in outer diameter, is used to minimize the effects of the plasma sheath on the ion current collector's effective collecting area. The electrodes are constructed from molybdenum. A 0.56 $[mm]$ wide gap exists between the outer wall of the collector and the inner wall of the guard ring. The effective current collector area of the probe is calculated by adding a portion of the gap surface area to the collector surface area as suggested in Ref. 135. Ion charge flux is determined by dividing the current to the collector by its effective surface area. The disk and guard ring are biased to -30.0 $[V]$ with respect to chamber ground so that ion saturation is achieved. The effects of secondary electron emission are assumed to be less than a few percent[136] and are neglected in this analysis.

Figures 2-5(a) through 2-5(c) show Faraday cup data for the DCFT operating in HC-mode at various operating conditions. The uncertainties in the anode currents and keeper potentials represent the limits over which the values drifted during measurements, due in part to the thruster warming or cooling with changes in operating condition. The basic plume structure is

Table 2.1: HC-mode operating condition for which measurements are presented in Figs. 2-2(a) and 2-5(a).

V_a=	300 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	8.5 [sccm Xe]	I_k=	0.23 [A]
I_a=	0.53 ± 0.03 [A]	V_k=	16.8 ± 0.6 [V]
$R = 30$ [cm]			
I_b=	0.38 [A]	$\Theta_{90\%}$ =	72°
$R = 60$ [cm]			
I_b=	0.43 [A]	$\Theta_{90\%}$ =	72°

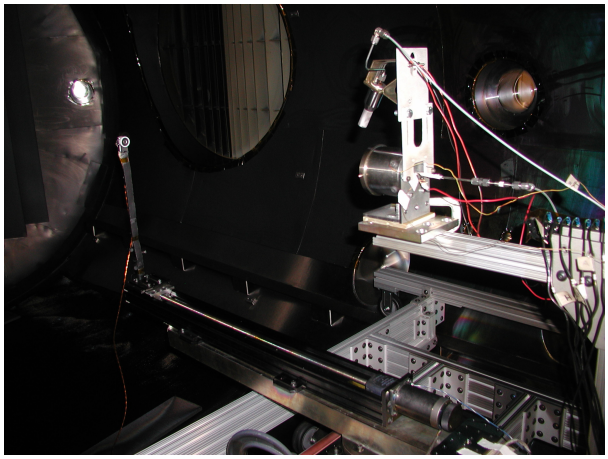

(a) Faraday cup mounted and aligned to measure current density within the DCFT plume.

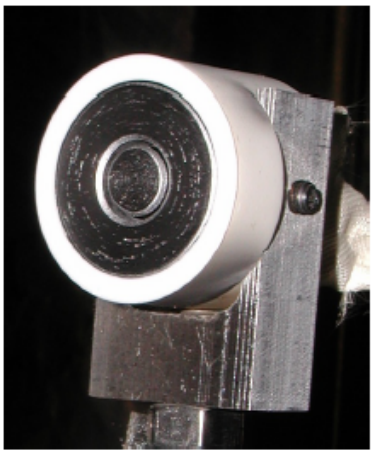

(b) Detail of the Faraday cup.

Figure 2-4: The Faraday cup used to record plume measurements.

Table 2.2: HC-mode operating condition for which plume measurements are presented in Figs. 2-2(c) and 2-5(b).

V_a=	450 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	8.5 [sccm Xe]	I_k=	0.23 [A]
I_a=	0.58 ± 0.03 [A]	V_k=	16.4 ± 0.5 [V]
	$\boldsymbol{R = 30\ [cm]}$		
I_b=	0.45 [A]	$\Theta_{90\%}$ =	70°
	$\boldsymbol{R = 60\ [cm]}$		
I_b=	0.40 [A]	$\Theta_{90\%}$ =	70°

Table 2.3: HC-mode operating condition for which plume measurements are presented in Figs. 2-2(e), 2-3(a), and 2-5(c).

V_a=	450 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	8.5 [sccm Xe]	I_k=	0.48 [A]
I_a=	0.617 ± 0.005 [A]	V_k=	16.1 ± 0.2 [V]
	$\boldsymbol{R = 30\ [cm]}$		
I_b=	0.46 [A]	$\Theta_{90\%}$ =	70°
	$\boldsymbol{R = 60\ [cm]}$		
I_b=	0.43 [A]	$\Theta_{90\%}$ =	69°

Table 2.4: LC-mode operating condition for which plume measurements are presented in Figs. 2-3(e) and 2-5(d).

V_a=	475.7 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	5.8 [sccm Xe]	I_k=	1.48 [A]
I_a=	0.1415 ± 0.0015 [A]	V_k=	17.1 ± 0.5 [V]
	$R = 30$ [cm]		
I_b=	0.133 [A]	$\Theta_{90\%}$ =	61°
	$R = 60$ [cm]		
I_b=	0.135 [A]	$\Theta_{90\%}$ =	61°

identical in all three plots, showing a hollow plume with the maximum current density located between $30° \leq |\Theta| \leq 35°$. The current density on the centerline of the thruster is about a factor of two lower than in the densest regions. Note the logarithmic scale of the vertical axes in all the figures shown in Fig. 2-5. The current density also begins to increase past $|\Theta| > 80°$ before decreasing again at larger angles. This is most likely due to charge exchange collisions with neutral particles. Repeating these measurements at varied chamber pressures would determine whether the ions are colliding with unionized xenon in the thruster flow, or with the background gas population of neutral particles within the vacuum chamber.

Figure 2-5(d) shows the current density per steradian for the DCFT operating in the low-current mode. The anode current is much lower here than for the high-current mode, though the basic structure of the current density appears identical over a wide range of angles. The one exception is on the thruster centerline, where the current density is less than the maximum value seen between $30° \leq |\Theta| \leq 35°$ by an order of magnitude, compared to the high-current mode, where the centerline current density dropped by only a factor of 2 compared to the maximum value. Although the data suggest that current densities on the centerline are reduced in the low-current mode relative to the high-current mode, the dissimilarities in the operating conditions between the two modes suggest that more plume probing of the thruster operating in the low-current mode is needed.

Values for the beam current I_b and 90%-current half-angle $\Theta_{90\%}$ are shown in Tables 2.1–2.4 for each of the data sets shown in Fig. 2-5. The beam current is determined by integrating current density for $0° \leq |\Theta| \leq 90°$, which also allows $\Theta_{90\%}$ to be determined. High-current

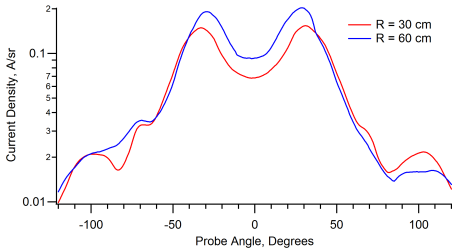

(a) The operating condition and integral plume parameters corresponding to these data are provided in Table 2.1.

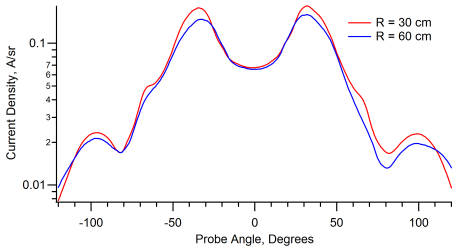

(b) The operating condition and integral plume parameters corresponding to these data are provided in Table 2.2.

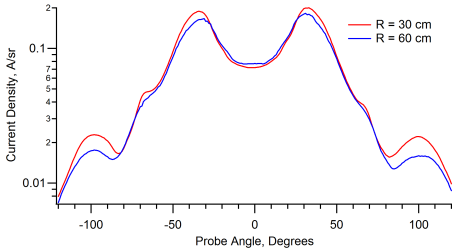

(c) The operating condition and integral plume parameters corresponding to these data are provided in Table 2.3.

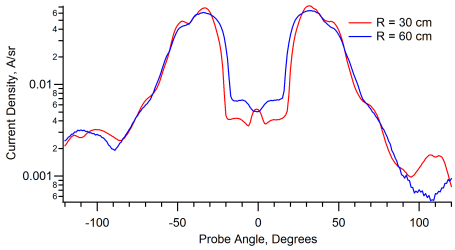

(d) The operating condition and integral plume parameters corresponding to these data are provided in Table 2.4.

Figure 2-5: Faraday cup measurements at different operating conditions and locations in the DCFT plume. Only Fig. 2-5(d) shows data from LC-mode — the other three data sets were recorded while operating in HC-mode.

conditions are characterized by lower ratios of beam to anode current, but higher ratios of beam to neutral current[a]. The ion current corresponding to the complete single ionization of 8.5 $[sccm\ Xe]$ is 0.614 $[A]$, and 0.419 $[A]$ for 5.8 $[sccm\ Xe]$. Additionally, for the LC-mode condition shown in Fig. 2-5(d), the ion beam in the plume is more focused than the HC-mode cases because it has the smallest $\Theta_{90\%}$ values. However, unlike in past studies[27,132], the low-current mode could only be maintained long enough to complete probe sweeps by increasing the cathode keeper current in all tests performed for this investigation. Previously, increasing the anode potential was often sufficient to initiate the transition from the high to low-current mode. Subsequently, it is difficult to justify drawing conclusions related to the benefit of one mode over the other based on integrated current density analyses.

[a]Here, "neutral current" represents what the beam current would be if all the incoming propellent were singly ionized and contributed to the beam leaving the thruster.

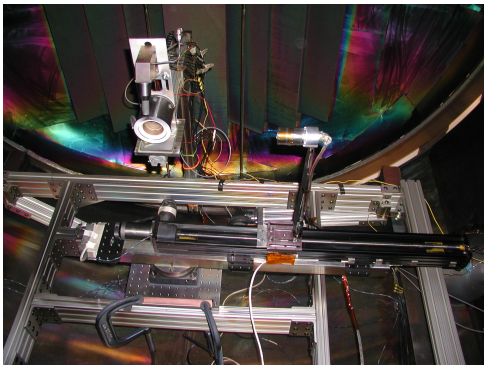

(a) RPA mounted and aligned to measure energy distributions of ion flux within the DCFT plume.

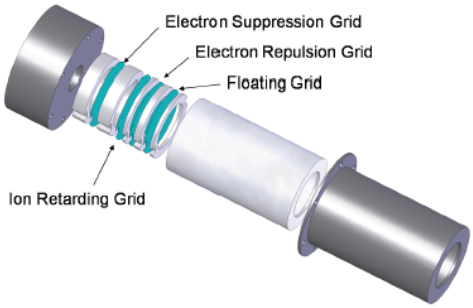

(b) Detail of the retarding potential analyzer.

Figure 2-6: The retarding potential analyzer used to record plume measurements.

This section presents RPA results for several different operating conditions, shown in Figs. 2-7–2-9. The RPA data are angularly resolved from $\pm 90^\circ$ or $\pm 120^\circ$ off axis in 5° intervals in both directions, taken at two different radii from the thruster exit plane at each condition, except where indicated. The RPA used in this experimental investigation is shown in Fig. 2-6(b) and employs four grids. The floating grid potential remains at the local plasma potential and is meant to minimize probe effects within the plume. The electron repulsion grid, biased negatively with respect to chamber ground, repels incoming electrons while allowing ions to pass through. The potential of the ion retarding grid is swept positively relative to chamber ground to allow selective passage of ions based on their kinetic energy. The electron suppression grid is placed in front of the collector to suppress the electron current induced by ion impacts on the retarding grid. Reference 135 provides a more complete description of the RPA used in this investigation.

As the ion retarding potential is increased from the ground potential to a specified limit, ions with lower energies are prevented from reaching the collector. The difference in ion currents measured between two distinct retarding potentials provides the portion ion flux comprised of ions with energies between the energies needed to overcome the retarding potentials. For appropriately small steps in retarding potential, the energy distribution of ion flux incident upon the RPA at a particular location is proportional to the negative of the derivative of the ion current with respect to the retarding potential, as seen in Eqs. 2.1–2.3. The minimum ion

energy allowed to reach the RPA collector (ε_{min}) is assumed equal to the energy associated with the ion retarding potential (eV_r) in order to arrive at Eq. 2.3 from Eq. 2.2. The quantity on the right-hand-side of Eq. 2.3 is plotted in Figs. 2-7 through 2-9 after being normalized such that the integral over the range of given retarding potentials is unity. Hence, the plots show the relative values of the energy distribution of ion flux at each specified location.

$$I = eA_p\Gamma_i \tag{2.1}$$

$$\Gamma_i\left(\varepsilon_{min}\right) \equiv \int_{\varepsilon_{min}}^{\infty} f_\Gamma\left(\varepsilon\right)\, d\varepsilon \tag{2.2}$$

$$f_\Gamma\left(eV_r\right) \propto -\frac{\mathrm{d}I}{\mathrm{d}V_r} \tag{2.3}$$

Figure 2-7 shows RPA data for the DCFT operating in HC-mode as shown in Table 2.5. Data in Fig. 2-8 correspond to the DCFT operating in the high-current once again, but at the condition shown in Table 2.6. Measurements at both operating conditions show energy distributions of ion flux for $0^\circ \leq |\Theta| \leq 20^\circ$ with a mixture of high and low energy ions on the center line, transitioning toward a distribution with less energy dispersion for $25^\circ \leq |\Theta| \leq 50^\circ$, where a sharp peak is evident near the applied anode voltage. The most energetically uniform ion flux distributions are located within regions of the plume that coincide with the greatest current densities, while centerline distributions are broader, comprised of both low and high energy ions. Between $55^\circ \leq |\Theta| \leq 75^\circ$ distributions remain peaked just below the anode voltage, though the number of low energy ions increases as compared to $25^\circ \leq |\Theta| \leq 50^\circ$. At larger angles, for $80^\circ \leq |\Theta| \leq 120^\circ$, the high energy peak disappears and is replaced by a broader, lower energy peak. Although the current density at these high angles is an order of magnitude less than the maximum, the energies of ions measured at high angles in the DCFT are greater than those measured in Hall thrusters [126]. Data identical to those shown in Fig. 2-8 were taken for the same operating condition, but with a keeper current of 0.50 $[A]$ instead of 0.25 $[A]$. Doubling the keeper current at this operating condition had little or no effect on the RPA probe results, which are not presented.

Table 2.5: HC-mode operating condition for which plume measurements are presented in Fig. 2-7.

V_a=	300 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	8.5 [sccm Xe]	I_k=	0.25 [A]
I_a=	0.525 ± 0.005 [A]	V_k=	15.4 ± 0.2 [V]

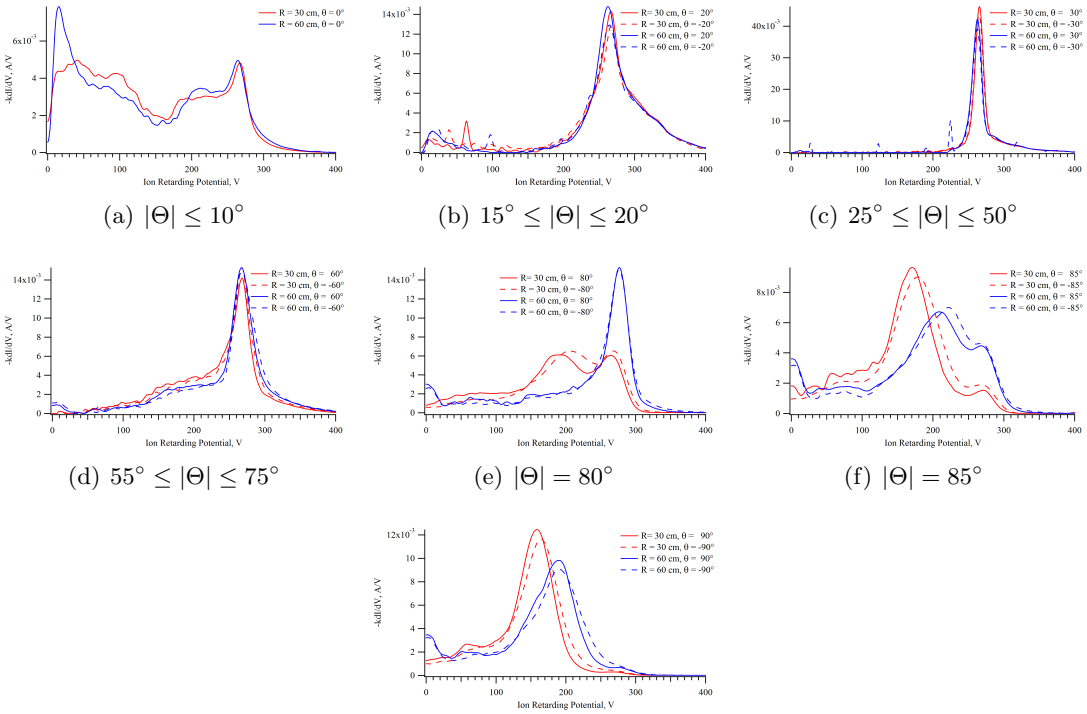

(a) $|\Theta| \leq 10°$ (b) $15° \leq |\Theta| \leq 20°$ (c) $25° \leq |\Theta| \leq 50°$

(d) $55° \leq |\Theta| \leq 75°$ (e) $|\Theta| = 80°$ (f) $|\Theta| = 85°$

(g) $|\Theta| = 90°$

Figure 2-7: Ion energy distributions representing various positions with respect to the DCFT exit plane. The thruster operated as shown in Table 2.5 while these data were recorded.

Table 2.6: HC-mode operating condition for which plume measurements are presented in Fig. 2-8.

V_a=	450 [V]	$\dot{m}_c$=	1.0 [sccm Xe]
$\dot{m}_a$=	8.5 [sccm Xe]	I_k=	0.25 [A]
I_a=	0.57 ± 0.05 [A]	V_k=	15.7 ± 0.1 [V]

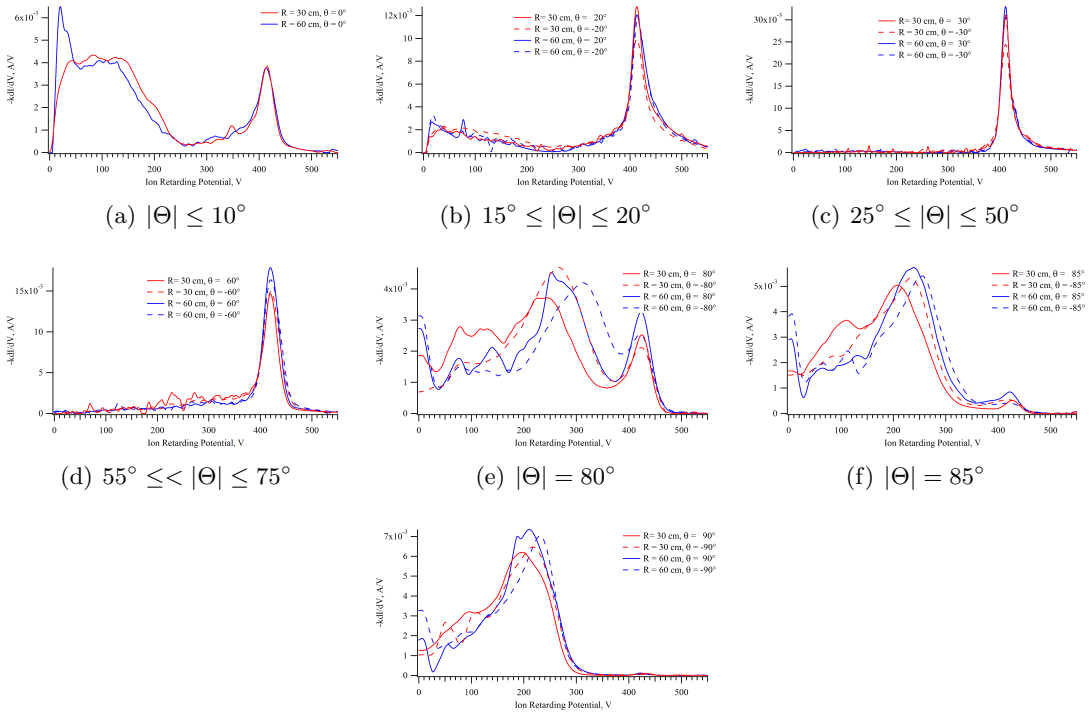

(a) $|\Theta| \leq 10°$ (b) $15° \leq |\Theta| \leq 20°$ (c) $25° \leq |\Theta| \leq 50°$

(d) $55° \leq< |\Theta| \leq 75°$ (e) $|\Theta| = 80°$ (f) $|\Theta| = 85°$

(g) $90° \leq |\Theta| \leq 120°$

Figure 2-8: Ion energy distributions representing various positions with respect to the DCFT exit plane. The thruster operated as shown in Table 2.6 while these data were recorded.

Figure 2-9 shows RPA data for the DCFT operating in LC-mode at the condition shown in Table 2.7. The operating condition for LC-mode RPA data in Fig. 2-9 is not the same as the condition for FC data in Fig. 2-5(d) — compare Table 2.4 and 2.7. The energy distributions of ion flux are much different at this LC-mode operating condition than for the HC-mode conditions shown in Figs. 2-7–2-8. The degree to which these differences are attributable to the different modes or the varied operating conditions is not addressed in the present investigation. Nonetheless, ions exhibit low energy dispersion between $30° \leq |\Theta| \leq 50°$, as in HC-mode. However, the relatively high energy ion populations present at large angles in the high-current mode are absent in the LC-mode examined here, except for $60° \leq |\Theta| \leq 70°$. Instead for $0° \leq |\Theta| \leq 10°$ and $80° \leq |\Theta| \leq 120°$ the energy peaks of the flux distribution are much closer to energies that could be expected of initially slow ions that have been accelerated sideways by a potential hump along the beam.

Table 2.7: LC-mode operating condition for which plume measurements are presented in Figs. 2-3(c) and 2-9.

V_a=	451.5 [V]	$\dot{m}_c$=	2.0 [sccm Xe]
$\dot{m}_a$=	4.3 [sccm Xe]	I_k=	1.98 [A]
I_a=	0.052 ± 0.009 [A]	V_k=	16.55 ± 0.05 [V]

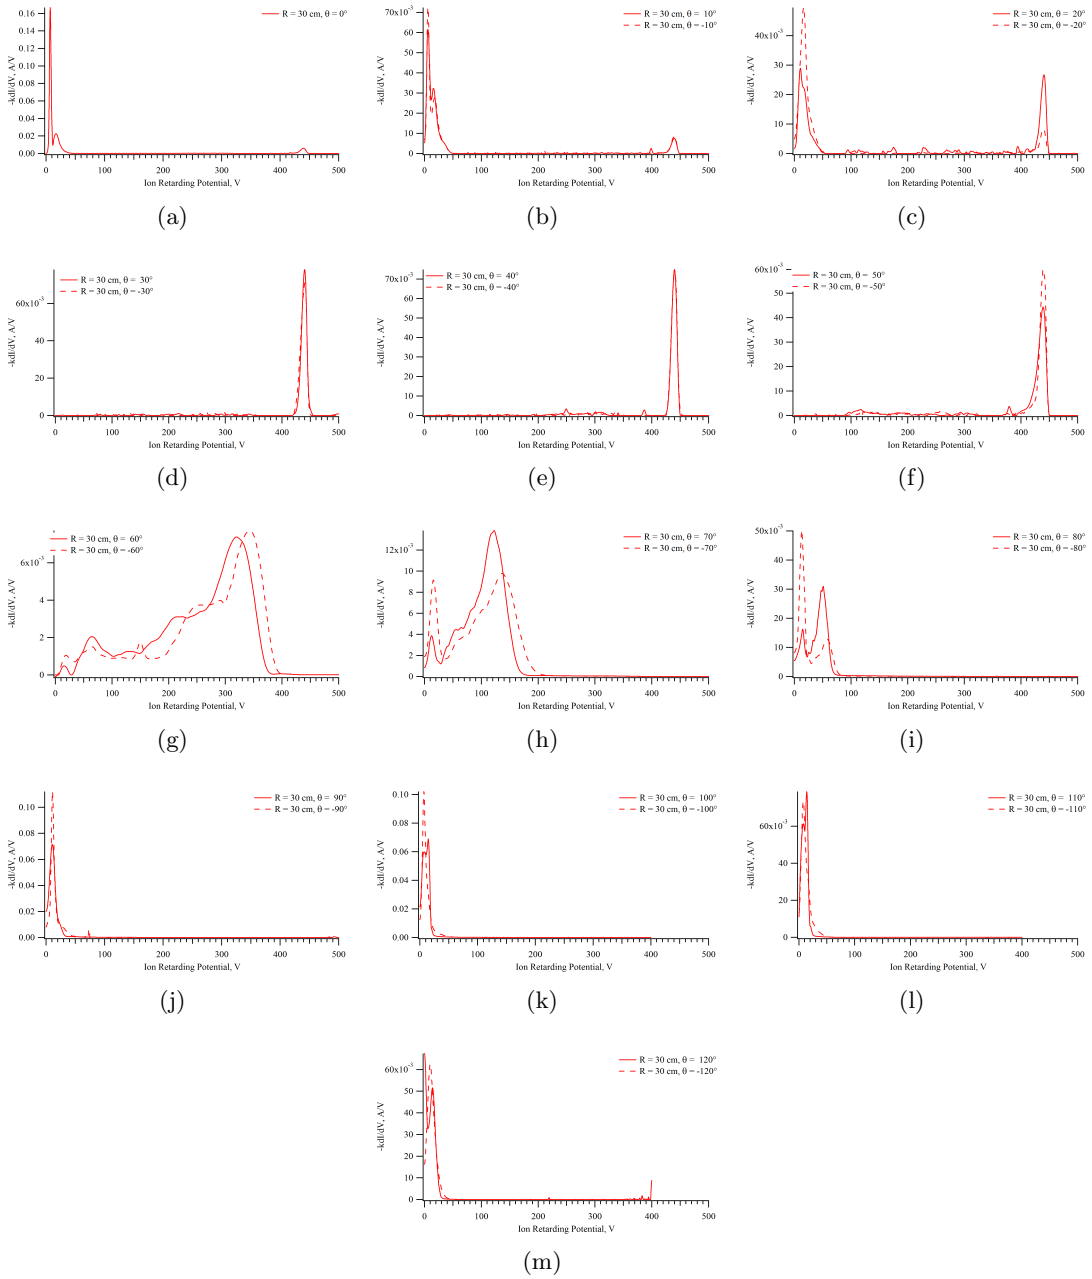

Figure 2-9: Ion energy distributions plotted at varying positions with respect to the DCFT exit plane. The thruster operated as shown in Table 2.7 while these data were recorded.

The plume measurements expand upon and confirm previous RPA and Faraday cup data[27,132]. A comparison of RPA and Faraday data for identical operating conditions shows that the region of the plume with the greatest current density contains the most energetically uniform ion populations. However, the RPA results indicate that highly energetic ions are expelled at large angles relative to the DCFT axis of symmetry in HC-mode, though the current densities at these angles are much lower. Time resolved anode current data reveal a distinct difference between the low and high-current modes. The data demonstrate that the anode current in the low-current mode is much less oscillatory than in the high-current mode. Identifying the cause of this difference in dynamic behavior will likely contribute toward the determination of basic DCFT operating principles in each mode.

While running the thruster for the purpose of this investigation, the operating conditions of the DCFT and cathode were found to be strongly coupled, motivating a future study that systematically examines the effect of cathode location and operating condition on the performance of the DCFT. A detailed look at the ion energy distributions at high angles in the low-current mode would also be beneficial, to determine if high energy ions are as prevalent at high angles in the low-current as in the high-current mode. Additionally, repeating some of the measurements made here at different chamber pressures would address the role that charge exchange collisions with background neutrals play in deflecting ions to high angles.

2.2 Identifying Trends in the Anode Current Waveform

Experimental Setup All results reported here were obtained from experiments conducted at the Massachusetts Institute of Technology (MIT) Space Propulsion Laboratory, and were originally reported in Ref. 30. The DCF thruster was operated in the ASTROVAC vacuum facility, with dimensions of 1.5 by 1.6 $[m]$ and an Argon pumping speed of 7500 $[^{L}/_{s}]$ obtained using two CTI-Cryogenics cryopumps (CT10 and OB400). Pressure is measured using a cold cathode gauge, corrected for xenon and recorded manually.

The body is electrically isolated from chamber ground through a dielectric pad, while anode and cathode keeper potentials are applied with respect to a common chamber and facility

ground. A Busek Inc. BHC-1500 hollow cathode is used in conjunction with the DCFT, and both are powered by Agilent N5772A Programmable Power Supplies. Two Omega FMA-A2400 series mass flow controllers, calibrated to provide $0 - 10$ $[sccm\ Xe]$ each, supply propellant to the anode and cathode. One collecting and one indicating oxygen trap are installed in the cathode flow line to avoid accidental exposure to a warm cathode emitter.

The operating condition of the thruster was monitored in LabVIEW, sampling anode and cathode voltages, currents and flow rates 100 times each minute. The chamber pressure was monitored manually at regular intervals, especially when flow rates into the chamber were varied. The oscillatory nature of the anode discharge current was measured using a Tektronix P6016 AC current probe attached to an Agilent Infinivision oscilloscope from which the waveform was extracted. The time that each current trace was made allowed each measurement to be synchronized with the thruster monitor data.

Each anode current trace was processed to determine the primary oscillation frequency and peak-to-peak amplitude. In cases where a sharp spike in current amplitude occurred, the current doubling and halving times were also measured. Oscillation measurements that were taken during a period of stable thruster operation were correlated to that condition, defined by the mean and standard deviation of anode and cathode currents, voltages and flow rates. In this way, the sensitivities of oscillation characteristics to set operating conditions were examined by varying one variable, such as anode voltage, while keeping other variables fixed. Data reduction was completed using Igor Pro6.04, a graphing and data analysis program.

The anode flow rate was swept from $3.5 - 10$ $[sccm\ Xe]$ with constant anode voltage and cathode operation. The effect on chamber pressure of varied flow raises the pressure linearly from $13.6 - 34.1$ $[\mu Torr]$. A plot of the variation of chamber pressure with total flow rate into the chamber is provided in Fig. 19 of Ref. 30. Other monitored quantities are provided in Table 2.8.

Anode voltage scans were performed for two flow rates. The flow rates were chosen so that the thruster operated primarily in HC-mode at the higher flow rate, and in the LC-mode for the lower flow rate. Other monitored quantities are provided in Table 2.9. The previous description of the pressure variation provides an adequate representation of the pressure variation in the

Table 2.8: Values held constant during the anode flow scan are presented along with their standard deviation if greater than 1% of mean. The duration of the scan was approximately 200 $[min]$ in order to allow the thruster to warm up as anode power increased. Temperature trends were monitored using a uncalibrated thermocouple inserted between the insulating cone and permanent magnets.

Anode Flow Scan Operating Condition	
Cathode Flow	1.5 ± 0.25 [sccm Xe]
Anode Voltage	250 [V]
Keeper Current	0.5 [A]
Keeper Voltage	14.9 ± 0.3 [V]

Table 2.9: Values held constant during the anode voltage scan are presented along with their standard deviation if greater than 1% of mean. The duration of the scan was approximately 130 $[min]$ in order to allow the thruster to warm up as anode power increased. Temperature trends were monitored using a uncalibrated thermocouple inserted between the insulating cone and permanent magnets.

Anode Voltage Scan Operating Conditions		
Anode Flow [sccm Xe]	6.5	8.5
Cathode Flow [sccm Xe]	1.5 ± 0.18	1.5 ± 0.14
Keeper Current [A]	0.5	0.5
Keeper Voltage [V]	15.0 ± 0.15	14.7 ± 0.5
Pressure [μTorr]	21.0	27.7 ± 0.5

chamber due to the difference in anode xenon flow rates for each scan.

Chamber pressure variations from $18.4 - 41.7$ $[\mu Torr]$ were studied at two operating conditions. The lower flow, higher voltage case was intended to provide data in the LC-mode and examine the transition from high to low current that can be induced by changes in pressure. The higher flow, lower voltage case was intended to examine the sensitivity of the HC-mode to pressure changes. Other monitored quantities are provided in Table 2.10 for each operating condition.

Results from these three scans are shown in Sections 2.2.1, 2.2.2 and 2.2.3, showing oscillation frequencies, peak-to-peak amplitudes, doubling times and halving times. The variation of the average anode current for each scan is given in Figure 2-10. Aside from the easily discernible trends, some effects of mode changes are visible in these data.

In Figure 2-10(a), the current-flow plot sharply changes slope below 5 $[sccm\ Xe]$. This is due

Table 2.10: Values held constant during the chamber pressure scan are presented along with their standard deviation if greater than 1% of mean. The duration of the scan was approximately 150 [min] in order to allow the thruster to warm up as anode power increased. Temperature trends were monitored using a uncalibrated thermocouple inserted between the insulating cone and permanent magnets.

Pressure Scan Operating Conditions		
Anode Flow [sccm Xe]	5.0	4.0
Anode Voltage [V]	200	250
Cathode Flow [sccm Xe]	1.5 ± 0.10	1.5 ± 0.09
Keeper Current [A]	0.5	0.5
Keeper Voltage [V]	15.2 ± 0.19	15.9 ± 0.12

to the mode change that was observed while increasing the flow from 3.5 to 4.5 [$sccm\ Xe$]. This is confirmed by the sharp change in oscillation magnitude at this point in the scan, visible in Figure 2-13(a). Another mode change occurred during the voltage scan, whose fixed parameters are given in Table 2.9. Figure 2-13(c) shows the peak-to-peak size of the oscillations dropping from over 2 [A] to less than 60 [mA] when the anode voltage is raised from 300 to 350 [V] for the 6.5 [$sccm\ Xe$] condition. The average value of the anode current also decreases. Although the anode current continues to rise during the pressure scan shown in Figure 2-10(b), excessive chamber pressures caused the thruster to switch into a third mode. This mode is accompanied by a visual change in the plume appearance from the typical blue-green color to a a tint of red. This *red*-mode, is sometimes accompanied by a drop in oscillation amplitude similar to what happens during a typical HC to LC-mode transition. However, the thruster often is not stable in this mode, and will frequently jump between the HC, LC and red-modes, preventing the oscillation magnitude from dropping. This was the case for the 5 [$sccm\ Xe$], 200 [V] pressure scan operating condition. Here, the plume was tinted red while the average anode current and oscillation magnitude were maintained after the mode transition. The results of the pressure scan suggest that thruster operation in the red-mode may be an indication that facility effects are having a dominant effect on thruster operation.

Uncertainties of plotted quantities are displayed when available. In Figures 2-10, 2-11, 2-13 and 2-14, uncertainties representing a 95% confidence interval are shown by horizontal and vertical bars where their span exceeds the approximate size of the marker used for each data point. No statistical variation in the primary frequency, peak-to-peak amplitude or rise and fall

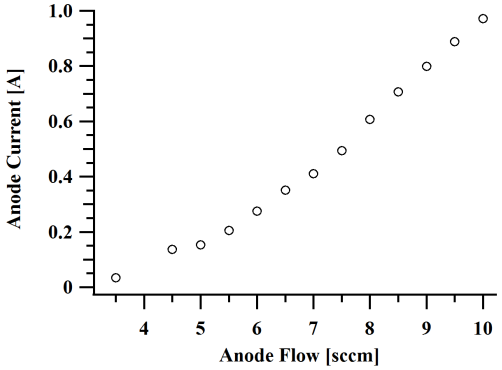

(a) Variation of anode current during the anode flow scan. The anode voltage was fixed at 250 $[V]$. The cathode operating condition is described in Table 2.8

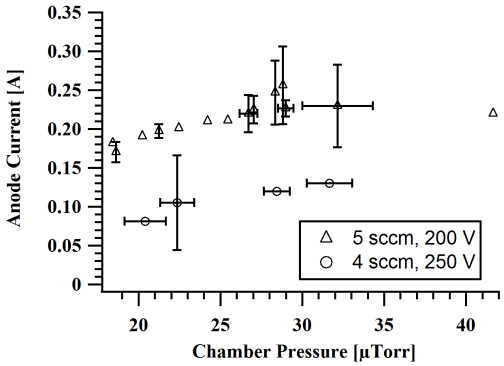

(b) Effect of chamber pressure scan on anode current, accomplished using an auxiliary xenon flow line

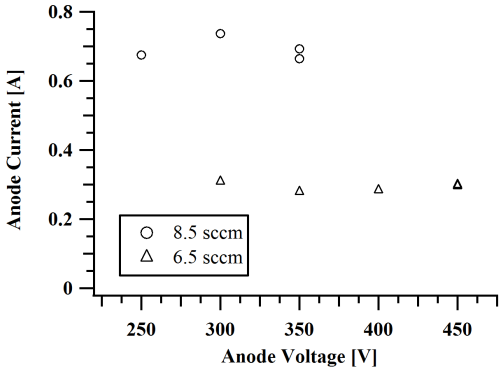

(c) Anode voltage scan.

Figure 2-10: Summary of varied thruster operating conditions. The variable on the horizontal axis of each plot was varied to measure the effect on anode current oscillations. The anode current for each of the varied conditions is displayed.

times are given as there were not enough samples of such data points to provide a measure of the uncertainty. The primary frequency is defined as the maximum discernible peak in the discrete Fourier transform (DFT) of a current trace. The peak-to-peak values for a given trace were calculated simply as the difference between the maximum and minimum recorded probe values for a given current trace in time. The doubling (rise) and halving (fall) times were calculated by fitting a growing or decaying exponential around a single peak in appropriate traces, using the peak itself and half-maximum values on either side. In the event that the wave form did not arrive at the half-maximum value, the time constant for exponential growth or decay was calculate from points on either side of the peak, and related to the doubling and halving times through the factor $\ln 2$. The variation of rise and fall times for a given trace varied by 10% or less.

Simultaneous, internal, temporally resolved floating potential measurements were attempted at different axial and azimuthal positions using flush electrostatic probes installed in the insulating cone of the thruster. However, the signals between the probes were discovered to be leaking between the different channels of our multi-channel DAQ device, a phenomenon sometimes referred to as *ghosting*. As a result, measurements were limited solely to discharge current oscillations. The presence of the internal probes is not thought to have significantly perturbed thruster operation because the currents drawn by the probes were negligible due to the large resistors between the probes and chamber ground. A more complete description of the instrumented thruster is given in another report [44].

2.2.1 Oscillation Frequency

The effects of the anode flow, chamber pressure and anode voltage scan are present in Figure 2-11. Interpreting the plots of frequency requires an emphasis on how the frequency for each point was determined. The selected frequency corresponds to the peak magnitude in the DFT of the anode current trace. This method is vulnerable to misleading data when the zero frequency of the DFT dominates, or if a lower frequency develops in the signal with a slowly varying magnitude. It was decided to plot the frequency for which the Fourier transform had the maximum amplitude, but this does not mean that that frequency is the fundamental harmonic of the signal. This can make it appear that the dominant frequency of the signal has been halved

or doubled, when really the difference between the two situations is similar to the differences between the current traces in Figures 2-12(a) and 2-12(b). All the anode flow scan points with flow rates of 8.5 [$sccm\ Xe$] or greater have strongly peaked current traces like the one shown in Figure 2-12(a). Then in going from 8.5 down to 8 [$sccm\ Xe$], the wave form becomes transitionary, with a mixture of strong peaks and stunted peaks, such as in Figure 2-12(b). When the flow reaches 7 [$sccm\ Xe$], the current trace takes the form of Figure 2-12(b). The mixing of the peak shapes brings about the lower fundamental frequency due to the fact that every other peak is stunted. With this lower harmonic, it is possible for the new maximum of the DFT to suddenly decrease by factor of two, which is the case of the outlier point in the anode flow scan at 6.5 [$sccm\ Xe$]. A peak in the DFT remains near 3 [kHz], but the new fundamental frequency overcomes it by a slim margin. It is necessary to realize that the dominant frequency component can be halved without a dramatic or particularly noticeable change in the current trace. Four data points in the chamber pressure scan are subject to this criterion in Figure 2-11(b). Doubling the frequencies for the two lowest values in each series changes the interpretation of the data to indicate that the oscillation frequency is not strongly dependent on chamber pressure over the range of pressures tested here. Cases for which no clear frequency was dominant or could be discerned were not assigned a frequency value. This most often occurs when the thruster is in the LC-mode. Perhaps a more robust indicator of frequency would be to record the first peak in frequency space greater than zero, taking notice each time the fundamental frequency drops by nearly a factor of two. Within each mode, the frequencies in Figure 2-11(c) do not change significantly, similar to the expected behavior of large amplitude oscillations in Hall thrusters[133]. Of the anode current oscillation regimes in the SPT-100, categorized by Gascon et al. oscillations in HC-mode most resemble Region 5[48].

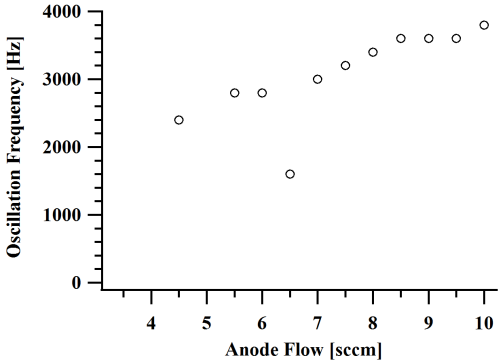

(a) The oscillation frequency increases with anode flow rate. The frequency is selected corresponds to the peak magnitude of the DFT of the current current trace.

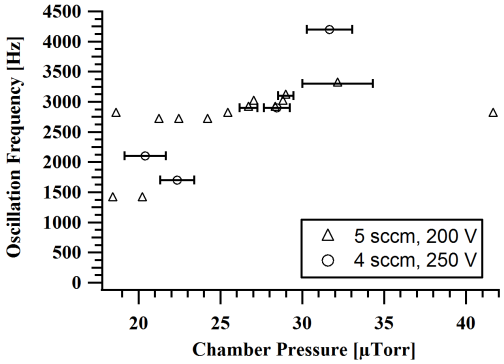

(b) Taking into account that the two lowest points in each series correspond to halved fundamental frequencies, chamber pressure has no significant effect on the oscillation frequency in the range of pressures tested here.

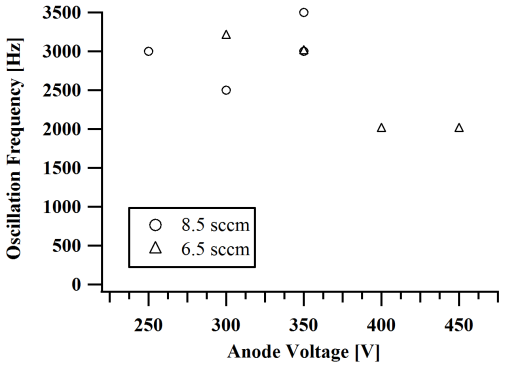

(c) The mode in which the thruster operates can be determined by the peak-to-peak magnitude of the discharge oscillation, referring to Fig. 2-13(c).

Figure 2-11: Dependence of discharge oscillation frequency on varied thruster operating conditions.

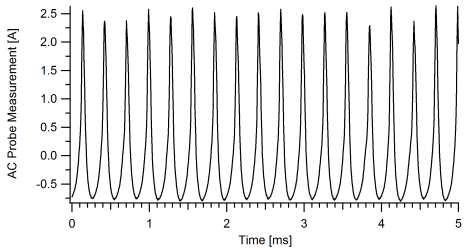

(a) Example of a common high-current AC waveform, marked by regular large amplitude current peaks. This trace was taken at the 8.5 [$sccm\ Xe$] data point for the anode flow scan.

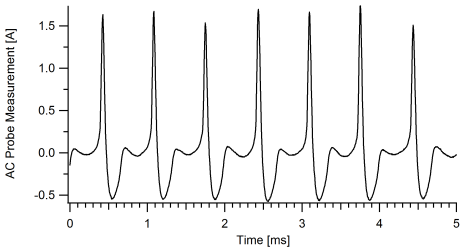

(b) Example of another commonly encountered high-current AC waveform. Here, every other peak is suppressed. This trace was taken at the 7.0 [$sccm\ Xe$] data point for the anode flow scan.

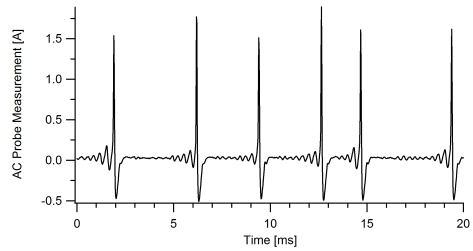

(c) This is a HC-mode trace at lower flow, closer to a transition to the LC-mode. When the flow is lowered further, the oscillations become smaller to the point that the large peak no longer occurs, and the trace takes on the appearance seen in Figure 2-12(d). This data point was taken at the 5.5 [$sccm\ Xe$] data point for the anode flow scan.

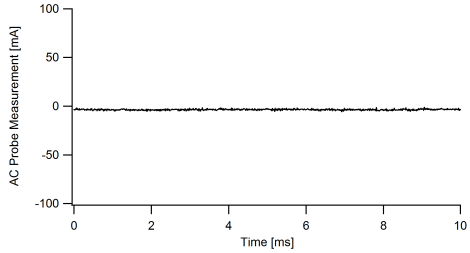

(d) Example a LC-mode AC current trace. Peak-to-peak amplitudes are typically on the order of several milliamperes. This trace was taken at the 3.5 [$sccm\ Xe$] data point for the anode flow scan.

Figure 2-12: Different AC waveforms encountered during anode flow scan.

2.2.2 Strength of Oscillations

The data in Figure 2-13 reveal that the peak-to-peak oscillation amplitude varies weakly with thruster operating condition, but differ dramatically between modes. In Figure 2-13(a), the amplitude increases sharply when going from the LC to the HC-mode, and again with voltage scan in Figure 2-13(c). The amplitude continues to rise slowly as pressure increases. Pressure scan data in Figure 2-13(b) show no discernible frequency trend within modes. The two low amplitude points are the result of the thruster operating in a mode that may only occur when chamber effects dominate.

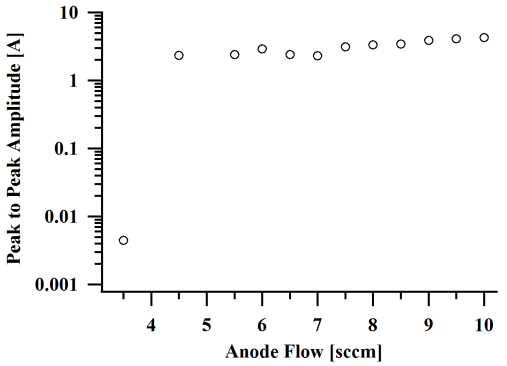

(a) Amplitude increases by nearly two orders of magnitude after the transition from the LC to HC-mode. Once in the HC-mode, the amplitude slowly increases with increased flow.

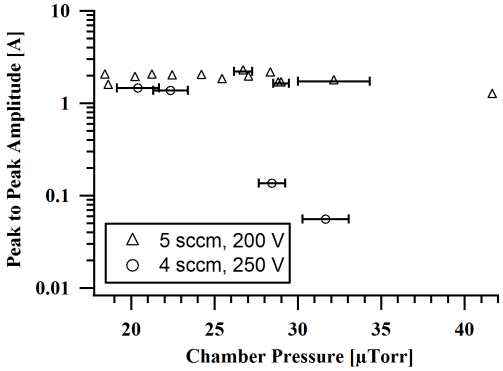

(b) The peak-to-peak amplitude appears constant within a given mode. The two operating points with amplitudes closer to 0.1 $[A]$ correspond to a red-mode that was only observed at elevated chamber pressures in this study.

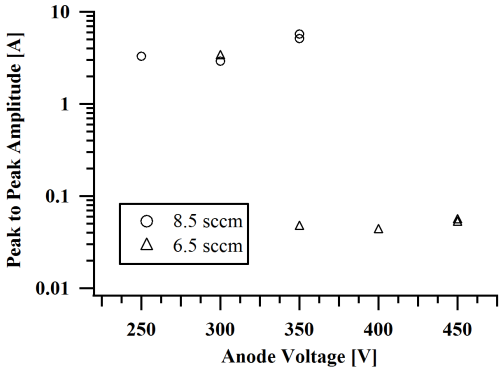

(c) No trend can be assigned to the oscillation amplitude within either the low or high-current modes. However, the drop in magnitude when changing from HC to LC-mode is evident.

Figure 2-13: Dependence of discharge oscillation peak-to-peak magnitude on varied thruster operating conditions.

2.2.3 Oscillation Feature Time Scales

Figures 2-14(a), 2-14(b) and 2-14(e) all show shorter rise (doubling) times than fall (halving) times. Increasing the anode flow decreases the rise time in a seemingly linear fashion, while the fall time remains effectively constant. The doubling and halving times have nearly identical values in measurements made during the pressure scan, shown in Figures 2-14(c) and 2-14(d). The HC-mode condition in the pressure scan had steady timescale values over all pressure ranges tested here. However, in the red-mode, the doubling and halving rates both increased during the pressure scan. No definitive trend between the rise and fall times and anode voltage can be made due to the large spread in data points.

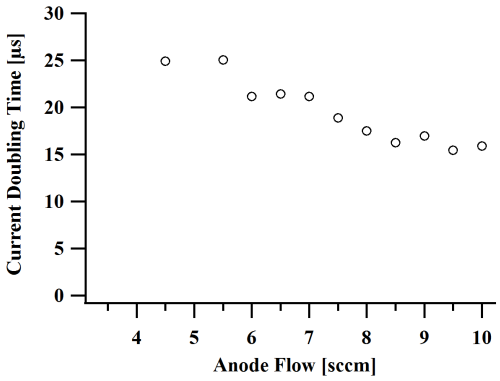

(a) Current doubling time plotted against anode flow.

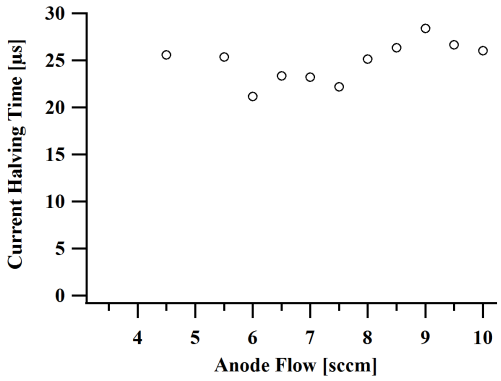

(b) Current halving time plotted against anode flow.

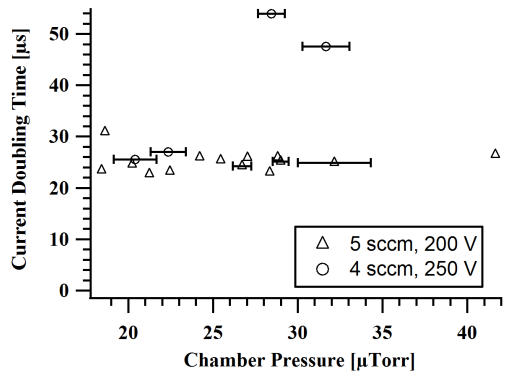

(c) Current doubling time plotted against chamber pressure.

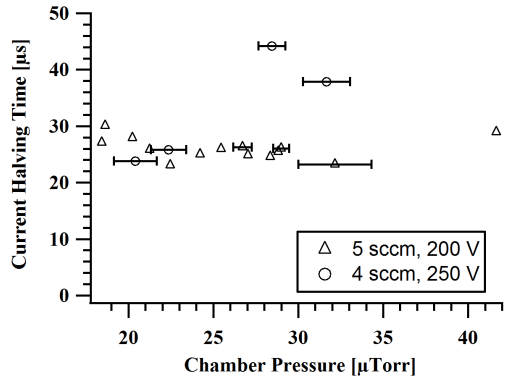

(d) Current halving time plotted against chamber pressure.

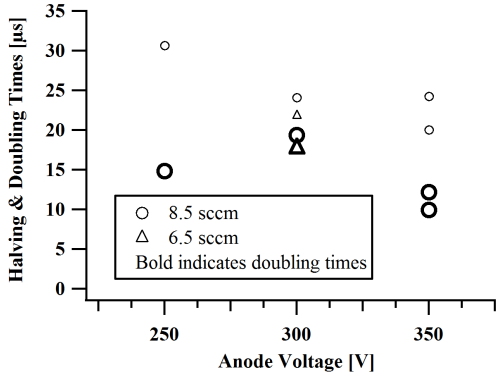

(e) Current doubling and halving times plotted against anode voltage

Figure 2-14: Dependence of discharge oscillation doubling and halving times on varied thruster operating conditions.

The DCF thruster was viewed at oblique angles and recorded during operation with a high speed digital video camera. Still frame images of these videos are presented in Ref. [30]. High and low-current modes are characterized by concentrations of luminosity corresponding to the number and location of four magnetic cusps. Through the course of a HC-mode oscillation, luminosity measured in the discharge chamber is dominated by a bright spot near the anode and C1 cusps. As the overall intensity of the image grows, weaker peaks materialize downstream in correlation with the C2 cusp. This sequence is independent of the different current trace shapes identified in Figure 2-12, with the shapes of luminosity oscillations matching the shapes of AC discharge current measurements. The correlation of luminosity oscillations to anode current oscillations suggests that areas of greatest luminosity may correlate well with ionization regions within the thruster. More direct and rigorous proof, via simulations or measurements, are necessary to firmly establish the regions of ionization within the chamber. Measurements and simulations are, however, indicating that the largest electric fields exist near or past the exit plane of the thruster.

The visual appearance of the LC-mode is distinguishable from the HC-mode within the thruster as well within the plume. The LC-mode has compact luminosity features near the anode and in C1, as well as an annular feature in C2. In addition to these features seen in the HC-mode, the LC-mode also has a ring shaped luminous area in the C3 magnetic cusp near the thruster exit. Transition from the LC to HC-mode is characterized by the onset of slowly growing low frequency disturbances in the discharge current, as well as intermittent current spikes. The transition from a stable HC-mode to LC-mode is observed to occur as the strong current peaks shown in Figure 2-12(a) give way to alternating stunted peaks, as in Figure 2-12(b). The transition to LC-mode is completed as the stunted peaks, seen in Figure 2-12(c), oscillate with progressively smaller magnitudes until the large current spikes disappear altogether. The result is a steady current trace, visible in Figure 2-12(d).

Thruster and cathode operating conditions were systematically varied to examine their effects on characteristics of the large amplitude, low frequency discharge oscillations. As anode flow is increased, the discharge current frequency and peak-to-peak amplitude grow steadily also. Doubling times for bursts of current decrease with increased flow, while halving times

remain level. Varying the background chamber pressure does not result in significant changes to oscillation frequency, current growth and decay rates, except when causing a mode transition from the standard HC-mode to a mode with a reddish plume and a lower amplitude oscillations. This mode should not be confused with the LC-mode. Repeating these measurements in a chamber capable of operating at lower pressures could determine if any changes between the HC and LC-modes can be initiated in repeatable manner. There are not enough data points for the voltage scans in this study to offer a clear picture of their effects on discharge oscillations, but they are most evident when changes in anode voltage result in thruster mode changes. No data are given for variations in cathode flow rate, which does not significantly alter either the average anode current value, or the oscillation characteristics. In some instances, increasing the keeper current can cause the thruster to switch from HC to LC-mode, but the currents required to do so are large enough to inhibit extensive testing of this sort, out of concern for maintaining the cathode in an operational state. In videos taken at low enough frequencies to allow the inclusion of the cathode orifice in the image, brightening of the cathode plume also coincides with the strong peaks in luminosity in the thruster chamber. Also, operating conditions were selected in this study in an attempt to avoid operating the thruster much above a 200 $[W]$ anode power threshold. Magnet temperature recordings indicate that operating lower than 200 $[W]$ is a good indicator that the magnet temperature is likely less than the manufacturer quoted maximum operating temperature of 300 $[^{\circ}C]$.

2.3 DCFT Long Duration Test & Erosion Measurements

2.3.1 Long Test Setup

Chamber 1 was again used for the long duration testing described in this section, which has also been presented in Refs. 128 and 45. All thruster and cathode electrical power was provided by commercially available devices. The thruster anode and cathode keeper were each driven by a Sorensen DCS600-1.7E, while the cathode heater power was delivered by a Sorensen DLM 40-15. The anode was biased in reference to cathode potential, which was allowed to float with respect to chamber ground. Outputs from power supplies were calibrated using an Agilent 34410A Digital Multimeter. Purified xenon (99.995%) flow rates to the anode and cathode were set

manually and delivered using a UNIT Instruments power supply and mass flow controllers. The cathode position, which has been shown to significantly affect thruster operation, was located identically to past studies[27]: 73 [mm] from the thruster axis of symmetry and 31 [mm] from the exit plane. During thruster operation, power supply voltages and currents, along with the cathode floating potential, were digitally recorded twice per second. Chamber pressure was recorded by hand until hour 60 of testing, at which point modifications to data acquisition software enabled automated recordings. After this point, pressure was recorded twice per second as well. A Busek BHC-1500 hollow cathode was used to initiate the discharge and neutralize the plume for all tests described herein.

Throughout all tests, the cathode heater current was fixed at 3.0 [A] to decrease the likelihood of the keeper voltage rising substantially and causing a shutdown during unsupervised operation. The keeper current was held constant at 0.5 [A]. Maintaining operation in LC-mode was not possible for more than ten minutes. As a result, the thruster operated in HC-mode at a single operating condition for nearly the entire test. An anode voltage of 300 [V] and mass flow rate of 8.5 [$sccm\ Xe$] were selected to avoid heating the permanent magnets above 300 [$^{\circ}C$] — the manufacturer listed maximum operating temperature[27]. The curie temperature for the magnets is higher, at 825 [$^{\circ}C$]. Unpublished tests with a thermocouple placed between the middle magnet and chamber insulation suggest an upper limit of $185 - 200$ [W] to the anode power to avoid overheating the magnets beyond 300 [$^{\circ}C$]. Additionally, laser-induced fluorescence velocimetry and thrust data are available at this operating condition from other campaigns[27,43].

The main purpose of this long duration test was to quantify erosion rates in the DCF thruster and evaluate the potential of this concept for extending available operational lifetimes of low-power plasma thrusters. Before and after the long duration test in Chamber 1, the profile of the insulating channel was measured using the Micromeasure system from Micro Photonics, Inc. This system consists of a STIL CHR150 chromatic confocal spectrometer, a CHR Contactless Optical Sensor, three linear stages and a power supply. The hardware and data acquisition are controlled by a desktop computer using a program called Surface Map. Specifications of the optical sensor and stepper motors are provided in Table 2.11, and a schematic of the profilometer is shown in Figure 2-15. The optical sensor was positioned at the angle allowing for the greatest

Table 2.11: Specifications of the optical sensor, digital level, and stepper motors.

Optical Sensor		Flexbar PRO 360 Digital Level		Linear Stages	
Working Distance	66.9 [mm]	Accuracy, ±10° from level/vertical	±0.1°	Z-axis precision	0.01 [μm]
Measuring Range	10.0 [mm]	Maximum error, other angles	±0.2°	X & Y axes precisions	0.1 [μm]
Lateral Resolution	25.5 [μm]	Range	360°		
Static Noise	556 [nm]				

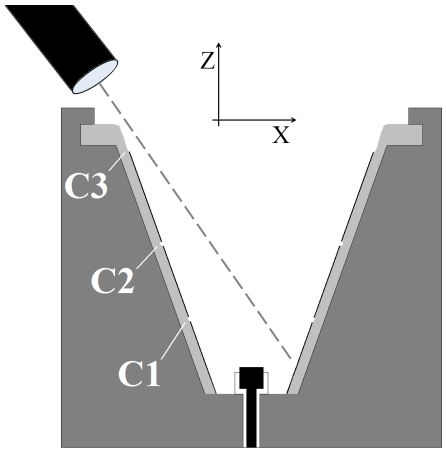

Figure 2-15: In this schematic, the optical sensor is illustrated measuring the surface position at a point upstream of C1.

signal strength a short distance upstream of C1. This was necessary because the signal intensity was not sufficient to allow measurements over a wide range of incidence angles. Once a suitable orientation was found, the position of the sensor was secured and measured with the digital level described in Table 2.11. An additional requirement of the optical sensor position was line-of-sight accessibility to the reference region (described below), without having to change the sensor's orientation. This condition was imposed to avoid additional uncertainties related to the relative orientations of the optical sensor between scans.

Using the profilometer, the relative distances between two points in a scan can be determined to an accuracy better than 1 [μm]. However, quantifying erosion requires the determination of distances between points in different scans. This system is not configured to provide measurements while the thruster is under vacuum, so a method relating the coordinate systems of different scans was implemented.

A digital level was used along the X and Y axes to ensure that the exit plane of the thruster was parallel with the XY measurement plane. The zero-coordinate of the Y-axis

for each scan was estimated as the average extrema location of parabolic fits to measured profiles of the distance from the optical sensor to the thruster surface. These centering data were obtained from sweeps of the Y-coordinate at constant values of X. Then, the detailed topology of a small feature carved out of a protected part of the thruster was measured with the profilometer. Scans of the reference feature were made before and after operating the thruster. The coordinate systems of the erosion scans were then related to one another through the coordinate transformations necessary to match the two reference scans. Figure 2-26 shows the "before" scan matched to the "after" scan of the reference feature, to within ± 10 $[\mu m]$. The matching scheme could be improved if the coordinate transformation were obtained through the constrained minimization of an error function. However, in this case, plots similar to Figure 2-26 were inspected visually to select the coordinate transformation that most closely matched the contours. The quoted uncertainty from the technique is the smallest change in each coordinate that caused a noticeable difference in the similarity of the two scans.

After testing was complete and post-test profiles were measured, a sample of the surface was taken for analysis using Energy Dispersive X-ray Spectroscopy (EDS). The sample consists of wall material removed from the entire interior of the surface, so no definitive statements about the distribution of materials on the surface can be made. Also, some of the material removed was clearly magnetized, gathering in the cusps of the thruster.

2.3.2 Long Duration Test: Thruster Operation

The thruster and cathode ran at a single operating condition for almost the entirety of the test, given in Table 2.12. Throughout testing, the chamber pressure was $2 - 3$ $[\mu Torr]$. The anode current is plotted in Figure 2-16(a), and remained between $0.53 - 0.56$ $[A]$ from 5.5 hours until the end of testing. Over the same interval, the floating cathode potential, relative to chamber ground, began at -16 $[V]$ and slowly increased to -12 $[V]$. The cathode keeper voltage remained between $11 - 13$ $[V]$ with the aid of a constant 3.0 $[A]$ heater current. The repeated drops in anode current correspond with intentional thruster shutdowns to allow for thrust stand calibration and subsequent measurements. These measurements are not described in detail because the thrust stand was not configured properly before testing. After analyzing the data, the standard deviations of measured anode efficiencies were larger than $\pm 10\%$. The

Table 2.12: Thruster and cathode operating condition during the long duration test.

Anode Flow	8.5 [sccm Xe]
Anode Voltage	300 [V]
Cathode Flow	1.5 [sccm Xe]
Keeper Current	0.5 [A]
Heater Current	3.0 [A]

average values of thrust, specific impulse, and anode efficiency were significantly and consistently lower than other DCFT performance measurements[28,131]. Standard deviations in thrust and specific impulse approached or exceeded 50% of average values in most cases. The average thrust values also decreased as the thruster operated, though any trends in the average were within uncertainty margins. The poor quality of these data does not warrant further discussion.

The anode current took on a wider range of values over the first 5.5 hours of testing. The thruster ignited near the 2.5 $[h]$ mark after cathode conditioning. Operation in LC-mode was attempted during the first ten minutes of testing. The anode current crept up with time in conjunction with the sporadic onset of current pulses — a conspicuous precursor to the transition from LC to HC-mode. After this was observed, the decision was made to operate in HC-mode, and the operating parameters were set to the values listed in Table 2.12. From here, the current rose above 600 $[mA]$, very similar to the anode current measured during initial performance measurements performed at Busek in 2007[27]. Then, as seen in Figure 2-16(b), during a period slightly less than half an hour, the current dropped below 500 $[mA]$, and then rose gently and leveled off at 510 $[mA]$. All operating conditions were held constant over this time, and the thruster was left unattended between hours 3.5 and 60. During this time, the anode current rose from 510 $[mA]$ to 550 $[mA]$ just after hour 5, where it remained above 538 $[mA]$ until hour 55.38, when it dropped suddenly again to 532 $[mA]$. After this, the anode current varied smoothly until the end of the test. However, the onset of very low frequency anode current oscillations is evident starting between hours 90 and 100. These had a period of several minutes and their amplitude increased with time until testing concluded. As shown in Figure 2-16(c), the amplitude near hour 200 was $10 - 20$ $[mA]$. Changes in the cathode potential, also shown in Figure 2-16(c), coincide with fluctuations in the anode current, though

they are negatively correlated.

As mentioned earlier, the anode current settled to about 550 $[mA]$ for the majority of the test, lower than the maximum $610 - 620$ $[mA]$ encountered when HC-mode operation was initiated here, or in the first measurements in Ref. 27. However, in two previous tests, the anode current was 530 $[mA]$ at the operating condition used in this study. One test [29] was done in the same chamber as this study, using the same thruster as in Ref. 27, while the other [43] used the same thruster as the test described here, but in a different chamber. Although the thrusters used in this test and in Ref. 43 are dimensionally identical to the original thruster used in Ref. 27, the anode was made of 316 stainless steel for the first performance measurements described in Ref. 27. All subsequent tests have used a graphite anode. However, this alone does not explain the behavior of the anode current during the first hour of HC-mode operation, an interval in which the anode power varied by as much as 35 $[W]$. Transients in the thruster, injected xenon, or cathode temperatures may all have contributed, as well as surface conditioning of the insulator. Another important difference may be that the cathode was allowed to float with respect to chamber ground in these tests[b], while the cathode was grounded in Ref. 27. Chamber pressure is not suspected in this case, as the xenon corrected background pressure was between $2 - 5$ $[\mu Torr]$ for the tests described in Refs. 27, 29, and 43.

Finally, accumulated deposition between cusps might also have had an effect. The visual appearance of the thruster is documented in Figure 2-17. When performance measurements were first made [27], the insulator was closer in appearance to Fig. 2-17(a) than Fig. 2-17(b). Before any testing, the BN cone is clean and white, as seen in Figure 2-17(a). Contrastingly, Figure 2-17(b) shows the insulator after $65 - 70$ $[h]$ of operation, but prior to the experiments described in this report. The walls are dark and black, except for white rings visible in each cusp. Finally, after an additional 204 $[h]$ of firing, the accumulation of more deposition on the walls is visible in Figure 2-17(c), where the aluminum cap holding the cone in place has been removed in preparation for profilometry. Parts of the wall appear lighter than surrounding areas, and with shades of grey or white. The cause of this coloration is unknown, but an elemental analysis of the deposition is discussed in Section 2.3.3.

[b] As a reminder, the floating cathode potential varied between -16 $[V]$ and -12 $[V]$ with respect to chamber ground.

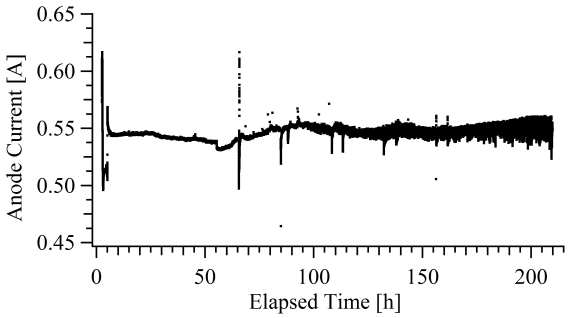

(a) Anode current during the entire test.

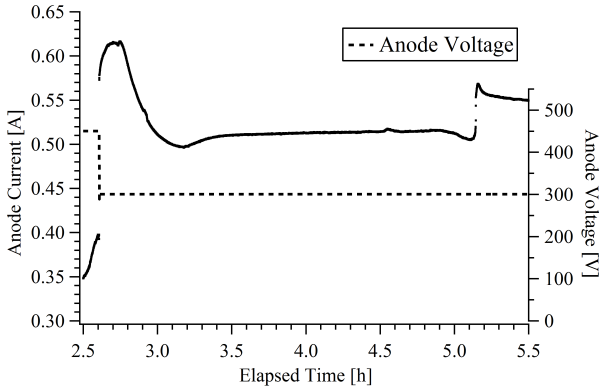

(b) Anode current and voltage during the first several hours of operation. The current (voltage) scale is on the left (right) axis.

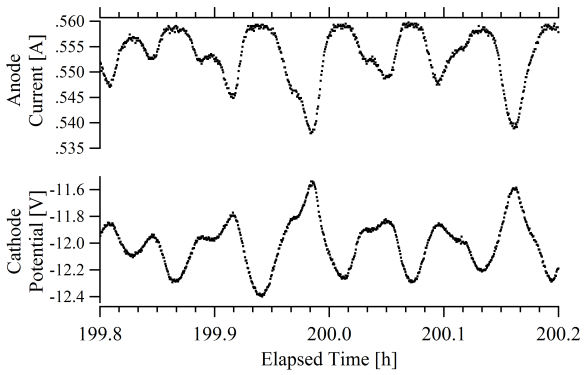

(c) Anode current and cathode potential near hour 200. Axial increments in time are 1 $[min]$.

Figure 2-16: Anode current during testing.

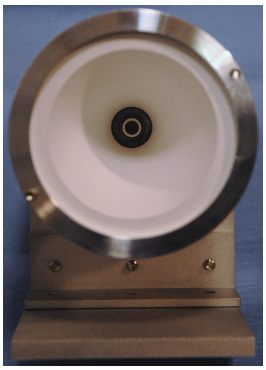

(a) The thruster before operation.

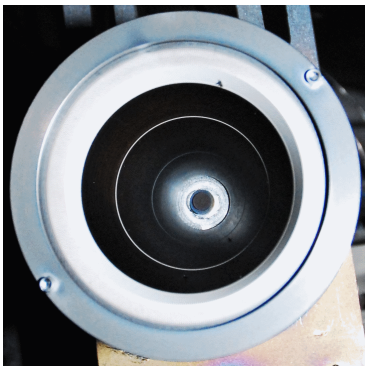

(b) The thruster after $65 - 70$ $[h]$ of operation, but before the start of the long duration test.

(c) The thruster after 204 $[h]$ of operation at the conditions given in Table 2.12, for a total run time of $269 - 274$ $[h]$.

Figure 2-17: Photographs of the HP BN insulator at different stages of its operational history. The thruster is oriented with the cathode above the thruster in Figs. 2-17(b) and 2-17(c).

2.3.3 Erosion Quantification & Deposition Analysis

The profile of the insulating HP BN cone was measured before and after the long duration test. An overview of the measurements is given in Figure 2-18, showing the locations of maximum erosion relative to the permanent magnets and insulator. Each scan sampled the surface height along a 0.5 $[mm]$ wide strip along the Y-axis, with a 25 $[\mu m]$ step size. The data points plotted in Figures 2-18 and 2-19 represent the average coordinates of data points along the width of the strip. The offset of points from the center of the strip was included in calculating the average radial coordinate. The systematic error introduced to this average by the curvature of the measurement surface is less than 1 $[\mu m]$, much less than the errors due to coordinate system matching. Erosion is limited to the immediate vicinity of the cusp locations in C1 and C2. Closer to the exit plane, erosion begins near C3 and continues toward the exit plane. Detailed plots of the measured profiles in the cusp regions are given in Figure 2-19. The reason for less resolution on the downstream side of the erosion feature visible in Figure 2-19(c) is due to the slope of the surface relative to the incident beam of light. The dependence of measurement resolution on the relative orientations of the surface and light beam is illustrated in Figure 2-20(a). If the surface slope goes past parallel with the beam, a portion of the surface will be inaccessible for measurement unless the beam angle is changed.

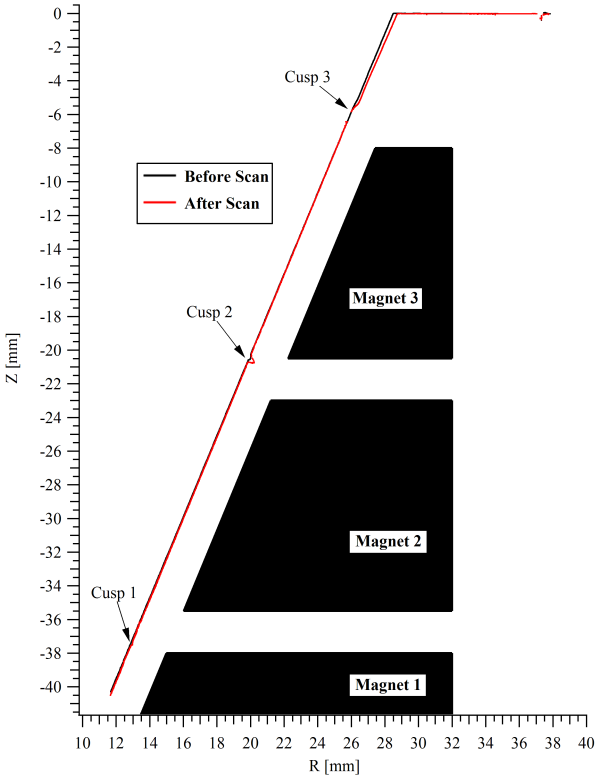

Figure 2-18: Profilometer scans, before and after 204 $[h]$ of operation at the operating condition given in Table 2.12. Magnets are numbered in their order downstream of the anode. The exit plane of the thruster corresponds to $z = 0$ $[cm]$. Measurements were limited to just upstream of C1, due to the angle of the incident beam of light and to more specular reflection closer to the anode.

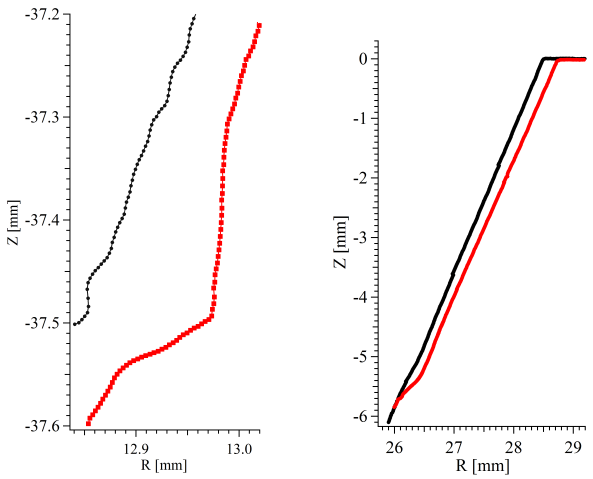

(a) Cusp 1 detail.

(b) Cusp 3 detail.

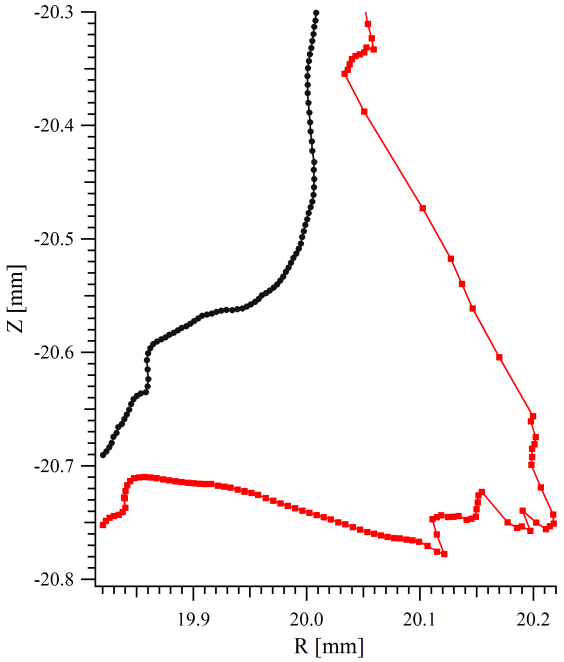

(c) Cusp 2 detail.

Figure 2-19: The "before" ("after") profile consists of circles (squares) connected by the darker (lighter) line. Note the different scales in each figure.

From these two profiles, average erosion rates are calculated based on the distances between them, perpendicular to the original cone surface — which has a 22.5° half-angle with respect to the centerline. Referring to Figure 2-20(b), this is done by fitting a line perpendicular to the nominal wall angle through a point in the "after" scan, referred to as point "A". Next, the point of intersection between this line and a segment joining two points in the "before" scan is found, referred to as point "B". The distance between A and B (d_{AB}) is taken as the eroded distance corresponding to point B. The results are shown in Figures 2-21 and 2-22, where the erosion rate is plotted in terms of the radial coordinates of interpolated points from the "before" scan. Comparing Figures 2-22(b) and 2-19(c), this explains why the plotted erosion rate is double valued in areas where the "before" profile is sufficiently curved from preexisting erosion. It also explains why the shapes of the cusps do not match those visible in Figure 2-19. The slope of erosion data between cusps in Figure 2-21 is due to a 0.25° misalignment, very close to the precision of the level used to position the thruster for measurements. The alignment error is smallest where the alignment was made, and increases closer to C1. The effect of the error is to *overestimate* the maximum erosion rate of the thruster by about 0.15 [$\mu m/h$]. Aside from this systematic error, the statistical uncertainty of eroded depths between cusps is estimated to be ± 20 [μm], so the uncertainty in erosion rate is ± 0.1 [$\mu m/h$]. In the second cusp, uncertainties are larger because of the orientation of the "after" surface. The maximum erosion rate resulting from a single measurement in this region is 2.52 [$\mu m/h$], providing a lifetime estimate for the DCFT prototype between $920 - 1220$ [h]. A detailed description of erosion rate uncertainties in C2 is provided in Ref. 45.

The peak erosion rate is in C2, rather than at the exit plane as observed in Hall thrusters, with a maximum average value of 1.9 [$\mu m/h$]. The first and second ring-cusps coincide with regions of increased erosion, while C3 is roughly where erosion at the exit begins in the downstream direction along the thruster wall. Erosion rates increase toward the exit, but never exceed the erosion rate measured in the second cusp. The initial thickness of the insulating cone is 2.5 [mm][c] in the radial direction, allowing for an effective insulation lifetime of 1220 [h]. This lifetime is calculated assuming a constant erosion rate perpendicular to the nominal insulator surface in the second cusp. Many Hall thrusters are predicted to continue operating

[c]The shortest distance through the insulator is less than this, ≈ 2.31 [mm], because the nominal insulator surface makes a 22.5° angle with the axis of symmetry.

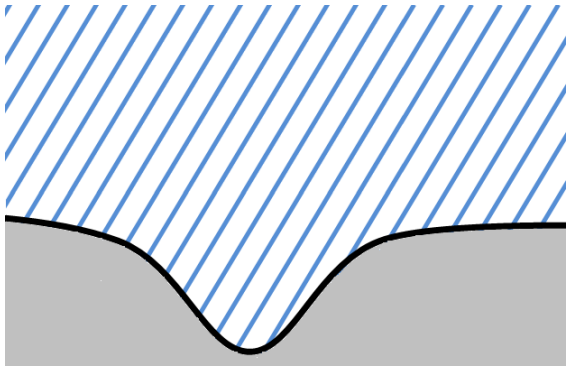

(a) The dependence of measurement resolution on surface slope relative to the incidence angle of the beam. As the beam and surface become more parallel, the local measurement resolution deteriorates.

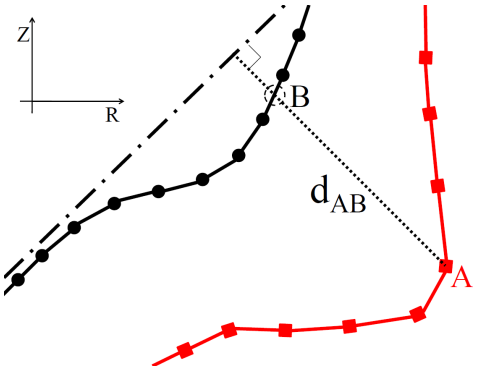

(b) The eroded distance is taken as the distance between points A and B. Point A is measured in the "after" scan, while point B is interpolated from two points in the "before" scan.

Figure 2-20: These illustrations are meant to demonstrate how measurement resolution is affected by surface slope, and how the eroded distance is calculated.

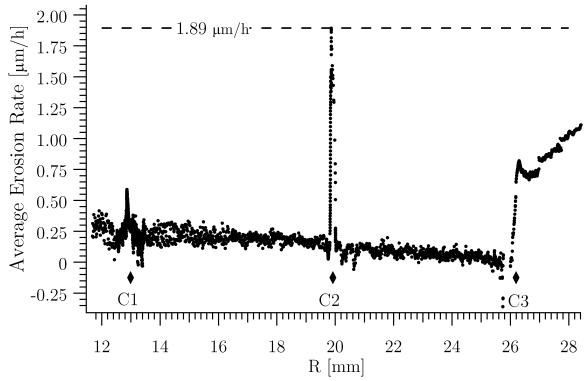

Figure 2-21: Average erosion rate during the 204 hours of operation. The locations of the cusps along the wall are indicated by the black diamonds.

after the insulating ceramic is unable to completely shield the magnetic circuit — see Table 4 in Ref. 45. This is appropriately referred to as a "soft failure" [8] in situations where a satellite and thruster can continue to function within specifications after the insulator has been compromised. An extensive comparison of beginning-of-life (BOL) erosion rates and thruster lifetimes determined from measurements in Hall thrusters is again provided in Table 4 of Ref. 45. A distinction is made between the predicted lifetime of a thruster and the time at which soft failure is reported or predicted to occur. Also, in the Appendix of Ref. 45, a description of insulator materials used in many different thrusters is provided.

Deposition on the cone surface during operation was also significant. Figure 2-23 shows photographs of different areas of the cone surface after the test was completed. These photos should be compared to Figure 2-17(b) to observe the difference in the appearance before and after testing. The coloring of the deposition is lightest near the porous 316 stainless steel diffuser behind the anode up until the first ring-cusp. Downstream of C1, the deposition is mostly darker, with the exceptions of some areas that are lighter grey or white in between cusps. The direct source of these previously unseen patterns is not known, but they exhibit some symmetry with respect to the cathode location. A sample of the deposition was scrapped off with a teflon coated spatula and analyzed using EDS. Localized samples were not taken, preventing the examination of compositional variation between different areas of the chamber. A sample of deposited materials on an external surface was also analyzed.

The composition of the internal deposition contained copper, constituents of 316 stainless steel and additives commonly used to lower the work function of tungsten in hollow cathode emitters. The exterior sample was composed entirely of copper, tin and other components common to copper alloys. The portion of the thrust stand visible in Figure 2-24 is made of copper to allow for more efficient active cooling using an exterior cold water supply, and is likely the source of copper identified in both deposition samples. The large amount of iron on the interior is likely sputtered from the annular porous stainless steel disc located upstream of the anode. The diffuser disc is visible in Figures 2-17 and 2-23. The source of lighter deposition on the upper and side surfaces is not known, but the symmetry of the pattern suggests that it originates from the cathode.

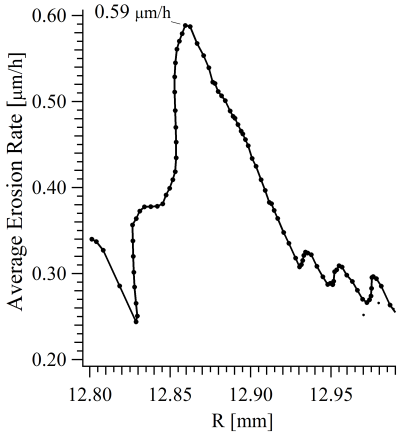

(a) Cusp 1 average erosion rates.

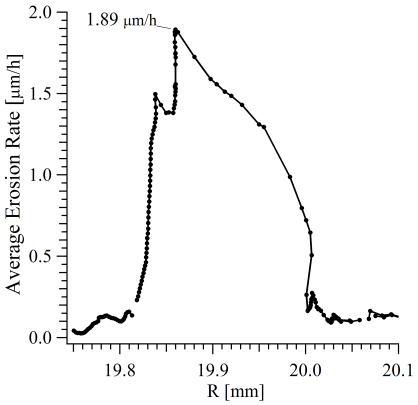

(b) Cusp 2 average erosion rates.

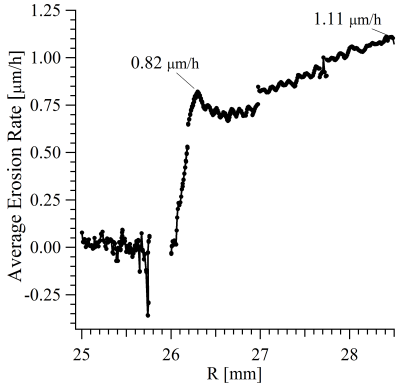

(c) Cusp 3 average erosion rates.

Figure 2-22: Average erosion rate measurements in each ring cusp. The coordinate on the horizontal axis is the radial coordinate of an interpolated point in the "before" scan. Note the different scales in each of these figures.

(a) Photo of the left half cone surface, facing the anode.

(b) Photo of the right half cone surface, facing the anode.

Figure 2-23: Photos of the DCFT insulator after testing was completed. Eroded regions are evident as white rings in each cusp. The uneven, but somewhat symmetrical, pattern of deposition is also visible. The location of the cathode is above the thruster for the orientations shown here.

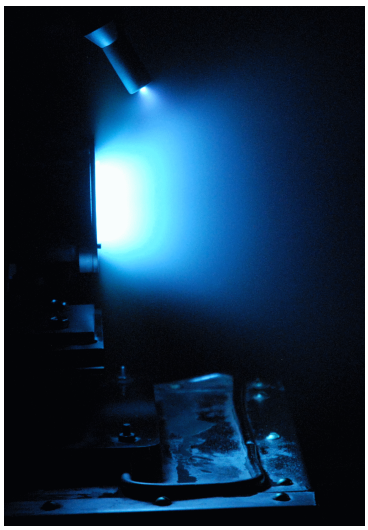

Figure 2-24: The plume, overexposed in this photograph, illuminates the thrust stand.

2.3.4 Near-Wall Magnetic Field & Coordinate System Matching

Figure 2-25 contains detailed information regarding the magnetic topology along the insulator in the DCF thruster. The magnetic stream function φ is proportional to the product of the radial coordinate and the azimuthal component of the magnetic vector potential, so lines of constant φ represent magnetic field lines. The dividing surfaces between magnetic cells are called separatrices, and are defined as those surfaces whose stream function value is zero. Thus, the rigorous locations of magnetic cusps along the wall are where the stream function value is zero. The values plotted here were calculated using the MAXWELL SV software package. Field strength measurements along the thruster axis validate the simulated results [27].

Figure 2-26 shows data used to relate the "before" and "after" profilometry scans to one another. In each case, with the optical pen orientation fixed, the topology of a reference feature was measured. Identical coordinate transformations used to create the plot in Figure 2-26 were then used to relate the coordinate systems of the "before" and "after" profilometry scans. In Figure 2-26, the thicker set of lines represents data from the "after" reference scan, while the thin lines show measured heights from the "before" reference scan. The square outlines the domain of data for the "before" reference scan. The uncertainty in alignment resulting from this approach is estimated to be 10 $[\mu m]$.

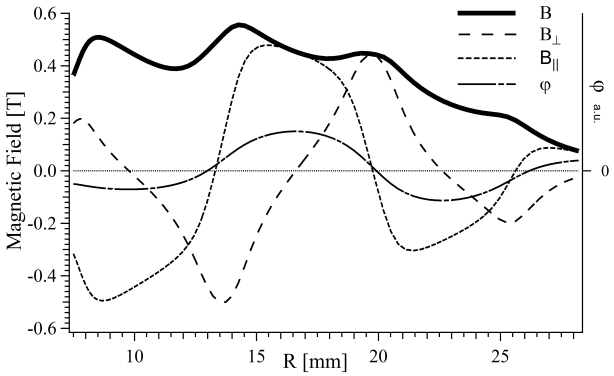

Figure 2-25: The parallel and perpendicular components of the magnetic field along the thruster wall are plotted, as well as the magnitude of the field and value of the magnetic stream function (φ).

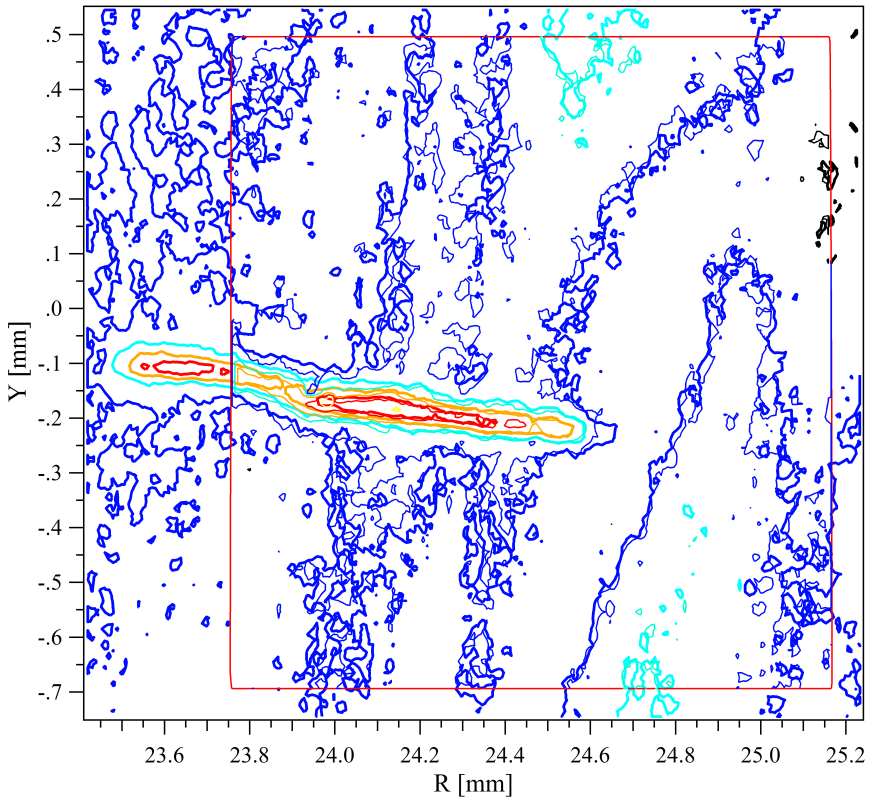

Figure 2-26: Contour plots of height for two scans of area surrounding the reference feature. Identically shaded lines represent equivalent surface heights, and appear in divisions of 10 [μm]. Note that the scales of the Y and R axes are not equal.

2.3.5 Conclusions From Long Test

The DCF thruster was subjected to a long duration test to measure average erosion rates of the insulating ceramic. The thruster operated consistently in HC-mode throughout, with the average anode current remaining within a range of 30 $[mA]$ for the last 98.5% of the 204 $[h]$ test. Very low frequency oscillations in the anode current were observed, increasing in amplitude as the test continued. The deposition of conductive, ferromagnetic material on the BN insulator was also identified. This motivates avoiding further use of stainless steel in the interior of laboratory thrusters.

Profilometry in each cusp was completed, as well as in between cusps. Erosion features in the cusps and downstream of C3 firmly establish the locations of wall losses in cusped-field thrusters. The peak average erosion rate was 1.9 $[\mu m/h]$, measured in the second cusp. This is lower than any limiting erosion rate reported in the reviewed Hall thruster literature, even amongst thrusters with anode powers less than 200 $[W]$. A lifetime estimate based on the average erosion rate and measurement uncertainty is $920 - 1220$ $[h]$. The design examined here is a laboratory prototype, originally intended for exploratory testing of the cusped-field thruster concept.

The observation of higher erosion in C2 than in C3 was unexpected and may offer insight into important differences between the HC and LC-modes. The anode current in HC-mode is characterized by ~ 3 $[kHz]$ oscillations with magnitudes several times larger than the average value. Therefore, the electron temperature estimates shown in Table 5 of Ref. 128 must be interpreted as time averages over the oscillation period. High-speed digital videos show that the HC-mode instability is accompanied by an increase in luminosity from the plasma upstream of the exit plane[30]. One possible interpretation of these observations is that an ionization avalanche is initiated throughout the thruster, consuming a large fraction of the available neutrals. Ions created during this burst of ionization are more easily lost to the second cusp than ions created near the exit because of the potential valley present in the bulk within ring-cusps[34,44,137]. This particular cusp drains from the largest volume in the interior of the thruster chamber, as shown in Figure 1-1(b). Rather than moderate, steady amounts of ionization that must be present in the LC-mode, the emerging view of the HC-mode is one of near complete ionization of the neutral population as some instability is triggered, leading to spikes in the anode current. The

time between bursts is set by the neutral replenishment time scale. As such, second cusp erosion in the LC-mode, where ionization is steady and presumably localized near the third cusp, is expected to be lower than in the HC-mode mode. Confirming this hypothesis, and the extent to which the erosion is lessened, will require further experimental and numerical investigation. Quantifying erosion in the low-current mode will be an important goal of continued studies involving the DCFT. Predictions from PTpic simulations for erosion in LC-mode are presented in Chapter 5.

Part I

PTpic: Plasma Thruster Particle-in-Cell Simulator

An outline of tasks executed by PTpic in order to simulate a given plasma thruster is provided in Fig. 2-27. PTpic shares many of the same steps that were present in earlier version written by Szabo and modified by others. However, several important changes in the way these steps are carried out have been made. The biggest change in the higher level procedure is the ordering in which species are moved. Now, electrons are moved first, followed by ions, then neutrals. Prior to this work, they were moved in the opposite order.

Before the start of the main loop, the Message Passing Interface (MPI) and Basic Linear Algebra Communication Subprograms (BLACS)[d] are initialized – along with all the necessary variables for PTpic to operate. Next, PTpic loads input parameters from files provided by the user. Some of these inputs include the simulation mesh, applied magnetic field, boundary locations, the thruster operating condition, mass and permittivity factors, the desired number of iterations, and many more. Full descriptions of what inputs can be provided, or are required, are provided in the respective input files. Once the mesh and boundary information are given, important constant quantities are calculated – including all the matrix elements for the linear system of equations for the node potentials. Finally, particles are initialized within the domain. Particle positions and velocities can be read from files or generated pseudo-randomly, depending on whether a simulation is starting from scratch, starting a plasma simulation from a neutrals-only simulation, or restarting a full plasma simulation.

When the previously described tasks are complete, the main simulation loop begins. The timestep is constant and specified by the user in an input file, though PTpic still calculates and saves the minimum timestep needed to properly resolve cyclotron or plasma frequencies – modified slightly by Eq. 3.11 – allowing the user to assess how appropriate the fixed step size is for a given simulation. In the event that the timestep Δt over-resolves the γ-modified plasma-frequency-timestep Δt_p, PTpic delays updating the potential solution for a number of iterations determined by the value of $\lfloor \frac{\Delta t_p}{\Delta t} - 1 \rfloor$, where these brackets contain arguments of the "floor" function — indicating that the value of the interior expression is replaced by the greatest integer less than or equal to its value. The remaining procedures listed as part of the main loop in Fig. 2-27 are discussed in the following Chapters. Chapter 3 describes how particle trajectories are integrated forward in time, and how particles interact with each other and boundaries.

[d]The BLACS support necessary operations within the ScaLAPACK libraries, used by PTpic to solve for potential distributions. These are only needed if the direct potential solver is used – the current default setting.

Chapter 4 discusses how potential solutions are obtained in detail. Algorithms carried over from the works of Szabo, et al. will not be described in detail, except where significant changes were made in the development of PTpic.

PTpic Adaptability & Error Messages Initially, the SPL-PIC Hall thruster model was not well suited to simulate the DCFT. In fact, adapting the model to new configurations was an arduous task – requiring a large number of source code modifications. An important consequence of changes made to accommodate the DCFT has been the generalization of many algorithms in the model that have made PTpic much more adaptable than its direct predecessor. Possibly more importantly, over 300 error checks have been added that cause PTpic to abort execution if satisfied. Additionally, an error message, error code, or both are provided to inform users of the error, allowing them to look up the cause in a specific portion of the source code, if needed. Prior to PTpic, "silent" errors were the norm, which made locating sources of known errors difficult – and also allowed many errors to go unnoticed, initially. Great effort was undertaken to notify the user of unsupported or invalid inputs to PTpic, so that they can quickly discern what input or source code changes are required for their application. It is not practical to describe all the changes made in the development of PTpic, though the most important adjustments will be presented in the following Chapters. For a more detailed discussion of PTpic's adaptability and limitations, see Chapter 6.

Pre-Loop Procedures

- Initialize variables, MPI, BLACS
- Load inputs & data from files
- Calculate geometry dependent parameters
- Initialize particles within domain

Main Loop Procedures

1. If required: solve for ϕ & calculate $\underline{\mathbf{E}}$-field
2. If desired: contribute to (e)/(i) time-avg.
3. Advance electrons by Δt
4. Advance ions by Δt
5. Start neutral sub-cycle procedure
   a) Accumulate tallies of (n) consumption
   b) If sub-cycle completed
      - Inject neutrals
      - If desired: contribute to (n) time-avg.
      - Advance neutrals
   c) If sub-cycle not completed: do nothing
6. Normalize species moments
7. Update diagnostic data buffer
8. If desired: save data to files

Post-Loop Procedures

- Save final, instant particle moments
- Free allocated memory
- Finalize BLACS, then MPI

Figure 2-27: Overview of main procedures in PTpic.

Chapter 3

Particle Trajectories & Interactions

3.1 Time Integration of Particle Trajectories

PTpic numerically integrates particle trajectories using the leapfrog scheme as suggested by Boris[87] (pp.58-59). This is identical to previous versions of the code (see Section 3.14 of Szabo's PhD thesis[97]). The leapfrog method is typically implemented by updating particle positions at integer time steps and velocities halfway between them. However, it is possible to advance positions and velocities simultaneously with the same order of accuracy[138] using Eqs. 3.1 and 3.4 – referred to here as the synchronous leapfrog method (SLM). In this section, derivations of two leapfrog schemes will be presented in preparation for a discussion of an error identified in the previous integration method, as well as a demonstration of the resulting energy error magnitudes for a simple problem.

Expanding a future particle position $\underline{r}(t + \Delta t)$ to second order about t using a Taylor expansion provides Eq. 3.1. Adding the expansions of $\underline{v}(t + \Delta t)$ and $\underline{v}(t)$ about $t + \frac{\Delta t}{2}$ gives Eq. 3.2, which can be combined with Eq. 3.3 to arrive at Eq. 3.4. Together, Equations 3.1 and 3.4 allow particle trajectories to be integrated forward in time with second order accuracy. The more common "offset" leapfrog method (LM) can be obtained just as easily by shifting the time axis in Eq. 3.2 with $t \rightarrow t - \frac{\Delta t}{2}$, giving the finite difference relation for velocities in Eq. 3.5. Finally, the leading staggered velocity in Eq. 3.6 can be substituted into Eq. 3.1, providing the equation used to update particle positions in the LM. The standard leapfrog equations are outlined with double boxes for emphasis – see Eq. 3.5 and 3.7. Both algorithms accurate to

second order – though the SLM requires more floating point operations.

$$\boxed{\underline{r}(t+\Delta t) = \underline{r}(t) + \left[\underline{v}(t) + \underline{a}(t)\frac{\Delta t}{2}\right]\Delta t + \vartheta(\Delta t^3)} \tag{3.1}$$

$$\underline{v}(t+\Delta t) = \underline{v}(t) + \underline{a}(t+\frac{\Delta t}{2})\Delta t + \vartheta(\Delta t^3) \tag{3.2}$$

$$\underline{a}(t+\frac{\Delta t}{2})\Delta t = [\underline{a}(t) + \underline{a}(t+\Delta t)]\,\frac{\Delta t}{2} + \vartheta(\Delta t^3) \tag{3.3}$$

$$\boxed{\underline{v}(t+\Delta t) = \underline{v}(t) + [\underline{a}(t) + \underline{a}(t+\Delta t)]\,\frac{\Delta t}{2} + \vartheta(\Delta t^3)} \tag{3.4}$$

$$\boxed{\boxed{\underline{v}(t+\frac{\Delta t}{2}) = \underline{v}(t-\frac{\Delta t}{2}) + \underline{a}(t)\Delta t + \vartheta(\Delta t^3)}} \tag{3.5}$$

$$\left[\underline{v}(t) + \underline{a}(t)\frac{\Delta t}{2}\right]\Delta t = \underline{v}(t+\frac{\Delta t}{2})\Delta t + \vartheta(\Delta t^3) \tag{3.6}$$

$$\boxed{\boxed{\underline{r}(t+\Delta t) = \underline{r}(t) + \underline{v}(t+\frac{\Delta t}{2})\Delta t + \vartheta(\Delta t^3)}} \tag{3.7}$$

The acceleration terms $\underline{a}$ in Eqs. 3.1-3.7 depend entirely on the physical system being modeled. In Hall and cusped-field thrusters, aside from boundary and inter-particle interactions, all accelerations are calculated according to the Lorentz force. The electric field $\underline{E}(\underline{r},t)$ is taken to be electroquasistatic[a], while the magnetic field $\underline{B}(\underline{r})$ is assumed to be static – induced magnetic fields from plasma currents or time-varying electric fields are neglected. Boundary conditions and particle positions at t provide $\underline{E}(\underline{r},t)$, while the magnetic force is given by $q\underline{v}(t) \times \underline{B}(\underline{r})$.

When applying the SLM to advance from initial conditions $[\underline{r}(t_0), \underline{v}(t_0)]$, particle positions are updated first using Eq. 3.1– providing enough information to calculate $\underline{E}(t_0+\Delta t)$. At this point, the only remaining unknown is $\underline{v}(t+\Delta t)$, which appears explicitly on the left side of Eqn. 3.4, and implicitly on the right side within $\underline{a}(t+\Delta t)$. Fortunately, analytical solutions for each component of the new velocity are easily found, especially with $B_\theta \equiv 0$[b]. Most plasma PIC models, including PTpic, do *not* use the SLM – despite the fact that it provides new positions

[a] "Electroquasistatic" refers to an approximation applied to Maxwell's equations for electricity & magnetism where the electric field is assumed to be irrotational [139] (Ch. 3). In other words, time-variations in the magnetic field $\underline{B}$ are ignored in the equation expressing Faraday's law of induction.

[b] Rather than applying analytical equations within source code, it may be quicker to express each vector equation in terms of 3x3 matrices and 3x1 vectors to exploit preexisting hardware and software optimizations related to solving linear systems of equations

and velocities at identical instants in time. The offset leapfrog is ubiquitous to the point that "leapfrog method" (LM) is commonly understood to specify the time-staggered variant. From this point forward, "leapfrog" should be interpreted as referring to the offset method – unless stated otherwise.

When PTpic creates a particle at a time t, its position $\underline{r}(t)$ and/or velocity $\underline{v}(t - \frac{\Delta_t}{2})$ are pseudo-randomly generated and then integrated forward in time using the LM. A method more inline with a purist's leapfrog implementation might generate $\underline{v}(t)$ and then send it back half a timestep, but the precise initial conditions of each particle are unimportant for this application. The first step is to update the particle velocity using Eq. 3.5. The velocity change $\underline{a}(t)\Delta t$ is evaluated as in Eq. 3.8, where $\underline{r}$ is understood to represent a particle position at time t. Here, the average of the offset particle velocities can be used because $\underline{v}(t) = \frac{1}{2}\left[\underline{v}(t + \frac{\Delta t}{2}) + \underline{v}(t - \frac{\Delta t}{2})\right]$, accurate to first order in Δt. Multiplication by another Δt factor makes the entire equation accurate to second order, as desired.

$$\underline{a}(t)\Delta t = \frac{q}{m}\left\{\underline{E}\left(\underline{r}, t\right) + \frac{1}{2}\left[\underline{v}(t + \frac{\Delta t}{2}) + \underline{v}(t - \frac{\Delta t}{2})\right] \times \underline{B}\left(\underline{r}\right)\right\}\Delta t + \vartheta(\Delta t^3) \tag{3.8}$$

The method of Boris reduces the computational cost of the algorithm by using the intermediate variables $\underline{v}^+$ and $\underline{v}^-$, as defined in Eq. 3.9 — where $K \equiv \frac{q\Delta t}{2m}$. The full expression for $\underline{v}^+$ is provided in Eq. 3.10. The final step is to obtain $\underline{v}(t + \frac{\Delta t}{2})$ from the definition of $\underline{v}^+$ in Eq. 3.9. Once the new velocity is known, each particle's position is updated using Eq. 3.7.

$$\underline{v}^+ \equiv \underline{v}(t + \frac{\Delta t}{2}) - K\underline{E}(\underline{r}, t) \quad , \quad \underline{v}^- \equiv \underline{v}(t - \frac{\Delta t}{2}) + K\underline{E}(\underline{r}, t) \tag{3.9}$$
$$\underline{c} \equiv \underline{v}^- + K\left(\underline{v}^- \times \underline{B}(\underline{r})\right)$$

$$\begin{aligned} v_\theta^+ &= \frac{c_\theta + K\left(B_r c_z - B_z c_r\right)}{1 + K^2\left(B_z^2 + B_r^2\right)} \\ v_z^+ &= c_z - KB_r v_\theta^+ \\ v_r^+ &= c_r + KB_z v_\theta^+ \end{aligned} \tag{3.10}$$

The main motivation for providing the details of the LM here comes from how it was applied in previous versions of PTpic. The first, and ultimately harmless detail has to do with the coor-

dinate systems in which particles are moved. Comparing Eq. 3.10 to its counter part in Szabo's PhD thesis[c], we see that there is disagreement in the equations for v_θ^+: terms proportional to K in the numerator have opposite signs. Looking at his derivation, it is clear that this difference can be attributed entirely to how the positive azimuthal direction is defined. Szabo used a non-standard coordinate system with a right-handed ordering of polar coordinates (r, z, θ), whereas (r, θ, z) is far more common[d]. This can cause confusion when interpreting output from earlier versions of the code, because the positive azimuthal direction is opposite what would normally be expected. The more standard coordinate system is in place in the final PTpic revision resulting from this work. Again, this did not appear to cause any errors in prior simulations. The second reason for this discussion involves errors that are introduced when timesteps are allowed to vary.

3.1.1 Timestep: Variable Step Sizes Destabilize the Leapfrog Method

Prior to PTpic, the SPL-PIC model selected a time-step each iteration from two choices. The first, Δt_p defined in Eq. 3.11, balances the need to manage numerical plasma heating and maximize the timestep [88] (p. 321). The same authors describe this choice as optimal because smaller timesteps offer diminishing reductions of numerical heating, and larger choices allow for excessive heating. Previously, Δt_p was chosen by finding the minimum value of $\frac{H}{2\lambda^*_{D,e}}$ over the entire domain, where H represents the size for each mesh cell, $\lambda^*_{D,e}$ is the γ-modified Debye length at each node on the mesh, and $\omega^*_{p,e}$ is the global maximum value of the γ-modified plasma frequency[e]. The timestep was then chosen as the smallest value between Δt_p and $\Delta t_c = {}^{0.3}/_{\omega_{c,MAX}}$, where $\omega_{c,MAX}$ is the maximum resolvable gyrofrequency in the simulation domain.

$$\omega^*_{p,e}\Delta t_p = \min\left(\frac{H}{2\lambda^*_{D,e}}, 1.0\right) \tag{3.11}$$

Unfortunately, by attempting to pick the best timestep for each iteration, significant errors were likely introduced to past simulation results — including my own [34,114]. In retrospect, the error is obvious for the offset LM when Eq. 3.5 is examined closely. Assuming $\Delta t = \Delta t_1$ before

[c] See Eq. 3.152 in Ref. 97

[d] See sections 3.14.3 and 3.14.4 in Ref. 97

[e] Using the *maximum* modified plasma frequency is more stringent than necessary. A better method would pick the minimum timestep at each node using local values of $\omega^*_{p,e}$, and then select the timestep from the global minimum amongst the nodes.

time t, and that $\Delta t = \Delta t_2$ after time t, we see that $\underline{v}(t + \frac{\Delta t_2}{2})$ depends on $\underline{v}(t - \frac{\Delta t_2}{2})$. However, the only known velocity at this stage in the algorithm is $\underline{v}(t - \frac{\Delta t_1}{2})$. Clearly, significant errors should be expected to accrue by repeatedly misusing $\underline{v}(t - \frac{\Delta t_2}{2})$ in place of $\underline{v}(t - \frac{\Delta t_1}{2})$.

At this stage, it is possible to wonder if the timestep could be varied slowly or smoothly enough to allow at least some flexibility in Δt when applying the LM. Going even further, at first glance it appears that the SLM discussed earlier should allow variable timesteps because positions and velocities are always known at, and advanced to, identical instances in time. The central differences used to derive Eq. 3.4 only involve the time coordinates t, $t + \frac{\Delta t}{2}$, and $t + \Delta t$. The spacing of these terms has no dependence on the previous timestep. Despite the fact that no symmetries are broken using the SLM in place of the standard LM, using adaptive timesteps has been shown to destabilize leapfrog algorithms [140]. Even if Δt values are bounded above by a value that would normally guarantee stability, the error magnitudes introduced to particle energies cannot be meaningfully bounded. According to Richardson and Finn [141], "parametric instabilities may arise because of resonance between oscillations in $\Delta t(t)$ and the orbital motion". The same authors, and others, have developed and applied adaptive symplectic integrators[f] in some circumstances [138,141,143], though the author was unable to identify a suitable alternative to fixed-step leapfrog algorithms for the present application.

Previous benchmarks of particle orbits in the SPL-PIC model used fixed timestep sizes, so variable timestep effects could not have been noticed in those tests [97] (Sections 3.14.7 & 4.1.4). If variable timesteps were used in validating tests, then they probably were not run long enough to observe noticeable errors in orbiting particle energies. As a simple demonstration, a comparison of results obtained using variable and fixed timesteps is presented in the following discussion for a particle subjected to a one-dimensional harmonic oscillator potential. The governing equations expressed in Hamiltonian form are in Eq. 3.12, initial conditions and the exact solution are given in Eqs. 3.13 and 3.14, respectively. The exact solution is used to provide boundary conditions for the SLM and LM, and allows errors in numerical solutions to

[f] Leapfrog schemes fall within the category of symplectic integrators, for which the symplectic 2-form is exactly conserved — preventing energy errors from increasing with time. The symplectic 2-form can be thought of as an area in phase space related to the governing Hamiltonian of a system. The basic theory of symplectic integrators is described in Ref. 142.

be calculated accurately.

$$\dot{v} = -x \tag{3.12}$$

$$\dot{x} = v$$

$$x(0) \equiv 1 \quad , \quad v\left(\frac{-\Delta t_0}{2}\right) \equiv 0 \tag{3.13}$$

$$x(t) = \cos(t) - \tan\left(\frac{\Delta t_0}{2}\right)\sin(t) \tag{3.14}$$

Error data from four different integration schemes applied to a single particle's trajectory are summarized in Fig. 3-1(a). In cases where the timestep varied, Δt values for each iteration were generated pseudo-randomly and uniformly between $0-30\%$ of the system's inverse angular frequency. However, the first step size Δt_0 was specified as $.3/\omega$, matching initial conditions amongst all cases. A step size of $.3/\omega$ was utilized when timesteps were held constant. Therefore, on average, step sizes were half as large in cases where timesteps were varied as when they were not. Values for Δt were generated randomly in order to emulate the behavior of timesteps in PTpic if they were to be selected based on the previously described criteria. For this example, it is worth emphasizing that the characteristic frequency of the system ω was held constant, and that all time-steps in the varied or constant cases satisfied $\Delta t \leq .3/\omega$

The absolute energy error is defined as the difference in the total energy (kinetic plus potential) between the exact and numerical solutions, normalized by the exact value[g]. From Fig. 3-1(a), it is clear that varying the timestep of the LM causes errors several factors larger than all the other cases. Errors in the SLM are lower than with the LM, as expected based on symmetry considerations, but errors are observed to grow to large amplitudes in a seemingly random fashion. The maximum error when the timestep is fixed is never more than about 2% for both the SLM and LM – in agreement with previous statements that both methods are accurate to the same order in Δt.

The errors for the adaptive SLM and LM tests in Fig. 3-1(a) are representative of error magnitudes encountered when generating many sequences of timesteps, but Fig. 3-1(b) serves

[g]Calculating the total system energy is trivial when state variables are known at common times. For the fixed-step LM, velocities at integer timesteps were estimated using an analogue of Eq. 3.3. Adaptive LM errors presented here were calculated ignoring offsets in time, though they did not differ significantly from estimates made by substituting analytical values of velocity at integer timesteps

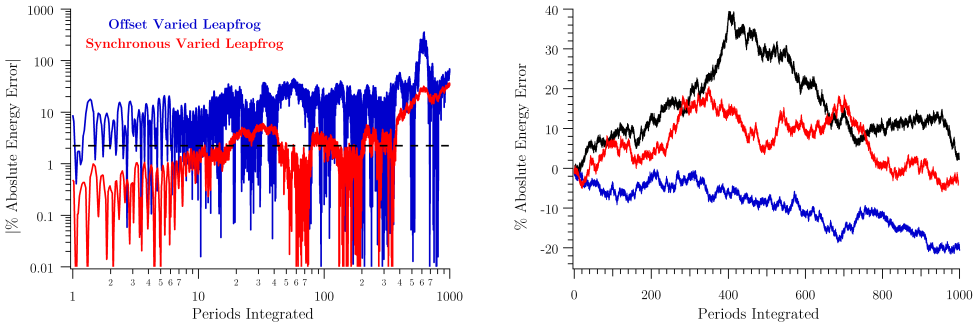

(a) Results from four different methods are summarized in this figure. The dashed horizontal black line represents the maximum error from the offset and synchronous leapfrog methods using a *constant* timestep.

(b) Energy error histories using the synchronous leapfrog method with variable timesteps. These trajectories come from three different sequences of pseudo-randomly generated timesteps – over the same range as in Fig. 3-1(a).

Figure 3-1: Energy error for a particle advanced using variations of the leapfrog method. In all cases in this figure, the particle is subjected to a 1D harmonic oscillator potential, with initial conditions supplied to each integrator from the analytical solution for the trajectory.

to illustrate that maximum errors and error-trajectories vary significantly for different sequences of step sizes. Similar error histories are shown in Fig. 4 of a report summarizing results from the application of adaptive symplectic integrators to several test problems related to celestial mechanics[138]. However, their method is in general iterative and may require several evaluations of forcing terms per step[141].

Offset vs. Synchronous Leapfrog Method: The synchronous and standard leapfrog methods have identical stability and error characteristics when applied to the one-dimensional harmonic oscillator. They probably also have the same accuracy and stability with a full Lorentz type forcing term. However, the procedure for the SLM is more complex than the LM because of the need to determine the acceleration at the new time $t+\Delta t$ *before* calculating the new velocity. In order to do this for all particles, a program would first loop through each particle to update their positions. Then, the potential distribution would be updated based on the new particle positions. Next, a second loop through the particles would be necessary to calculate their new velocities. Contrastingly, the LM is simpler: requiring that the potential distribution be updated after new positions are found, from which the next velocity and positions can be calculated together in a single particle loop for each species. Therefore, in addition to being marginally

more computationally expensive, the code required to implement the SLM is more complex than with the LM because it requires two separate loops through all particles. The tradeoff with the LM is not knowing particle velocities at integer timesteps. In order to calculate them, the previous velocity must be stored in memory, from which $\underline{v}(t) \approx \frac{1}{2}\left[\underline{v}(t-\frac{\Delta t}{2}) + \underline{v}(t+\frac{\Delta t}{2})\right]$, an estimate that is accurate only to first order in Δt, whereas the values at half-steps are accurate to second order. Alternatively, higher-order algorithms such as Runge-Kutta methods might be used, but these are more computationally demanding and do not preserve constants of motion — particularly energy — unlike symplectic integrators [144]. However, the advantages of using variable timesteps do invite the consideration of non-symplectic integrators if errors can be acceptably bounded. This might be facilitated by choosing a maximum electron Hall parameter value Ω_{max} for which a specified error tolerance is maintained — where Ω_{max} simply specifies the number of orbital periods over which the energy error of a single particle should remain below the given threshold. This approach would need to be compared against test cases using the more established LM in order to establish its validity.

3.1.2 Tracking Particles as they Move

The change in position for each particle is determined according to Eq. 3.7 after the new velocity $\underline{v}(t+\frac{\Delta t}{2})$ is calculated, which will generally result in displacements out of the R-Z simulation plane. However, $\underline{v}(t+\frac{\Delta t}{2})$ and the displacement are calculated in terms of the unit vectors for the *old* position vector $\underline{r}(t)$, so a simple rotation transformation is required to account for the rotation of the polar coordinate system as the particle moves. The new cylindrical radius is determined directly from the radial and azimuthal displacements, and the velocity components are transformed into the rotated frame. Details are well described by Szabo in Section 3.14.6 of Ref. 97 — specifically in Fig. 3-21 and Eq. 3.160 of that document. PTpic does not attempt to resolve any azimuthal asymmetries, so the angle of rotation is not calculated or stored in memory. The non-transformed velocity vector is retained until it is determined whether or not the particle has intercepted any objects defined within the domain. Particle collisions with boundaries are addressed in Section 3.2.

The new coordinates (z, r) for the particle must now be translated to identify its location on the structured mesh. In other words, the *computational coordinates* $[\xi, \eta]$ of the particle

must be determined – these were previously illustrated in Fig.1-8. Computational coordinates are used to weight particles characteristics to nodes on the mesh, and also to weight quantities calculated at nodes back to each particle. They are used to calculate particle moments such as densities, bulk velocities, temperatures, pressure tensor components, and heat flux vectors. They also translate particle positions into charge densities that constitute the inputs to the potential solver. In turn, electric and magnetic fields are weighted to each particle's position based on the same coordinates. This is a complicated task in PTpic because the mesh is not required to be orthogonal and because of particle collisions with predefined thruster surfaces.

The first step in calculating $[\xi, \eta]$ from (z, r) is to identify which cell $[k, j]$, if any, a particle is located within. Particles leaving the domain or entering predefined objects are dealt with separately. Figure 3-2 shows a particle p located within cell $[k, j]$ with physical coordinates (z, r). The cell coordinates $[\alpha_k, \alpha_j]$ are defined within the context of Eq. 3.15, where all functions in PTpic are interpolated in this fashion. For the linear functions $z(\alpha_k, \alpha_j)$ and $r(\alpha_k, \alpha_j)$, the interpolation is exact [145]. Based on the stored node coordinates and other values calculated as the simulation starts, it is sufficient at this point to state that the cell coordinates $[\alpha_k, \alpha_j]$ can be calculated for a given position (z, r) once the cell containing the point is known. The details of the calculation are easy to derive, and are presented in Ref. 145 and in Section 3.8.3 of Ref. 97. Once the cell coordinates are known, the weighting factors $P_{\Delta k \Delta j}$ and computational coordinates $[\xi, \eta] = [k + \alpha_k, j + \alpha_j]$ are easily calculated. Additionally, rounding the computational coordinates to their nearest integer value provides the mesh coordinates of the node that a particle is closest to $[k_{NGP}, j_{NGP}]$. This is used to count the number of particles within regions around each node — vital information when calculating ionization collision frequencies that are dependent on the background density of ions or neutrals. When particles are weighted to their nearest neighbor, this is referred to as "zero-order" or "nearest-grid-point (NGP)" weighting. When the weighting factors $P_{\Delta k \Delta j}$ are used, this is called "first-order" or "cloud-in-cell (CIC)" weighting. These terms refer to the effective shape the computational particle assumes under each scheme [87] (Sec. 4-6).

$$\begin{aligned} f(\alpha_k, \alpha_j) &\approx (1-\alpha_k)(1-\alpha_j) f_{0,0} + (1-\alpha_k)\alpha_j f_{0,1} + \alpha_k(1-\alpha_j) f_{1,0} + \alpha_k \alpha_j f_{1,1} \\ &= P_{00} f_{0,0} + P_{01} f_{0,1} + P_{10} f_{1,0} + P_{11} f_{1,1} \end{aligned} \tag{3.15}$$

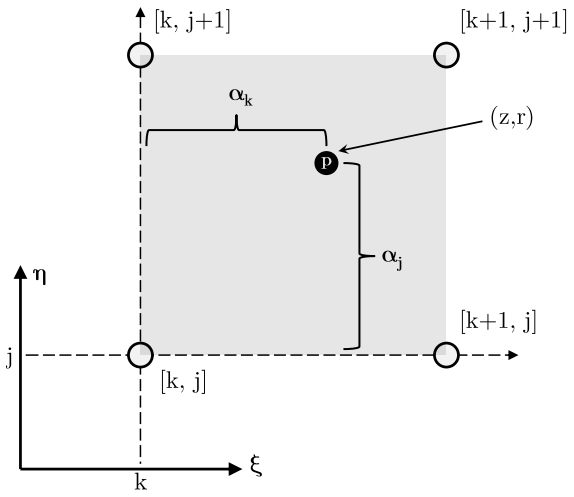

Figure 3-2: Schematic illustrating the relationships between computational and physical coordinates. Cell $[k, j]$ is colored grey in this figure. The dashed vertical and horizontal rays are defined by $\xi = k$ and $\eta = j$, respectively.

Originally, the SPL-PIC model estimated the computational coordinates of a particle from Eq. 3.16 and then checked whether the particle was within cell $[k, j] = [\lfloor \xi \rfloor, \lfloor \eta \rfloor]$. If the particle's cell was not identified by this process a computationally expensive spiral search, starting from the first cell to be checked, was initiated outward until the cell containing the particle was located. Further details are in Section 3.14.7 in Ref. 97, where it is said that each particle's weighting constants should be recalculated once the proper cell is identified. However, when the model the author inherited at the start of this project from J. Fox was first applied to the DCFT geometry, a common and recurring error occurred where particles were lost by the model. Running the same code for the P5 Hall thruster geometry did not produce the same errors, likely because it lacked an angled wall like the DCFT — necessitating greater mesh element distortion. Eventually, after combing through the SPL-PIC source code, it was determined that computational coordinates for some particles were not being updated properly. In some cases, the only modifications to the computational coordinates came from the application of Eq. 3.16. This would cause the model to mistakenly assume particles were in the wrong cell as error from the estimates accrued. Previously, in the event that a particle was lost, a message would be printed to output and the particle was deleted from memory.

This serious flaw in the code was corrected for PTpic by only updating the computational coordinates of a particle after its resident cell has been located. Additionally, if a single particle is not properly located, an error message is issued and the program terminates execution immediately. The new procedure, which will now be summarized, appears to resolve the issue of the lost particles successfully. No particles have been misplaced after several hundred million iterations applied to three different thruster geometries. Most of the tasks listed in the remainder of this section are performed as part of the `check_location_and_external_boundaries()` subroutine, defined in `tracking.c`.

$$\begin{pmatrix} \xi \\ \\ \eta \end{pmatrix} \approx \begin{pmatrix} \xi_0 + \frac{\partial \xi}{\partial \hat{z}} \Delta \hat{z} + \frac{\partial \xi}{\partial \hat{r}} \Delta \hat{r} \\ \\ \eta_0 + \frac{\partial \eta}{\partial \hat{z}} \Delta \hat{z} + \frac{\partial \eta}{\partial \hat{r}} \Delta \hat{r} \end{pmatrix} \tag{3.16}$$

New Procedure for Updating Computational Coordinates: First, the new particle location is checked to see if it is located within the same cell it started in. Most particles will fall into this category. Next, some of the surrounding cells are checked based on which boundaries in the original cell the particle was closest to when its computational coordinates were last known. The idea depends on the fact that most particles will leave the cell through the boundary they are closest to because timesteps should be small enough to force most particles, especially ions and neutrals, to take steps that are small compared to cell dimensions. Based on the old values of $[k, j]$ and $[k_{NGP}, j_{NGP}]$ the previous location can be classified into one of four quadrants, visualized in Fig. 3-3(a). Within each quadrant, the closest boundary is selected by comparing the previous value of η to either $j + \alpha_k$ or $j + 1 - \alpha_k$. The division of the quadrants by these criteria are illustrated in Fig. 3-3(b).

Once the closest surface is selected, it is checked for an intersection with the line segment joining $\underline{r}(t)$ and $\underline{r}(t + \Delta t)$. If an intersection is found, the cell sharing that boundary is checked to see if it contains the new location. If the particle is not in this cell, the cells above and below (or to the left and right) of the most recently checked cell are tested to see if they contain the new location – in the order of their proximity to the old location. If the correct cell has not been found at this stage, then Eq. 3.16 is used to determine a guess $[\xi_g, \eta_g]$ from which a possibly new cell is suggested from $[k, j] = [\lfloor \xi_g \rfloor, \lfloor \eta_g \rfloor]$. Finally, if the proper cell still has not been identified, then the dreaded spiral search is called. The spiral search succeeds by brute

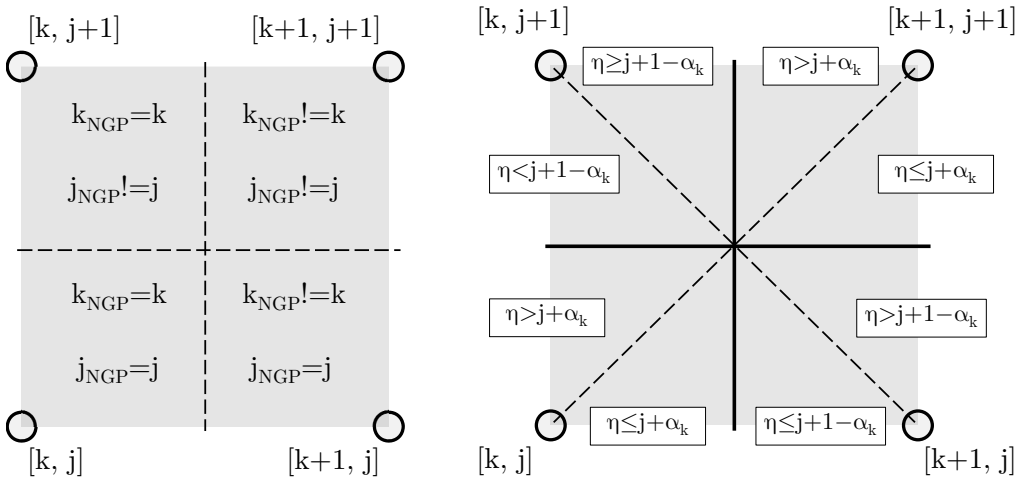

(a) First, the particle's starting point is classified as belonging to one of four quadrants. Here, "!=" means "not equal."

(b) Within a given quadrant, the old value of η is compared to one of dashed lines in computational space. The coordinate system present in Fig. 3-2 applies in this figure as well.

Figure 3-3: Schematic illustrating how cells are partitioned when identifying which boundary a particle was likely to leave a cell from.

force, searching every cell on the mesh if needed — starting near the original cell. It typically finds particles after searching less than 10 cells, and is not invoked often because the preceding attempts are normally successful. However, if too large a step size is specified, allowing many electrons to advance distances larger than one cell per iteration, the frequency of calls to the spiral search will increase. This will degrade the performance of PTpic noticeably. It is always a good idea to compare the specified timestep with the step size PTpic would pick if variable timesteps were allowed — this value is calculated and written to the diagnostic data file for each iteration that PTpic executes (the values are saved in a memory buffer and written to file after a specified number of iterations has passed).

Absent from the previous discussion is any mention of what happens when a particle encounters a thruster surface or leaves the computational domain. If the particle is found within a cell, the cell's "type"[h] is checked to see if it is an `INTERIOR` cell or not. If not, then it has been found within an interior object — such as the anode, dielectric, or other component — and will be subjected to the procedures outlined in Section 3.2. Otherwise, it is treated as

[h]The cell type is an attribute assigned when the simulation is initialized, based on user inputs.

normal: the cell's computational coordinates are updated now that a suitable cell containing the new coordinates (z, r) has been located by PTpic. Throughout the effort to find a particle's new cell, if a cell suggested by PTpic for searching is not part of the domain, appropriate mesh boundaries are checked for intersections with the segment joining the particle's new and original locations. If an intersection is identified, and the particle is cleared of having intersected any internal boundaries (or external boundaries representing an objects surface) on its way out of the domain, the particle is deleted from memory after appropriate data related to particle, momentum, and energy fluxes have been recorded.

3.2 Boundary Interactions

When a particle's new position $\underline{r}(t + \Delta t)$ is located within a cell, or found to have collided with a surface on its way out of the domain, its treatment depends on whether the particle is an electron, ion, or neutral. In PTpic, electrons are never bounced from surfaces[i], though secondary electron emission is included. At all surfaces, electrons are deleted from memory after their charges are taken into account. When impacting a `DIELECTRIC`[j] surface, super-electron charges are distributed amongst the two nearest nodes defining the surface, based on the distance from each node. If a `FLOATING_WALL` is encountered, the super-electron charge contributes to that body's total charge. At `ANODE` and `FREE_SPACE` boundaries, the super-electron is deleted after it contributes to the anode current and other data. Ion super-charges, whether singly or doubly charged, are counted the same way as electron charges at all boundaries, except that a super-ion is converted to a neutral particle after any boundary collision — excluding `FREE_SPACE` boundaries. The surviving neutral particle is then bounced diffusively according to the algorithm described later in this section. Neutral particles are bounced off all surface they collide with, except `FREE_SPACE` boundaries. As a quick note about the deposition of ion charges, an idea for a new sheath patch was tested which only counted ion charges at `DIELECTRIC` and `FLOATING_WALL` surfaces a fraction $\frac{1}{\sqrt{r}}$ of the time. However, this

[i] Before PTpic, V. Blateau implemented a sheath patch in SPL-PIC which reflected some low energy electrons and later selected high energy electrons for removal. This was meant to correct changes in the energy–composition of electrons lost at walls – caused by using a mass factor r greater than unity [110]. This has been removed for PTpic.

[j] The boundary types in all caps are macro names used in the PTpic source-code to distinguish between different boundaries.

was found to prevent the floating body's potential from recovering to higher values after electrons initially drove it downward. As a result, this sheath patch was deactivated for all results presented in this thesis, unless otherwise noted. It could be reactivated by setting the macro `NO__MASSFACTOR_SHEATH_CORRECTION` to zero in `pic.h`, and then recompiling PTpic.

Particle Bouncing Algorithm: The routine for bouncing particles was rewritten from scratch for PTpic because the version the author inherited was overly complex and did not handle the case of a sloped boundary satisfactorily. It may have been modified from the original version, because the thruster Szabo first simulated had an angled wall [97] (Fig. 4-1). The new routine is defined in `boundaries.c` as `bounce_diffusely()`. The logic for this algorithm is divided into three steps. First, the time of intersection t_c is found within the interval from t to $t + \Delta t$. Second, a particle's velocity after bouncing is calculated with a specified accommodation factor f_a determining its speed. Finally, the particle is advanced forward in time from t_c to $t + \Delta t$, using the new velocity, so that the bounced particle is re-synchronized with all other particles at the start of the next iteration. Bouncing is done within a loop to accommodate the unlikely event that a particle bounces multiple times per step.

A detailed description of each step is not given here, but some of the most important aspects are given in the order in which they are useful in the process outlined in the paragraph above. When determining t_c, the newest velocity $\underline{v}(t + \frac{\Delta t}{2})$ available for the particle is used. This is entirely justifiable, despite the fact that positions and velocities are not known simultaneously, because $\underline{v}(t)$ and $\underline{v}(t + \frac{\Delta t}{2})$ are equivalent to first order in Δt — compare Eqs. 3.1 and 3.7 to see why this is true. The speed of a particle after colliding with the wall w_a is determined based on three things: the particle's speed before hitting the wall w_b, a randomly generated speed w_w defined in Eq. 3.18, and the accommodation factor f_a. Equation 3.17 completely defines how w_a is calculated. The randomly generated speed is intended to be selected from a Maxwellian distribution at the specified wall temperature T_w using a uniformly distributed, pseudo-randomly generated number $R \in (0, 1)$[k]. The unit vector for the new velocity is pseudo-randomly chosen amongst all directions facing away from the surface, but will be directed along the surface normal $\hat{n}$ on average. With the new velocity $\underline{w}_a$, the particle's position is advanced

[k]This is the method inherited from Szabo for generating speeds for particles distributed as Maxwellian — see Eq. 3.40 in Ref. 97. However, Szabo states that this is the Box-Muller method for generating a 2D distribution, but PTpic requires speeds compatible with a 3D distribution.

from t_c to $t + \Delta t$ using Eq. 3.19, which is accurate to first order in the step size. In this equation, $\underline{r}_c^*$ is the modified location where the particle collided with the wall, and Δt_{rem} is the time remaining from the original timestep. The wall collision position $\underline{r}_c^*$ is equal to the true collision location plus a small offset $(\delta r)\hat{n}$, where δr is supplied by the user in the variables input file – an example value used in DCFT simulations is about $1\mu m$. This was needed to remove any ambiguity between the newest and prior wall collisions when the new collision location is located for multiply-bounced particles. The distinction between Δt_{rem} and Δt is only meaningful when a particle experiences more than one bounce per timestep.

$$w_a = (1 - f_a)w_b + f_a w_w \tag{3.17}$$

$$w_w = \left(\frac{8kT_w}{\pi m}\right)^{1/2} \sqrt{-\ln R} \tag{3.18}$$

$$\underline{r}(t + \Delta t) = \underline{r}_c^* + \underline{w}_a \left(\Delta t_{rem} - t_c\right) \tag{3.19}$$

3.2.1 Secondary Electron Emission (SEE) Model

The secondary yield Σ for a given electron super-particle incident upon a dielectric surface, in this case boron nitride, is calculated according to Eq. 3.20, as listed in Sec. 7.3.3 of Ref. 22. Though Goebel & Katz list E_{eV} as the electron temperature, in PTpic each incident electron's kinetic energy is calculated in electron volts and then used to determine the yield for the single particle. A more accurate model of the secondary electron yield would specify the yield for individual electrons of known energies. With the yield σ known, the number of secondaries N_s emitted by the wall is calculated according to Eqn. 3.21, where $\lfloor x \rfloor$ signifies the floor operator applied to the number x, and r is a uniformly distributed pseudo-random number on the interval $(0, 1)$. No upper bound is placed on the number of secondaries that may be released for a single electron because the model should be able to resolve the effects of charge saturation within the simulated sheath structure. Also, no dependence on the incidence angle of the primary electron is assumed. Finally, the primary electron is always deleted from memory after its charge has been deposited onto the insulator wall. Secondary electrons are injected at thermal speeds associated with the boundary temperature specified by the user in the `IN__CONSTANTS` input file, which means secondary electrons are injected with very low energies — though they quickly

gain energy if they make it into the bulk.

$$\Sigma = (1.38)(.15)(E_{eV})^{.549} \tag{3.20}$$

$$N_s = \lfloor \sigma \rfloor + \{r < (\sigma - \lfloor \sigma \rfloor)\} \tag{3.21}$$

3.2.2 Dielectric Erosion Rate Calculations in PTpic

A module has been added to PTpic that will calculate erosion rates at each mesh node located on a dielectric surface. This module can be turned on or off by manipulating the appropriate parameter within the `IN__VARIABLES` input file. If the erosion module is active, PTpic reads parameters to calculate sputter yields from a separate input file whose name is stored in the character array `IN__Sputter`, which must reside in a folder named `sputter` — the `sputter` folder must be located in the same directory as the PTpic executable file. The name for `IN__Sputter` must be specified in the file named `MasterPTpic.in`.

The erosion module calculates contributions to erosion from singly and doubly charged ions, as well as charge-exchange neutral particles. When one of these particles strikes a dielectric wall, the volume eroded $\mathbb{V}$ from the impact of a single super-particle — no matter the charge — is calculated according to Eqn. 3.22, where W is the super-particle weight, e is the elementary charge, and the sputter yield S_ξ is a function of the incident particle energy E and incidence angle Θ relative to a unit vector $\hat{n}$ normal to the surface. Here, Θ is defined as zero when the particle strikes the surface with a velocity parallel to the normal direction. A redeposition model is not included in PTpic.

The sputter yield is typically given in units of volume eroded per charge, with the erosion rate ξ calculated as $\xi = jS_\xi$, where j is a xenon ion current density. However, absent separate yields for neutrals and doubly charged ions, PTpic converts the yield per charge to a yield per particle via multiplication by the factor e. Therefore, with equivalent impact energies and trajectories, neutrals and ions of any charge will contribute identical amounts of erosion with this model. However, doubly charged ions will typically arrive at walls with more energy than singly charged ions, so greater yields can be expected due to the energy dependence of S_ξ. The details of how an atom interacts with the surface may depend on its charge, but no effort to differentiate between the yields for differently charged species is made in this erosion model.

Furthermore, experimental and computational data, that curve-fits for S_ξ are based upon, are measured using beams that are expected or assumed to be comprised entirely of singly charged ions[50,146]. In fact, in the molecular dynamics simulations performed by J. Yim, each impact ion is simulated as a neutral particle because it is assumed that an electron from the surface neutralizes the ion before impact[50] (Sec. 4.1.2.2).

$$\mathbb{V} = W \cdot e \cdot S_\xi\left(E, \Theta\right) \tag{3.22}$$

$$\delta = \frac{\mathbb{V}}{A} \tag{3.23}$$

Following the method used by S. Cheng in her erosion simulations using a modified version of HP-Hall, the eroded depth δ is calculated using Eqn. 3.23, and this displacement is weighted linearly to the nearest two nodes based on the collision location's distance from each node[147] (Sec. 4.2.5). In Eqn. 3.23, A is the area of the boundary segment rotated about the thruster axis of symmetry. The vector displacement of a node is given by $\delta\hat{n}$, and accumulates additively as individual super-particles collide with each segment. The accumulation of erosion can continue from one simulation to the next, or be reset by the user using a switch in the `IN__VARIABLES` input file.

The sputter yield S_ξ is assumed to have the form shown in Eqn. 3.24, where Y_n is the sputter yield when $\Theta = 0$ and $\hat{Y}$ is the normalized angular variation in the yield—which may also depend on the energy of the incident particle. PTpic currently allows choosing one of three different functions in each of two energy ranges to specify the sputter yield at a constant angle Θ_o. The energy ranges are defined by $E \in (0, E_{limit}]$ and $E > E_{limit}$, where E_{limit} and Θ_o must be specified in the `IN__SPUTTER` input file. The available fits for $S_\xi(E, \Theta_o)$ are shown in Eqn. 3.25–3.27, giving the fits proposed by Wilhelm, Zhang & Zhang, and Bohdansky, respectively. The background behind these fits, and associated references, are well summarized in Sec. 2.4.2 in J. Yim's PhD thesis[50]. Note that when $\Theta_o \equiv 0$, $S(E, \Theta_o) = Y_n(E)$. Also, each curve fit has a threshold energy, so over the range that $S < 0$, the yield is set to zero. The factor $s_n(E)$ is defined in Eqns. 3.28 — identical to Eqn. 2.14 in Yim's PhD thesis. The fits are specified by providing a string of characters to identify the two fits, and two sets of coefficients: $\{A_1, E_{th,1}, E_{o,1}\}$ and $\{A_2, E_{th,2}, E_{o,2}\}$. This should all be specified in `IN__SPUTTER`, though

the third coefficient E_o need only be specified when the Zhang fit is used. The first group of coefficients is used for $E \leq E_{limit}$, and the second is used for $E > E_{limit}$. In her simulations, S. Cheng used a modified Yamamura fit for the normal yield that is not provided here, but is given in Eqn. 3.4 in her PhD thesis. It would not be difficult to add or change the curve fits by modifying a small portion of the PTpic source code.

$$S_\xi\left(E,\Theta\right) \equiv Y_n(E)\hat{Y}(E,\Theta) \tag{3.24}$$

$$S(E,\Theta_o)_W = A\left(E - E_{th}\right)^2 \tag{3.25}$$

$$S(E,\Theta_o)_Z = As_n(E)\left[1 - \sqrt{\frac{E_{th}}{E}}\cos\Theta_o\right] \tag{3.26}$$

$$S(E,\Theta_o)_B = A\left[1 - \left(\frac{E_{th}}{E}\right)^{2/3}\right]\left(1 - \frac{E_{th}}{E}\right)^2 \tag{3.27}$$

$$s_n(E) = \frac{\ln\left(1 + 1.1383^{E/E_o}\right)}{2\left[{}^{E}\!/_{E_o} + .01321\left({}^{E}\!/_{E_o}\right)^{0.21226} + .19593\sqrt{{}^{E}\!/_{E_o}}\right]} \tag{3.28}$$

Two different curve fits for the angular yield $\hat{Y}$ are given in Eqns. 3.29–3.30, where $x \equiv \frac{1}{\cos\Theta}$. These are referred to, in order, as Yamamura's light-ion (LI) and heavy-ion (HI) fits[148]. The exponent f is defined in Eqn. 3.31. However, f has a pole at $E = a^2$, which can be avoided by using f_C in place of f with $a < 0$, as suggested by S. Cheng[147] (Sec. 3.3.2). Cheng's modified definition of the exponent is given in Eqn. 3.32. Note that these equations have some slight differences than those presented in the listed references. For the Yamamura HI fit in Eqn. 3.30, the exponent has a negative sign in front of it in Eq. 17 of Ref. 148 — though the same result can be recovered by assigning a negative value to f_s in Eqn. 3.31. In Sec. 3.3.2 of Cheng's PhD thesis, a minus sign is present in front of f_s in Eqn. 3.32, while there is not one in this report. This appears to be a typographical error, because the angular yield has the incorrect shape when calculated as printed in Ref. 147. In PTpic, the angular fit is selected by providing one of three strings in the `IN__SPUTTER` input file, and the adjustable parameters within PTpic are the factors f_s, a, and Σ.

$$\hat{Y}(E,\Theta)_{LI} = x^f \exp\left[-\Sigma\left(x-1\right)\right] \tag{3.29}$$

$$\hat{Y}(E,\Theta)_{HI} = x^f \exp\left[-\Sigma\left(x-1\right)\right]\left[\frac{1-{}^{\left(a^2/E\right)^{1/2}}/_{x}}{1-\left({}^{a^2}/_{E}\right)^{1/2}}\right] \tag{3.30}$$

$$f \equiv f_s\left(1+2.5\frac{\left({}^{a^2}/_{E}\right)^{1/2}}{1-\left({}^{a^2}/_{E}\right)^{1/2}}\right) \tag{3.31}$$

$$f_C \equiv f_s\left(1+2.5\frac{a\sqrt{{}^{1}/_{E}}}{1-a\sqrt{{}^{1}/_{E}}}\right) \tag{3.32}$$

Coefficient Selection

To calculated erosion rates during simulations of the DCFT, curve fits for the yield at a constant angle, discussed above, are used with coefficients suggested by J. Yim in Sec. 5.2.1.2 of his PhD thesis. Yim provides coefficients to represent $S_o \equiv S_\xi\left(E,45^\circ\right)$ for $0 \leq E \leq 50$ $[eV]$ and $E > 50$ $[eV]$. The normal yield may then be determined from $Y_n(E) = \frac{S_o}{\hat{Y}(E,45^\circ)}$. Therefore, Y_n — in this case — will depend on the fit that is used to describe the angular yield $\hat{Y}$, a consequence of Eqn. 3.24. One option would be to use the fit suggested by S. Cheng, and also used by J. Yim, which takes the form of Eqn 3.29, but uses f_C in place of f for the exponent of x. Cheng's fit for $\hat{Y}$ is based on Yalin's measurements[146] at ion energies of 250 $[eV]$ and 350 $[eV]$ — shown together in Fig. 3-4(a). The data point at $(E,\Theta) = (250, 75^\circ)$ was below the detection limit of the measurement technique, and the coefficients used to generate Cheng's fit are: $\{f_s, a, \Sigma\} = \{5.97563, -3.63786, 1.41355\}$. However, this fit was not chosen for this work because of the resulting energy dependence of Y_n, shown as the finely dotted line in Fig. 3-4(c). The curve provided by Yim for $S_\xi\left(E,45^\circ\right)$ is given in the same figure for a visual reference. Using Yim's fit for $\Theta = 45^\circ$ in conjunction with Cheng's fit for $\hat{Y}$ causes the normal yield to exceed oblique yields for E below about 80 $[eV]$, and for Y_n to decrease as E increases beyond approximately 160 $[eV]$. These are not desirable trends, though the author is not aware of existing data that would contradict the trend in the angular yield at low energies. In an effort to double check Cheng's coefficients, the author performed a new two-dimensional, nonlinear

regression[l] for Cheng's proposed fit using Mathematica[150], providing very similar curves. The new curves are plotted with Yim's data in Fig. 3-4(b), and use similar coefficients as those given by Cheng: $\{f_s, a, \Sigma\} = \{6.13315, -3.55112, 1.49095\}$. This example also verifies the functionality of the nonlinear regression procedure.

A separate fit, this time using an unmodified Yamamura LI fit, is also plotted in Fig. 3-4(b) as the solid curve. The same regression technique used to verify Cheng's coefficients was applied to Eqn. 3.29, except this time using the unaltered exponent f. Constraining $a \geq 0$, this fit is specified by $\{f_s, a, \Sigma\} = \{3.54165, 0, 1.50406\}$. With $a = 0$, $\hat{Y}$ is no longer dependent on energy. Using this form of $\hat{Y}$ results in a different energy dependence for Y_n, plotted as the solid curve in Fig. 3-4(c). A regression using the Yamamura HI equation provided the same fit because the heavy and light-ion fits are interchangeable when $a = 0$. Performing a regression with more data points in $E - \Theta$ space, and perhaps for the total yield S_ξ, might provide better curve fits.

To summarize all of this, the fit types and coefficients used to calculate erosion estimates in simulations of the DCFT are all provided in Table 3.1. Note that Yim's fit for $S_\xi\,(E, 45^\circ)$ is based on simulated xenon sputtering of pure BN at a temperature of 423 $[K]$. The DCFT insulator is made of HP Combat® BN, which is not pure BN because it is prepared using a boric acid binder[151]. Finally, rather than erosion displacements, adjusted erosion rates $\xi = {}^{\xi^*}/_{\sqrt{r}}$ are saved to file by PTpic. The simulated erosion rate ξ^* is interpreted as too large by a factor $\sqrt{r}$ because ion current densities are increased by this factor when using a reduced heavy-particle mass.

[l]This was done using the fuction `NonlinearModelFit` with the following optional arguments[149]: `Method→{NMinimize, Method→{``DifferentialEvolution", ``ScalingFactor"→.9, ``CrossProbability"→.1, ``PostProcess"→{FindMinimum, Method→``QuasiNewton"}}, MaxIterations→150}`

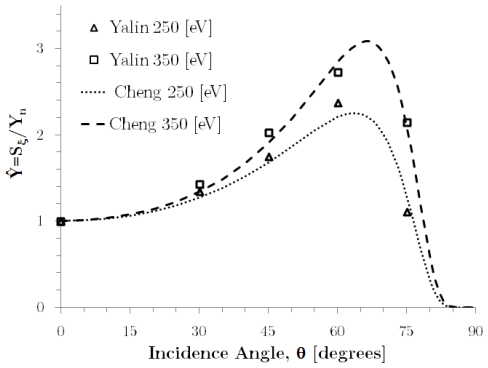

(a) Yalin's measurements with S. Cheng's fits to the data.

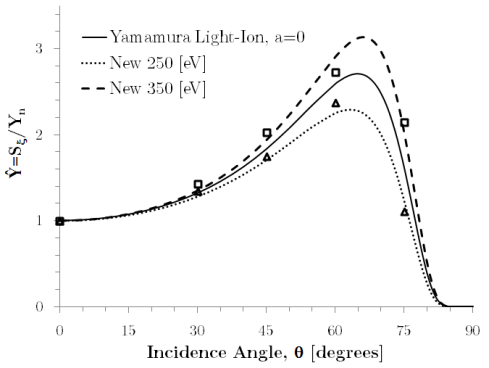

(b) Yalin's measurements presented with new coefficients for Cheng's curve fit, as well as an unmodified Yamamura LI fit.

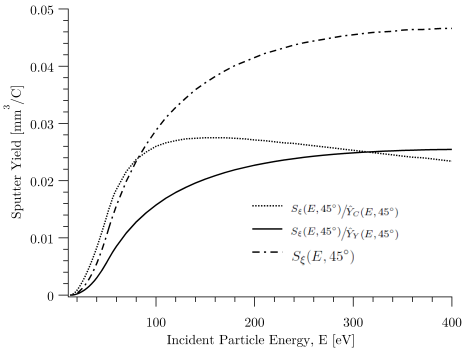

(c) Two different normal yields are plotted, as well as Yim's suggested fit for representing $S_\xi(E, 45°)$

Figure 3-4: Plots comparing fits and data for different angular, normal, and total sputter yields.

Table 3.1: Coefficients and fit–types used to calculate erosion in DCFT simulations.

Constant Angle Yield: $\Theta_o = 45^\circ$	
$0 < E \leq 50$ $[eV]$*: Wilhelm fit*	
A_1	7.62×10^{-6}
$E_{th,1}$	13.0
$E_{o,1}$	-
$E > 50$ $[eV]$*: Zhang fit*	
A_2	.205
$E_{th,2}$	70.8
$E_{o,2}$	465.0
Angular Yield: Unmodified Yamamura LI fit	
f_s	3.54165
a	0.0
Σ	1.50406

3.3 Ionization & Charge-Exchange Collisions

First, a brief review of common terminology related to how hardware and software are organized in parallel computing may help the reader avoid confusion in the following discussion. The terms "process", "processor", "core", and "thread" are all used here with distinct meanings. Typically, a *processor* refers to a single hardware component on a computer's motherboard housing 1 or more *cores*, or central processing units. Hence the term "multi-core processor". A *process* is a set of programming instructions assigned to run using one or more cores. Assigning multiple processes to execute a computer program is an example of *distributed*-memory parallelism, because each process can only access memory that it has been assigned. With distributed memory parallelism, communication between processes is required to share data amongst different processes. If a process is run on multiple cores, subsets of a process' instructions, called *threads*, may be executed in parallel on the available cores to exploit *shared*-memory parallelization. Currently, PTpic only utilizes distributed-memory parallelism using the openMPI library. This means that each process is run on a single core. In order to take advantage of shared-memory parallelism, openMP being a common implementation, each process would need to be run using multiple cores. The acronym openMPI stands for "Open Message Passing Interface", while openMP is short for "Open Multiprocessing". Of course, benefits of multi-core computing can only be realized if processes are assigned to run on one or more core.

Working within the constraints of how SPL-PIC was structured by J. Fox[111] to run with multiple processes, extensive modifications to the algorithms for ionization and charge-exchange (CEX) collisions were necessary to ensure that particles and charges are conserved exactly and locally (within each cell) in *all* circumstances. Previously, the parallel version of SPL-PIC could only ensure global conservation in a statistical sense. This meant that, on average, for every ionized neutral, an ion and an electron would be created, and a neutral particle would be destroyed. However, there was no guarantee that these actions would occur in the same cell, or that a neutral would be destroyed every time a pair of charged particles was created. The situation was not as imprecise as the previous statement might indicate, because neutrals were selected for deletion based on a probability distribution formed as ions and electrons were created. So, just because exact conservation at all times could not be promised, this does not mean approximate conservation was not provided most of the time.

The main impetus for rewriting this portion of the SPL-PIC model for PTpic was the behavior of the model during transients accompanied by strong levels of ionization — especially during the initial phases of plasma simulations. In order to ensure that the code provided a physically accurate mechanism to quench these instabilities, it was necessary to make sure that neutrals were destroyed when and where any ionization event occurred. The overall picture is complicated further by neutral sub-cycling, where neutrals are advanced using a much larger timestep than ions and electrons[m]. This necessitates counting super-neutrals that have been consumed between neutral steps — allowing the consumption of neutrals to feedback properly into calculations of ionization event probabilities. These changes were expected to improve the model's capability to simulate breathing mode oscillations and unsteady ionization phenomena.

The current method of parallelization necessitates an additional layer of complexity in order to locally conserve particles and charges. Particles are distributed in memory amongst all active processes in a roughly even manner. Therefore, particles in one cell may belong to any of the processes assigned to run PTpic. Consequently, whenever any aggregate quantity related to particles in a given cell is calculated — for instance, something as simple as counting particles in the cell — a communication operation amongst all processes is required because they only

[m]The number of timesteps to wait before advancing neutrals is provided by the user in the variable input file. A typical value used has been 1000. Before PTpic was modified to more carefully track neutral consumption, a value of 10 was used.

have access to their own memory. The same idea applies in PTpic when electrons are pseudo-randomly selected to ionize neutrals and ions: each process loops through its list of electrons and independently chooses which of its electrons, if any, will participate in ionization reactions. At this stage, the processes have no information about how many ionization events have been selected in each cell by all the other processes, so there is no guarantee that there are enough neutrals or ions present to support all the reactions that have been statistically chosen to occur. The use of superparticles, weighted to represent W real particles, makes this a more important consideration because each simulated ionization event represents W real ionizations.

One solution is to sum the number of reactions each process plans on executing in each cell — requiring a collective communication operation. Next, each event can be approved or canceled after checking if there are enough background superparticles to support the event's occurrence. However, this step must be done sequentially on a single process because each approval could conceivably influence the approval or denial of the next event. Once this is completed, the decisions regarding each processes' list of events are returned, allowing that process to cancel or enforce each event. Only now, for example, would a process be allowed to create a new electron-ion pair resulting from an ionization event. In between neutral iterations, the neutral densities in each cell are updated to reflect changes from ionization and CEX reactions. When it is time to update neutral positions and velocities, the required number of neutrals are deleted from memory in each cell. If discrepancies are detected between expected and needed numbers of neutral superparticles, PTpic should exit with a diagnostic error code and message. This procedure has been debugged extensively in hopes of removing any errors. Most of the examples given in this section refer only to electrons ionizing a background species, specifically neutrals and ions. However, this discussion is equally applicable to CEX collisions involving a primary ion and a background neutral population. In this case, candidate CEX events for ions are identified statistically, approved or denied based on updated super-neutral counts, and then implemented only when approved. For CEX collisions, the ion is converted to a neutral particle, and an ion is created with the bulk neutral velocity because there is no particle paring in this algorithm.

3.3.1 Collisions as a Performance Bottleneck

As could reasonably be expected, the serial event-approval process just described represents a barrier to running PTpic with a large number of processes. The current method of distributing particle data amongst processors is not well suited for calculations requiring data from all particles in a cell — the most notable examples being the calculation of particle moments, and the ionization/CEX procedure outline earlier in this section. All other aspects of the particle moving algorithm scale efficaciously over the range of processors used (< 100), but the collision procedure execution time increases as more processors are used.

An alternative data structure, that would probably scale favorably using larger numbers of processors, entails distributing particles based on their physical locations. For example, groups of co-located cells could be assigned to each process, and all particles within a process' region would belong to that process. Using this concept, each process would be free to move and collide its particles without the need to communicate with other processes. Calculating particle moments in each sub-region could also be done without any communication requirement. This idea is certainly not new, and was used in DSMC simulations using hundreds of millions of neutral particles[152] over 15 years ago. It is likely that the communication costs of exchanging particle data as they move from one process' region to another can be managed so that the overall simulation scales favorably. Incorporating an efficient and scalable method for obtaining potential solutions into a model such as this would also be necessary.

Balancing computational loads between processes will be more complicated if particles are distributed in memory according to their physical location. Presently, each process takes turns injecting neutrals and electrons, ensuring an acceptable load distribution in a trivial way. If the memory distribution of particles is altered as suggested, a method for adjusting each process' domain may be desirable, or necessary. This will prevent regions with significantly higher number densities from slowing down the entire simulation. Load re-distribution will probably have a non-negligible computational cost, so a compromise between tolerated levels of load imbalance and the allowed time between re-distribution actions would be necessary. This topic is also discussed in Ref. 152.

3.4 Particle Injection

In PTpic, neutral particles are injected on the left-hand[n] computational boundary only, between nodes specified in the `IN__BOUNDARIES` input file. They are injected with half of an isotropic distribution function at a temperature specified by the user in the `IN__VARIABLES` input file. Electrons are injected through either the top or right-hand side boundaries, also through nodes provided by the user in `IN__BOUNDARIES`. The temperature at which electrons are injected is specified within the `IN__CONSTANTS` input file. The enforcement of a background neutral pressure is not currently support in PTpic.

The number of electrons injected is calculated based upon the net charge collected by the anode each iteration. If the net charge is positive, meaning more ions hit the anode than electrons, no electrons are injected at the boundaries. The case where ions outnumber electrons is very rare, and has only been noticed a small handful of times out of over 100 million iterations of PTpic operation.

The cathode is not explicitly modeled in PTpic, so depending on the specifications provided by the user, the locations of electron injection may or may not represent their likely injection point from the cathode into the domain. At the start of a plasma simulation, the user has the option to seed a specified density of particles plasma with given electron and ion temperatures. Additionally, or alternatively, the user may elect to seed no particles and rely on electron supplied at the boundaries to start the discharge. For this reason, the user may also specify a sort of "keeper current" in the `IN__VARIABLES` input file that supplies a fixed excess (with respect to the electron current collected at the anode) electron current to the boundary. There are inputs in the same file that tell PTPic to shut-off the extra supply when the anode current exceeds a certain value. However, due to simulation noise — depending on the super-particle size, a super-electron may not reach the anode each timestep — there are also switches that tell PTpic to shutoff the extra electron supply once the simulated counts of ion and/or electron super-particles exceed user-specified thresholds.

[n] These boundaries are described assuming that the z-axis is on the "bottom", and that positive z points to the right.

3.5 Quench Model of for Anomalous Electron Collisions

In PTpic, the quench model can be turned on and off based on user-set variables in the `IN__VARIABLES` input file. The simulated anomalous collision frequency ν_a^* is calculated in Eq. 3.33, where ω_c is the local cyclotron frequency, σ is the shear rate, and the other terms Ω_B, τ_{DC}, and c are also read from the `IN__VARIABLES` input file. In the absence of any shear, the relationship between the Bohm diffusion coefficient and the fit parameter Ω_B is given in Eq. 3.34, and we see that Ω_B is simply the Hall parameter due to Bohm diffusion in the absence of any shear. The parameter τ_{DC} is referred to as the turbulent decorrelation time by Scharfe[58]. Currently, in PTpic, the shear is calculated as the magnitude of the gradient of the azimuthal electron speed. As a point of reference, M. Scharfe uses the fit parameters $\{\Omega_B, \tau_{DC}, c\} = \{8, 16 \times 10^{-8}\ [s], 2\}$ for some of her simulations[58]. This is Scharfe's implementation of the quench model, which is borrowed from the magnetically confined fusion literature. No temporal or spatial averaging of the shear is currently done. One potentially interesting option for smoothing out time-variations of shear in the quench model is suggested and applied by T. Matlock in Sec. 5.3.9.1 of Ref. 35.

$$\nu_a^* = \frac{\omega_c/\Omega_B\sqrt{r}}{1 + (\sigma\tau_{DC})^c} \tag{3.33}$$

$$D_B = \frac{1}{\Omega_B}\frac{T_e}{eB} \tag{3.34}$$

3.6 Ion Scattering Off of Background Neutrals

Representing collisions between ion and neutrals is not an ideal application for the MCC procedure because ions and neutrals have effectively identical masses. Furthermore, low-energy ions may have speeds similar to neutrals, another indicator that DSMC would be the better algorithm to use for representing ion-neutral elastic scattering collisions. Nonetheless, this section describes an attempt to implement these collisions using the MCC procedure.

Between iterations in the simulation that call for neutrals to be moved, elastic scattering events for ions with background neutrals are determined within in each cell, during each iteration. The relative speed between an ion and a representative neutral is calculated from the

difference between the ion super-particle's velocity $\underline{w}_i$ and the bulk neutral velocity $\underline{u}_n$ — note that no individual neutral particles are selected and used to calculate a collision probability, as that would constitute an implementation of the DSMC method. At this point, an impact parameter for each collision that occurs is pseudo-randomly generated and the ion velocity is scattered according to a variable hard-sphere model of the ion-neutral interaction[o]. The resulting momentum and energy changes for each ion superparticle undergoing these collisions contributes to a net change of momentum and energy in each cell. Thus, for each cell, total amounts of velocity $\Delta\underline{V}$ and energy ΔE that should be imparted to neutral particles in that same cell — in order to conserve momentum and energy — are stored until the neutral particles are iterated forward in time.

When neutrals are moved, the velocity of each neutral particle $\underline{w}_n$ in a given cell is altered as shown in Eq. 3.35 to represent the transfer of momentum and energy between ions and neutrals resulting from their elastic interactions. Here, $\Delta\underline{u} \equiv {}^{\Delta\underline{V}}/_{N}$, and N is the number of neutral superparticles remaining in the cell. The unit vector $\hat{n}$ is formed from two pseudo-randomly generated angles, and c represents the magnitude of the random velocity change provided to each neutral superparticle in the cell — it is defined in Eq. 3.36. Therefore, momentum and energy are conserved in an average sense: to the extent that the average value of $c\hat{n}$ amongst all neutrals in the cell is equal to zero.

$$\underline{w}_{n,NEW} = \underline{w}_{n,OLD} + \Delta\underline{u} + c\hat{n} \tag{3.35}$$

$$c^2 = \frac{2\Delta E}{mN} - \left[|\underline{u}|^2 + 2\left(\underline{u}\cdot\underline{u}_n\right)\right] \tag{3.36}$$

The main drawback of this method, aside from the fact that these collisions would be better represented by a DSMC procedure, is that the quantity c^2 is not *always* greater than zero. The simplest example of when this occurs is when neutrals transfer a net positive amount of energy to ions via collisions, meaning that $\Delta E < 0$. In practice, this does not happen often, but when it does occur, PTpic does not scatter neutrals in the cell for which $c^2 < 0$, so momentum and energy are not conserved over the relevant time interval, in that particular cell. From observing the output of the code, which displays a warning message when this situation

[o]The total cross section used for ion-neutral scattering collisions is described in Appendix D

arises, this scenario occurs infrequently relative to the total number of times it is applied. Additionally, the magnitude of momentum and energy conservation violation, resulting from cases where $c^2 < 0$, is summed for the entire simulation duration — over the whole domain — and displayed after the last iteration has finished. This information can help users understand if this procedure is appropriate for their application, and aid in future decisions to alter how ion-neutral scattering events are implemented in PTpic.

Chapter 4

Potential Solver

In order to calculate the electric field $\underline{\mathrm{E}}$, the potential ϕ at each node point $[k, j]$ is determined first. Once all ϕ values are known, the radial component of $\underline{\mathrm{E}}$ can be determined using Eq. 4.1. The expression for the axial component is obtained by replacing r with z, while the azimuthal component is zero by assumption. Shorthand notation is defined as well. The derivatives of the computational coordinates $[\xi, \eta]$ with respect to r and z are pre-computed on the structured mesh when the simulation is initialized. The derivatives of ϕ with respect to ξ and η are approximated using finite differences of ϕ values at surrounding nodes.

$$E_r = -\frac{\delta\xi}{\delta r}\frac{\delta\phi}{\delta\xi} - \frac{\delta\eta}{\delta r}\frac{\delta\phi}{\delta\eta} \equiv -\xi_r\phi_\xi - \eta_r\phi_\eta \tag{4.1}$$

To actually solve for ϕ, linear algebraic equations relating potential values at adjacent nodes are obtained by applying the integral form of Gauss' law to each node, provided in Eq. 4.2 — expressed in SI units. Here, S_V is a closed surface bounding a volume V, both of which contain or enclose a net charge Q. The permittivity factor *gamma* must be≥ 1 — as determined by computational resources — and has been left on the RHS in Eq. 4.2 to emphasize the sensitivity to net space charge that is lost by using large values of γ.

$$\int_{S_V} \epsilon\underline{\mathrm{E}} \cdot \mathrm{d}\underline{\mathrm{S}} = \frac{Q}{\gamma^2} \tag{4.2}$$

In practice, the code uses the CGS expression of Eq. 4.2 as a starting point. From there, it is reexpressed in a normalized form, given in Eq. 4.3. As with all quantities throughout this

document, quantities expressed in SI units are the default, CGS quantities are denoted by an overbar, and normalized quantities are indicated by the use of a caret. The normalized Gauss' law is applied to each node, with the control surface illustrated in Fig.4-1. This is referred to as the 9-point method, because when node $[k, j]$ is not on a boundary, the potentials at 9 points end up contributing to the LHS of Eq. 4.3, referred to as $\hat{I}_{k,j}$ in the following discussion.

$$\hat{I}_{k,j} \equiv \int_{S_V} \hat{\epsilon}\hat{\underline{E}} \cdot \underline{d\hat{S}} = \hat{Q}_{k,j} \tag{4.3}$$

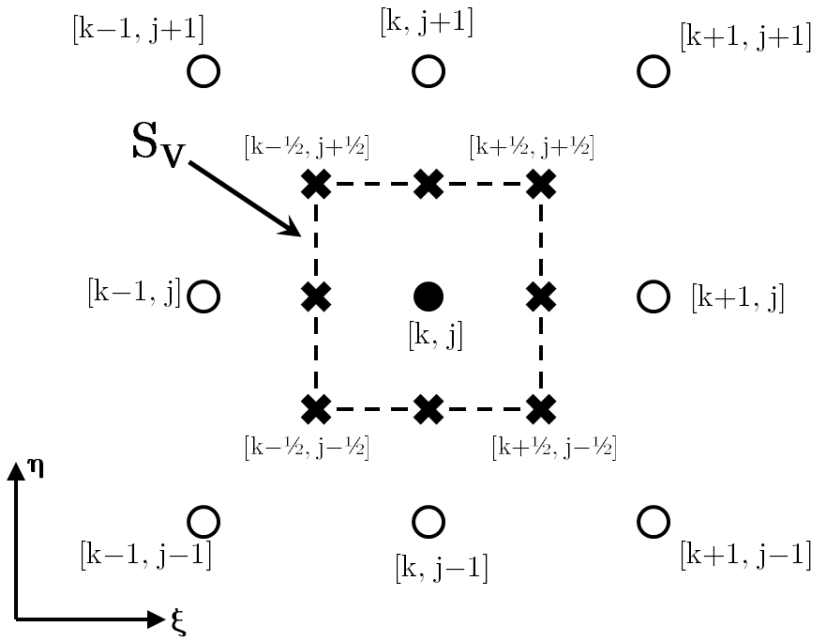

Figure 4-1: Illustration of the control surface S_V cross section in computational space [ξ,η].

Disregarding nodes on boundaries for now, S_V is comprised of 8 smaller surfaces, with 2 residing in each quadrilateral mesh element, or cell, surrounding node $[k, j]$. Using a more systematic notation simplifies the expressions for each integral contribution and allows boundary conditions to be handled in a straightforward manner. The descriptors for nodes and cells surrounding each node are summarized in Tables 4.1 and 4.2. Labels with all capital letters are reserved for mesh elements, while nodes are represented by lowercase letters. As an example, the NE mesh element is specified by the values $[\Delta k, \Delta j] = [+1, +1]$. With this in mind, the

integral $\hat{I}_{k,j}$ is divided into contributions from the four surrounding mesh elements, as shown explicitly in Eq. 4.4. Now, the problem has been reduced to finding an expression for the simpler, but more general, quantity: $\hat{I}_{\Delta k,\Delta j}$. This involves an integral over only two surfaces in a single cell, rather than eight. The definition of $\hat{I}_{\Delta k,\Delta j}$ is provided in Eq. 4.5, and illustrated in Fig. 4-2 — where H is the computationally horizontal surface, and V is the computationally vertical surface. Here, $\hat{\epsilon}_{\Delta k,\Delta j}$ is relative permittivity of the region contained in the mesh element. This allows the potential solver to extend into regions with different permittivities, such as insulators, as long as charges trapped on surfaces of the insulator contribute to $\hat{Q}_{k,j}$.

Table 4.1: Labels for nodes in terms of the changes in k and j, relative to the "center" node at $[k, j]$.

Node Labels			
+1	nw	n	ne
0	w	c	e
-1	sw	s	se
$\Delta j/\Delta k$	-1	0	+1

Table 4.2: Labels for mesh elements in terms of the changes in k and j, relative to the "center" node at $[k, j]$.

Cell Labels		
+1	NW	NE
-1	SW	SE
$\Delta j/\Delta k$	-1	+1

$$\hat{I}_{k,j} \equiv \sum_{\Delta k=-1,+1} \sum_{\Delta j=-1,+1} \hat{I}_{\Delta k,\Delta j} = \hat{I}_{NE} + \hat{I}_{SE} + \hat{I}_{SW} + \hat{I}_{NW} \tag{4.4}$$

$$\hat{I}_{\Delta k,\Delta j} \equiv \int_{H+V} \hat{\epsilon}_{\Delta k,\Delta j} \underline{\hat{E}} \cdot \underline{d\hat{S}} \tag{4.5}$$

Now, a relationship between the node potentials and space charge can be developed by expressing $\hat{I}_{\Delta k,\Delta j}$ in terms of the area-averaged components of $\underline{\hat{E}}$ and the projections of the normalized surface areas, as shown in Eq. 4.6. Here, the subscripts H and V indicate which

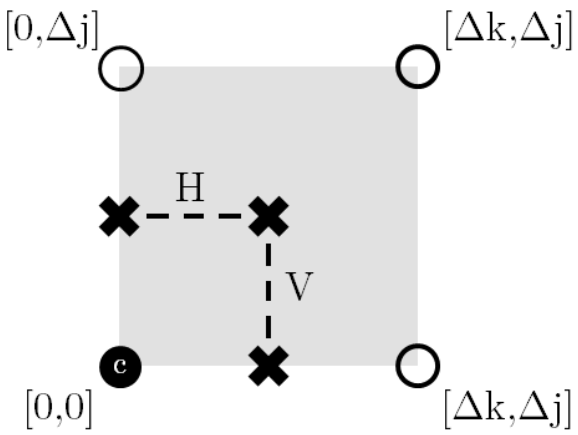

Figure 4-2: Canonical mesh element used to evaluate $\hat{I}_{\Delta k,\Delta j}$.

surface the radial or axial component has been averaged over. These average components of $\underline{\mathrm{E}}$ are evaluated in terms of their values at the four surrounding nodes using bilinear interpolation. PTpic interpolates function values within cells using Eq. 4.7. A cell is called out by the computational coordinates of the node in its lower-left corner (a pair of numbers $[k, j]$). Thus, *cell* $[k, j]$ is bounded by the surfaces connecting nodes $[k, j]$, $[k, j+1]$, $[k+1, j+1]$, and $[k+1, j]$. In Eq. 4.7, α_k and α_j are defined as $\xi - k$ and $\eta - j$, respectively, and have values $\in [0, 1]$. The subscripts on the function values are chosen such that $f_{0,1} \equiv f(0, 1)$, which is also equivalent to sampling f at node $[k + 0, j + 1]$.

$$\hat{I}_{\Delta k,\Delta j} = \hat{\epsilon}_{\Delta k,\Delta j} \left(\hat{\underline{\mathrm{E}}}_{r,H} \hat{A}_{r,H} + \hat{\underline{\mathrm{E}}}_{z,H} \hat{A}_{z,H} + \hat{\underline{\mathrm{E}}}_{r,V} \hat{A}_{r,V} + \hat{\underline{\mathrm{E}}}_{z,V} \hat{A}_{z,V} \right) \tag{4.6}$$

$$\begin{aligned} f(\alpha_k, \alpha_j) &\approx (1 - \alpha_k)(1 - \alpha_j) f_{0,0} + (1 - \alpha_k)\alpha_j f_{0,1} + \alpha_k (1 - \alpha_j) f_{1,0} + \alpha_k \alpha_j f_{1,1} \\ &= f_{0,0} + \alpha_k (f_{1,0} - f_{0,0}) + \alpha_j (f_{0,1} - f_{0,0}) + \alpha_k \alpha_j \left[f_{0,0} + f_{1,1} - (f_{0,1} + f_{1,0}) \right] \end{aligned} \tag{4.7}$$

$$\begin{aligned} \langle f \rangle_{H,cart} &= 2 \int_0^{1/2} f\left(\alpha_k, 1/2\right) \, d\alpha_k \\ &= \frac{3}{8} \left(f_c + f_{0,\Delta j} \right) + \frac{1}{8} \left(f_{\Delta k,0} + f_{\Delta k,\Delta j} \right) \end{aligned} \tag{4.8}$$

Using the assumed form of functions within a cell, averages of any function whose values are

known at the four surrounding nodes can be evaluated once a surface is specified. Originally, these averages were evaluated using Eq. 4.8. The equation for $\langle f \rangle_{V,cart}$ is obtained by exchanging the subscripts $0, \Delta j$ and $\Delta k, 0$ in Eq. 4.8. However, this did not fully account for cylindrical geometry because values further from the z-axis are not given more weight in the average. The cylindrical average is provided in Eq. 4.9. Again, the expression for $\langle f \rangle_V$ can be obtained by exchanging the subscripts $0, \Delta j$ and $\Delta k, 0$ for the values of the function f and the radial coordinate r. Notice that Eq. 4.9 reduces to 4.8 when all radial coordinates are identical.

$$\begin{aligned} \langle f \rangle_H &= \frac{\int_0^{1/2} f\left(\alpha_k, 1/2\right) \hat{r}\left(\alpha_k, 1/2\right) \; d\alpha_k}{\int_0^{1/2} \hat{r}\left(\alpha_k, 1/2\right) \; d\alpha_k} \\ &= \frac{a_H}{c_H}\left(f_c + f_{0,\Delta j}\right) + \frac{b_H}{c_H}\left(f_{\Delta k,0} + f_{\Delta k,\Delta j}\right) \end{aligned} \tag{4.9}$$

$$a_H \equiv 7\left(\hat{r}_c + \hat{r}_{0,\Delta j}\right) + 2\left(\hat{r}_{\Delta k,0} + \hat{r}_{\Delta k,\Delta j}\right)$$

$$b_H \equiv 2\left(\hat{r}_c + \hat{r}_{0,\Delta j}\right) + \hat{r}_{\Delta k,0} + \hat{r}_{\Delta k,\Delta j}$$

$$c_H \equiv 6\left[3\left(\hat{r}_c + \hat{r}_{0,\Delta j}\right) + \hat{r}_{\Delta k,0} + \hat{r}_{\Delta k,\Delta j}\right]$$

Although Eq. 4.9, combined with Eq. 4.6 — where either component of $\hat{\underline{\mathrm{E}}} \rightarrow f$ and including the surface V — accounts fully for cylindrical geometry, the dominant cylindrical correction on a fine mesh is contained in the components of the surface areas, as opposed to the non-Cartesian function averages. This is because the mesh is typically more uniform and orthogonal near the axis of symmetry. These corrections would have been less important in previous applications — which dealt exclusively with Hall thrusters, whose chambers are annular. The correction provided by the new expressions can be evaluated directly, and is provided in Eq. 4.10, again for the H surface. The equation for sufrace V is obtained by exchanging appropriate indices, as described previously. As an example, for cells on the axis of symmetry — meaning $\hat{r}_c = \hat{r}_{\Delta k,0} = 0$ — a good approximation describing the shape of the mesh can be expressed as: $r_{0,\Delta j} \approx r_{\Delta k,\Delta j} \approx \Delta r$, where Δr just represents the mesh spacing near the axis of symmetry. With these simplifications, the approximate differences between the cartesian and cylindrical averages is provided in Eq 4.11, showing that the cylindrical averaging correction can be small — even where it might have been thought to be the most significant. However, the accuracy of the 9-point method, now accounting fully for cylindrical geometry, has been improved overall —

which is important when applied near the $r = 0$ axis, particularly in the DCFT. No additional computational penalty is incurred because the geometry dependent constants are all evaluated before the main simulation loop begins.

$$\langle f \rangle_H = \langle f \rangle_{H,cart} + \frac{(f_{\Delta k,0} - f_c) + (f_{\Delta k,\Delta j} - f_{0,\Delta j})}{24} \cdot \frac{(\hat{r}_{\Delta k,0} - \hat{r}_c) + (\hat{r}_{\Delta k,\Delta j} - \hat{r}_{0,\Delta j})}{3(\hat{r}_c + \hat{r}_{0,\Delta j}) + \hat{r}_{\Delta k,0} + \hat{r}_{\Delta k,\Delta j}} \tag{4.10}$$

$$\begin{aligned} \langle f \rangle_H - \langle f \rangle_{H,cart} &\approx 0 \\ \langle f \rangle_V - \langle f \rangle_{V,cart} &\approx \frac{(f_{0,\Delta j} - f_c) + (f_{\Delta k,\Delta j} - f_{\Delta k,0})}{24} \end{aligned} \tag{4.11}$$

Up to this point, the only approximation introduced has been the interpolation scheme. Now, to represent the fields at each node, the derivatives of potential on the mesh, present in Eq. 4.1, are represented in terms of finite differences between values at the four nodes defining the cell corners. The finite differences utilized at each node are provided in Eq. 4.12 — where all subscripts specify which node the normalized potential $\hat{\phi}$, or gradient in computational space ($\hat{\phi}_\xi$ or $\hat{\phi}_\eta$), is sampled at. Higher-order finite difference schemes could be used, but would make the method more computationally expensive and complicate the handling of boundary conditions. This is especially true for nodes near boundaries where the free-space permittivity changes, as gradients in potential are discontinuous at these points. Furthermore, the present scheme has demonstrated satisfactory levels of accuracy. The partial derivatives of the computational coordinates $[\xi, \eta]$, also present in Eq. 4.1, are each evaluated with respect to the two spatial

coordinates (z, r) using second-order finite difference schemes.

$$\begin{aligned}
\hat{\phi}_{\xi,c} &\approx \Delta k(\hat{\phi}_{\Delta k,0} - \hat{\phi}_c) \\
\hat{\phi}_{\xi,0,\Delta j} &\approx \Delta k(\hat{\phi}_{\Delta k,\Delta j} - \hat{\phi}_{0,\Delta j}) \\
\hat{\phi}_{\eta,c} &\approx \Delta j(\hat{\phi}_{0,\Delta j} - \hat{\phi}_c) \\
\hat{\phi}_{\eta,\Delta k,0} &\approx \Delta j(\hat{\phi}_{\Delta k,\Delta j} - \hat{\phi}_{\Delta k,0}) \\
\hat{\phi}_{\xi,\Delta k,0} &\equiv \hat{\phi}_{\xi,c} \\
\hat{\phi}_{\xi,\Delta k,\Delta j} &\equiv \hat{\phi}_{\xi,0,\Delta j} \\
\hat{\phi}_{\eta,0,\Delta j} &\equiv \hat{\phi}_{\eta,c} \\
\hat{\phi}_{\eta,\Delta k,\Delta j} &\equiv \hat{\phi}_{\eta,\Delta k,0}
\end{aligned} \tag{4.12}$$

Finally, an expression for $\hat{I}_{\Delta k,\Delta j}$ can be obtained. A simplified form is provided in Eq. 4.13, with each coefficient defined in Eqs. 4.15-4.18. Note the shorthand notation for the partial derivatives, first demonstrated in Eq. 4.1, and that the partial derivatives are with respect to the *normalized* spatial coordinates ($\hat{z}$,$\hat{r}$) — not the dimensional coordinates (z,r). Mathematica[150] was used to derive these expressions due to the large amount of algebra involved. The previously used cartesian constants are obtained by again substituting $\frac{a_H}{c_H} = \frac{a_V}{c_V} = \frac{3}{8}$ and $\frac{b_H}{c_H} = \frac{b_V}{c_V} = \frac{1}{8}$. The choice of sign for non-"center" node coefficients was made to have the terms be additive when used in the iterative scheme for determining the potential distribution — see Eq. 4.26. The meaning of the subscripts are different for the coefficients $[C, N, NE, E]$ and node potentials $[\phi]$. For coefficients (and also for $\hat{I}_{\Delta k,\Delta j}$), subscripts specify which *cell* is contributing to the total integral $\hat{I}_{k,j}$ — just as in in Table 4.2. Subscripts for node potentials and computational coordinate derivatives have the same meaning as in Table 4.1, where they specify the sampling position relative to node $[k, j]$.

$$\hat{I}_{\Delta k,\Delta j} = C_{\Delta k,\Delta j}\hat{\phi}_c - N_{\Delta k,\Delta j}\hat{\phi}_{0,\Delta j} - (NE)_{\Delta k,\Delta j}\hat{\phi}_{\Delta k,\Delta j} - E_{\Delta k,\Delta j}\hat{\phi}_{\Delta k,0} \tag{4.13}$$

Summing the contributions from each surrounding cell, we arrive at Eq. 4.14 — giving an approximate relationship between the potentials at each node and the charge contained in the control volume around it. In the definitions of the total coefficients (distinguished by their lack

of subscripts, though they might be better identified with a subscript $[k,j]$, to associate them with node $[k,j]$, as with $\hat{Q}_{k,j}$), subscripts for terms on the RHS again specify which cell each contribution comes from. An equation similar to Eq. 4.14 applies to each node on the mesh for which the potential is unknown. However, before discussing the linear system as a whole, or methods for solving it, the treatment of boundary conditions must be addressed.

$$C\hat{\phi}_c - N\hat{\phi}_n - (NE)\hat{\phi}_{ne} - E\hat{\phi}_e - (SE)\hat{\phi}_{se} - S\hat{\phi}_s - (SW)\hat{\phi}_{sw} - W\hat{\phi}_w - (NW)\hat{\phi}_{nw} = \hat{Q}_{k,j} \tag{4.14}$$

$$\begin{aligned}
C &= C_{+1,+1} + C_{+1,-1} + C_{-1,-1} + C_{-1,+1} \\
N &= N_{+1,+1} + N_{-1,+1} \\
(NE) &= (NE)_{+1,+1} \\
E &= E_{+1,+1} + E_{+1,-1} \\
(SE) &= (NE)_{+1,-1} \\
S &= N_{+1,-1} + N_{-1,-1} \\
(SW) &= (NE)_{-1,-1} \\
W &= E_{-1,-1} + E_{-1,+1} \\
(NW) &= (NE)_{-1,+1}
\end{aligned}$$

$$\begin{aligned}
C_{\Delta k,\Delta j} = \hat{\epsilon}_{\Delta k,\Delta j} \Bigg\{ & \hat{A}_{r,H} \left[\frac{1}{c_H} \Delta k \left(a_H \xi_{\hat{r},c} + b_H \xi_{\hat{r},\Delta k,0} \right) + \frac{a_H}{c_H} \Delta j \left(\eta_{\hat{r},c} + \eta_{\hat{r},0,\Delta j} \right) \right] \\
& + \hat{A}_{z,H} \left[\frac{1}{c_H} \Delta k \left(a_H \xi_{\hat{z},c} + b_H \xi_{\hat{z},\Delta k,0} \right) + \frac{a_H}{c_H} \Delta j \left(\eta_{\hat{z},c} + \eta_{\hat{z},0,\Delta j} \right) \right] \\
& + \hat{A}_{r,V} \left[\frac{a_V}{c_V} \Delta k \left(\xi_{\hat{r},c} + \xi_{\hat{r},\Delta k,0} \right) + \frac{1}{c_V} \Delta j \left(a_V \eta_{\hat{r},c} + b_V \eta_{\hat{r},0,\Delta j} \right) \right] \\
& + \hat{A}_{z,V} \left[\frac{a_V}{c_V} \Delta k \left(\xi_{\hat{z},c} + \xi_{\hat{z},\Delta k,0} \right) + \frac{1}{c_V} \Delta j \left(a_V \eta_{\hat{z},c} + b_V \eta_{\hat{z},0,\Delta j} \right) \right] \Bigg\}
\end{aligned} \tag{4.15}$$

$$N_{\Delta k,\Delta j} = -\hat{\epsilon}_{\Delta k,\Delta j}\Bigg\{ \hat{A}_{r,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{r},0,\Delta j}+b_H\xi_{\hat{r},\Delta k,\Delta j}\right)-\frac{a_H}{c_H}\Delta j\left(\eta_{\hat{r},c}+\eta_{\hat{r},0,\Delta j}\right)\right] \tag{4.16}$$
$$+\hat{A}_{z,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{z},0,\Delta j}+b_H\xi_{\hat{z},\Delta k,\Delta j}\right)-\frac{a_H}{c_H}\Delta j\left(\eta_{\hat{z},c}+\eta_{\hat{z},0,\Delta j}\right)\right]$$
$$+\hat{A}_{r,V}\left[\frac{b_V}{c_V}\Delta k\left(\xi_{\hat{r},0,\Delta j}+\xi_{\hat{r},\Delta k,\Delta j}\right)-\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{r},c}+b_V\eta_{\hat{r},0,\Delta j}\right)\right]$$
$$+\hat{A}_{z,V}\left[\frac{b_V}{c_V}\Delta k\left(\xi_{\hat{z},0,\Delta j}+\xi_{\hat{z},\Delta k,\Delta j}\right)-\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{z},c}+b_V\eta_{\hat{z},0,\Delta j}\right)\right]\Bigg\}$$

$$(NE)_{\Delta k,\Delta j} = \hat{\epsilon}_{\Delta k,\Delta j}\Bigg\{ \hat{A}_{r,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{r},0,\Delta j}+b_H\xi_{\hat{r},\Delta k,\Delta j}\right)+\frac{b_H}{c_H}\Delta j\left(\eta_{\hat{r},\Delta k,0}+\eta_{\hat{r},\Delta k,\Delta j}\right)\right] \tag{4.17}$$
$$+\hat{A}_{z,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{z},0,\Delta j}+b_H\xi_{\hat{z},\Delta k,\Delta j}\right)+\frac{b_H}{c_H}\Delta j\left(\eta_{\hat{z},\Delta k,0}+\eta_{\hat{z},\Delta k,\Delta j}\right)\right]$$
$$+\hat{A}_{r,V}\left[\frac{b_V}{c_V}\Delta k\left(\xi_{\hat{r},0,\Delta j}+\xi_{\hat{r},\Delta k,\Delta j}\right)+\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{r},\Delta k,0}+b_V\eta_{\hat{r},\Delta k,\Delta j}\right)\right]$$
$$+\hat{A}_{z,V}\left[\frac{b_V}{c_V}\Delta k\left(\xi_{\hat{z},0,\Delta j}+\xi_{\hat{z},\Delta k,\Delta j}\right)+\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{z},\Delta k,0}+b_V\eta_{\hat{z},\Delta k,\Delta j}\right)\right]\Bigg\}$$

$$E_{\Delta k,\Delta j} = \hat{\epsilon}_{\Delta k,\Delta j}\Bigg\{ \hat{A}_{r,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{r},c}+b_H\xi_{\hat{r},\Delta k,0}\right)-\frac{b_H}{c_H}\Delta j\left(\eta_{\hat{r},\Delta k,0}+\eta_{\hat{r},\Delta k,\Delta j}\right)\right] \tag{4.18}$$
$$+\hat{A}_{z,H}\left[\frac{1}{c_H}\Delta k\left(a_H\xi_{\hat{z},c}+b_H\xi_{\hat{z},\Delta k,0}\right)-\frac{b_H}{c_H}\Delta j\left(\eta_{\hat{z},\Delta k,0}+\eta_{\hat{z},\Delta k,\Delta j}\right)\right]$$
$$+\hat{A}_{r,V}\left[\frac{a_V}{c_V}\Delta k\left(\xi_{\hat{r},c}+\xi_{\hat{r},\Delta k,0}\right)-\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{r},\Delta k,0}+b_V\eta_{\hat{r},\Delta k,\Delta j}\right)\right]$$
$$+\hat{A}_{z,V}\left[\frac{a_V}{c_V}\Delta k\left(\xi_{\hat{z},c}+\xi_{\hat{z},\Delta k,0}\right)-\frac{1}{c_V}\Delta j\left(a_V\eta_{\hat{z},\Delta k,0}+b_V\eta_{\hat{z},\Delta k,\Delta j}\right)\right]\Bigg\}$$

4.1 Boundary Conditions

Nodes on the mesh can be divided into two broad categories, "knowns" and "unknowns". PTpic allows for the potential at a node to be specified as one of three known values: the anode potential ϕ_a, cathode potential ϕ_{cath}, or a floating body potential ϕ_f. Section 4.1.2 discusses floating body potentials in more detail. Aside from being classified based on whether or not the potential at a node is specified, it is also important to alter the treatment of nodes located on external boundaries of the computational mesh. In these cases, the number of terms in the

governing equation of the node must be altered because some nodes identified in Eq. 4.14 will not be part of the mesh.

In addition to nodes falling into different categories, cells are distinguished by their permittivity values, and by whether they are within a body modeled as a perfect conductor (i.e. the anode or floating body). In its present state, PTpic allows for a single value of permittivity different than in free space, reserved to represent the permittivity of dielectric bodies. The total charge accumulated on all dielectric surfaces is tracked and included in the normalized node charge $Q_{k,j}$. On borders between cells with differing permittivity values, potential gradients are approximated by finite difference schemes that do not mix values on either side of the boundary surface. The electric field is only used to calculate forces on charged particles, so finite differencing is only performed amongst nodes contacting cells with the free-space value of permittivity. This is especially important for field calculations near the interface of a dielectric body and a plasma.

4.1.1 Governing Potential Equations for Nodes Near Boundaries

Boundary conditions for nodes on external boundaries of the mesh were previously provided by the omission of appropriate terms in Eq. 4.14. For example, a node on the LHS computational boundary has no cells to the left (west) of it, so the contributions from the SW or NW cells are *assumed* to be zero. This corresponds to the pairs of $[\Delta k, \Delta j]$ values in Table 4.2 of $[-1,-1]$ and $[-1,+1]$. As a result, the governing equation for this examples lack terms proportional to $\hat{\phi}_{sw}$, $\hat{\phi}_{w}$, and $\hat{\phi}_{nw}$.

Physically, this method for handling the boundary conditions forces values of displacement ($\epsilon \underline{\mathrm{E}}$) flux through external mesh boundaries to zero. The result for the electric field is that no field lines intersect the mesh boundary, except where the cathode potential is specified. In other words, selecting this boundary condition forces field lines tangent to the mesh boundary, except where the potential is specified.

The same method for handling external mesh boundaries is applied to cells inside bodies modeled as perfect conductors (i.e. the anode and floating body). This is appropriate because the contribution to $\hat{I}_{k,j}$ is zero in these bodies, as there can be no electric field within a perfect conductor. The important difference is that node potentials on these body surfaces

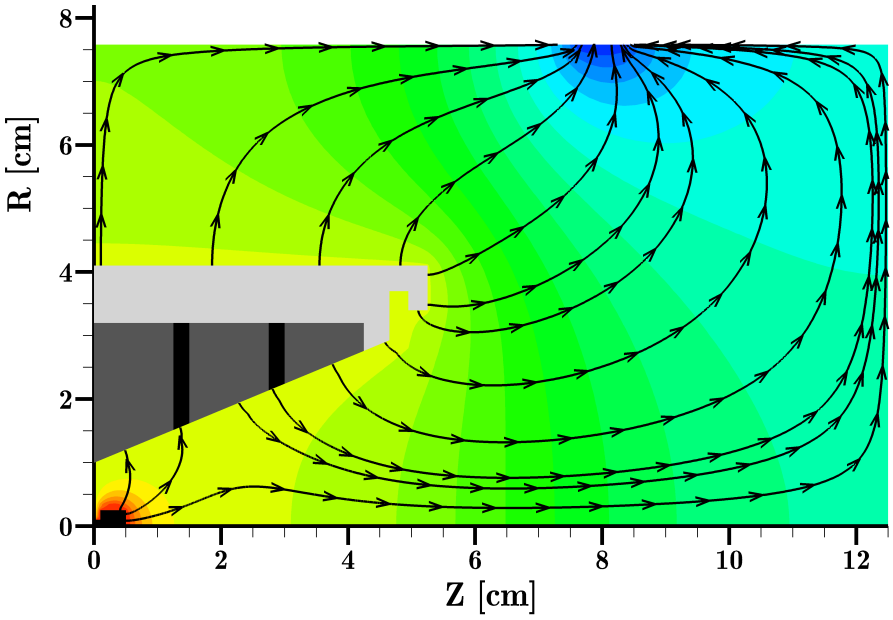

Figure 4-3: Potential solution with electric filed lines overlayed. This plot serves to demonstrate visually the boundary conditions of the potential solver — excluding the modifications described in the next section. The lines outlining the two dielectric bodies have been removed from this plot.

are not determined within the system of equations for each node. The anode potential is a given parameter, and the floating body potential is determined based on its accumulated net charge (see Sec. 4.1.2). Therefore, the anode, cathode and floating body potentials appear only as *known* quantities in equations for surrounding nodes with unknown potentials. Additionally, electric fields must be perpendicular to outer surfaces of perfect conductors because all nodes on their surfaces are set to equal potential values. Figure 4-3 shows an example potential solution with no charged particles presents, illustrating the expected behavior of E-field lines. The boundary condition extension described in the following section was not applied in this case.

Boundary Condition Extension for PTpic

Rather than assuming the electric field normal to the mesh boundaries is zero, a feature new to PTpic uses an analytical estimate of the finite differences in potential beyond the mesh to allow for an oblique electric field with respect to mesh boundaries. The estimate is based entirely upon a superposition of the electrostatic potential ϕ due to the presence of two conducting bodies. To simplify expressions, the bodies are assumed spherical, and distances between a

given point and the body are approximated by the distance between a point and the center of the body.

The starting point of this model for the boundary conditions is provided in Eq. 4.19. Here, $\hat{\phi}(\underline{\mathrm{r}})$ is the potential at some point on or beyond the boundary mesh at a specified location $\underline{\mathrm{r}}$. The first term describes a Coulomb potential profile beyond a sphere of radius L_f, approximating the presence of a floating body at a potential $\hat{\phi}_f$ centered about $\underline{R}_f$. The next term represents the cathode located at $\underline{R}_{cath}$ with a size L_{cath}. The potential $\hat{\phi}^*$ is chosen such that $\hat{\phi}(\underline{R}_{cath}) = \hat{\phi}_{cath}$, where $\hat{\phi}_{cath}$ is the dimensionless cathode potential. Currently, PTpic requires $\hat{\phi}_{cath} = 0$ in order to avoid considering the mutual capacitance of the floating body and cathode, though this extended boundary condition is valid for $\hat{\phi}_{cath} \neq 0$. For brevity, $d_f \equiv |\underline{r} - \underline{R}_f|$, $d_{cath} \equiv |\underline{r} - \underline{R}_{cath}|$, and $d_{fc} \equiv |\underline{R}_{cath} - \underline{R}_f|$, allowing $\hat{\phi}(\underline{\mathrm{r}})$ to be expressed as in Eq. 4.20 — using the previously mentioned choice of $\hat{\phi}^*$. From here, finite differences between potential values *beyond* the mesh and on the mesh boundary are estimated with $\Delta\hat{\phi} \approx \hat{\nabla}\hat{\phi} \cdot (\Delta\underline{\hat{r}})_{mesh}$ — where $\Delta\underline{\hat{r}}_{mesh}$ is the estimated dimensionless displacement from a node located on an external boundary to the next node in the ξ or η direction. For example, the z-component of the dimensionless displacement vector between a node on the RHS boundary of the mesh and the "next" node beyond the mesh is found by extrapolation: $\Delta\hat{z} = \frac{\partial\hat{z}}{\partial\xi}$ — by setting $\Delta\xi \equiv 1$, as it is between all neighboring nodes with identical η coordinates. The partial derivatives of dimensionless spatial coordinates on the mesh are evaluated and stored prior to the main simulation loop of the program, and are utilized in many important algorithms in PTpic. Finally, Eq. 4.21 provides estimates of finite potential differences at nodes located on external boundaries in terms of $\hat{\phi}_f$ and $\hat{\phi}_{cath}$.

$$\hat{\phi}(\underline{\mathrm{r}}) \approx \frac{\hat{\phi}_f L_f}{L_f + |\underline{r} - \underline{R}_f|} - \frac{\hat{\phi}^* L_{cath}}{L_{cath} + |\underline{r} - \underline{R}_{cath}|} \tag{4.19}$$

$$\hat{\phi}(\underline{\mathrm{r}}) \approx \hat{\phi}_f \left[\left(1 + \frac{d_f}{L_f}\right)^{-1} - \left(1 + \frac{d_{fc}}{L_f}\right)^{-1} \left(1 + \frac{d_{cath}}{L_{cath}}\right)^{-1} \right] + \hat{\phi}_{cath} \left(1 + \frac{d_{cath}}{L_{cath}}\right)^{-1} \tag{4.20}$$

$$\Delta\hat{\phi} \approx \left[\left(1 + \frac{d_{fc}}{L_f}\right)^{-1} \left(1 + \frac{d_{cath}}{L_{cath}}\right)^{-2} \frac{\Delta d_{cath}}{L_{cath}} - \left(1 + \frac{d_f}{L_f}\right)^{-2} \frac{\Delta d_f}{L_f} \right] \hat{\phi}_f - \left(1 + \frac{d_{cath}}{L_{cath}}\right)^{-2} \left(\frac{\Delta d_{cath}}{L_{cath}}\right) \hat{\phi}_{cath} \tag{4.21}$$

As with all other terms that depend on geometry, the coefficients for $\hat{\phi}_f$ and $\hat{\phi}_{cath}$ are

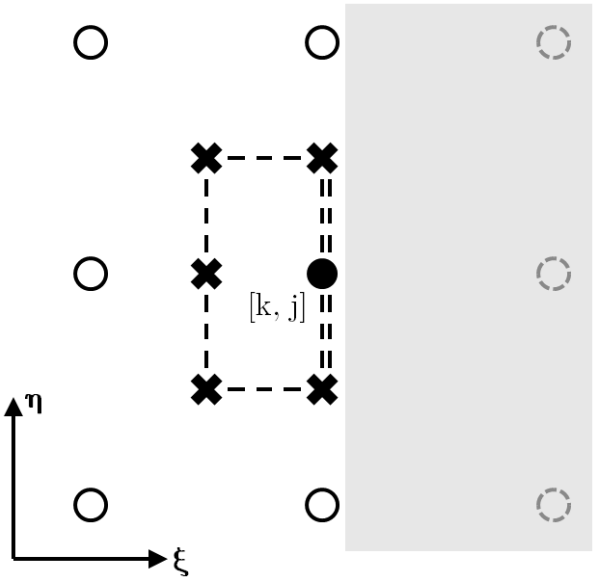

Figure 4-4: Example showing new surfaces of integration, allowing for a non-zero electric field beyond the mesh. The grey shaded area represents the region beyond the mesh. The double-dotted lines represent surfaces not present prior to this work. The potential differences between the nodes on the RHS boundary and the dotted nodes beyond the mesh are estimated using equation Eq. 4.21.

evaluated prior to the start of the simulation, allowing $\Delta\hat{\phi}$ in Eq. 4.21 to be expressed as $\Delta\hat{\phi} = a_f\hat{\phi}_f - a_{cath}\hat{\phi}_{cath}$. With the coefficients a_f and a_{cath}, the displacement field flux can be estimated through open surfaces on or beyond the mesh. Figure 4-4 illustrates an example case, showing a node on the RHS computational boundary.

Prior to this new boundary condition, only the singly dotted surfaces contributed to the integral in Gauss' law, so the flux through all other surfaces was implicitly zero. With this extension, the doubly dotted surfaces also contribute to the integral, supplementing the terms present in Eq. 4.14. Repeating the process for expressing the displacement flux through each new surface in terms of the relevant potentials, Eq. 4.22 is obtained for nodes lying on external mesh boundaries. Some terms are not shown explicitly because the ones that do contribute are dependent upon whether the node is on the left, top, or right boundary (nodes at corners have even fewer terms). The potential equation for nodes on the bottom boundary ($r = 0$) is unaffected by this boundary condition extension. Also, some coefficients for potentials surrounding a node on a left, top, or right boundary are also modified by the allowance of electric fields beyond the mesh — as described in the supplemental definitions in Eq. 4.22.

Complete descriptions of each terms are not necessary to convey the functionality of this expanded boundary condition. This equation for the potential of a node located on the mesh boundary has the same form as Eq. 4.14, which applies to internal nodes. However, the net effect of the new boundary conditions is to add terms dependent on the floating body and cathode potentials — which arise from estimates of electric fields beyond the mesh. If further details are needed, the brief list of items below may be helpful, but the full details are present in the source code of PTpic.

(a) Subscripts after the subscripted comma in a_f and a_{cath} specify which node the coefficients apply to, just as they do for the partial derivatives of the computational coordinates $[\xi, \eta]$.

(b) Here, the surfaces V and H refer only to the surfaces on the exterior of the mesh, and their orientation in *computational* space. Additionally, the changes in the node index (Δk and Δj) select which external boundary the evaluated term contributes toward. For example, $C[\Delta j = +1]$ is the contribution from the surface extending from $[\xi, \eta] = [k, j]$ to $[\xi, \eta] = [k, j + 1/2]$. If either of the two nodes bounding the surface are not part of the external boundary of the mesh, it should be understood that no contribution to the flux integral is made.

(c) Nodes at the top-right and top-left corners of the mesh have contributions from both the side and top boundary expressions. As an example, the top-left node on the mesh has $\beta_{f,TOT} = \beta_f[\Delta j = -1] + \beta_f[\Delta k = +1]$. The other combinations of Δk and Δj connect this node to surfaces beyond the boundary, so they are ignored.

$$C\hat{\phi}_c - (\ldots regular\ terms \ldots) - \beta_{f,Tot}\hat{\phi}_f - \beta_{cath,Tot}\hat{\phi}_{cath} = \hat{Q}_{k,j} \tag{4.22}$$

Left & Right Boundaries:

$$\begin{aligned}
N &\equiv C[\Delta j = +1] && + (\ldots regular\ terms \ldots) \\
C &\equiv C[\Delta j = +1] && + C[\Delta j = -1] + (\ldots regular\ terms \ldots) \\
S &\equiv C[\Delta j = -1] && + (\ldots regular\ terms \ldots) \\
\beta_{f,Tot} &\equiv \beta_f[\Delta j = +1] && + \beta_f[\Delta j = -1] \\
\beta_{cath,Tot} &\equiv \beta_{cath}[\Delta j = +1] && + \beta_{cath}[\Delta j = -1]
\end{aligned}$$

$$C[\Delta j] \equiv \frac{\Delta j}{\hat{R}_{3,\Delta j}} \left[\hat{R}_{1,\Delta j}\left(\hat{A}_{r,V}\eta_{\hat{r},c} + \hat{A}_{z,V}\eta_{\hat{z},c}\right) + \hat{R}_{2,\Delta j}\left(\hat{A}_{r,V}\eta_{\hat{r},0,\Delta j} + \hat{A}_{z,V}\eta_{\hat{z},0,\Delta j}\right)\right]$$

$$\beta_f[\Delta j] \equiv \frac{1}{\hat{R}_{3,\Delta j}} \left[a_{f,c}\hat{R}_{1,\Delta j}\left(\hat{A}_{r,V}\xi_{\hat{r},c} + \hat{A}_{z,V}\xi_{\hat{z},c}\right) + a_{f,0,\Delta j}\hat{R}_{2,\Delta j}\left(\hat{A}_{r,V}\xi_{\hat{r},0,\Delta j} + \hat{A}_{z,V}\xi_{\hat{z},0,\Delta j}\right)\right]$$

$$\beta_{cath}[\Delta j] \equiv \frac{-1}{\hat{R}_{3,\Delta j}} \left[a_{cath,c}\hat{R}_{1,\Delta j}\left(\hat{A}_{r,V}\xi_{\hat{r},c} + \hat{A}_{z,V}\xi_{\hat{z},c}\right) + a_{cath,0,\Delta j}\hat{R}_{2,\Delta j}\left(\hat{A}_{r,V}\xi_{\hat{r},0,\Delta j} + \hat{A}_{z,V}\xi_{\hat{z},0,\Delta j}\right)\right]$$

$$\hat{R}_{1,\Delta j} \equiv \frac{7\hat{r}_c}{24} + \frac{\hat{r}_{0,\Delta j}}{12}\ ,\ \ \hat{R}_{2,\Delta j} \equiv \frac{\hat{r}_c}{12} + \frac{\hat{r}_{0,\Delta j}}{24}\ ,\ \ \hat{R}_{3,\Delta j} \equiv \frac{3\hat{r}_c}{8} + \frac{\hat{r}_{0,\Delta j}}{8}$$

Top Boundaries:

$$\begin{aligned}
E &\equiv C[\Delta k = +1] && + (\ldots regular\ terms \ldots) \\
C &\equiv C[\Delta k = +1] && + C[\Delta k = -1] + (\ldots regular\ terms \ldots) \\
W &\equiv C[\Delta k = -1] && + (\ldots regular\ terms \ldots) \\
\beta_{f,Tot} &\equiv \beta_f[\Delta k = +1] && + \beta_f[\Delta k = -1] \\
\beta_{cath,Tot} &\equiv \beta_{cath}[\Delta k = +1] && + \beta_{cath}[\Delta k = -1]
\end{aligned}$$

$$C[\Delta k] \equiv \frac{\Delta k}{\hat{R}_{3,\Delta k}} \left[\hat{R}_{1,\Delta k}\left(\hat{A}_{r,H}\xi_{\hat{r},c} + \hat{A}_{z,H}\xi_{\hat{z},c}\right) + \hat{R}_{2,\Delta k}\left(\hat{A}_{r,H}\xi_{\hat{r},\Delta k,0} + \hat{A}_{z,H}\xi_{\hat{z},\Delta k,0}\right)\right]$$

$$\beta_f[\Delta k] \equiv \frac{1}{\hat{R}_{3,\Delta k}} \left[a_{f,c}\hat{R}_{1,\Delta k}\left(\hat{A}_{r,H}\eta_{\hat{r},c} + \hat{A}_{z,H}\eta_{\hat{z},c}\right) + a_{f,\Delta k,0}\hat{R}_{2,\Delta k}\left(\hat{A}_{r,H}\eta_{\hat{r},\Delta k,0} + \hat{A}_{z,H}\eta_{\hat{z},\Delta k,0}\right)\right]$$

$$\beta_{cath}[\Delta k] \equiv \frac{-1}{\hat{R}_{3,\Delta k}} \left[a_{cath,c}\hat{R}_{1,\Delta k}\left(\hat{A}_{r,H}\eta_{\hat{r},c} + \hat{A}_{z,H}\eta_{\hat{z},c}\right) + a_{cath,\Delta k,0}\hat{R}_{2,\Delta k}\left(\hat{A}_{r,H}\eta_{\hat{r},\Delta k,0} + \hat{A}_{z,H}\eta_{\hat{z},\Delta k,0}\right)\right]$$

$$\hat{R}_{1,\Delta k} \equiv \frac{7\hat{r}_c}{24} + \frac{\hat{r}_{\Delta k,0}}{12}\ ,\ \ \hat{R}_{2,\Delta k} \equiv \frac{\hat{r}_c}{12} + \frac{\hat{r}_{\Delta k,0}}{24}\ ,\ \ \hat{R}_{3,\Delta k} \equiv \frac{3\hat{r}_c}{8} + \frac{\hat{r}_{\Delta k,0}}{8}$$

Limitations & Alternatives: Estimating the electric field beyond the mesh provides more realistic potential distributions in the absence of any plasma, as shown in Section 4.2.3. However, if the cathode is grounded and ϕ_f is also close to zero, the magnitudes of the estimated potential differences become small. As a result, the boundary conditions approach the field-free exterior condition as the floating and cathode potentials approach zero — see Eq. 4.21. Additionally, there is no guarantee these estimates will always be desirable in the presence of a robust plasma at or beyond the exit of a simulated thruster.

This boundary condition extension is *not* currently the default setting for PTpic. However, it can be enabled by setting the compiler option `ALLOW_EXTERNAL_FIELD` to one in the header file `pic.h`. Also in `pic.h`, when using the direct solver — requiring that `DIRECT_SL_SOLVER` also be non-zero — the `FREESPACE_CATHODE` option is selected by default. This causes the direct solver to set all `FREE_SPACE` nodes to the cathode potential. For any changes to take effect, PTpic must be recompiled.

4.1.2 The Floating Body Boundary Condition

Prior to PTpic the model calculated ϕ_f from the net charge Q_f accumulated by the floating body and a specified capacitance. Before the potential solver was invoked, the change in the floating body potential was calculated simply in terms of the change in the floating body charge: $\Delta\phi_f = {}^{\Delta Q_f}/_{C_f}$. As appealing as this simple approach may be, it is not rigorously accurate because it fails to account for the cross-capacitances between the floating body and other conductors within the domain (i.e. the anode, cathode, and plasma). However, as discussed at the end of this section, it may be adequate in certain cases.

Ignoring the effect of the plasma on the overall system capacitance, PTpic was modified in this work to include the cross-capacitance of the anode and cathode. The correct relationship between ϕ_f and the potentials of other conductors in the domain — again ignoring the plasma — is provided in Eq. 4.23[153] (p. 43). Stipulating that the cathode must always be grounded, and fixing ground to 0 $[V]$ removes the last two terms from this equation.

$$Q_f = C_f\phi_f + C_{f,a}\phi_a + C_{f,cath}\phi_{cath} + C_{f,gnd}\phi_{gnd} \tag{4.23}$$

In the case of plasma thruster simulations, ϕ_a is held constant, leaving the two unknown

capacitances and the desired value of Q_f as a constraint. One way to solve for the floating potential is to determine the potential distribution on the domain for two different values of ϕ_f, and then integrate the displacement field over the surface of the conductor to determine the total charge of the floating body in each case. This would provide the values C_f and $C_{f,a}$ needed to determine the floating body potential corresponding to any specified net charge Q_f. PTpic reads these two capacitances from the input file `variables.in`.

Alternatively, the Maxwell Student Version (SV) software package can determine the entire capacitance matrix, using a slightly different method — according the technical notes provided its documentation. In the presence of N floating bodies, N potential distributions are determined. In each case, one perfect conductor is set to 1 $[V]$ with respect to ground while all other floating bodies are grounded. Then, each element of the capacitance matrix is evaluated using Eq. 4.24, where V is the volume of the entire solution domain. The electric field subscripts specify which conductor is not grounded when the given field is calculated.

In either case, the net result is simply a voltage offset when the floating body has zero net charge: $\phi_f = \frac{Q_f}{C_f} + V_o$, $V_o \equiv -\frac{C_{f,a}\phi_a}{C_f}$. The only advantage to using $C_{f,a}$ rather than V_o is that the correct, charge-neutral value of the floating body potential is easily calculated when the anode potential is changed. It is worth emphasizing that capacitance values are used to determine floating body potentials, and thus provide boundary conditions to the potential solver. In order to avoid any direct dependencies of boundary conditions on γ, capacitances are left unchanged with respect to γ, even though capacitances are proportional to ϵ_o. This emphasizes the physical effect of γ, which is to reduce the sensitivity of potential distributions to space charge, rather than a wholesale alteration of the free-space permittivity in all circumstances.

$$C_{i,j} = \int_V \epsilon \underline{\mathrm{E}}_i \cdot \underline{\mathrm{E}}_j \; dV \tag{4.24}$$

The discussion of the floating body potential now shifts to how the presence of a plasma alters the capacitance of the entire system. A simple example demonstrating the magnitude of this effect involves a spherical floating body immersed in a steady-state plasma. In the limit of a small plasma-screening-length relative to a spherical radius R, the potential distribution beyond the sphere can be approximated as: $\phi(r \geq R) \approx \phi_f \; \exp\left(-\frac{(r-R)}{\lambda_{De}}\right)$ — with the screening length taken as the electron Debye length λ_{De} and $\phi(r \to \infty) \equiv 0$. From here, it is easy to show

that $Q_f = 4\pi\epsilon_o R\left(\frac{R}{\lambda_{De}}\right)\phi_f$, predicting that the presence of the plasma increases the capacitance of the sphere by a factor of $\frac{R}{\lambda_{De}}$ over its isolated, free-space value. This expression represents the small Debye length limit of Eq. 4.25, whose derivation was motivated by investigating forces on charged objects exposed to plasmas, such as contaminants in magnetically confined fusion experiments or asteroids in the solar system[154]. The modifying term comes from using a screened Coulomb potential rather than the purely exponential one from the previous example. Hutchinson[155] compares charge values on a conducting sphere, calculated using Eq. 4.25, to results from the PIC model SCEPTIC. Excellent agreement is observed when $\lambda_{De} \gg R$, though differences on the order of 10% appear at $\lambda_{DE} = R$. However, the discrepancy also grows as the relative ion bulk speed increases, and the opposite limit where $\lambda_{De} \ll R$ was not examined in that study. At least part of the inaccuracy arises from the derivation of Eq. 4.25, which relies on the linearization of a Boltzmann factor for the electron density in Poisson's equation. This approximation is only accurate when $|\phi_f| \ll T_e$, where T_e has units of electron volts. When the potential drop in the sheath is comparable to or greater than T_e, as is often the case in Hall thrusters, the approximate equations discussed here may not be appropriate. For example, as $|\phi_f|$ increases, the sheath thickness can increase, which will reduce C_f if the electric field at the floating body surface also decreases.

$$C_f = 4\pi\epsilon_o R\left(\frac{R}{\lambda_{De}} + 1\right) \tag{4.25}$$

Seeing that the plasma should have a dominant effect on the capacitance matrix of the entire system, it could be argued that calculating self-consistent capacitances and floating potentials each time the potential solution is desired should be a necessary feature in PTpic. In this case, a typical problem with an anode, floating body and grounded cathode would require three potential solves[a]. The first two solutions would provide ϕ_f, as outlined previously, and the third solution would obtain the potential distribution consistent with the properly determined value of ϕ_f. The main challenges associated with implementing such an approach, aside from adding to the already high computational cost of each iteration, would be quickly evaluating the displacement flux on the surface of the floating body, or computing an expression similar

[a]Multiple potential solutions are necessary when the constraint on the floating body charge is not included in the system of linear equations. The inclusion of this condition is discussed later in this section.

to Eq. 4.24 (as long as it is verified that such an expression is valid in the presence of a plasma). Also, in cases where the conducting body does not lie entirely within the simulation domain, good estimates of the charge contribution from the exterior surfaces would also be necessary if the charge on each floating body was desired.

Allowing for multiple floating bodies in the simulation domain is *not* currently supported in PTpic. In order to address such a case, with N floating bodies and one or more perfectly conducting electrodes at specified potentials in the domain, $N + 1$ potential solutions would have to be determined from $N + 1$ linearly independent, known combinations of floating body potentials. Using the information in these $N+1$ solution sets, $N(N+1)/2$ evaluations of integrals similar to Eq. 4.24 — or, alternatively, floating body charges — would provide the needed set of capacitance coefficients. At this point, the small $(N \times N)$ capacitance matrix could be inverted directly to provide the instantaneous, self-consistent potential for each floating body. After all of this, an additional potential solution would be needed using the correct floating body potential values as boundary conditions. This procedure would only work if the capacitance of the system remains linear in the presence of a plasma. If the system's capacitance is non-linear, then this method could only provide approximate information that might be usable as part of a root finding method to determine the self-consistent value of ϕ_f.

J. Szabo discusses ideas related to the wall potential and how electron injection and the behavior of the model are coupled in Secs. 4.3.2–4.3.4 of Ref. 97. It is also worth pointing out here that using a permittivity factor should weaken fields at the floating body surface, and that effects of dielectrics bordering floating bodies may complicate any desired corrections for the permittivity factor. Hypothetically, if C_f were calculated self-consistently, then it might be best not to adjust C_f for γ unless problems related to C_f values being too low are identified.

Determining floating potentials in a single step: Using $N+2$ different potential solutions would not provide the correct floating potential values if the capacitance is non-linear, and the computational cost of this method cannot be taken lightly considering the nearly inhibitive time required to run PTpic in its current state — even if the capacitance were linear. Rather than solving for floating potentials by calculating each element of the capacitance matrix, the charge constraint on each body can be expressed as a single algebraic equation. Supplementing the system of node potential equations with an additional equation for each floating body keeps

the number of equations and unknowns in balance — allowing this new system to be solved by whichever method is desired.

Specifying the charge on a conducting body is mathematically expressed as an integral boundary condition using Eq. 4.2 — where in this context S_V specifies the surface of the conductor, Q is the known body charge, and γ is set to unity[b]. Continuing with the same line of reasoning used to derive the node potential and boundary condition extension equations leads to the following insight: each sub-surface comprising S_V will make an additive contribution to the integral equation, providing a constraint through the value of Q. Applying the finite differences of potential given in Eq. 4.12 ensures that the constraining equation will depend linearly on the potentials of all nodes bordering the floating body.

Implementing this approach would be straightforward, though tedious, and would significantly alter the structure of the underlying coefficient matrix by introducing direct coupling between floating body potentials and all bordering unknowns. The direct solver added to PTpic is described in Section 4.2.2, and its scalability to larger problems depends on exploiting the matrix structure to reduce memory consumption and execution time for each solve. Therefore, adding the capability to solve directly for all floating potentials in a single step will require using a direct solver meant for a full matrix — such as the `ScaLAPACK` routines `PDGETRF` and `PDGETRS`. However, the memory required to store a full LU decomposition scales as the square of the number of unknowns, as apposed to the ⅔ power for matrices solved in the current version of PTpic. Therefore, it will probably be necessary to use a new iterative solver, or return the parallel SOR solver, if this modification is pursued in future work with larger meshes. Alternatively, a sparse linear solver might be well suited to determine floating body potentials in a single step. The Watson Sparse Matrix Package (WSMP) is one example that has been used in a 3D Hall thruster model[117], and can be obtained online[156].

Why the rigorous solution may not always be needed: In the presence of a steady state plasma, any floating conductor (implicitly assumed to have a finite capacitance) will collect charge until it reaches a potential that drives the total current reaching the body to zero. This floating potential is then a characteristic voltage of the surrounding plasma. In this sense, on

[b]Modifying γ in the plasma domain serves to relax length scale resolution requirements, while doing so at the surface of the conductor does not serve any rational purpose. Having $\gamma = 1$ for this application prevents any direct influence on the boundary conditions from γ.

a long enough timescale, a floating conductor will reach a plasma's floating potential *regardless* of the body's capacitance. Therefore, in applications where a steady-state, or even a "slowly" evolving plasma is of interest, the use of a single, approximate capacitance for the floating body appears justifiable. Accounting for the voltage offset V_o that has been added for PTpic is likely not even necessary for simulating certain problems[c]. However, if the dynamics of sheath formation or the floating body potential are of interest or otherwise deemed important, then the self-consistent method for calculating the floating body potential should be employed.

In its current state, PTpic reads the capacitances C_f and $C_{f,a}$ from file, and treats them as constants throughout the simulation. As such, the self-capacitance C_f provides a mechanism for the code to adjust the floating body potential according to the charge it accumulates — which will drive ϕ_f toward the plasma floating potential when the overall system evolves slowly enough. What will qualify as "slow" is difficult to bound in a meaningful way, especially if the capacitance and floating body current vary over several orders of magnitude between the time of plasma ignition to when the body reaches the plasma floating potential. Tracking the floating body potential during preliminary DCFT simulations, with constant values of $C_f = 23.5\ [pF]$ and $C_{f,a} = -16.7\ [pF]$[d], suggests that ϕ_f changes on a time scale of several tens of microseconds when corrected for the mass factor value. This may be tolerable compared to a neutral convergence time between $1 - 4\ ms$, but may be an underestimate of the real time scale due to the use of a constant C_f value.

A counter to this argument comes from considering the following: if the incorrect capacitances affect the transient behavior of the floating body potential, it may be possible for this to feed back and ultimately distort the bulk properties and long-term behavior of the plasma. This could be accentuated in PIC models, because initial conditions are typically very different from how the plasma is found to evolve afterward. Additionally, fluctuation amplitudes about mean floating body potentials would probably be smaller using the self-consistent approach, because capacitance values will be larger than their free-space counterparts. The only sure way to test the importance of this omission from the model would be to compare simulations run with and without the proper calculation of floating potentials.

[c] To remove the voltage offset, simply specify a value of 0.0 for $C_{f,a}$

[d] These values are provided by Maxwell SV solutions, in the complete absence of plasma

4.2 Solving the Linear System

All the potential solver details presented thus far have dealt only with forming the desired set of coupled, linear, algebraic equations in an automated way using a systematic notation. The direct predecessors of PTpic used iterative methods (serial and parallel) to solve the resulting system based on Eq. 4.26. From some initially specified set of values for the unknown potentials, new values were calculated based on the surrounding node potentials, until all values changed by an amount smaller than some specified tolerance. Boundary conditions were enforced by substituting the desired values in the equations where appropriate. In the case of the extended PTpic boundary conditions, the terms proportional to $\hat{\phi}_f$ and $\hat{\phi}_{cath}$ in Eq. 4.22 are added to the numerator on the RHS of Eq. 4.26.

$$\hat{\phi}_{c,new} = \frac{\hat{Q}_{k,j} + N\hat{\phi}_n + (NE)\hat{\phi}_{ne} + E\hat{\phi}_e + (SE)\hat{\phi}_{se} + S\hat{\phi}_s + (SW)\hat{\phi}_{sw} + W\hat{\phi}_w + (NW)\hat{\phi}_{nw}}{C} \tag{4.26}$$

4.2.1 Parallel Iterative Potential Solver: Successive Over-Relaxation(SOR) with Chebyshev Acceleration

PTpic has the option to use one of two solvers: an iterative method inherited from previous work, and a direct solver added during this investigation. The iterative solver was initially written by Szabo [97], improved by Blateau [110], and parallelized by Fox [111][e]. In PTpic, when the compiler variable `DIRECT_SL_SOLVER` has a value of zero, the iterative solver is used. Otherwise, the direct solver is used.

Prior to modifications for PTpic, the solver was in the state left by Fox. At this point, when initially applied to the DCFT, the floating body boundary condition was not functional, and the potential solver did not extend into dielectrics. For this work, both of these features were added to the model — though Szabo did allow for a floating body boundary in his initial version of the code. Fox parallelized the iterative solver by dividing the domain into P rows, where P is the number of processes[f] running in parallel. Each process updates its own potential

[e] For clarity, this discussion refers to the work done by Fox extending Szabo's and Blateau's work, and is separate from Fox's work on a different plasma model [112,113]

[f] The first paragraph of Sec. 3.3 defines the intended distinction between the terms "process", "processor" and "core" in this discussion, and throughout this thesis. In its current state, PTpic can use only one core per process.

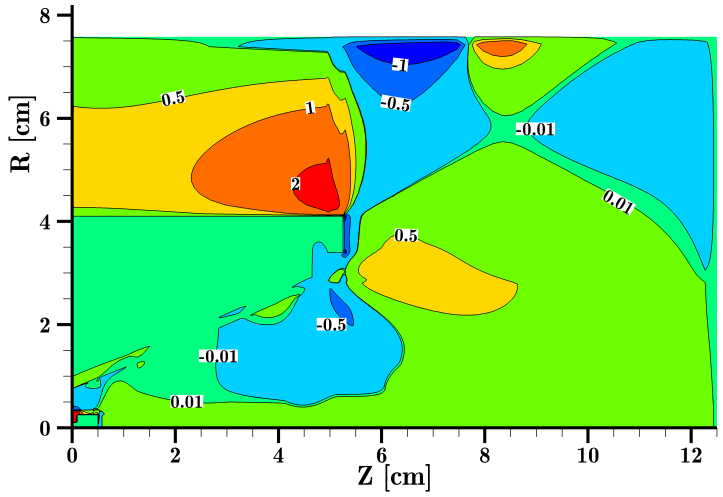

Figure 4-5: Error in potential, expressed in volts, after iterative solver convergence. These values are characteristic of results for $\epsilon_\phi \in [10^{-5} - 2.7 \times 10^{-12}]$ $[V]$.

values, sharing border values with other processes as needed. However, several undesirable characteristics identified during this work will now be discussed.

The convergence criterion for the iterative solver requires that the largest absolute change in a potential amongst all nodes must be smaller than some tolerance ϵ_ϕ. Using a DCFT mesh as an example, the potential value at each node was checked for consistency with the local space charge and surrounding node potentials. More precisely, the value of the expression on the RHS of Eq. 4.26 was compared to ϕ_c as determined by the converged iterative solver. Unexpectedly, differences between these two values were found with magnitudes much larger than ϵ_ϕ — Figure 4-5 shows an example. Curiously, the size of the error was insensitive to ϵ_ϕ over a range from $2.7 \times 10^{-12} - 1 \times 10^{-5}[V]$, changes to the permittivity factor, and changes in the number of processes used. The only parameter that was observed to affect this error was the floating body potential — though the values did not differ much from those shown here. Perhaps the order in which the node potentials are updated, or rounding errors, ultimately limit the achievable tolerances of the solver — though the actual reason remains unknown.

Dividing the domain into rows also limits the scalability of the potential solver. In practice, the hard limit for the number of processes that can work together within the iterative solver depends on the number of rows in the mesh, where a row is defined by all nodes with equal computational coordinates η, or equivalently j. The iterative solver is written such that the

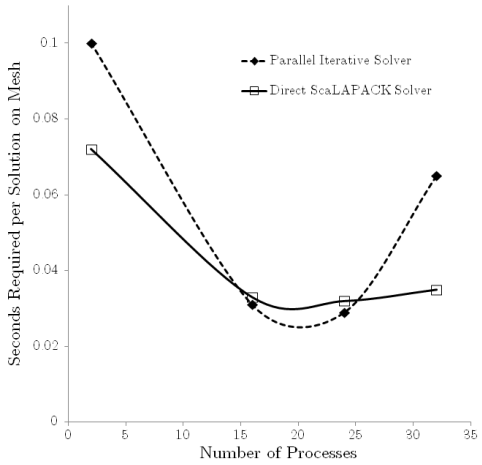

Figure 4-6: Comparison of times to solve one iteration for the parallel iterative and direct (based on PDGBTRS) solvers. The iterative solver is only competitive with $\epsilon_\phi \geq 10^{-5}$ $[V]$, whereas the direct solver produces solutions that are accurate to the order of 10^{-12} $[V]$. The mesh on which these solutions were obtained is identical to other example problems discussed in this section.

solver fails completely when any process is assigned 1 or fewer rows. However, this is rarely an issue because the inherited potential solver does not scale favorably when run with more than 30 cores[], and so a practical limit of 24 cores was used on simulation runs. Measurements of the time consumed by the iterative solver as the number of processes[g] varied are provided in Fig. 4-6. For 16 and 24 processes, the iterative solver appears only slightly faster when $\epsilon_\phi = 10^{-5}$ $[V]$ — though they are seemingly within the range of statistical variation for each solver. Of course, a solver's scalability often depends on the size of the problem being attacked, so it is important to mention that the optimal core numbers discussed here apply only within the context of a mesh with 244 columns and 156 rows[h].

4.2.2 Direct Potential Solver: Repeated Use of LU Factorization for a Banded Linear System

The direct solver employs Scalable LAPACK (ScaLAPACK) [157], a computational library that extends selected Linear Algebra PACKage (LAPACK) functions to utilize the capabilities of

[g] For all simulations described in this thesis, the number of processes and cores are equivalent because each process is assigned to run on a single core.

[h] The actual number of unknowns determined by the solver is less than the number of nodes in the mesh, due to specified potentials at the cathode and within the anode and floating body. For the example problem discussed here, the number of unknowns is closer to $30,000$ than $156 \times 244 = 38,064$

distributed memory parallel computers. ScaLAPACK is freely available, portable, documented in the literature[158] and supports the MPI standard for interprocess communications. The underlying LAPACK library is also reliable — it is used by MATLAB for tasks such as matrix multiplication and LU decomposition. The following discussion will outline how two specific functions in ScaLAPACK are used to solve for potential distributions in PTpic, namely `PDGBSV` and `PDGBTRS`. It is worth mentioning here that Fox implemented a direct solver from the MUltifrontal Massively Parallel sparse direct Solver (MUMPS) library in a different plasma simulating program using an adaptive mesh[112,113]. However, the MUMPS solver did not scale as well as hoped for in that application, used an excessive amount of random access memory (RAM), and exhibited a memory leak that preventing the code from running for long intervals.

In PTpic, direct solvers offer a competitive alternative to the iterative solver for one main reason: the matrix of coefficients $\underline{\underline{A}}$ for the linear system is constant, evaluated before the simulation loop begins. This can be seen in Eq. 4.14, where the node potential coefficients for a given row of the matrix depend only on the fixed mesh geometry. Put into the form of a matrix equation, $\underline{\underline{A}}\,\underline{\hat{\phi}} = \underline{\hat{Q}}$ represents the linear system of equations whose solution provides unknown potential values on the mesh. In this equation, $\underline{\hat{\phi}}$ is the vector of all unknown, normalized, node potentials, and $\underline{\hat{Q}}$ is the vector of normalized space charge values at each node. Whenever a given node $[k, j]$ is coupled to a known potential value, this information is moved into $\underline{\hat{Q}}$.

The structure of the coefficient matrix is determined by how the unknowns are labeled on the mesh. Ordering nodes with unknown potentials consecutively in the η-direction (see Fig. 4-1 for a visual aid), and placing the n^{th} node-equation in the n^{th} row of $\underline{\underline{A}}$, produces a matrix with the structure illustrated on the left in Fig. 4-7. For a given node $[k, j]$ labeled as the n^{th} unknown, coupling to nodes "north" and "south" of $[k, j]$ contribute to columns $n-1$ and $n+1$ within row n. The coefficient C occupies the diagonal entry of the same row. Predictably, nodes "west" of $[k, j]$ fill columns $n-(N_j+1)$, $n-N_j$, and $n-(N_j-1)$ while nodes to the "east" fill columns $n+(N_j-1)$, $n+N_j$, and $n+(N_j+1)$[i]. If any of the coupling nodes are not part of the mesh, then no corresponding matrix element is present. The same is true for nodes whose potential value is specified. A matrix with this structure is called a *band matrix*,

[i] N_j refers to the number of rows in the mesh, while N_k gives the number of mesh columns. These should not be confused with row/column indices in the coefficient matrix, which must have equal numbers of rows and columns — the total number of which can not exceed the product of N_j and N_k. In the PTpic source code, N_j and N_k are represented by the variables "N_R" and "N_Z", respectively.

and the bandwidth w is defined as the maximum index displacement of any matrix element away from the main diagonal entry in the same row. Hence, the bandwidth of $\underline{\underline{A}}$ is $N_j + 1$ — which partly determines the amount of memory and calculation required to factor and solve the system of equations. In a typical PTpic mesh, $N_k > N_j$, so this numbering scheme should minimize the computational expense of the potentials solver in most cases.

Computationally efficient LU decomposition algorithms often overwrite each element of $\underline{\underline{A}}$ with the entries needed to store the entire factorization when the matrix entries are no longer useful. For a band matrix LU factorization stored in-place, only entries located within the bandwidth will be non-zero (the LU decomposition is dense within the bandwidth). Therefore, the amount of memory required to store the LU factorization of a banded matrix is proportional to the product of the total number of unknowns N and the bandwidth ($w \approx \sqrt{N}$ for a square mesh). For large systems, this represents significant savings compared to the N^2 scaling for a general $N \times N$ matrix.

ScaLAPACK and LAPACK take advantage of band structure by only requiring allocation of enough memory to store the entries within the banded portion of $\underline{\underline{A}}$ — plus a manageable amount of supplemental memory required for extra fill-ins resulting from the details of how the factorization process is optimized within supporting libraries. A visualization of the band storage scheme is provided on the right in Fig. 4-7 — where vertical lines show how the entries are divided amongst processes in memory. Each row in the lower section of the compact representation of $\underline{\underline{A}}$ is reserved for an entire diagonal of the matrix, while coefficients are stored in the same column for either representation. Thus, the entry at row R, column C of $\underline{\underline{A}}$ is stored in row $3w + (R - C)$, column C, in band storage. The vector $\underline{\hat{Q}}$ is also spread amongst processes so that a processes p owns the same rows of $\underline{\hat{Q}}$ as columns of $\underline{\underline{A}}$. This storage method prevents the matrix and subsequent factorization from ever having to be stored entirely within a single process' RAM.

After filling the distributed portions of $\underline{\underline{A}}$ and $\underline{\hat{Q}}$ before the simulation loop, the first potential solution is obtained by invoking `PDGBSV`, which computes the LU factorization of $\underline{\underline{A}}$ and a solution $\underline{\hat{\phi}}$, stored in the same memory as the inputs $\underline{\underline{A}}$ and $\underline{\hat{Q}}$. Internally, `PDGBSV` calls `PDGBTRF` to compute the LU decomposition, followed by `PDGBTRS` to generate the solution using forward and backward substitution[j]. Then, within the main simulation loop, `PDGBTRS` is called to

[j]ScaLAPACK routines extend the naming conventions of LAPACK. The `P` distinguishes ScaLAPACK from

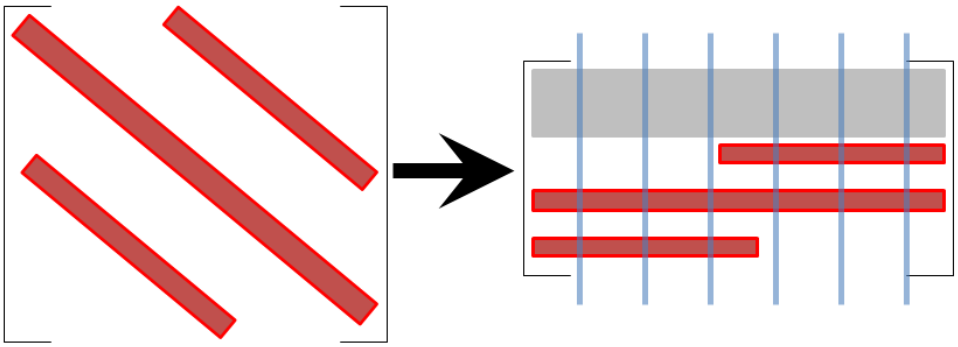

Figure 4-7: **Left:** Non-zero elements of matrix $\underline{\underline{A}}$ occur in groups of three diagonals at a time — one near the main diagonal, and right/left of the main diagonal. **Right:** Band storage scheme used by ScaLAPACK. The first group of rows stores additional fill-ins for the LU factorization, and remaining rows store all diagonals within the bandwidth. The column index of each entry in $\underline{\underline{A}}$ is preserved in the band storage representation. Vertical lines represent how the banded representation is divided amongst processes.

generate potential solutions as needed. Full details are provided in the ScaLAPACK manual [157], installation files, or online. A few additional details are discussed at the end of this section.

Direct solver execution times during a DCFT simulation are given in Fig. 4-6. The scaling appears to be minimally degraded (especially considering statistical uncertainties of these measurements) by required ancillary tasks — such as reforming $\underline{\hat{Q}}$ on each process and then distributing the solution $\underline{\hat{\phi}}$ back to all processes. Execution times for `PDGBTRS` measured for a sample problem with 4×10^4 unknowns and $w = 200$ demonstrated no deterioration of performance using up to 80 cores. This was the maximum number of cores used in test runs. However, returns were diminishing: a parallel speed up by a factor of 8 was measured using 24 cores, versus a speed up of 9.8 using 80 cores. This behavior is consistent with the `PDGBTRS` documentation, which states that the algorithm is best applied to cases where the ratio of unknowns N to processes P is much greater than the bandwidth w. In this example, $\frac{N}{P\,w} = 10$ for $P = 20$, suggesting correctly that gains in performance from using additional processes should increase more slowly beyond this point. Again, PTpic has used only one core per process for the simulations described in this thesis.

With the incorporation of the direct solver into PTpic, the potential solver performance improved from consuming 75% or more of each iteration, to between $10-15\%$ when using $16-24$

LAPACK routines, `DGB` specifies a Double precision General Band matrix, and the suffix describes the operation to be performed.

processes on the SPL cluster[k]. Additionally, $\underline{\hat{\phi}}$ values are more consistent with the governing equations by many orders of magnitude — on the order of 10^{-12} $[V]$, in comparison the values shown in Fig. 4-5.

In principle, this direct solver could scale up to allow a permittivity factor γ close to unity in simulations run on large computing facilities. Doing so would increase the number of unknowns by γ^2 — which would scale the number of unknowns in the DCFT mesh up to about 7.5×10^7 (the mesh currently has about 30,000 unknowns using $\gamma = 50$). A total of about $19,360$ GB RAM would then be needed to store the banded version of the LU decomposition, requiring the use of $10 - 20$ thousand cores if each has access to $1 - 2$ GB RAM. This would necessitate operating with $\frac{N}{P\ w} \approx 0.4 - 0.9$. In the best case, the speed up would probably not significantly exceed $5 - 10\%$ of ideal $1/P$ scaling — though even these parallel efficiencies would reduce execution times by a factor of $1,000$ over the serial case. In the worst case, using too many cores will slow the solver down as the time spent communicating increases relative to time spent computing.

The resources for this large-scale calculation are available at various super-computing facilities[l]. For example, the Kraken Cray XT5 computing system at the National Institute for Computational Sciences provides a total of 147 TB of RAM distributed evenly amongst $9,408$ nodes with 12 cores each. Running with the minimum 2 processes per node, each process would then have 8 GB of RAM to work with, allowing the job to be run with $2,420$ processes on half as many nodes. At the other extreme, using 12 processes per node would require running with $14,520$ processes on the same number of nodes[m]. Wall clock limitations on Kraken jobs would limit the number of iterations that could be completed in a single run, but these estimates motivate modifying PTpic to run efficiently on thousands, rather than tens, of cores. Furthermore, incorporating Open Multiprocessing (OpenMP) into PTpic would allow taking advantage of shared-memory parallelization via multi-threading, in addition to the distributed-memory advantages provided by the Open Message Passing Interface (OpenMPI).

[k] This comparison refers to inherited iterative solver settings, which used $\epsilon_\phi = 2.7 \times 10^{-12}$, expressed in normalized voltage units.

[l] This feasibility estimate ignores the relatively poor scaling of other PTpic sub-algorithms

[m] By adding the capability for PTpic to exploit shared-memory parallelism, such as with openMP, all 12 cores per node could be utilized by only 2 processes per node.

How ScaLAPACK routines can simplify solving for capacitance matrices: This discussion is in reference to the ideas described in Sec. 4.1.2. Although one method of solving for the N floating body potentials in the domain requires a total of $N+2$ different solutions $\underline{\hat{\phi}}$, they could be obtained with only two calls to `PDGBTRS`[n]. This is not meant to suggest that the computational work required to generate these solutions will be significantly reduced. However, the procedure itself would be much simpler. The first $N+1$ solutions could be found after providing the same number of RHS vectors $\underline{\hat{Q}}$ as inputs to `PDGBTRS`. Processing this first set of solutions would provide N floating potentials that would then be used to form the last RHS $\underline{\hat{Q}}$, and in turn, the updated potential distribution from a second call to `PDGBTRS`. This procedure avoids having to call the same function $N+2$ times.

Some Important Details:

(a) The general banded solvers in ScaLAPACK divide the band representation of $\underline{\underline{A}}$ using a one-dimensional block column distribution, and the RHS vector using a one-dimensional block row distribution. The two-dimensional block-cyclic distribution is *not* used in this case, which has implications for how array descriptors must be initialized (see ScaLAPACK manual[157] for more details).

(b) When calling the ScaLAPACK LU factorization routine and using as many MPI processes as there are cores on a node, it is necessary to ensure that the environment variable `OMP_NUM_THREADS` is set to 1 when openMP is also available. Otherwise, the factorization runs prohibitively slow. With this setting, LU factorization times are reasonable, and scale favorably. Running with fewer processes than cores on a given machine may allow for improved performance using different values of `OMP_NUM_THREADS`, however.

(c) Memory for the RHS vector is allocated slightly differently than the band storage array. Each process designated as a member of the communication group for calls to the BLACS library from ScaLAPACK must allocate an array of the same leading dimension (number of rows) to store its portion of the RHS vector. The appropriate value for this shared leading dimension is given in Eq. 4.27, where P is the number of processes, N is the number of

[n] N is now used to represent the number of floating bodies in the domain, not the number of unknowns in the linear system governing the potential distribution.

unknowns, and LLD_B is the leading dimension for each process' portion of the RHS vector. This value of LLD_B should be used even when $N < (P \times \text{LLD\_B})$.

$$\text{LLD\_B} = \begin{cases} \frac{N}{P} & \text{if } N\%P = 0 \\ \frac{N+(P-N\%P)}{P} & \text{otherwise} \end{cases} \tag{4.27}$$

(d) An integer array referred to as `IPIV` in the ScaLAPACK documentation must be provided when calling either `PDGBSV` or `PDGBTRS`. Although the documentation states that this array must have a length equivalent to the value of LLD_B in Eq. 4.27, a *different* value, $(\text{LLD\_B} + 2w)$, was needed for the routines to function properly. This bug, and the suggested solution, was found on the LAPACK/ScaLAPACK Development online forum. The `IPIV` array is used to store information about pivoting that is done during the LU factorization.

4.2.3 Potential Solver Benchmarks

Comparing the direct PTpic solver to Maxwell SV: In order to benchmark the PTpic solver against a trusted electrostatic solver, boundary conditions that could be reproduced in both solvers were needed. For this purpose, the option to set all nodes on exterior boundaries (that are not already specified as part of other boundary conditions) to the cathode potential was added. This boundary condition can currently only be used for the direct solver, and is turned on or off by setting the value of the compiler option `FREESPACE_CATHODE`. This boundary condition was replicated in Maxwell SV by placing a perfect conductor at the same locations as the edge of the PTpic mesh and setting its potential to 0 $[V]$. The potential distributions calculated by PTpic and Maxwell SV with this common boundary condition are shown in Fig. 4-8. The domain in Maxwell extends beyond the PTpic mesh in the $\pm z$ and r directions. The presence of perfect conductors on the boundaries, however, means that the only differences in boundary conditions are on the $z = 0$ plane: Between the anode stem and floating body from $r \in [0.1 - 0.15]$ $[cm]$, and the gap between the floating body and grounded surface from $r \in [4.1 - 4.15]$ $[cm]$.

Over the full voltage range, the solutions look identical, and more precise comparisons demonstrate that their agreement is within justifiable uncertainties. The PTpic solution is

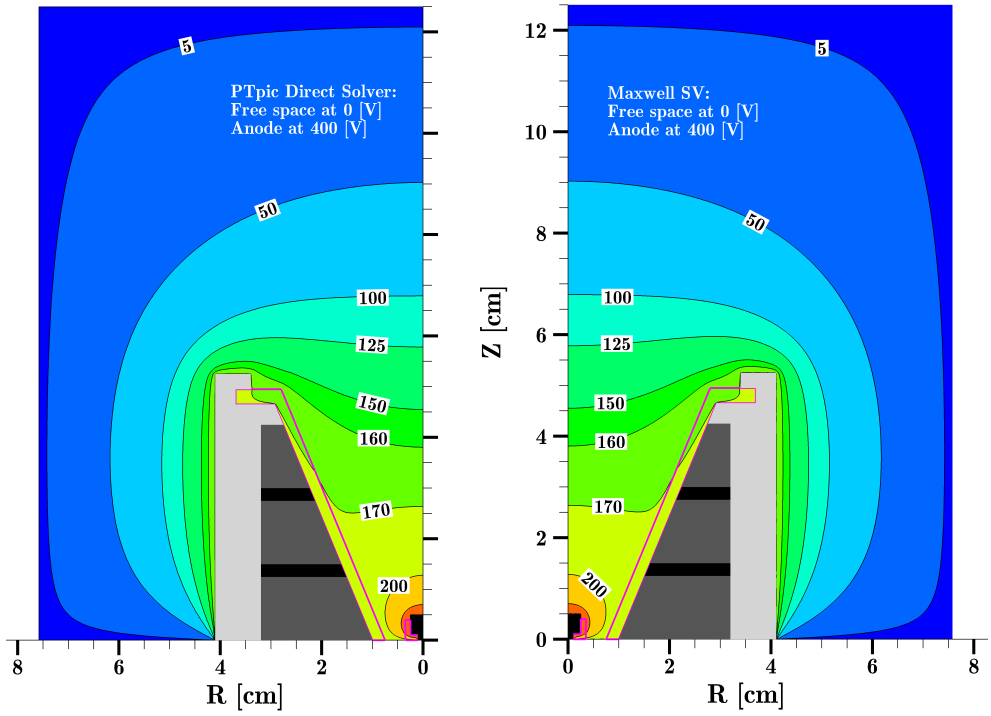

(a) PTpic solution with boundary nodes grounded.

(b) Maxwell SV solution with nearly identical boundary conditions.

Figure 4-8: Comparison of PTpic and Maxwell SV potential solutions with no space charge, and similar boundary conditions. Potential values are given in volts.

saved on its native mesh in a format readable by Tecplot. The Maxwell solution was determined on an unstructured triangular mesh with over 1.3×10^5 elements, and then interpolated to a structured mesh and written to a file that was laster converted to the desired Tecplot format. The result is plotted in Fig. 4-8(b).

More detailed comparisons were made by extracting values along several lines from each data set and calculating the absolute values of differences at common sampling points. These differences are shown in Fig. 4-9. For the majority of data points, the differences are smaller than .1 $[V]$, though differences in some regions are larger. On the line running from a corner of the floating body to where the cathode would be located, the maximum difference is slightly larger than 1 $[V]$.

There are several justifications for differences of these sizes. First, each data point is interpolated by Tecplot from surrounding data points to the points on each line from two different

meshes. Second, the Maxwell data are subjected to an additional interpolation layer when going between the raw Maxwell solution to a Tecplot readable format. Third, the floating body potential in PTpic is calculated based on capacitances obtained from Maxwell — for reasons discussed in Section 4.1.2. Based upon the sensitivities of the capacitance values on the resolution of the mesh in Maxwell, their accuracy is estimated to be about 1%. This level of precision is consistent with the maximum relative errors shown in the table portion of Fig. 4-9. Finally, a small handful of points on the surface of the anode in the Maxwell solution had potentials in excess of 400 $[V]$ despite the boundary condition specifying the potential there. Considering all of these factors, the agreement between solutions serves to validate the PTpic potential solver.

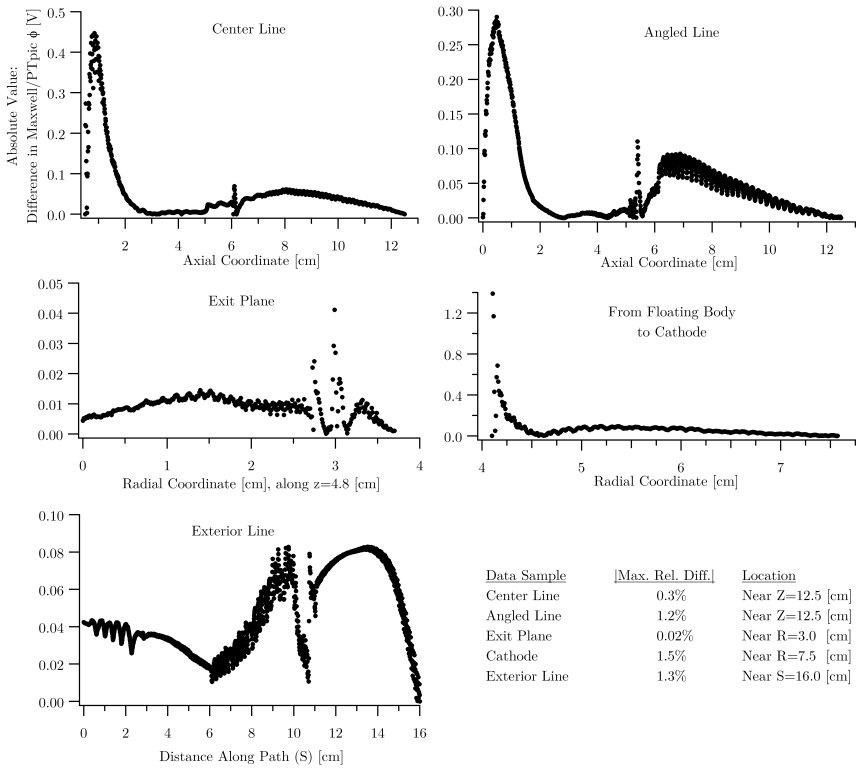

Figure 4-9: **Plots -** Absolute values of differences between the two solutions on selected lines. Data was extracted along the following paths (all (z, r) coordinates expressed in $[cm]$): "Center":$(.5, 0) \rightarrow (12.5, 0)$;"Angled":$(0, .55) \rightarrow (12.5, 5.18)$; "Exit":$(4.8, 0) \rightarrow (4.8, 3.7)$; "Cathode":$(5.25, 4.1) \rightarrow (8, 7.58)$;"Exterior":$(10, 0) \rightarrow (10, 6) \rightarrow (0, 6)$. **Table -** The absolute value of the maximum relative difference between the solutions on each line is given, along with its approximate location on the line.

A detailed benchmark accounting for the presence of space charge has not been completed for the PTpic solver because no third party solvers were readily available to replicate such conditions. However, previous versions of the iterative solver were tested against analytical charge density distributions [97,111], and the direct solver is based on a slightly modified set of equations. A good validating test in the presence of space charge would compare results for two solutions with identical boundary conditions and a known space charge distribution. One candidate for providing this capability may be the COMSOL Multiphysics software environment, or perhaps another analytical test. However, the numerical solution is preferred, because it would enable comparisons on meshes used in "real-world" problems.

Assessing PTpic boundary conditions: In addition to testing the validity of the PTpic potential solvers, the influences of two different boundary conditions were tested. The original boundary conditions essentially required an electric field parallel to the mesh boundary at the edges of the domain, except where the cathode potential is enforced. For comparison, an electrostatic model of the thruster — absent any plasma — was created and solved using Maxwell SV on a much larger domain. This enabled comparisons between PTpic solutions on smaller meshes to solutions less influenced by boundary conditions. In an attempt to provide more physical boundary conditions, the extension described in Section 4.1.1 was implemented. A visual comparison of solutions employing these three different boundary conditions is provided in Fig. 4-10. The PTpic solution obtained using electric field estimates at boundaries provides a closer match to the Maxwell solution, especially beyond $z \approx 5$ $[cm]$. Potentials along two of the same lines described in Fig. 4-9 are given in Fig. 4-11 for a clearer comparison of the three cases. Larger differences shown in Fig. 4-11(b) past $S = 10$ $[cm]$, corresponding to $(z \in [0, 6], r = 6)$ $[cm]$, occur — though not in the primary region of interest.

Although the external field boundary condition is more realistic in this simple case, its utility as a good boundary condition in the presence of a plasma has not been firmly established. Nonetheless, allowing for perpendicular electric field components beyond the mesh is appealing, so this boundary condition has been left as the default setting in PTpic. The input parameters for this condition are read from the input file `potentialBC.in` at run-time. Currently, two different sets of parameters, described in Table 4.3, are used to reflect different orientations with respect to the thruster. From the "TOP" boundary, the closest part of the floating body

is offset from the axis of symmetry, while on the "RHS" it is centered about that axis. Also on the RHS, the characteristic length of the floating body is its outer radius of 4.1 $[cm]$, while from the "TOP" boundary L_f was selected as $\sqrt{4.1^2 + 7.5^2}$ $[cm]$ — where the larger dimension is the total floating body length (extending back to approximately $z = -2.5[cm]$).

Table 4.3: The two sets of parameters used in some DCFT simulations and the potential comparisons shown in this section. Section 4.1.1 describes how each of these terms are used, if at all. "TOP/LHS" refers to boundaries defined by the $z = 0$ $[cm]$ and $r = 7.58$ $[cm]$ surfaces, respectively. The "RHS" boundary is on the $z = 12.5$ $[cm]$ surface.

	(z,r) [cm]		[cm]	
Boundaries	$\underline{R}_f$	$\underline{R}_{cath}$	L_f	L_{cath}
TOP/LHS	(5.3,4.1)	(8.0,7.58)	8.55	0.75
RHS	(5.3,0)	(8.0,7.58)	4.1	0.75

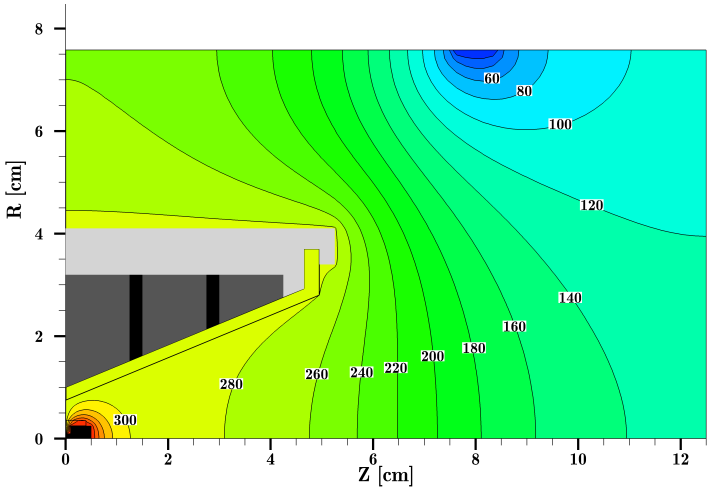

(a) PTpic solution with no external field (original boundary conditions).

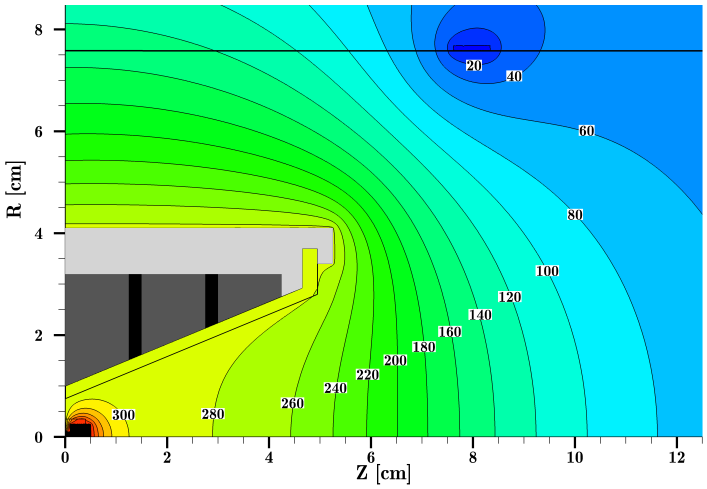

(b) Maxwell SV solution on a larger domain.

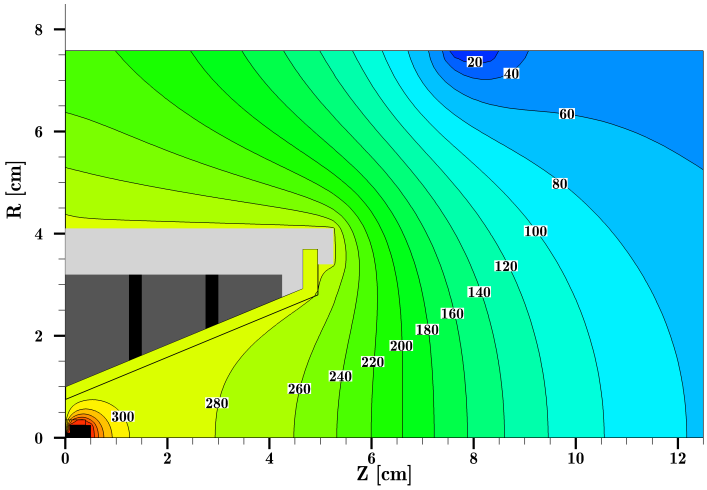

(c) PTpic solution using the boundary condition extension described in Sec. 4.1.1

Figure 4-10: Comparison of potential solutions with different boundary conditions. Potential values are given in volts.

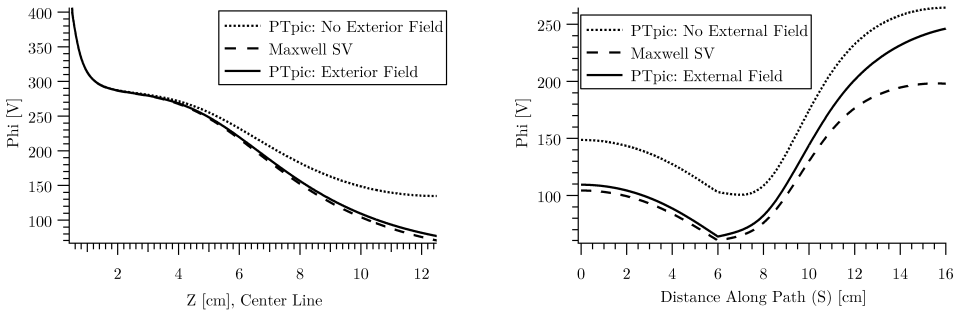

(a) Data extracted on "Center" line from solutions using different boundary conditions.

(b) Data along "Exterior" line using different boundary conditions.

Figure 4-11: Representative data comparing potential solutions obtained using three different boundary conditions.

4.3 Varying the Permittivity Factor

To illustrate the consequences of dampening the sensitivity of the potential to space charge, the potential distribution shown in Fig. 4-12(b) was calculated. This solution was obtained by restarting a simulation that had been running with $\gamma = 50$, and then reducing the permittivity factor to its physical value of 1. The potential solution in Fig. 4-12(a) was obtained from an identical charge distribution, but with $\gamma = 50$. The mesh resolution, meant to be adequate only for the $\gamma = 50$ case, is identical in each case. Clearly, the potential distribution for the more realistic case shows that larger γ values permit a degree of charge separation that would otherwise be more strongly counteracted by electric fields. If it were feasible to continue the simulation with $\gamma = 1$, charged particles should be expected to rearrange themselves in response to these large fields, eventually shielding them out in the bulk. At that point, the details of the potential topology would likely be different than in Fig. 4-12(a), but the values would certainly be on the same scale. Therefore, the differences between the distributions in Fig. 4-12 can not be directly interpreted as an indicator of error magnitudes in the model. Instantaneous views of PIC model data do not typically reflect every aspect of what should be expected from idealized, steady state plasmas. However, these data do starkly illustrate that the tradeoff for a linear increase in plasma shielding lengths is a quadratic reduction in sensitivity to space charge.

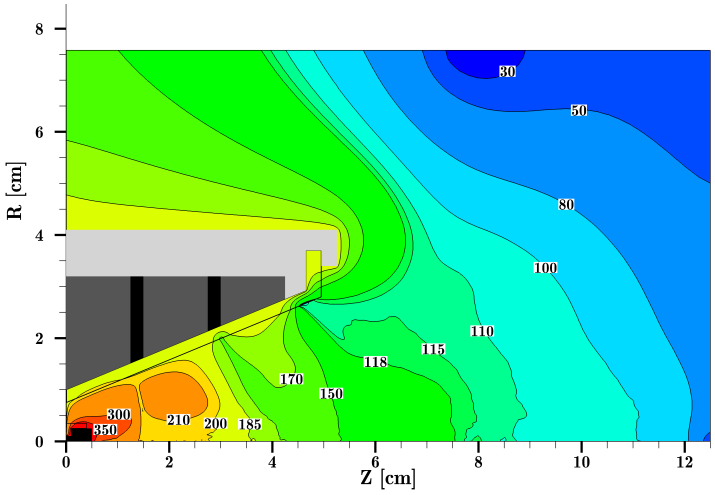

(a) Snapshot of a preliminary PTpic simulation of the DCFT with $\gamma = 50$.

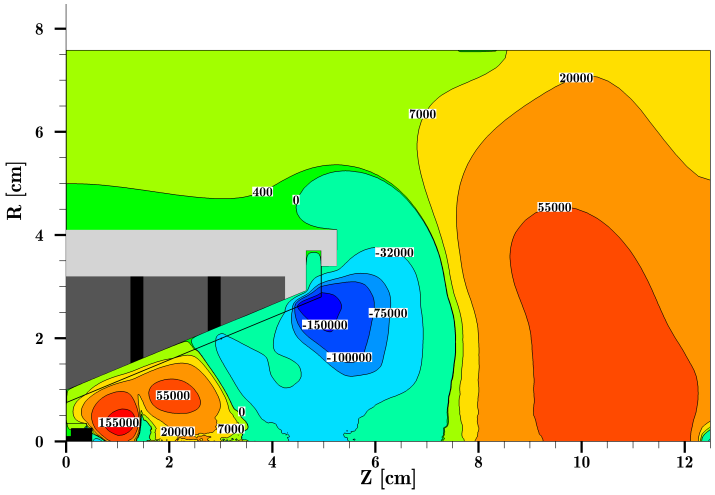

(b) Snapshot of a different solution with the same boundary conditions and charge distribution, but with $\gamma = 1$.

Figure 4-12: Comparison of two potential solutions, given in volts, with different permittivity factor values. The anode voltage is 400 $[V]$, and $\phi_{cath} = 0$ $[V]$ in both cases. The floating body potential is also the same in these two examples.

4.4 Potential Solver Summary

The essence of the PTpic potential solver has remained unchanged: the integral form of Gauss' law is applied to each node on the mesh, nominally coupling its potential value to the 8 surrounding values when the appropriate finite difference approximations are used. However, the method was modified to account fully for cylindrical coordinates — though the dominant cylindrical correction was included in previous versions of this model. As a result, large portions of the solver were entirely rewritten. In the process, a boundary condition extension was added to allow a perpendicular $\underline{E}$-field component on the mesh exterior by making analytical estimates of potential gradients. This boundary condition, and the new solver, were then validated against the commercially developed Maxwell SV electrostatic solver.

While checking the validity of the parallel iterative solver, systematic discrepancies were noticed that could not be corrected. Though the inherent errors do not render the solver useless, greater accuracy was desired. As a result, a direct solver was incorporated into PTpic, making use of ScaLAPACK functions to perform an LU factorization which is utilized to repeatedly provide new potential distributions as the charge distribution evolves. The direct solver provides solutions as quickly as the iterative solver, and should scale more favorably to using a larger number of cores. A detailed comparison between potential distributions calculated by the direct solver and Maxwell SV was performed, validating the new solver to within expected error bounds.

Furthermore, two algorithms for calculating floating body potentials are identified and proposed, but neither is implemented. Currently, PTpic oversimplifies the calculation, determining the floating body potential of up to one conductor using constant capacitance values and the known charge on the body. The presence of the plasma, a very good electrical conductor, is *not* taken into account. However, for some applications, simply providing a mechanism for the floating body to reach the surrounding plasma's floating potential may be adequate. The only way to decisively settle this issue seems to require implementing the extensions necessary to rigorously satisfy the integral boundary condition on the floating body. Doing so would quantify the effect of the plasma on the capacitance matrix, and possibly identify powerful approximations that would allow for a quicker, but still sufficiently accurate, treatment of the floating body boundary condition. For instance, it may be that the cross-capacitances of multiple bodies

immersed in a plasma will be greatly reduced compared to their values in vacuo — assuming a Debye length much smaller than the spacing between bodies.

Chapter 5

Simulations of Original DCFT Design

By far, the greatest obstacle toward obtaining meaningful simulation results with PTpic is the disparity in time scales between electrons and neutrals. In the absence of any plasma, an acceptably converged neutral simulation can be obtained in a half-hour or less of wall time[a] for $r \geq 1000$. Simulated on a mesh extending beyond the exit plane of the DCFT, neutral particles approached a convergent state in the times shown in Table 5.1, where the degree of convergence (DoC) is defined as $1 - \exp\left(\frac{-t^*}{\tau_r}\right)$—where $\tau_r \equiv \tau_1/\sqrt{r}$. It is convenient to express intervals of simulated time in terms of the convergence time scale $\tau_r \equiv \tau_1/\sqrt{r}$, especially amongst simulations with different mass factor values. The factor τ_1 depends on all factors related to the size and geometric configuration of the domain, as well as the neutral species temperature and flow rate. Though sensitive to many inputs, τ_1 can be measured during the neutrals only simulations that precede the full plasma simulations. Interpreting simulated time intervals as relating to physical times through the factor $\sqrt{r}$ is appropriate when examining timescales set by neutrals and ions. However, on the electron timescales, it is better to interpret simulated time intervals through a factor of γ. In these results, the timestep is on the order of the electron timescale, but time-averaging smooths over variations that might occur that quickly.

For the DCFT simulations described here, a timestep of 1×10^{-12} $[s]$ was used, which is

[a]The term "wall time" refers to the real time elapsed from when a simulation begins to when the simulation ends. This has nothing to do with intervals of time within the simulation.

approximately $1.73\times10^{-9}\tau_1$ for $\tau_1 = 578$ $[\mu s]$. The required number of iterations to reach $t^* = \tau_r$ can now be estimated as $\frac{5.8\times10^8}{\sqrt{r}}$. Depending on the number of particles at a given instant, an iteration typically requires $.1 - .2$ $[s]$, on average. From here, we are able to estimate that reaching τ_r should take at least $^{669}/_{\sqrt{r}}$ *days*. Furthermore, substantial increases in the number of plasma super-particles can drive a single iteration's wall time above 1 $[s]$, so it is necessary to use a large enough super-particle size W to maintain tractability. In all the simulations discussed in this chapter, neutral temperatures are elevated to be on the order of 0.1 $[eV]$, reducing the number of iterations required to reach a given degree of convergence. This undoubtedly has the undesirable effect of lowering the simulated neutral utilization efficiency because neutrals are probably no hotter than about 0.06 $[eV]$ in experiments.

Table 5.1: Estimated amount of simulated time t^*, in microseconds, required to reach a given degree of DoC. These numbers are representative of the simulation domain, with an injected xenon flow rate between $4 - 6[sccm]$.

		Mass Factor (r)				
DoC	t^*	1	1,000	5,000	10,000	40,000
63%	τ_r	578	18.3	8.17	5.78	2.89
87%	$2\tau_r$	1155	36.5	16.3	11.6	5.78
95%	$3\tau_r$	1733	54.8	24.5	17.3	8.67

In the following sections, simulations of the original DCFT design built by D. Courtney [131] are presented. Sec. 5.1 contains results of a steady state simulation, representing the low-current mode of operation. Next, results from an oscillatory simulation, more closely resembling the high-current mode are given in Sec. 5.2. The effects of varying parameters such as the mass factor, floating body capacitance, and limited results of an almost identical geometry represented on a different mesh are provided in Sec. 5.4. Finally, in Sec. 5.3, simulated erosion data are presented.

5.1 Converged, Steady State Simulation

A steady state simulation was achieved using the parameters listed in Table 5.2. The most distinct parameter, relative to other simulations, is the floating body capacitance C_f—it is 1000 times larger than the value used in Sec. 5.2. Physically, this enlarged capacitance is meant

to mimic the effect of the plasma on C_f, as discussed in Sec. 4.1.2. The actual factor by which C_f is increased here is only based on estimates, but it is clear that this parameter has a strong influence on the behavior of the model. Also, electron transport is not increased above classical levels, except of course for increases in electron collisional cross-sections by $\sqrt{r}$ that are meant to keep simulated neutral mean-free-paths for ionization in agreement with physical values.

Selected diagnostic data from the simulation are given in Fig. 5-1 for the simulated time interval $t^* = 13.20 - 36.30$ $[\mu s]$. The super-particle weight W was decreased from 5×10^8 to 1×10^8 at $t^* = 21.17$ $[\mu s]$, and the magnitude of statistical fluctuations are noticeably reduced in Fig. 5-1 as a result. The maximum relative neutral size M was decreased from 200 to 100 at the same time the super-particle size was decreased. Finally, the wall and injected neutral temperatures were about .1 $[eV]$ in this simulation.

As a quick aside on computational particle weighings, ions and electrons all have the same weighting W, and neutral weights vary from W to $(M \times W)$. In all plots presented in this chapter, particle counts are converted into units of W-sized particles. The converting factor is just one for electron and ions, and somewhere between 1 and M for each neutral super-particle. Whenever W or M are changed, mass and particles are conserved up to the point that an integer number of computational particles, constrained by the new and old values of W and M, can represent the same number of physical particles. Decreasing W and M by integer factors conserves mass very well and can increase the statistical resolution of a simulation without affecting its behavior too strongly in some circumstances. Whenever W and M are decreased by a fractional factor, or increased any factor, particles are only conserved in an average sense. Increasing the weighting factors has been avoided in the simulation results discussed throughout this chapter.

Table 5.2: Summary of this simulation's parameters.

V_a=	250 [V]	$\dot{m}$=	6 [sccm Xe]
γ=	50	r=	1000
W=	5e8,1e8	C_f=	39.1 [nF]

The first 13 $[\mu s]$ of this simulation were run with a lower value of C_f and actually provide the results discussed in the next section. The larger capacitance of the floating body had

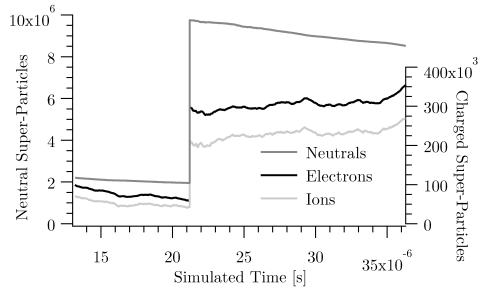

(a) Super-particle counts.

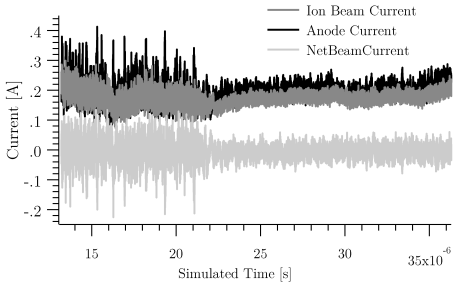

(b) Ion-beam, anode, and net beam current.

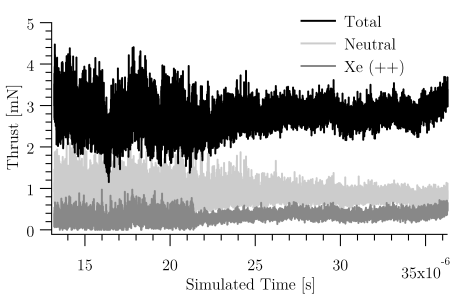

(c) Thrust and its components.

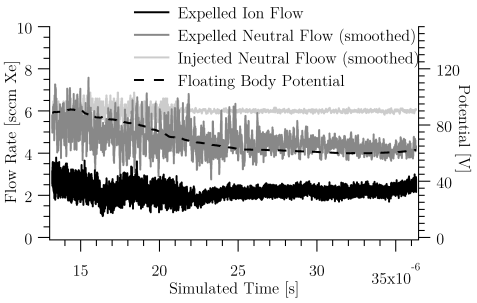

(d) Different flow rates and floating body potential.

Figure 5-1: Each quantity is time-averaged within groups of 1000 iterations. Values labeled as "smoothed" are subjected to an additional sliding average within a box containing 50 time-averaged values. The simulated time interval of $\Delta t^* \approx 1.26\tau_r$ spans 23.1 million iterations.

the effect of stabilizing the simulated discharge, despite the fact that the operating condition was not changed. The absence of large oscillations in the number of super-particles in the domain allowed W to be reduced, improving the statistical representation of the discharge. The abrupt increase in super-particle counts in Fig. 5-1(a) is the result this reduction. Through this transition, mass is almost perfectly conserved because each particle with weight W is transformed into 5 particles of weight $W/5$ — all with the same position and velocity. After the transition, all variations in super-particle counts are essentially magnified by the same factor by which W was decreased. After t^* reaches 25 $[\mu s]$, the currents, thrust, and flow rates have all reached steady states.

Important average quantities from the Fig. 5-1 are provided in Table 5.3. The mass flow rate given in the table is the net flow rate exiting the simulation domain, and is 11% larger than the input mass-flow rate of 6 $[sccm\ Xe]$. This difference reflects that the simulation is approximately 70% converged, which is properly reflected in the decreasing number of neutral super-particles represented in the domain. The predicted anode current is about two-thirds of measured values[28,30], though the lack of a neutral background pressure accounts for some of this discrepancy. Predicted thrust F and I_{sp} are also lower than measured values reported by R. Daspit—by approximately 50% in both cases[28]. However, the total efficiency is in excellent agreement with the 29% value measured at the same condition, though the simulated value is about 8 percentage points higher than Daspit's estimate of an efficiency corrected for a finite background pressure[159] (Sec. 4.4.2). The predicted *ion* specific impulse in Table 5.3 is much closer than the predicted *total* I_{sp} to the measured total value of ≈ 900 $[s]$ reported in Ref. 28.

With a simulated utilization efficiency of only 32%[b] and a useful energy per ion of 55 $[eV]$[c], the *simulated* performance of the thruster is unexpectedly bolstered by the neutral population. Here, neutrals comprise 4.54 $[sccm\ Xe]$ of the exiting mass flow, provide an average thrust of 0.82 $[mN]$, 24% of the useful beam energy, and 2.3 $[eV]$ per neutral of useful beam energy. This is well above the values from the neutrals-only simulations that precede the plasma simulations, where neutrals provided .194 $[mN]$ of thrust, and .07 $[eV]$ per neutral to the beam energy.

[b] Here, utilization efficiency is calculated as the ratio of exiting ion mass flow to nominal input mass flow, where the ion mass flow includes all ion charges. A more common definition is based on the total ion *current*, allowing efficiencies calculated in this way to exceed 100%[10].

[c] This energy is calculated according to $\frac{1}{2}m_i c_i^2$, where c_i is the effective ion exhaust speed. This effective speed is calculated according to $F_i = \dot{m}_i c_i$, where F_i & $\dot{m}_i$ represent the ion thrust and exit flow rate, respectively.

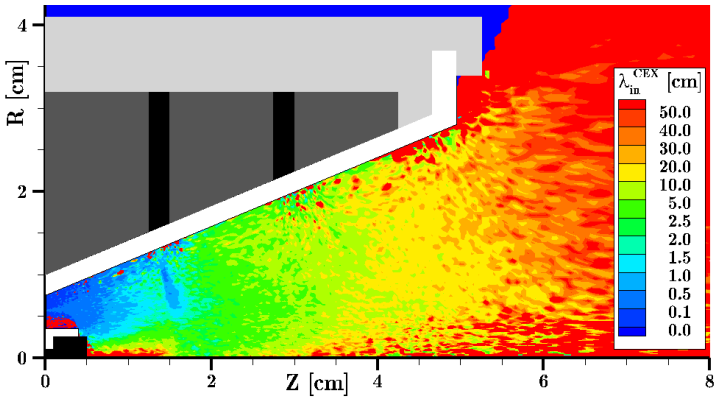

Figure 5-2: Calculated average ion MFP for CEX collisions. This is calculated from the following individually time-averaged quantities: the total ion-neutral CEX reaction rate, ion densities and velocities.

With a neutral injection temperature on the order of .1 $[eV]$, the lower energy contribution from the neutrals is very reasonable. The only mechanism available to explain the elevated neutral energies during the plasma simulation is energy exchange with ions. This is allowed through two avenues in PTpic: charge-exchange and elastic scattering collisions between ions and neutral particles — though the direct energy exchange accompanying a CEX interaction likely eliminates elastic scattering interactions as the primary suspect. Furthermore, although PTpic internally tracks which neutrals exist as a result of charge-exchange collisions, the model is not presently configured to directly extract data of this type. However, useful estimates can be made with the data that PTpic already provides.

Looking ahead to the plasma density and axial ion speeds shown in Fig. 5-3(b) and Fig. 5-4(b), the high density region in C1 and along the axis toward C2 is filled will slow ions. In order to become trapped in these small potential troughs, ions would either need to be created within the region, or lose energy via a collision with a neutral particle. Average ion CEX MFP λ_{in}^{CEX} values are given in Fig. 5-2, and are lower than 2 $[cm]$ in much of the first cusp. Therefore, in this simulation, ions upstream of C2 are likely to transfer their energy to a neutral particle, which may explain how the mass-averaged beam energy per neutral is able to attain a value over ten times larger than that afforded by the neutral temperature. Checking the time-history of neutral thrust also reveals that its starts very close to zero as the plasma simulation is initiated, and grows as the discharge becomes more robust.

Table 5.3: Time-averaged parameters from $t^* = 15 - 36.3$ [μs].

$\langle \dot{m}_{exit} \rangle =$	6.68 [sccm Xe]	$\langle I_a \rangle =$	0.198 [A]
$\langle F \rangle =$	2.72 [mN]	$\langle \eta_a \rangle =$	28.8%
$\langle I_{sp} \rangle =$	422 [s]	$\langle$Ion $I_{sp} \rangle =$	919 [s]

The time-averaged potential distribution in Fig. 5-3(a) shows a drop from 250 [V] to less than 70 [V] in the first ring cusp. This is contrary to inferences made from ion speed measurements [43] and emissive probe measurements [35] (Sec. 3.3.5), which suggest that the largest potential drop should occur past the exit plane, coinciding roughly with the separatrix surface in C3. Matlock's measurements provide a more direct comparison, at an operating condition with $V_a = 250$ [V] and $\dot{m} = 5$ [$sccm\ Xe$], which show potentials as high as 100 [V] at distances 5 [mm] or greater past the exit plane. The simulated potential then steps down in two comparably sized steps going from C1 to C2, and then from C2 to C3. Despite the larger potential drop upstream of C1, the average ion beam energy coincides with the potential difference between C2 and the cathode, rather than the drop that occurs between the anode and C1. Ion origins will be discussed later in this section when simulated ionization distributions are discussed.

Time-averaged electron density n_e is plotted in Fig. 5-3(b), showing densities greater than 10^{19} [m^{-3}] within C1 and in the region between C1 and C2. The greatest densities are on-axis, though rings of electrons are visible in C2 and C3 at density levels that are more in line with those typically found in Hall thrusters. These are an order of magnitude larger than the plasma densities used to initially estimate the feasibility of simulating the DCFT. The dark blue regions show where no super-particles are present, clearly indicating locations where densities fall below the resolvable threshold set by the super-particle size and local mesh resolution. Even with these high densities, the simulation timestep of 10^{-12} [s] is acceptable. Taking γ into account, this timestep is low enough to be on the threshold of stability for electron densities as high as 3×10^{24} [m^{-3}], though densities greater than about 7×10^{22} [m^{-3}] may unacceptably increase the growth of phase errors for particles experiencing Langmuir oscillations. Clearly these are acceptable values from the standpoint of resolving γ-modified Langmuir oscillations. This timestep was originally chosen to match values selected by SPL-PIC when it previously, and incorrectly, used an adaptive timestep. In addition to plasma frequency considerations,

SPL-PIC used the criterion suggested by Hockney and Eastwood [88] (Sec. 9-2-3) to mitigate numerical heating effects in the event that Debye lengths exceed twice the mesh resolution[d]. From these considerations, this choice of timestep may be overly conservative — something that should be exploited in future simulations if mesh spacing allows for it. Although the timestep can not be adaptive using the leapfrog algorithm, acceptable values could be calculated and saved as the simulation progresses to allow an after-the-fact assessment of the specified timestep.

Time-averaged charge neutrality is quantified in Fig. 5-3(c), but the normalized magnitude of the source term used in the potential solver is a factor of γ^2=2500 smaller than these numbers. Also, the coloring in the immediate vicinity of the white walls is meaningless in this figure, other than to indicate the sign of the net wall charge at a given location. From this we see that the walls on the insulating cone are only negatively charged within each ring cusp and past the exit plane, meaning that these are the only locations on the divergent insulator where ion-attracting sheaths can form. The third cusp has an expanded region of negative wall charging on the downstream face of the insulator, which is where many plume electrons arrive after traveling along magnetic field lines from their injection locations. Overall, the uncorrected normalized charge density is within $\pm 10\%$ in regions with the highest plasma densities. The regions with the most net negative charge are along the separatrices in C1 and C2, in the point cusp downstream of the anode, and the small pocket near the corner of the anode that secondary electrons likely play a role in sustaining. The charge density along the separatrices is expected because electrons should coalesce on these surfaces, being strongly magnetized, while ions are electrostatically linked to the electrons even if they are not magnetized. This concentration of negative charge is what provides the curvature in the potential distribution necessary to establish the potential troughs observed with flush wall-mounted Langmuir probes [44]. Another common aspect of the net charge distributions near each ring cusp is the region of positive net charge upstream of each cusp. These features are consistent with the stair-step nature of the simulated potential drop from the anode to the exit plane. In a simple 1D picture, where the electric field transitions from a plateau to a region of increasing electric field strength in the positive-z direction, a positive charge density must be present according to the Poisson's equation for electroquasistatics. The radial and axial electric field components are shown in

[d] This is meant to find a compromise between avoiding the over-resolution of plasma oscillations and not moving electrons in steps that are too large with respect to mesh spacing.

Figs. 5-7 and 5-8.

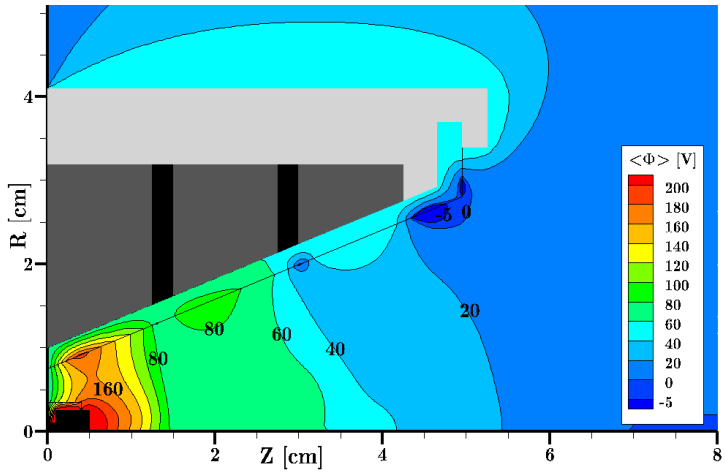

(a) Time-averaged potential.

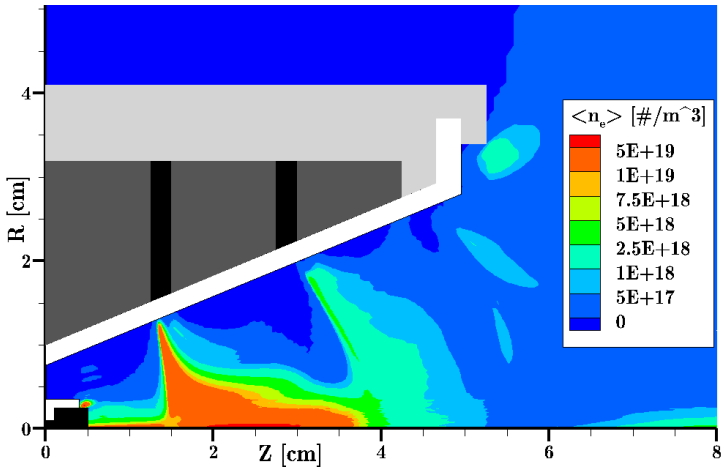

(b) Time-averaged electron density.

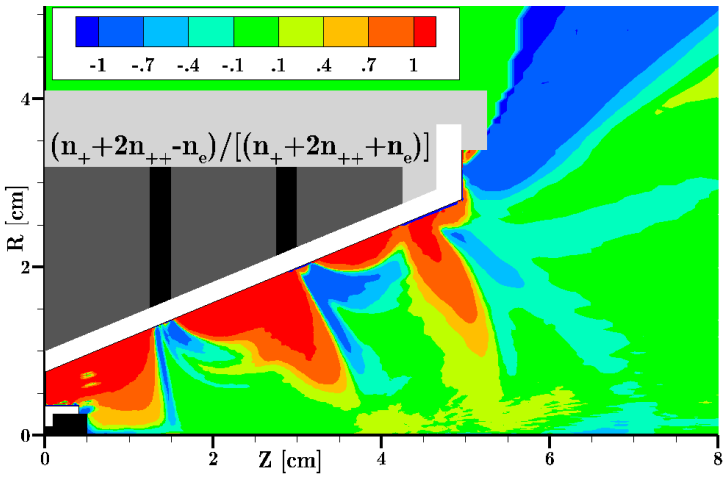

(c) Normalized, time-averaged ρ. Values of $+1(-1)$ indicate regions where only ion(electron) super-particles are present.

Figure 5-3: Time-averaged ϕ, n_e, and charge density ρ during the interval $t^* = 32.67 - 33.17$ $[\mu s]$.

Despite the potential structure not matching expectations based on interpretations of ion speed and plasma potential measurements, PTpic has been able to reproduce one of the main features of the DCFT discharge: the hollow conical plume structure. The hollow structure is apparent in each of the three figures grouped into Fig. 5-4. Fig. 5-4(a) shows the magnitude of an *estimate* of the ion current-density, expressed in Eq. 5.1, because PTpic is not currently configured to calculate the exact average. This should be improved if work with PTpic continues, because it would only be a minor adjustment. For this steady simulation, at least, this discrepancy is not expected to present significant obstacles in interpreting the data.

$$\langle \underline{J}_i \rangle \equiv e\left[\langle n_+ \underline{u}_+ \rangle + 2\langle n_{+2} \underline{u}_{+2} \rangle\right] \approx e\left[\langle n_+ \rangle \langle \underline{u}_+ \rangle + 2\langle n_{+2} \rangle \langle \underline{u}_{+2} \rangle\right] \tag{5.1}$$

Two main concentrations of current density are visible: one along the z-axis and another leaving the domain nearly parallel to the divergent insulator. Within the divergent beam, two beamlets are clearly visible. Typically, current density measurements of the DCFT plume have been interpreted to signify the presence of a single cone of current density [29,132]. However, T. Matlock has demonstrated that two beamlets emerge from what is normally considered a single beam when the magnetic field is altered using an electromagnet placed near the DCFT exit plane [35] (Sec. 3.3.3). Looking back at previous Faraday cup measurements taken in LC-mode, some measurements hint at the presence of two separate jets spaced closely together, such as Fig. 3-13 in Matlock's PhD thesis, and most distinctly in Fig. 5-5 of this document. In Fig. 5-5, the measurements taken closer to the thruster show two local maxima in current density on either side of the thruster, but the distinction is not as clear as the simulated data re-plotted in Fig. 5-4(c). This figure is an attempt to mimic what measurements of the simulation might look like[e]. The current density measurements taken closest to the DCFT, that the author is aware of, were taken by T. Matlock: a lateral sweep of a probe located 7.5 $[cm]$ from the exit plane, presented in Fig. 3-19 of Ref. 35. The measurement plane would be located at about $z = 12.5$ $[cm]$ in the 2D simulated data coordinate system. However, only one peak is discernible in this close scan.

Two other significant features of the simulated distribution of J_i are visible in Fig. 5-4(a),

[e]For Fig. 5-4(c), current density was sampled on a 7.5 $[cm]$ radius circle centered about $(z, r) = (5, 0)$ $[cm]$. Each data point is assigned the angle between the vector $\underline{J}_i$ and the z-axis, and the distance r is calculated between the sampling point and the center of the sampling circle.

though they seemingly do not contribute directly to the formation of the hollow beam. The first is the small stream of ions flowing back to the $z = 0$ $[cm]$ plane. This is one way in which the floating body potential is driven to higher potentials, counteracting electrons that charge the body negatively. The only other unexposed surfaces of the floating body are all downstream of the exit plane. Next, the only significant amount of ion flux impinging on the dielectric wall (colored white) appears in the first ring cusp. This differs from J_i distributions predicted by the simulation of an oscillatory operating condition, described in Sec. 5.2. These differences in current density will also be apparent when simulated erosion rate data from both cases are discussed in Sec. 5.3.

To ensure that the appearance of two beamlets in the simulation can not be explained by a quick shift in the beams position during the time-averaging window, a video was made where each frame displayed the instantaneous total current density sampled once every 25000 iterations — in this case the exact current density can be plotted. The result is visually similar to Fig. 5-4(a), again showing the presence of two separate beamlets within the conical portion of the discharge. Therefore, it is conceivable that this simulation reveals a feature of the discharge that went unnoticed or was unresolvable in some previous measurements of the DCFT operating with its baseline magnetic field. It is also possible that different operating conditions for the baseline DCFT, including background chamber pressure, contribute to the beam appearing as one main divergent peak, or two slightly separated peaks. This has not been an area of experimental emphasis. Rather than continue along these lines, it is more appropriate to keep in mind the extent to which these simulations do not agree with other aspects of existing data, rather than to speculate too much on new interpretations of the data.

First, the divergence angle of the jet is closer to the insulator angle of 22.5° than to typical experimental values slightly greater than 30° — see Fig. 5-5 as an example. Furthermore, the current density on axis is larger than off-axis, which is qualitatively much different than any measurements have indicated. Then there is the matter of ion axial speeds, shown in Fig. 5-4(b), showing ion acceleration well upstream of the exit plane, whereas measurements suggest that ions should not exceed Bohm speeds until reaching the vicinity of the separatrix in C3[43]. Starting with the simulated potential distribution, the discussion will now turn to inferring why the simulated thruster operation still resembles some aspects of a real thruster,

and the main causes of the observed differences. The presence of three different potential drops in Fig. 5-3(a) and the complex axial speed profiles do not provide an obvious picture of how the ion jet forms in the simulation. The plasma density distribution helps to explain why the average ion speed decreases, despite the absence of a strong retarding electric field. To help form a clearer picture of what is happening in the simulation, estimates of the time-averaged volumetric production rate of singly charged ions is plotted in Fig. 5-6(b). The color scale is linear above 5×10^{24} $[m^{-3}s^{-1}]$, showing the most intense ionization in the point cusp downstream of the anode, with significant ionization extending along the centerline to where the C2 separatrix intersects the centerline. Looking now at the electric field components plotted in Fig. 5-7, the ions created upstream of C1 are accelerated by axial electric fields exceeding 10 $[V/mm]$. However, the CEX rate is elevated on-axis downstream of the anode and within the vicinity of the C1 separatrix, visible in Fig. 5-6(c). Even though ions are produced on-axis before C1 where potentials exceed 100 $[V]$, the average energy per ion leaving the simulation is only about 60 $[eV]$. This corresponds much more closely to the potential values in the region between C1 and C2 than to the elevated potentials near the anode. It seems that the consequence of the strongest electric fields moving so far upstream is that ions transfer a significant amount of energy to a small fraction of the neutral population. The smaller cross-sectional area of the chamber results in a relatively high neutral density here, which is also a contributing cause of this scenario.

Now a plausible explanation for the formation of the simulated ion jet begins to materialize. Ions created near the anode experience strong electric fields but lose energy through CEX collisions. Following the streamline overlays in Fig. 5-7, the ions continue downstream from C1 into a week field region. Note that the scale in Fig. 5-7(a) is linear between $-.5 - 5$ $[V/mm]$ and separately in the range $10 - 17.5$ $[V/mm]$. In Fig. 5-7(b), the scale is again piecewise linear between $-1 - 4$ $[V/mm]$ and $4 - 8$ $[V/mm]$. As these ions encounter the C2 separatrix, ions located sufficiently off-axis are accelerated radially and axially by localized maxima in each component of the electric field. This is where the simulated beam gains most of its radial momentum. After traversing another weak-field region between C2 and C3, the ions are again accelerated as they approach a separatrix — this time in C3. Beyond the C3 separatrix, ions in the simulation encounter relatively weak fields, so their trajectories are essentially ballistic in the absence of

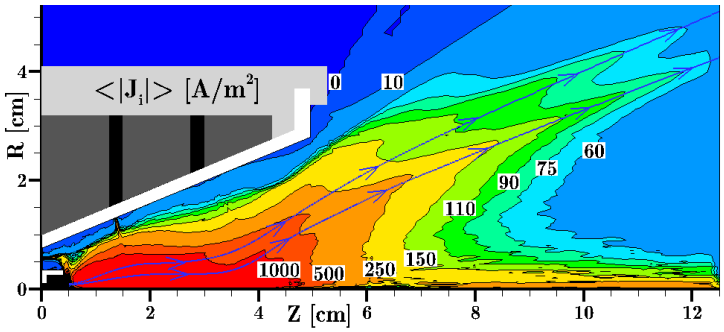

(a) Estimated magnitude of the time-averaged ion current density with two streamlines of the ion mass-averaged velocity.

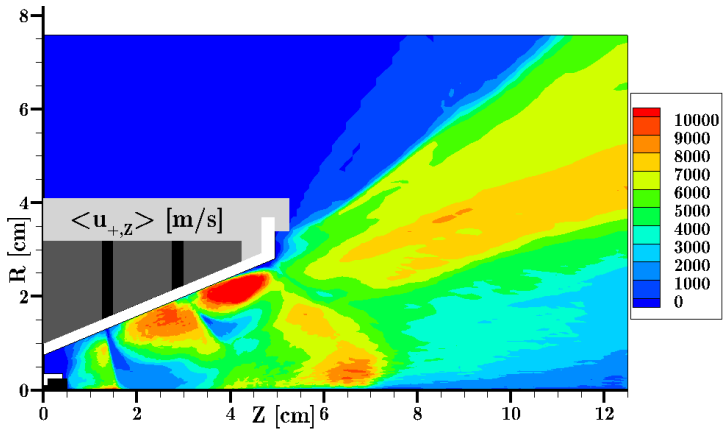

(b) Time-averaged axial ion speed.

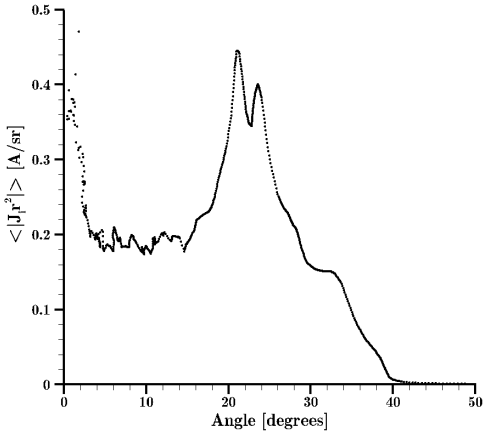

(c) A representation of the variation in ion current-density with the angle made between the bulk flow and symmetry axis.

Figure 5-4: Time-averaged $|J_i|$, $u_{+,Z}$, and an alternate presentation of the same data shown in Fig. 5-4(a) — all during the interval $t^* = 32.67 - 33.17$ $[\mu s]$.

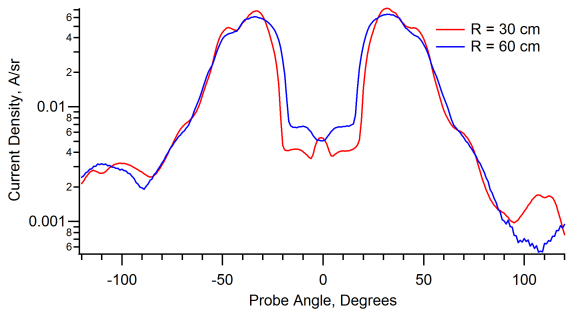

Figure 5-5: Measured angular ion current density distribution operating in LC-mode. Operating condition: $V_a = 476$ [V], $\dot{m}_a = 5.8$ [sccm Xe], $I_a = 0.14$ [A]. Originally presented in Ref. 29.

any scattering events with neutral particles.

A clearer picture of the relationship between the locations of each separatrix and the elevated fields is provided in Fig. 5-8, where the same contours from Fig. 5-7 are plotted with selected magnetic field lines. A closer look at E_z shows that the last two ring-cusps have several common features. Near the wall, the electric field acts to confine ions toward the separatrix. Away from the walls, moving out of the sheath in each cusp, the axial fields weaken in strength. Past these depressions in E_z, both cusps host localized elevations in E_z before the field components decrease again heading into null points in the magnetic field — where the separatrices intersect the centerline.

The radial electric field distribution shows that the simulated beam on the centerline may result from negative E_r values near the axis on either side of C1. This is best illustrated in Fig. 5-8(b), showing that some ions will be confined electrostatically within a tube that extends to roughly half the anode radius and up to about $z = 1.5$ [cm]. This is the same electrostatic ion confinement seen in each of the ring cusps, only this time for the anode point cusp. Ions that stay close to the centerline after crossing the C1-separatrix will again be confined by a much weaker radial field in between the C1 and C2 magnetic null points. Along the C2-separatrix, between the island of strong radial field and center axis, the ions that are not turned away from the z-axis do not experience as strong an axial E-field as those forming the divergent jet, accounting for the low ions speeds visible along the centerline in Fig. 5-4(b).

The radial field may play a similar role in the formation of the two beamlets comprising the main divergent beam, keeping in mind that only ions with energies less than a few electron volts could be magnetized within areas of interest to the present discussion. Figures 1-2 and B-

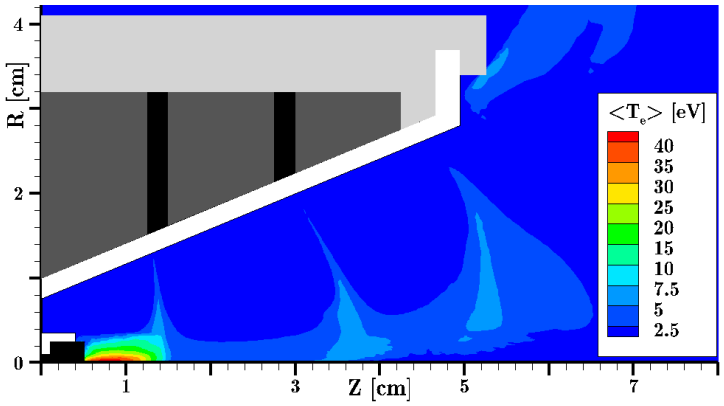

(a) Time averaged electron temperature.

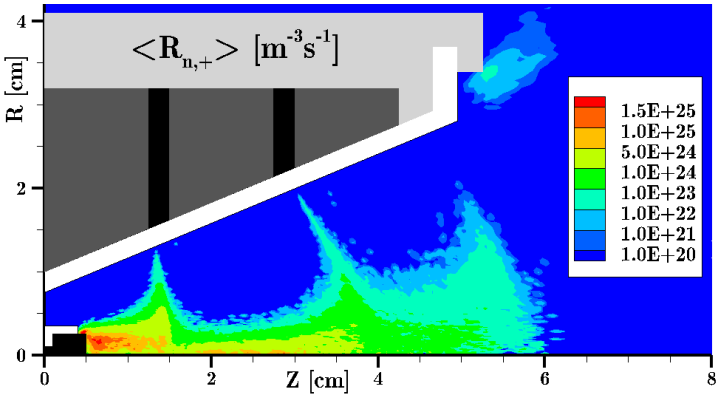

(b) Time averaged single-ion volumetric production rate $R_{n,+}$.

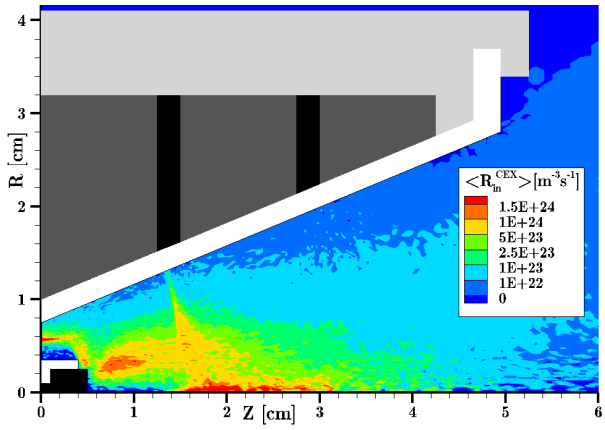

(c) Time-averaged volumetric rate of total ion charge-exchange collisions R_{in}^{cex}.

Figure 5-6: Time-averaged volumetric electron temperature, ionization and CEX rates during the interval $t^* = 32.67 - 33.17$ $[\mu s]$. These values have been reduced by $\sqrt{r}$ from the simulated values.

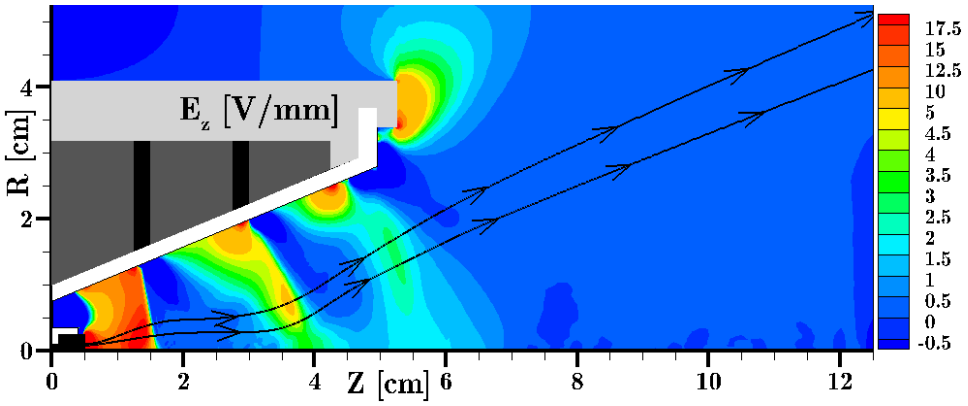

(a) Time-averaged axial electric field.

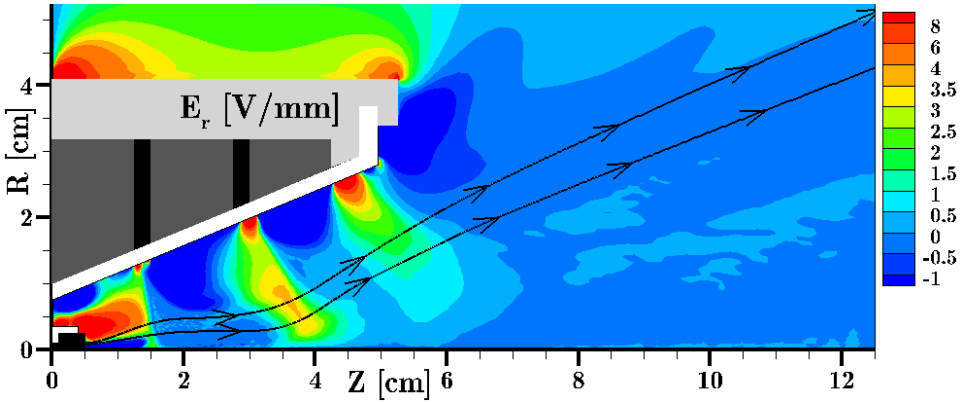

(b) Time-averaged radial electric field.

Figure 5-7: Time-averaged electric field components during the interval $t^* = 32.67 - 33.17$ $[\mu s]$. Note that the scales are different between figures, and irregular within each figure, to point out features at a variety of field strengths.

1, located respectively in Sec. 1.1 and Appendix B, should help convince the reader of this. Focusing again on Fig. 5-8(b), between C1 and C2 a region in which $E_r > 0$ lies between the two magnetic flux lines furthest from the axis downstream of $z = 2.35[cm]$. Following the topmost flux line until about $z = 2.7$ $[cm]$, there is a contour delineating a region where $E_r < 0$, further from the center axis than the previously describe region characterized by $E_r > 0$. This combination of fields is capable of stably confining some ions along the line where $E_r = 0$: for ions with low enough energies, excursions above and below it will be met with opposing electric fields. Shifting attention along the third flux line from the centerline, still between C1 and C2, the situation is reversed. Here, a second contour with $E_r = 0$ is present, but the arrangement of fields reinforces the motions of ions crossing into either surrounding region that does confine electrons. Looking again at Fig. 5-7(b), the streamlines are consistent with the picture just described. The outer beamlet may then be expected to form upstream of C2, off of the center axis within the region where $E_r > 0$. The inner beamlet forms closer to the centerline where $E_r < 0$.

Seemingly, the most important discrepancy between the results of this simulation and measurements is the location of the main potential drop — a common issue in Hall thruster modeling. In this case, the potential drop is suspected to occur too close to the anode, placing the strongest electric fields in the region with the greatest neutral density. The combination of energetic ions and a high-density of neutrals causes PTpic to predict a detrimental amount of CEX collisions within the chamber. This is reflected in the overprediction of neutral thrust and exit flow rate by a large factor. Even with an unrealistic potential distribution, the model still predicts the occurrence of a divergent ion beam, though the precise influences guiding the formation of the simulated beam may not apply if the strongest fields were at the exit. However, the similar distributions of E_z in C2 and C3 may be speculated to remain in the event that a more realistic potential distribution could be simulated. If this is the case, then the main concept applied in the recent design of a new cusped-field thruster may succeed in its primary objective: to reduce the divergence angle of the hollow plume by flattening the curvature of the separatrices in each cusp as much as possible.

Speculation aside, the results of this simulation acutely emphasize the need to improve the predicted potential distribution — namely to coax the largest drop in potential to occur in

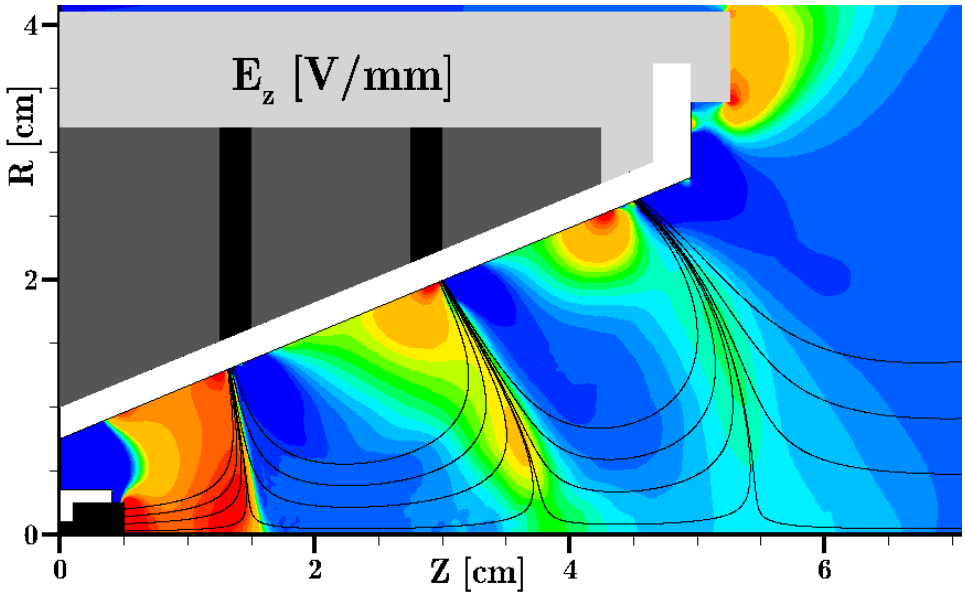

(a) Time-averaged axial electric field.

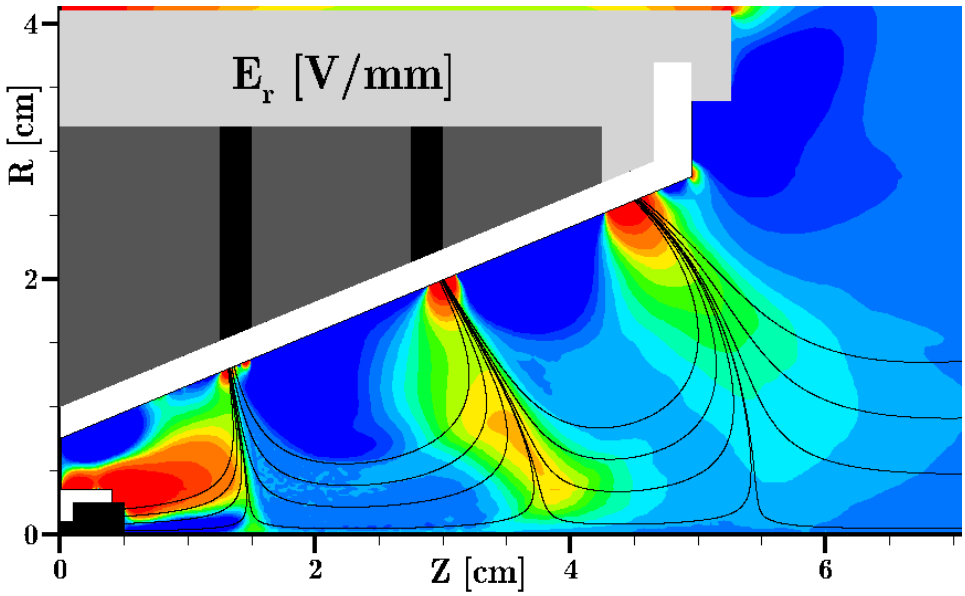

(b) Time-averaged radial electric field.

Figure 5-8: The contours in these figures are on identical color scales as in Fig. 5-7. The same set of selected magnetic field lines are shown in each figure to illustrate the spatial relationship between the separatrix in each cusp and areas with elevated electric fields.

the last cusp. Currently, with the potential drop mainly happening between the anode and C1 separatrix, it seems that electron mobility is too high in the simulation. One immediate suspect is the use of too large a mass factor value, because this requires all simulated classical electron collision frequencies to increase by the factor $\sqrt{r}$. Estimates of the *simulated*, time-averaged electron Hall parameter distribution is provided in Fig. 5-9, confirming that electrons remain magnetized in most places using $r = 1000$, though increasing electron magnetization uniformly, via a reduction in $\sqrt{r}$, might also decrease the relative mobility of electrons across the last cusp. This hypothesis can only be tested by performing simulations with lower mass factors, so it may be difficult to address directly.

A second way that the simulated electron mobility is artificially enhanced is slightly more subtle and not directly quantifiable, but is important to consider nonetheless. For this simulation, electrons were injected pseudo-randomly and uniformly over the entire right side and on the top boundary downstream of $z \approx 8$ $[cm]$[f]. The amount to inject each iteration is set equal to the number of electrons that reach the anode. This uniform distribution spreads electrons to all the field lines reaching allowed injection locations, whereas reality requires that they begin their trajectories on field lines intersecting the cathode orifice. So, it may be that by seeding electrons too close to the z-axis on the right boundary, electrons are able to bypass the cusps with much less effort than in reality.

If electron mobility is already too high because of the mass factor or electron injection procedure, then this precludes the useful application of anomalous collisionality. This has been confirmed in the sense that the use of anomalous collisions has not improved results or shown any hint of displacing the main potential drop further downstream. If a change were made that moved the potential drop too far downstream, or that made classical electron mobility insufficient to sustain a simulated discharge, then this might become an option to attempt once again. Using a large mass factor also limits the amount of anomalous collisionality that can be added to a particular simulation without overly reducing electron Hall parameters in the simulation. It may turn out that a helpful procedure would involve reducing r until the potential distribution moves downstream — or electrons cannot penetrate upstream of C3 — followed by the invocation anomalous diffusion, if needed. However, with very little empirical

[f]The locations to seed electrons on boundaries must be specified in the `IN__boundaries` input file.

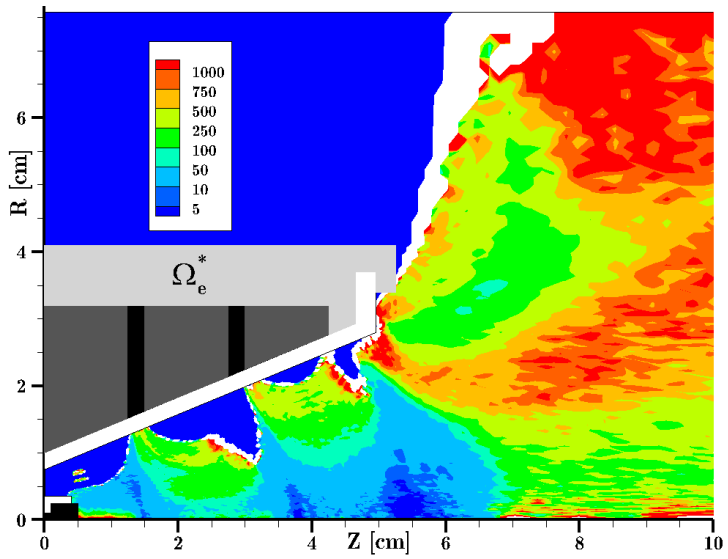

Figure 5-9: Time-averaged simulated electron Hall parameter Ω_e^* during the interval $t^* = 32.67 - 33.17$ $[\mu s]$. White areas, aside from the legend and insulator, indicate regions where the average electron collision frequency was zero-valued and $\langle n_e \rangle \neq 0$.

input regarding anomalous diffusion in cusped-field thrusters, nor the extent of their similarities to Hall thrusters in this aspect, specifying the correct amount of anomalous diffusion may prove challenging.

5.2 Oscillatory Simulation

All but two of the parameters used in the previous section, and shown in Table 5.2, were held constant for the simulation described here. The entire simulation uses a super-particle weighting $W = 5 \times 10^8$ and the floating body capacitance is 1000 times smaller, so $C_f = 39.1$ $[pF]$. Even though the thruster operating condition is the same as in Sec. 5.1, these results more closely resemble the HC-mode of DCFT operation. As a result, the data suggest important differences between the two modes that will be described in this section, and in Sec. 5.3.

The overall behavior of this simulation is summarized in Fig. 5-10. A fixed timestep was used after $t^* \approx 2.7 \mu s$, leading to a series of large amplitude, low frequency oscillations resembling the high-current mode of the DCFT. Reduced by $\sqrt{r}$, the estimated frequency of the slowest component of the oscillation is slightly greater than 6 $[kHz]$. This is consistent with an estimated neutral residency time for xenon particles injected at .1 $[eV]$ traveling a distance of

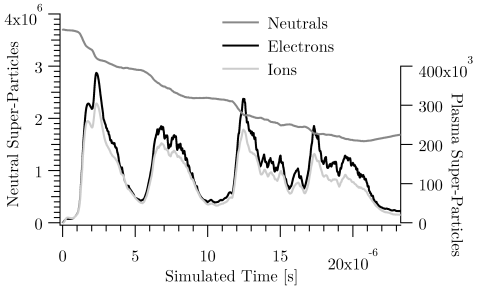

(a) Super-particle counts.

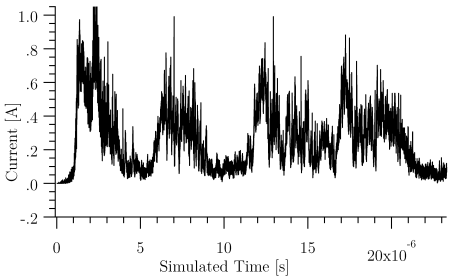

(b) Anode current.

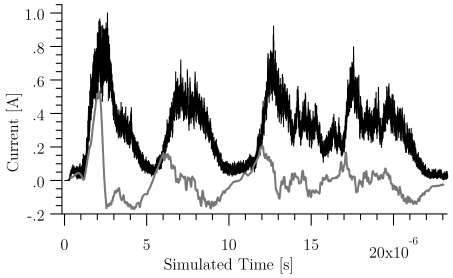

(c) Ion(black) and net beam(gray) currents. The net beam current is subjected to an additional boxcar average encompassing 1000 of the pre-averaged data points.

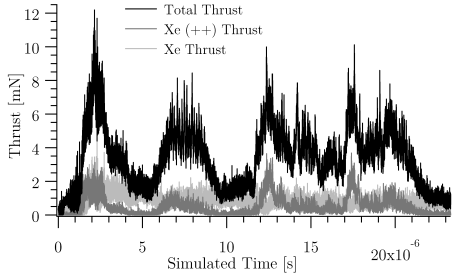

(d) Thrust and its components.

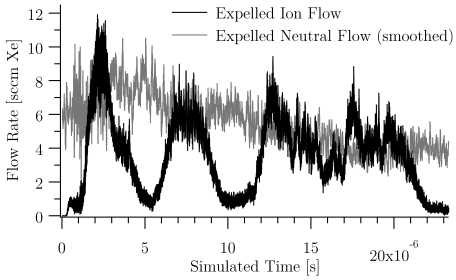

(e) Propellent flow rates. The expelled neutral flow is subjected to an additional boxcar average encompassing 50 of the pre-averaged data points.

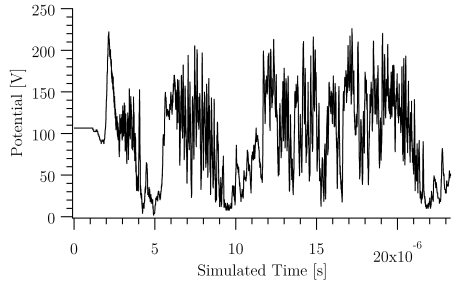

(f) Floating body potential.

Figure 5-10: Each quantity is time-averaged within groups of 1000 iterations. The simulated data shown here spans an interval $t^* = 1.27\tau_r$ after completing 25.13 million iterations. A constant timestep was used after $t^* = 2.7$ $[\mu s]$.

about 5 [cm]. As the simulation continues, however, the waveform changes as the total number of neutral super-particles levels off. Before its transition to a weaker state, the discharge does complete three oscillation cycles — not counting the first. The last two cycles show signs of the deterioration that eventually dominates the simulation, so this discussion of a simulated HC-mode operation is limited to the cycle occurring between $t^* \approx 5 - 10$ [μs].

Time averaged performance metrics are listed in Table 5.4. With the operating condition unchanged relative to the simulation in Sec. 5.1, these should be compared to the same data as before, indicating that many of the performance parameters are again underpredicted. Relative to the steady simulation, the anode current and thrust are 11% and 8.6% larger, respectively. In Fig. 5-10(e), the injected neutral flow is not shown, but it was fixed at 6 [$sccm\ Xe$] and a plot of the smoothed data is very similar in appearance to the injected flow rate plotted in Fig. 5-1(d). Similar to the steady case, the simulation also predicts too high a contribution from neutral particles to performance: neutral thrust is 0.83 [mN], beam energy per neutral is 1.3 [eV], providing 23.5% of the total useful beam energy, and the utilization efficiency is 34.5%. The average ion contribution to useful beam energy is 40% less in this case, at 33 [eV]. The same arguments supporting an excessive amount of CEX collisions near C1, as a result of the largest potential drop materializing in this region, apply to this case as well.

Table 5.4: Time-averaged parameters from $t^* = 5.87 - 11.70$ [μs].

$\langle \dot{m}_{exit} \rangle =$	9.29 [sccm Xe]	$\langle I_a \rangle =$	0.215 [A]
$\langle F \rangle =$	3.03 [mN]	$\langle \eta_a \rangle =$	24.8%
$\langle I_{sp} \rangle =$	339 [s]	$\langle$Ion $I_{sp}\rangle =$	713 [s]

From running a large number of preliminary simulations, one sign that a simulated discharge has been successfully initiated — when using values of C_f determined in the absence of plasma — is the onset of high-frequency, large amplitude oscillations in the floating body potential ϕ_f. This transition is visible in Fig. 5-10(f) between $t^* = 2 - 3$ [μs], as well as in each case discussed in Sec. 5.4 when $C_f = 39.1$ [pF]. The floating body potential also appears to play an important role in determining whether or not a simulated discharge is able to sustain itself. If C_f is driven too low by electrons, and ions do not compensate to bring the charge back up, then the voltage differences between where electrons are injected and the exit of the thruster may not be large

enough to drive electrons across the C3 separatrix in sufficient numbers. Since the injection of electrons follows the anode current[g], when ϕ_f and I_a take on low values at the same time, the simulation may be deprived of the electrons it needs to continue. This is discussed again in Sec. 5.3.

The high frequency components of anode current, present in all non-smoothed integral ion parameters, and ϕ_f — visible in Figs. 5-10(b)–5-10(f) — have simulated frequencies between $6.32 - 7.91$ $[MHz]$. If this is interpreted as a frequency limited by the mobility of the heavy species, then the predicted frequency range is $200-250$ $[kHz]$. This lies in the frequency range of ion transit-time oscillations observed in Hall thrusters[160]. Lower frequency components, aside from the primary HC-mode oscillation, have been measured in the range of tens of kilohertz[35] (Sec. 4.4) leading up to mode transitions in the DCFT. However, anode current measurements for signal components above 100 $[kHz]$ have not been presented in detail. Anecdotally, the author has observed low amplitude components of the anode current signal in the range of several hundred kilohertz during DCFT operation for the LIF measurements performed by N. MacDonald[43]. However, the data containing this information failed to write to disk. In any event, these simulated high-frequency oscillations coincide with plasma particles leaving the domain in clumps, rather than smooth streams. This is evident in the fluctuations in ion beam current, exit flow rate, thrust, and the net rate at which electrons leave the domain (not shown). Based on the differences observed in simulations using different values of C_f and super-particle weights, it appears that both of these parameters can influence the simulated amplitudes of high-frequency oscillations. See Fig. 5-1, and the results presented in Sec. 5.4, for instance.

For the remainder of this section, as well as in most Sec. 5.3, the discussion will focus on the oscillatory simulation in the interval of simulated time shown in Fig. 5-11. To a lesser degree than in the steady case, this simulation reproduces certain aspects of the hollow plume, though not during the entire oscillation period. This is easiest to see in plots of the axial Xe $(+)$ speeds $u_{+,z}$ shown at four different frames in Fig. 5-12. During the broad peak in anode current, $u_{+,z}$ appears as it does in Figs. 5-12(a) and 5-12(b), except in C3 where $u_{+,z}$ oscillates between the distributions shown in both figures. As in Sec. 5.1, elevated ion speeds first appear upstream

[g]To initiate a discharge, PTpic allows the user to set an injected electron current independent of the anode current within the `IN__VARIABLES` file. Other settings enable PTpic to shutoff this extra current automatically when certain adjustable criteria are met.

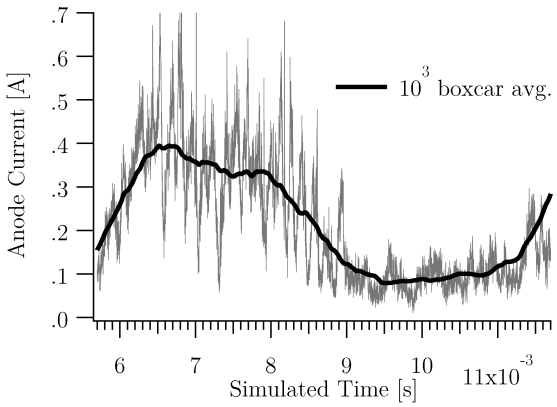

Figure 5-11: Detailed view of I_a for $t^* = 5.70 - 11.70$ $[\mu s]$ showing the best representation of an HC-mode oscillation cycle from the simulation summarized in Fig. 5-10.

of, and within C1. Aside from the low-density regions along the walls and between cusps, bulk ion speeds decrease past this point. The same explanations given in the previous section apply here as well, relying on low-energy ions created via ionization and CEX collisions.

The conical distribution of faster ions past the exit plane does not begin to materialize until about $t^* = 9.05$ $[\mu s]$ and remains until $t^* = 11.90$ $[\mu s]$. Therefore, the simulated beam persists while the anode current is depressed and also when it is rising up toward the broad peak. This is in qualitative agreement with time correlated LIF measurements of axial ion speeds downstream of C3, shown in Fig. 7.9 of Ref. 37. The data for the third point in MacDonald's measurements show an elevated most probable axial speed during the interval in which I_a is low, and the first significant decrease in speed is measured to occur at about the same time that the anode current has reached its maximum value. The location of the measurement, nor the magnitude of the speeds, are not in agreement with the simulated data. Attempts to compare measurements further upstream showed no clear agreement with the trends from the data, though the statistical noise from the simulation, and the presences of strong simulated gradients near measurement points do not allow for clear comparisons to be made.

Plots of total ion current density in Fig. 5-13 show that the simulated plume structure is less hollow than in the steady simulation. Before the divergent jet becomes visible, J_i is peaked on the axis. In fact, looking at Fig. 5-10(d) the thrust is greatest mostly during the anode current peak, including the times corresponding to the first two current density plots. In this sense, the existence of the divergent beam is less significant in this simulation because the thrust levels

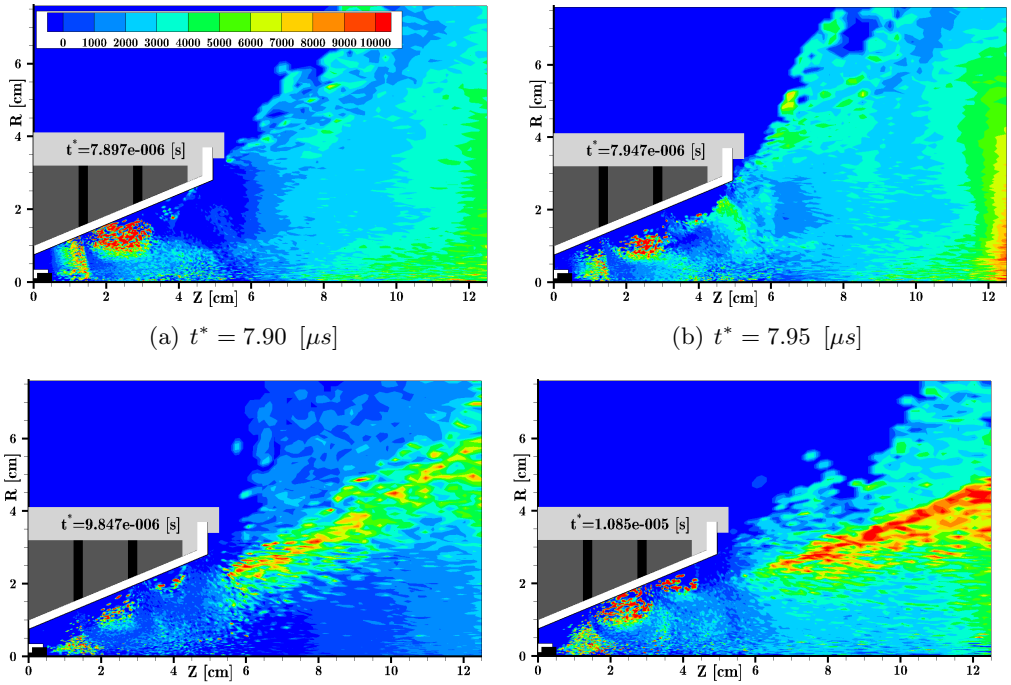

(a) $t^* = 7.90\ [\mu s]$ (b) $t^* = 7.95\ [\mu s]$

(c) $t^* = 9.85\ [\mu s]$ (d) $t^* = 10.85\ [\mu s]$

Figure 5-12: Axial, singly charged ion speed at selected frames in the interval shown in Fig. 5-11. The color scale is identical in each frame.

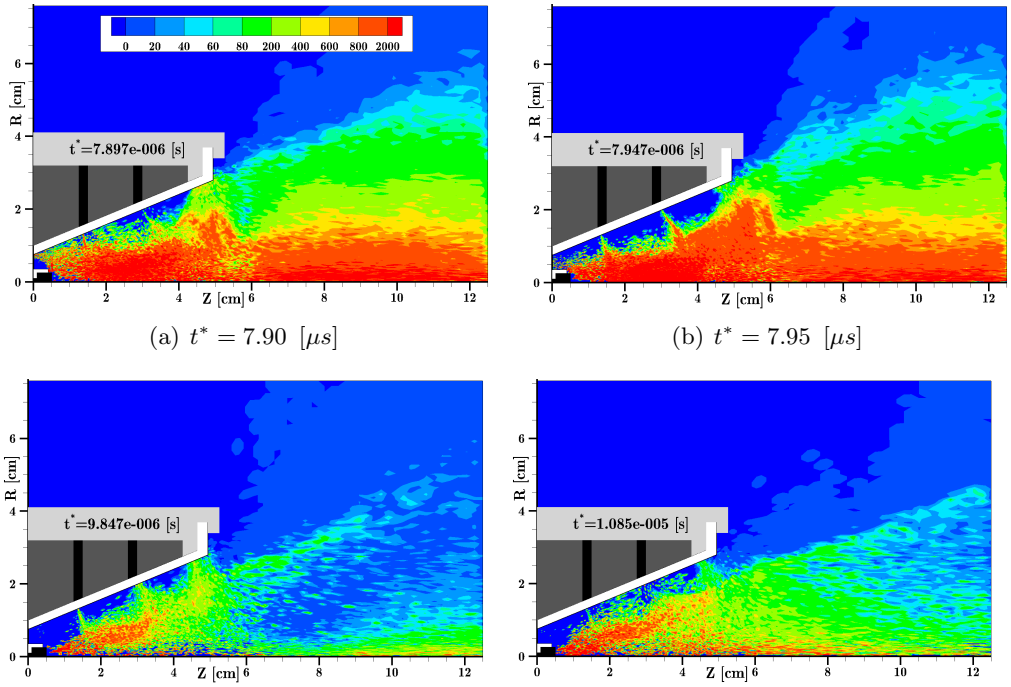

(a) $t^* = 7.90\ [\mu s]$ (b) $t^* = 7.95\ [\mu s]$

(c) $t^* = 9.85\ [\mu s]$ (d) $t^* = 10.85\ [\mu s]$

Figure 5-13: Total ion flux magnitude at selected frames in the interval shown in Fig. 5-11. The color scale is identical in each frame.

are lowest when the divergent beam is present.

In the context of the discussion in the next section, the most important differences between the steady and oscillatory distributions of J_i can be seen by comparing values between the two in each of the cusps. In the steady simulation, visible in Fig. 5-4(a), significant ions fluxes are present only in the first ring cusp. In Fig. 5-13, appreciable values of J_i can be seen in all three cusps except in Fig. 5-13(d), where the ion current density is lower in C2 than the other two. As expected, these differences result in very dissimilar predictions for the distribution and magnitude of erosion on the insulator wall.

5.3 Simulated Erosion Data

PTpic was used to calculated erosion rates on dielectric wall surfaces for the steady and oscillatory simulations presented in Secs. 5.1 and 5.2, respectively. A detailed description of how the erosion process is modeled has been provided in Sec. 3.2.2, and is based on the works of S. Cheng[15] and J. Yim[50]. Erosion data are accumulated as ions and CEX neutrals encounter dielectric surfaces on a fixed mesh, providing average erosion rates over specified simulated intervals of time. The mesh is not adapted to provide feedback into how the changing geometry might change erosion rates. Also, redeposition is not included as part of this calculation. Lastly, simulated erosion rates are interpreted as being larger than physical erosion rates by a factor $\sqrt{r}$ due to enhanced ion fluxes to all surfaces resulting from the artificially reduced xenon mass. All erosion rates presented here represent the physical erosion rates, in that they have been reduced by the appropriate factor relative to simulated erosion rates.

The only measurements of DCFT erosion available at the time of this writing are provided again here, in Fig. 5-14, a duplicate of Fig. 2-21. They were acquired by operating the DCFT for 204 $[h]$ in the HC-mode, using an anode voltage of 300 $[V]$ and an anode mass flow rate of 8.5 $[sccm\ Xe]$. The uncertainty is estimated at ± 0.1 $[\mu m/h]$ in the regions of peak erosion. The estimated time needed to erode through the insulator is 1220 $[h]$, based on extrapolation. Section 2.3 provides a detailed description of these measurements. As an example to provide some intuition for the magnitudes discussed in this section, an erosion rate of 1 $[\mu m/h]$ will consume 1 $[mm]$ of material in 1000 $[h]$, or 1 $[cm]$ of material in 10000 $[h]$. The starting minimum thickness of the conical DCFT insulator is ≈ 2.31 $[mm]$.

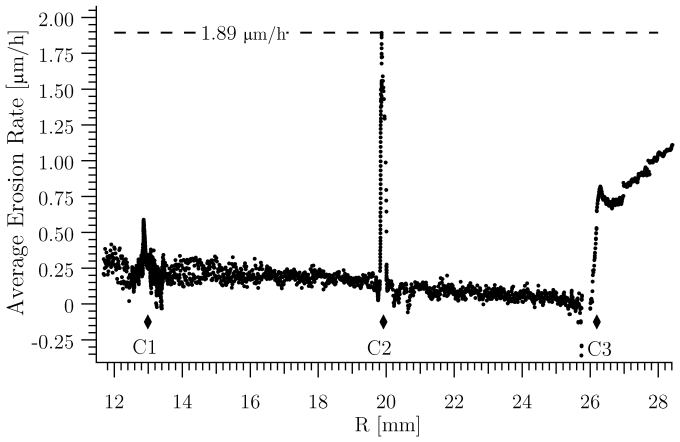

Figure 5-14: Measured values of average erosion rates in the DCFT operating with $V_a = 300$ $[V]$ and $\dot{m} = 8.5$ $[sccm\ Xe]$. These data are discussed in more detail in Sec. 2.3.

The simulated erosion data presented here are generated for a DCFT operating with $V_a = 250$ $[V]$ and $\dot{m}_a = 6$ $[sccm\ Xe]$. Therefore, these simulations were completed with the intent to predict whether or not significant differences in erosion should be expected based on the operational mode of the DCFT. Though erosion data has only been recorded in the HC-mode of operation, measurements of the DCFT in low-current mode may be performed by R. Sullivan et al. before the end of 2013. It has been speculated by the author and collaborators that erosion is greater in HC-mode, possibly caused by elevated ion current densities in cusps that may be present as a result of the characteristic large amplitude discharge oscillations [128].

5.3.1 Simulated Erosion for an Oscillatory Operating Condition

In order to examine average erosion rates over the course of an oscillation, data saved from the simulation described in Sec. 5.2 at $t^* = 5.70$ $[\mu s]$ was used to seed the first simulation during which erosion was calculated. Erosion was not calculated originally because this capability was added after the first simulation was performed. The only other significant change in PTpic made after the original simulation relates to how often the potential solution is updated in the event that the maximum electron cyclotron frequency exceeds the timestep based on the γ-modified plasma frequency and numerical heating considerations: the criteria that would allow PTpic to skip updating the potential solution, and for deciding how many iterations it could skip this important step, was made more stringent as a consequence of the transition to operating

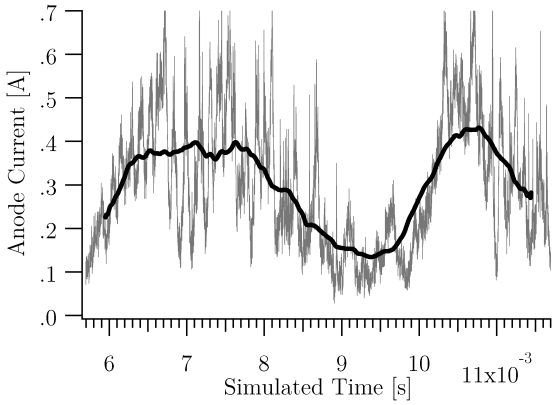

Figure 5-15: Detailed view of I_a for $t^* = 5.70 - 11.70$ $[\mu s]$. The black line shows a moving boxcar average using 10^3 data points.

exclusively with a fixed timestep. Aside from this change to the operational portion of the source code, no other changes to the model's input parameters were made for the simulated data presented in Figs. 5-15–5-17, except for the sequence of pseudo-random numbers generated throughout the course of this simulation and its predecessor.

The time-resolved anode current for the first oscillatory erosion simulation is visible in Fig. 5-15, and should be compared directly to Fig. 5-11. In this case, the simulated anode current is depressed for a shorter amount of time than in the original simulation, and enters a second peak that is narrower than the first. This simulation was not to continued past $t^* = 11.70$ $[\mu s]$, but in all likelihood it would eventually become steady as it continued, much like the original simulation. Other than this difference, the amplitude and duration of the first peak is similar to the original case, so the evolution of the average erosion rate, as time progresses though this single cycle, was selected for closer study. This provides the best representation of cycle-averaged erosion that can be estimated from the current set of results. The interval selected to represent a full oscillatory cycle is $\Delta t^* = 5.70 - 9.90$ $[\mu s]$.

The average erosion rate is plotted with respect to the radial coordinate along the divergent insulator wall in Figs. 5-16 and 5-17 for progressively longer intervals of time. Data was only extracted at nodes in the mesh, so the horizontal segments in these figures provide a visual representation of mesh resolution along the wall, which is finer within cusps. The three locations where the separatrix in each cusp intersects the wall are marked by the black diamonds in Fig. 5-14. The result from the first tenth of a simulated microsecond, encompassing only 10^5 iterations,

is shown in Fig. 5-16(a). At this point, a large amount of noise is present in the data because of the large statistical weights $W = 5 \times 10^8$ of the super-particles, and the relatively small number of particles that have hit the wall.

As the simulation continues, the magnitude of finite sampling effects diminish gradually. After 8×10^5 iterations and $\Delta t^* = .8$ $[\mu s]$, some features begin to materialize in the data, referring specifically to Fig. 5-16(b). Up to this point, the fastest erosion has taken place in the first cusp, where the contributions from CEX neutrals and ions are similar. Over the next 1.2 $[\mu s]$, bringing the total simulated interval to 2.0 $[\mu s]$, ion erosion in C3 nearly matches to total accrued in C1 in just 60% of the elapsed time. Erosion in the second cusp is observed to increase slightly during this interval as well.

Shifting to the final sequence of data selected for this simulation, shown in Fig. 5-17, previous trends are seen to continue as erosion in C3 surpasses the rate in C1 at $t^* = 8.40$ $[\mu s]$. After only a slight increase in leading up to Fig. 5-17(a), between $t^* = 8.4 - 9.1$ $[\mu s]$ the maximum average erosion rate in C2 is approximately 2.8 $[\mu m/h]$, increasing the overall average from about 0.2 $[\mu m/h]$ to above 0.55 $[\mu m/h]$ in less than a simulated microsecond. The average rate in C3 continues to increase as well.

Average erosion rates for the duration of the representative cycle are shown in Fig. 5-17(c). The peak in erosion due to CEX neutrals near and upstream of C1 is consistent with reasons offered in Secs. 5.1 and 5.2 to explain the unexpectedly large contribution to thrust from neutral particles, as well as low ion speeds downstream of C1. Erosion predictions in the steady simulation show CEX erosion in excess of ion erosion in this region — see Fig. 5-23. Perhaps the least expected prediction is the amount of erosion upstream of the first cusp. No measurements in this region are available, but there have been no visual indications that erosion might be important to consider in this area, though significant discoloration is observed — likely caused by redeposition of the porous stainless steel disc through which neutrals are injected. A photograph showing deposition on the DCFT after 204 $[h]$ of operation is shown in Fig. 2-23.

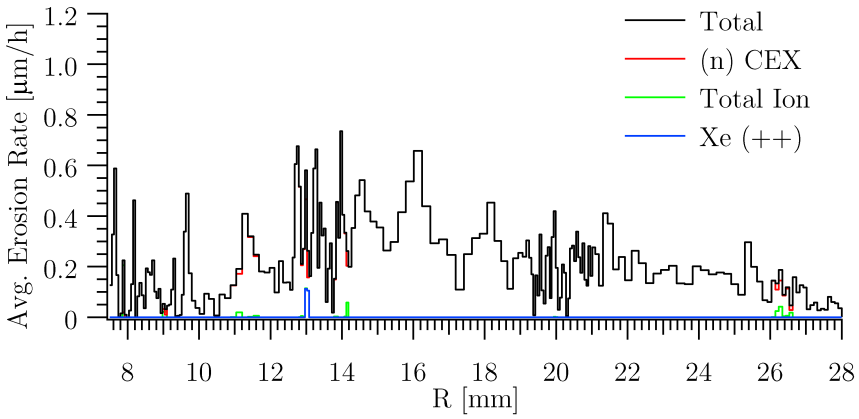

(a) Average erosion rate from $t^* = 5.70 - 5.80$ $[\mu s]$.

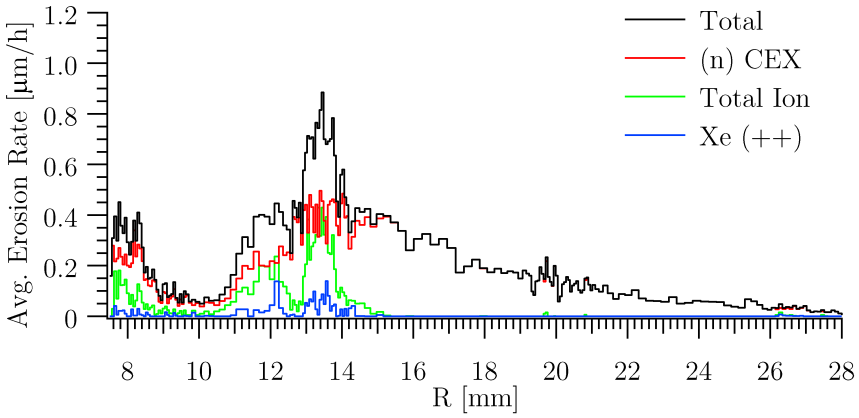

(b) Average erosion rate from $t^* = 5.70 - 6.50$ $[\mu s]$.

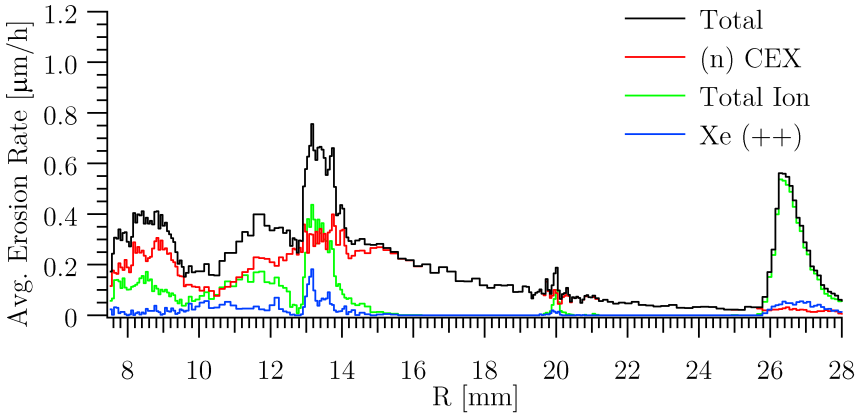

(c) Average erosion rate from $t^* = 5.70 - 7.70$ $[\mu s]$.

Figure 5-16: Average erosion rates for a rerun of the oscillatory simulation described in Sec. 5.2.

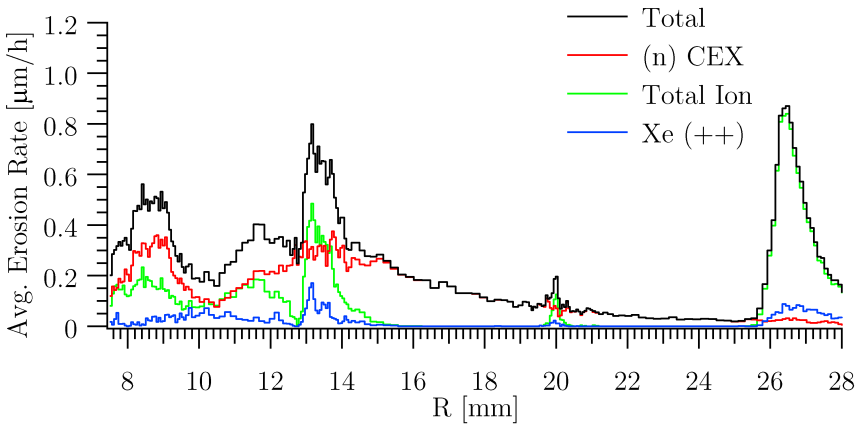

(a) Average erosion rate from $t^* = 5.70 - 8.40$ $[\mu s]$.

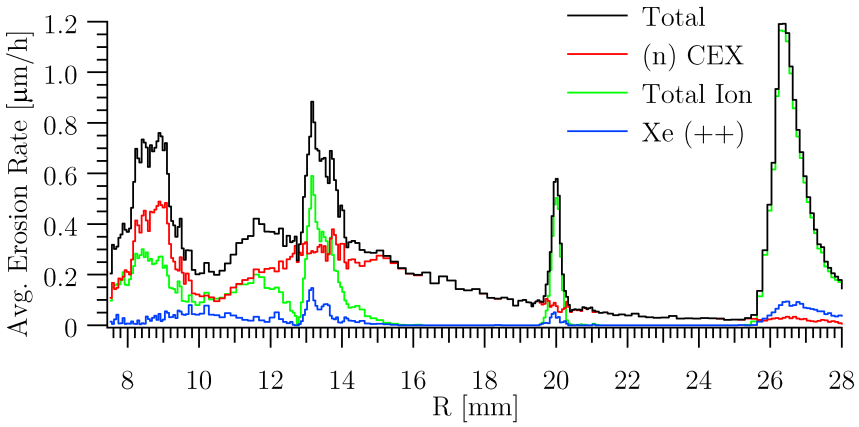

(b) Average erosion rate from $t^* = 5.70 - 9.10$ $[\mu s]$.

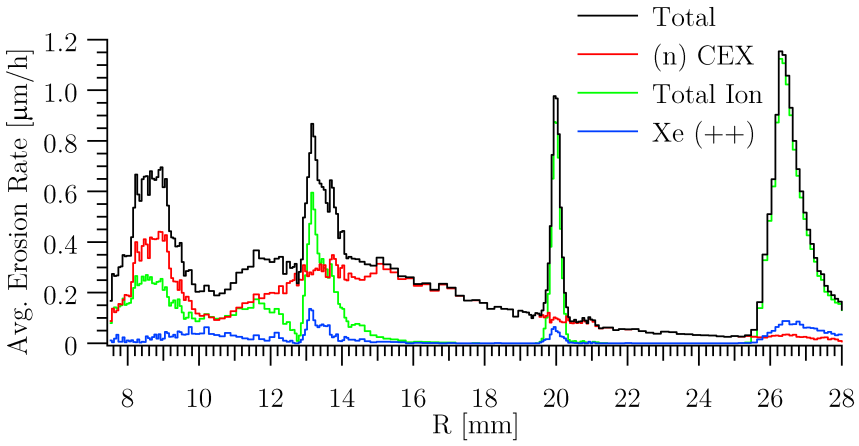

(c) Average erosion rate from $t^* = 5.70 - 9.90$ $[\mu s]$.

Figure 5-17: A continuation of Fig. 5-16, showing the temporal evolution of average erosion rates along the insulator wall.

Deferring further discussion of these results to Sec. 5.3.3, results from a similar simulation of the same period of time are now presented. The only simulation parameters that differed between these two simulations relate to super-particle weights. The baseline super-particle weight W was reduced by a factor of 4 from 5×10^8 in the previous simulation to 1.25×10^8. The maximum relative size of neutral super-particles was also decreased by a factor of 4. These effects combined to halve the expected noise level from charged particles, and the reduce statistical variations due to finite neutral statistics by a multiple of 4.

The waveform for I_a in this simulation is presented in Fig. 5-18, indicating a different behavior relative to the simulations summarized in Fig.5-11 and Fig. 5-15. Though one cycle is present, in the sense that the anode current rises and then falls, there is no second current, indicating that the change in super-particle weights has caused this simulation to decay to a more quiescent state sooner than its counterparts. The shape of the anode current cycle is also different, but the behavior of the discharge has many similarities to the previous simulations before $t^* = 10$ $[\mu s]$, such as a divergent beam of faster ions. As is typical of the DCFT simulations run for this study, the decay of a discharge is typically accompanied by a significant drop in the floating body potential ϕ_f. When this occurs, potentials near and past the exit of the thruster may be too low to maintain the flux of electrons into the chamber necessary to start the next cycle. This may be attributable to using too low a value of the floating body capacitance, the chosen method of electron injection[h], or some other aspect of the simulation. Most likely, it is a combination of several factors: keeping C_f constant or a lack of resolution for low plasma densities are two other possibilities. However, a persistently low simulated value of ϕ_f tends to be an indicator that a simulated plasma discharge has been effectively extinguished. Continuing simulations past similar points, in the author's experience, has not resulted in the reignition of any discharges.

Similar to the data shown in Fig. 5-16, in Fig. 5-19 the average erosion rate is dominated by statistical noise early on, but then acquires a more definite shape as the simulated interval of time increases. The benefit of using more computational particles to represent the same number of physical particles is discernible in the reduced noise in these data. Compare Figs. 5-19(a)

[h]When a depressed ϕ_f value is in phase with a depressed I_a value, a relatively small amount of electrons will be injected at the exit, and the field drawing them into the thruster will be weak because ϕ_f is depressed and the boundary is fixed to 0 $[V]$.

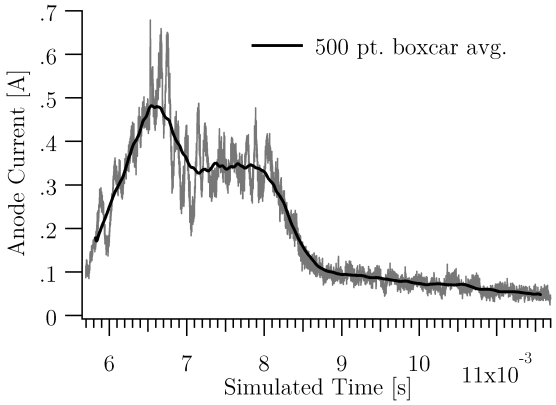

Figure 5-18: Detailed view of I_a for $t^* = 5.70 - 11.70$ $[\mu s]$ for the second oscillatory simulation run to generate erosion data.

and 5-19(b) to Figs. 5-16(a) and 5-16(b), for instance, keeping in mind the different scales in the two sets of figures. From Fig. 5-19(b), erosion first starts in and around C1, followed by the onset of erosion in C3 seen in Fig. 5-19(c).

Moving forward in time to the data shown in Fig. 5-20, at $t^* = 8.40$ $[\mu s]$ the most significant change relative to Fig. 5-19(c) is the near doubling of the CEX erosion rate upstream of C1. Then, similar to the previously presented simulation, between $t^* = 8.40 - 9.90$ $[\mu s]$ erosion in the second cusp increases sharply, and in this case exceeds erosion in C3. In order to do so, the average erosion rate between $t^* = 8.40 - 9.10$ $[\mu s]$ increases to 5.4 $[\mu m/h]$. As this happens, ion erosion in all three cusps increases simultaneously by varying amounts. In C3, the average erosion rate increases from .35 $[\mu m/h]$ at $t^* = 8.40$ $[\mu s]$ to about 1.0 $[\mu m/h]$ at $t^* = 9.90$ $[\mu s]$, giving an average rate of 2.2 $[\mu m/h]$ in between. As the end of the anode current pulse is reached, the average erosion rate in C2 decreases from its earlier maximum value to 2.5 $[\mu m/h]$. An example calculation showing an explicit example of what is meant by an average erosion rate between two times is provided in Eq. 5.2 for the case of the average erosion rate in C2 from $t^* = 8.40 - 9.10$ $[\mu s]$.

$$5.4\ [\mu m/h] \approx \frac{1.2 \times (9.1 - 5.7) - 0.1 \times (8.4 - 5.7)\ \ [\mu m/h][\mu s]}{(9.1 - 8.4)[\mu s]} \tag{5.2}$$

This second simulation makes similar predictions as the first: at an oscillatory operating condition, erosion should occur primarily in each cusp, and upstream of C1. The main difference

between the two is the ordering of each region in terms of increasing erosion rate. In Fig. 5-17(c), the first simulation predicts increasing erosion rates moving from upstream of C1 to the exit, whereas the second simulation, in Fig. 5-20(c), predicts that erosion upstream of C1 is greater than within C1, and that the overall peak in erosion should occur in C2. Clearly though, these predictions are sensitive to input parameters that have been specified, as are the anode current waveforms.

Beforehand, the time interval in which erosion rates were calculated was selected to emulate what the same data might represent if it had been calculated during a repeating sequence of anode current oscillations. In the absence of a better option, the two cases just described were used as substitutes. In order to show the effect of this choice on the predicted erosion rates, average erosion rates over the entire interval, $t^* = 5.70 - 11.70$ $[\mu s]$, are given in Fig. 5-21. Comparing Fig. 5-21(a) with 5-17(c), and Fig. 5-21(b) with 5-20(c), the effects of the last 1.8 μs on the averages can be observed.

In the first oscillatory simulation, the C1 erosion rate does not change significantly, while erosion rates in the other regions of interest decrease by small amounts — between $.05 - .15$ $[\mu m/h]$, approximately. However, the simulation did not complete the second, shorter cycle so the averages that might best represent erosion during a cycle for this simulation are presented in Fig. 5-17(c). The second simulation is different in the sense that the simulation does not recover into a second cycle at all, so lower erosion rates should be expected after $t^* = 9.9$ $[\mu s]$ because the discharge is less robust. If no erosion at all were to occur from $t^* = 9.9 - 11.7$ $[\mu s]$, each of the erosion rates would decrease by 30%. This is nearly the case, except in C2 where the average erosion rate decreases by only 15%. Therefore, the most relevant erosion rate predictions in the HC-mode obtained from these simulations are those shown in Figs. 5-17(c) and 5-20(c).

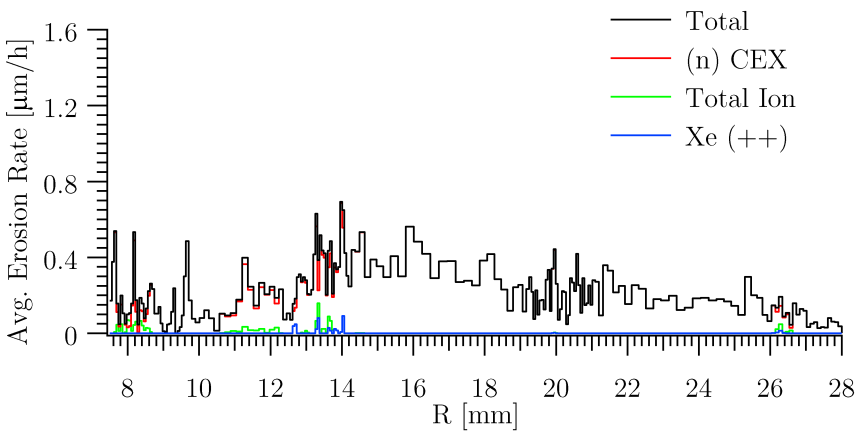

(a) Average erosion rate from $t^* = 5.70 - 5.80$ $[\mu s]$.

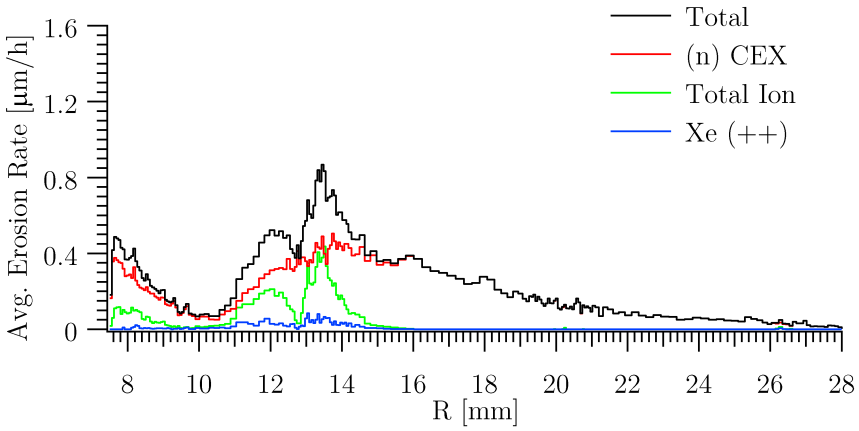

(b) Average erosion rate from $t^* = 5.70 - 6.50$ $[\mu s]$.

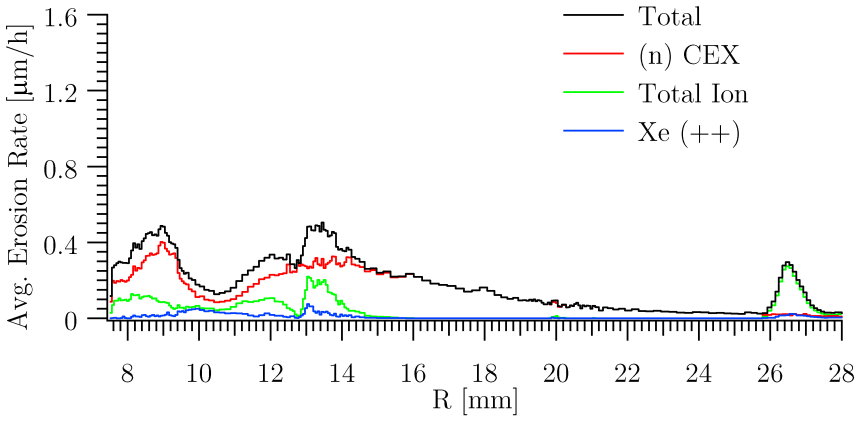

(c) Average erosion rate from $t^* = 5.70 - 7.70$ $[\mu s]$.

Figure 5-19: Average erosion rates for a rerun of the oscillatory simulation described in Sec. 5.2 using a smaller W value.

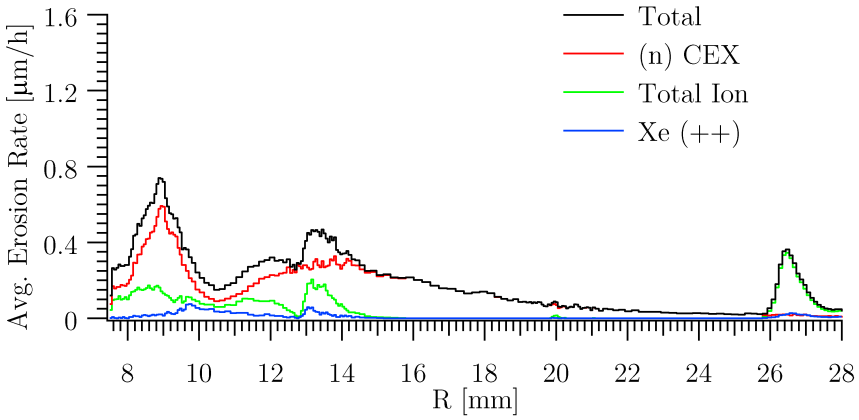

(a) Average erosion rate from $t^* = 5.70 - 8.40$ $[\mu s]$.

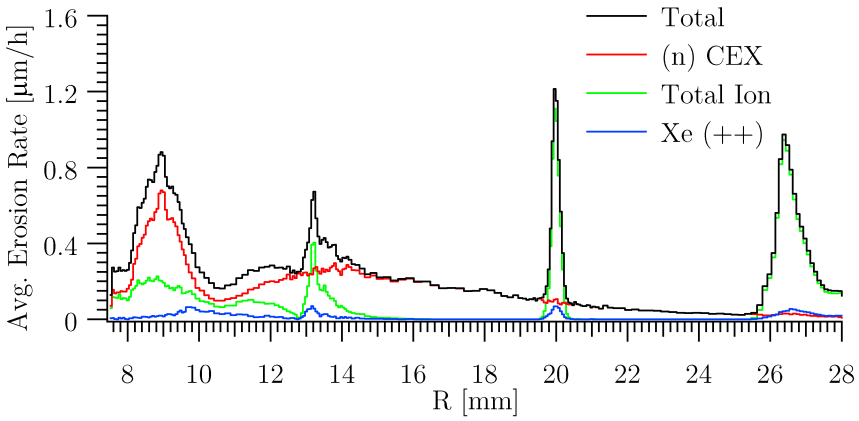

(b) Average erosion rate from $t^* = 5.70 - 9.10$ $[\mu s]$.

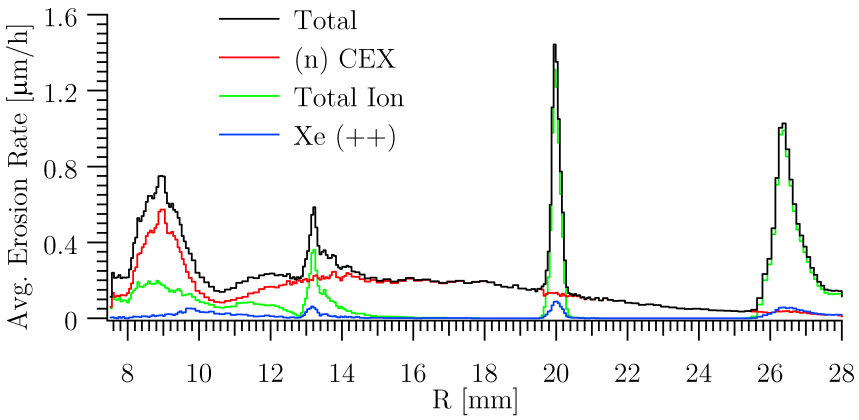

(c) Average erosion rate from $t^* = 5.70 - 9.90$ $[\mu s]$.

Figure 5-20: A continuation of Fig. 5-19, showing the temporal evolution of average erosion rates along the insulator wall.

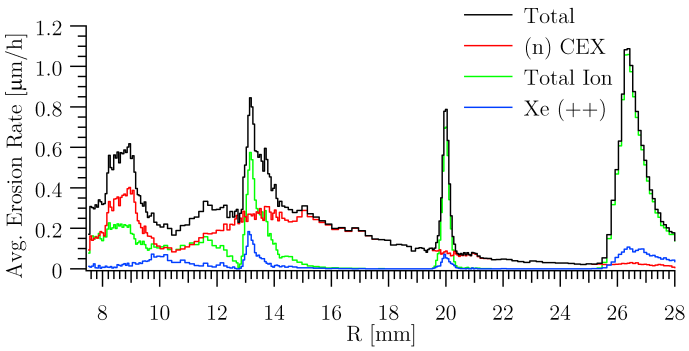

(a) Average erosion rate from $t^* = 5.70 - 11.70$ $[\mu s]$, for the simulation shown in Fig. 5-15.

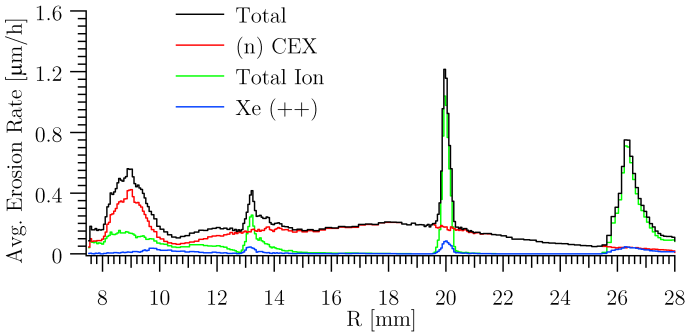

(b) Average erosion rate from $t^* = 5.70 - 11.70$ $[\mu s]$, for the simulation shown in Fig. 5-18.

Figure 5-21: Average erosion rates from beginning to end of both oscillatory simulations discussed in this section. Note the different scales in each figure.

5.3.2 Predicted Erosion for Steady Operation

Erosion data representing the LC-mode of DCFT operation were generated by continuing the steady simulation described in Sec. 5.1 for 4.8 [μs]. The anode current from this simulation is shown in Fig. 5-22 and is representative of the behavior of all the other diagnostic data recorded. The erosion profiles in this steady mode of operation are not similar to the erosion profiles representing the HC-mode of operation, presented in Fig. 5-23. Aside from small regions upstream of C1 and near the exit plane past where $R = 27.2$ [mm], CEX neutral contributions exceed total ion erosion over almost the entire insulator surface. This certainly supports claims made in Sec. 5.1 regarding the strong influence of CEX collisions on the simulation that are not expected to occur in experiments. Ion erosion is again confined to each cusp and upstream of C1.

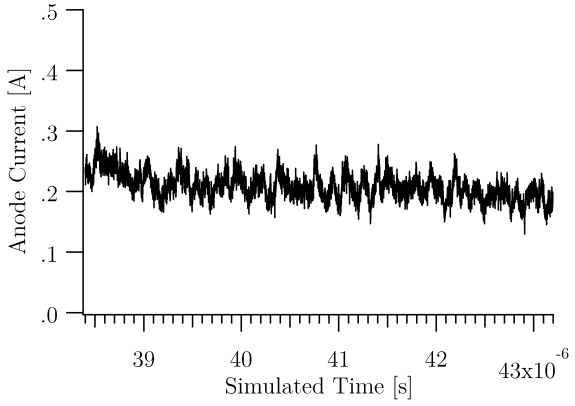

Figure 5-22: Detailed view of I_a the simulation described in Sec. 5.1, for $t^* = 38.40 - 43.20$ [μs].

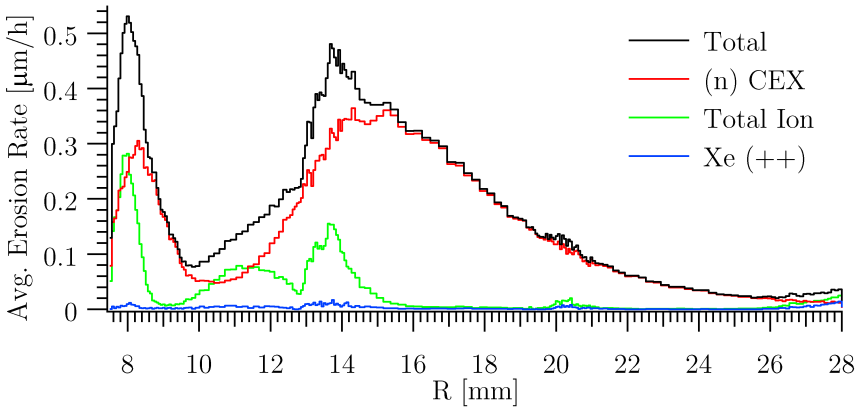

Figure 5-23: Avgerage erosion rate predicted by PTpic from $t^* = 38.40 - 43.20$ [μs] in the steady simulation presented in Sec. 5.1.

5.3.3 Summary of Erosion Predictions

Simulations were run with the intent to compare erosion characteristics in the high and low-current modes of DCFT operation. The use of two different values of C_f allowed these simulations to be run at the same thruster operating condition. Although the oscillatory simulations are somewhat representative of the HC-mode, a limit cycle representation of HC-mode has not yet been obtained with PTpic. The most appropriate comparisons with existing measurements can be made with simulated erosion data representing the HC-mode, but even then the operating conditions are not identical. Nonetheless, the three simulated erosion rate profiles are plotted together with the experimental measurements on the same absolute scale in Fig. 5-24. This provides a clear picture of which aspects of the simulated and experimental data most resemble one another. The simulated data are the same erosion rate predictions previously given in Figs. 5-17(c), 5-20(c), and 5-23.

The data sets representing the HC-mode of operation successfully reproduce the presence of concentrated erosion profiles in each cusp, and identify C1 as the cusp with the least amount of erosion. The second oscillatory simulation even predicts that erosion is greatest in C2, in agreement with one of the least expected aspects of the erosion measurements. However, the predicted levels upstream of C1 are unexpected. Even without measurements in this region, visual indicators that erosion might be significant in this region have not been noticed. It is possible, though, that redeposition occurs at a quicker rate upstream of C1 because of the proximity to the stainless steel diffuser in experiments. Based on materials analysis, the diffuser disc is suspected as one of the main sources of reposition between cusps [128]. As seen in Fig. 2-23, the wall upstream of C1 was lighter than other regions after 204 $[h]$ of operation. The simulated data also do not reproduce the observed increase in erosion downstream of C3, heading toward the exit plane.

Overall, considering that the simulations predict the largest potential drop to occur between the anode and C1, well upstream of the expected location from empirical observations, the degree to which the simulated data match the experimental data is surprising. Some of this agreement may be coincidental, especially between cusps. Originally, the author attributed trends in erosion measured outside of cusps to a slight misalignment during measurements [128]. Looking at the contributions to erosion between cusps, it is clear that the agreement between

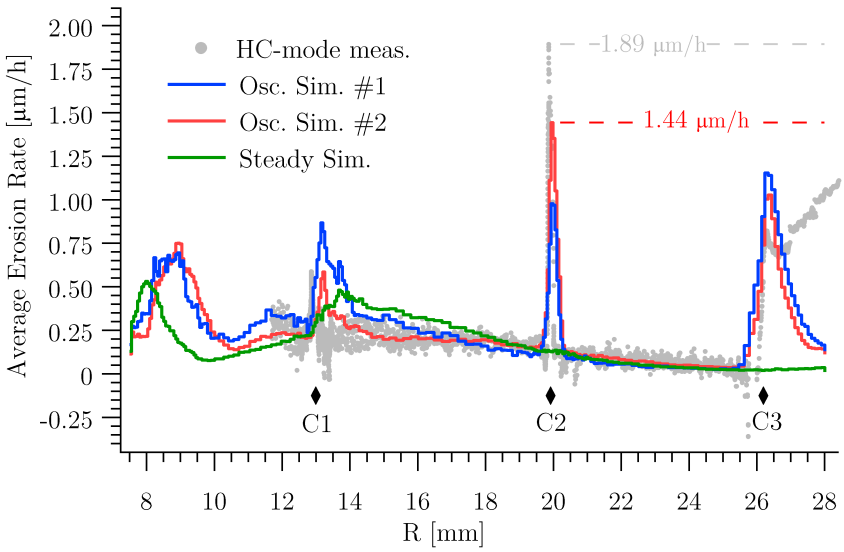

Figure 5-24: Comparison of average erosion rate measurements and simulations. The measured and simulated data do *not* correspond to the same operating condition.

predicted and measured values is due entirely to CEX neutrals. However, the ion CEX collision frequency increases monotonically as the ion speed relative to background neutrals increases — for a fixed neutral density. This means that the CEX erosion should contribute less to erosion in the chamber if the main potential drop is positioned at or past the C3 separatrix. The differences in the simulation would be amplified further by the non-linear energy dependence of the sputter yield on incident particle energy, especially considering the enforcement of a threshold energy. Without ion energies in excess of the sputtering threshold energy upstream of acceleration region — 13 $[eV]$ in these simulations — CEX neutral erosion would be significantly reduced in many places.

The preponderance of erosion caused by CEX neutrals in the LC-mode simulation is unlikely to exist in reality, or in a simulation with a properly located potential drop, for the same reasons discussed in the previous paragraph. The most specific prediction that might reasonably be expected to carry over to experiments might only be that erosion in the HC and LC-modes should be very different, and that erosion in LC-mode will be significantly lower than in HC-mode. The cause of these differences between modes can only be described within the context of these simulations, and the degree to which they might carry over to reality is difficult to quantify.

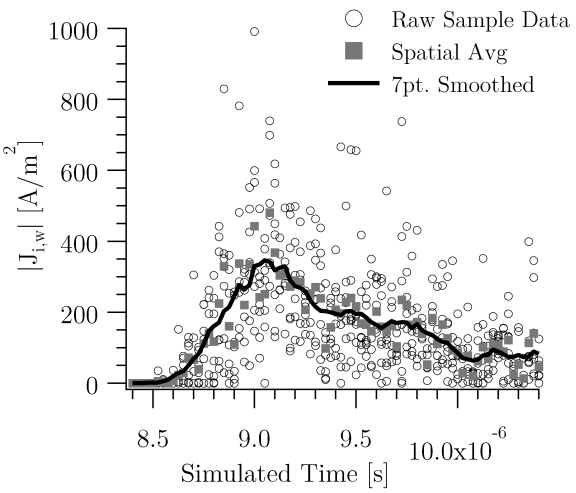

Figure 5-25: Time history of the spatially averaged magnitude of total ion current density in C2 for $R \in [19.76, 20.01]$ $[mm]$. The curve shows the variation of a moving, seven-point temporal average of the spatially averaged current density.

Considering lifetime limiting erosion rates, the most significant differences between erosion profiles in each mode are the rates present in C2 and C3. This is most clearly depicted in Fig. 5-24. In both oscillatory simulations discussed in Sec. 5.3.1, erosion profiles resemble the simulated LC-mode profile up until the times that erosion intensifies in C3 and C2. The most obvious example of this encountered so far occurs between $t^* = 8.40 - 9.10$ $[\mu s]$, evident through a comparison of Figs. 5-20(a) and 5-20(b), where the averaged erosion rate in C2 5.4 $[\mu m/h]$.

To examine what brings this about, the time-history of total ion current density J_i at the wall over a portion of C2 was examined in detail. The resulting simulated data are provided in Fig. 5-25. For this figure, J_i was sampled along 7 consecutive mesh nodes on the wall in C2 at each of about 80 instances in which instantaneous moment data were written to file between $t^* = 8.40 - 10.4$ $[\mu s]$. The raw sample data from each point are plotted as hollow circles, and the average of the 7 values at each instant are plotted as gray squares. Finally, a smoother representation of the data are provided by a moving time-average of these spatial averages. From this figure, it is clear that the strong peak in erosion is caused by the sudden influx of very high ion current densities in the second cusp. This can also be seen clearly in Fig. 5-26 showing a shift in the current density distribution from being peaked on the central axis at $t^* = 8.40$ $[\mu s]$, to a more complex distribution at later times. In Fig. 5-26(b), J_i is still peaked on axis near the anode, but downstream of C1 the region with elevated J_i values lies on a

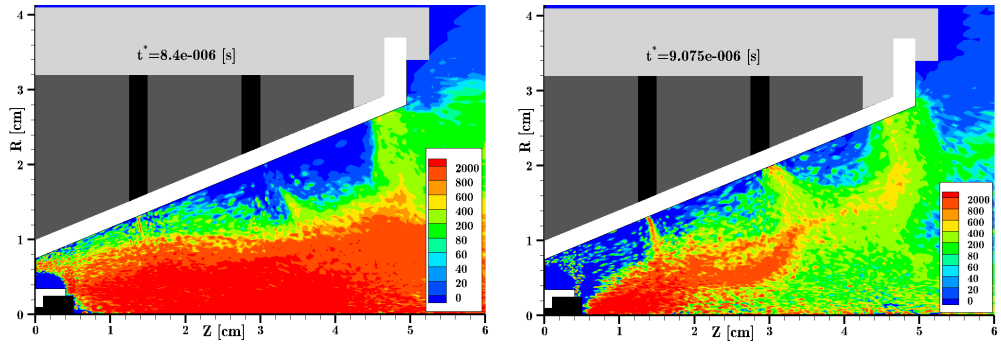

(a) $t^* = 8.40$ $[\mu s]$ (b) $t^* = 9.075$ $[\mu s]$

Figure 5-26: Snapshots showing the redistribution of the total ion current density magnitude during the onset of elevated erosion levels in C2.

contour resembling a magnetic field line connecting C1 to C2. Also visible is a thin, hollow jet extending from C1 out past the exit plane of the thruster. Looking back at Fig. 5-4(a) for a comparison with the LC-mode representation, ion current densities are insignificant in C2 and C3, just as they are insignificant in C2 as shown in Fig. 5-26(a). Of course, the energies of particles at the walls also determines the erosion rate, but it is this transition in the ion flux, shown in Figs. 5-25 and 5-26, that leads to the intense erosion in the oscillatory simulations. This phenomenon is not present in the LC-mode simulation because the J_i distribution is steady and low near the walls in the cusps downstream of C1.

5.4 Varied Simulation Parameters

This section presents a summary of successful simulations that were run to examine the effects of using different parameters on virtual thruster operation and to demonstrate that PTpic could converge to quasi-steady solutions. All the results presented here use an anode voltage of 250 $[V]$, a permittivity factor $\gamma = 50$, a super-particle size $W = 5 \times 10^8$ and specify a wall/anode/neutral injection temperature near 0.1 $[eV] \approx 1160$ $[K]$. Also, all free-space nodes are grounded, and no anomalous collisions for electrons are included. The main variations between simulations were mass flow rates, and the mass factor values. The lower flow rate conditions were intended to provide PTpic with a steady-state operating condition, whereas the higher flow condition was simulated in hopes of reproducing an oscillatory mode. Some

simulations were also run with the floating body capacitance C_f varied, or with the floating body potential fixed — this is equivalent to specifying an infinite capacitance. A summary of the parameters used for each simulation and their associated designations are provided in Table 5.5.

Table 5.5: Summary of simulation parameters for selected simulations.

Simulation Designation	r	$\dot{m}$ $[sccm\ Xe]$	C_f $[pF]$
S_1	5000	6.0	39.1×10^3
S_2	5000	6.0	39.1
S_3	5000	4.0	39.1
S_4	1000	4.0	∞, ϕ_f=75 [V]

Many more simulations were attempted than are mentioned here. Simulations utilizing anomalous electron collisions were attempted, but the use of large r values appears to preclude the examination of anomalous diffusion effects because the electron Hall parameter is reduced by $\sqrt{r}$. Furthermore, simulations with a different mesh and $W = 5 \times 10^7$ were run with the intent of better resolving the region upstream of C1. Using this different mesh, the anode structure within the domain was removed entirely and reduced to a portion of the left-hand domain. Some improvement was obtained, as the plasma potential remained near or above the anode potential immediately downstream of the anode with this configuration, but the large potential drop still occurred upstream of C1. Simulations using lower mass factors, combined with the quench model of anomalous diffusion have also been attempted, but have not yet yielded any significant results, and so are not included in this thesis.

During S_1, W was altered from 5×10^8 to 5×10^7 at $t^* = 30.13$ $[\mu s]$, and began using a fixed timestep after about $t^* = 4.1$ $[\mu s]$. In Fig. 5-27(a), the displayed particle counts prior to $t^* = 30.13$ $[\mu s]$ have been increased by a factor of 10 to allow the trend in simulated particles to be observed continuously, unlike in Fig. 5-1(a). In both cases, the number of computational particles representing an equivalent number of physical particles in increased when W is reduced. Also, the floating body capacitance value was increased from 39.1 $[pF]$ to 39.1 $[nF]$ at

$t^* = 10.13$ $[\mu s]$. The only difference between the long-term parameters of S_1 and the steady simulation from Sec. 5.1 is the mass factor r: here $r = 5000$, while the Sec. 5.1 simulation uses $r = 1000$.

S_2 began using a fixed timestep near $t^* \approx 4.1\mu s$. This simulation appears to have converged well, though the larger mass factor value appears to have resulted in different behavior than the "oscillatory" simulation run with identical parameters. Here, the floating body potential oscillates at a high frequency and an amplitude between $100 - 200$ $[V]$. Nonetheless, this was a very significant result for PTpic, demonstrating that convergence and sustainment of a plasma discharge can be obtained. S_1 was restarted from one of this simulations save states, using a larger capacitance value.S_2 uses a smaller value of C_f than S_1, and only differs in the use of $r = 5000$ relative to the oscillatory simulation discussed in Sec. 5.2.

In S_3, a fixed timestep was used after $t^* \approx 5.74\mu s$. Here, the number of neutrals in the simulation appears to be steadying, and other parameters appear to have converged. The large amplitude, high-frequency oscillations appear to be driven by the oscillating floating body potential. As a result, results from simulations such as this motivated the use of larger capacitance values.

In S_4, the floating body capacitance is effectively infinite because it was fixed at a potential of 75 $[V]$ to examine the effect of this condition on the simulation. This situation could be implemented in the laboratory by connecting a power supply to the floating body, but only if the power supply was capable of operating in the voltage limited regime in the presence of a plasma. So, in this simulation, the floating body is really another electrode, just like the anode. Rather than the floating body potential, the net charge collected by the "floating body" is shown in Fig. 5-30(f). The slope of this curve indicates the current collected by the body from the plasma. The average current magnitude collected after $t^* = 15$ $[\mu s]$ is .098 $[A]$ — larger than the anode current over the same interval. W was increased by a factor of 10 at

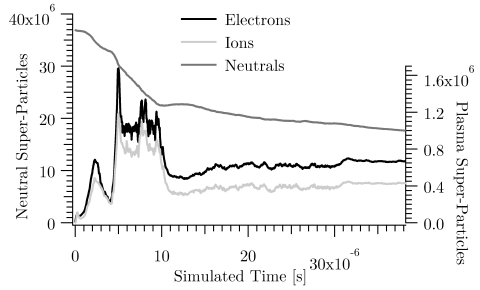

(a) Super-particle counts.

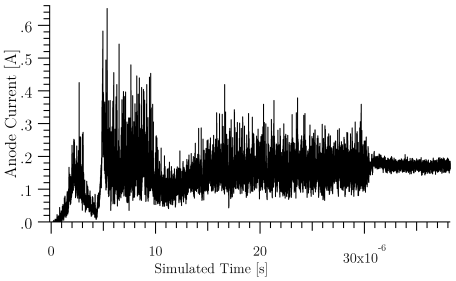

(b) Anode current.

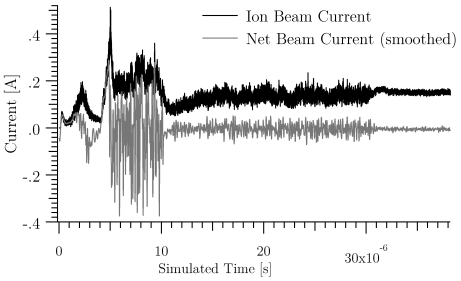

(c) Ion beam current and the net beam current. The beam current is further processed using a 50-point boxcar average.

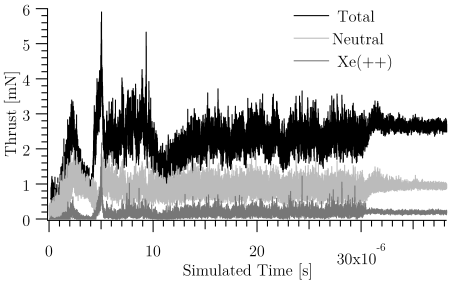

(d) Thrust and its components.

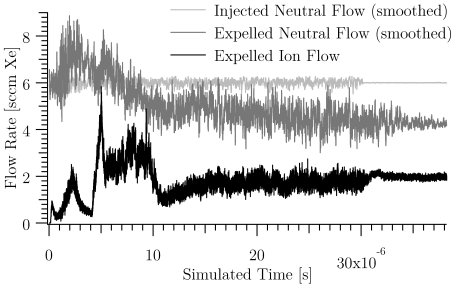

(e) Propellent mass-flow rates.

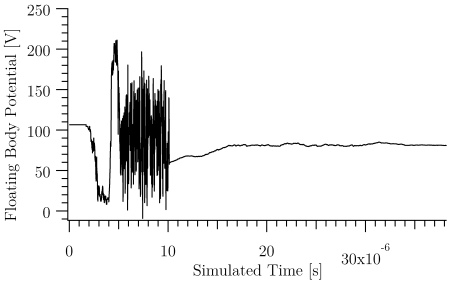

(f) Floating body potential.

Figure 5-27: Diagnostic data from S_1, where each quantity is time-averaged within groups of 1000 iterations. The simulated data shown here spans an interval $\Delta t^* = 1.27\tau_r$ after completing 25.13 million iterations. A constant timestep was used after $t^* = 2.7$ $[\mu s]$.

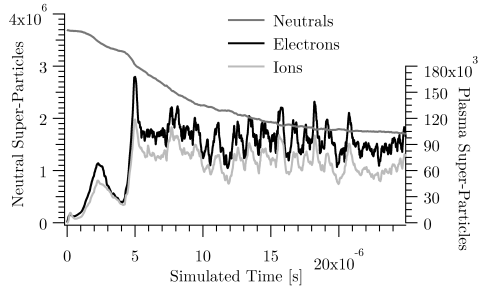

(a) Super-particle counts.

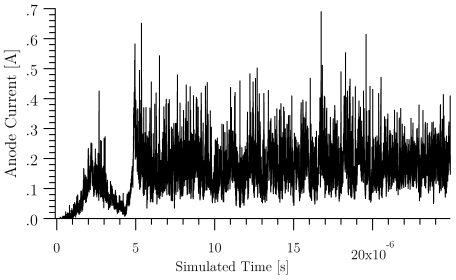

(b) Anode current.

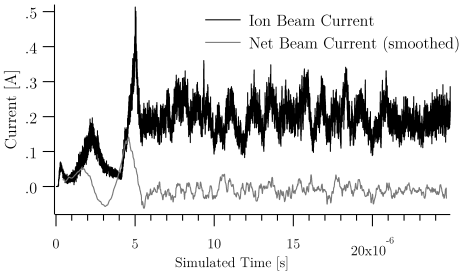

(c) Ion(black) and net beam(gray) currents. The beam current is subjected to an additional boxcar average encompassing 1000 of the pre-averaged data points.

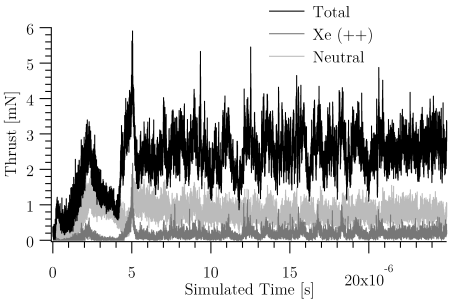

(d) Thrust components.

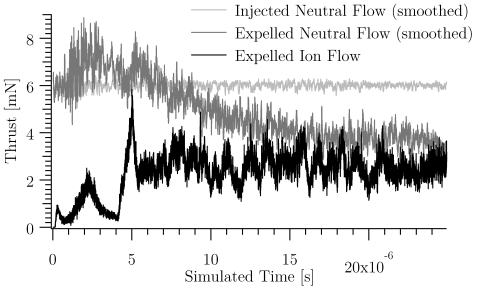

(e) Propellent mass-flow rates. The neutral flows were additionally smoothed with a boxcar average using 50 pre-averaged data points.

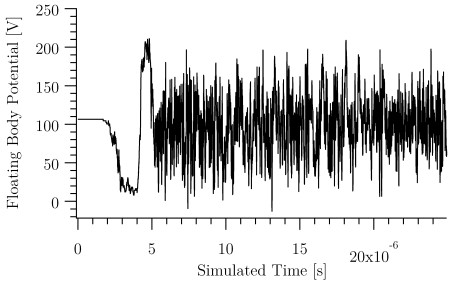

(f) Floating body potential.

Figure 5-28: Diagnostic data from S_2, where each quantity is time-averaged within groups of 1000 iterations. The simulated data shown here spans an interval $\Delta t^* = 3.05\tau_r$ after completing 24.8 million iterations.

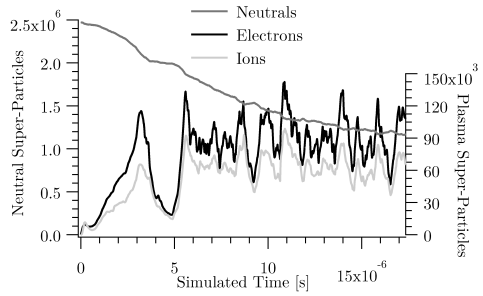

(a) Super-particle counts.

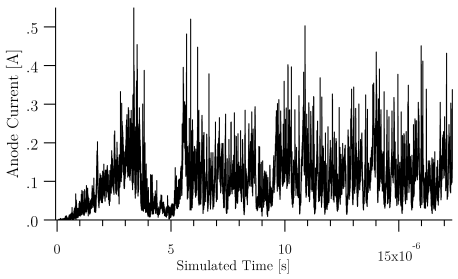

(b) Anode current.

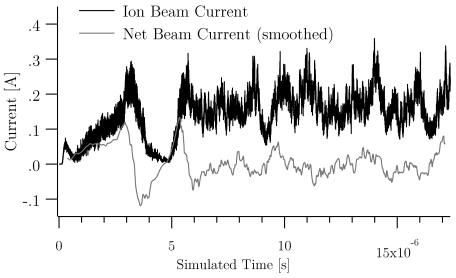

(c) Ion(black) and net beam(gray) currents. The beam current is subjected to an additional boxcar average encompassing 1000 of the pre-averaged data points.

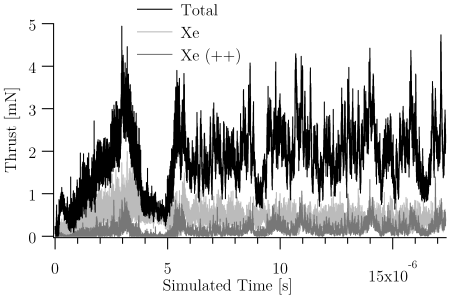

(d) Thrust components.

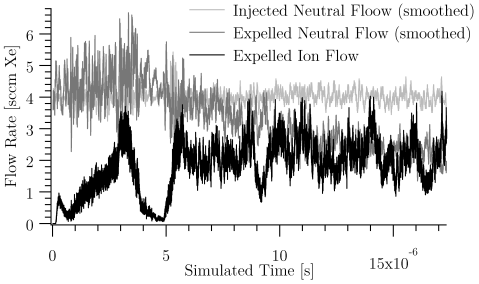

(e) Propellent mass-flow rates. The neutral flows were additionally smoothed with a boxcar average using 50 pre-averaged data points.

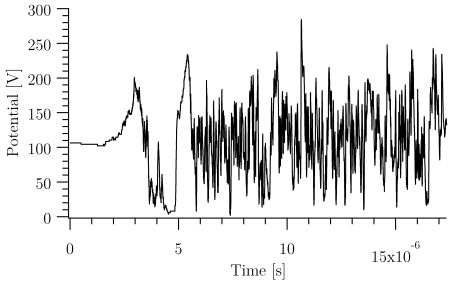

(f) Floating body potential.

Figure 5-29: Diagnostic data from S_3, where each quantity is time-averaged within groups of 1000 iterations. The simulated data shown here spans an interval $\Delta t^* = 2.13\tau_r$ after completing 19.25 million iterations.

$t^* = 15$ $[\mu s]$, but the values have been compensated to show trends on a single scale. A constant timestep was used for the entirety of this simulation.

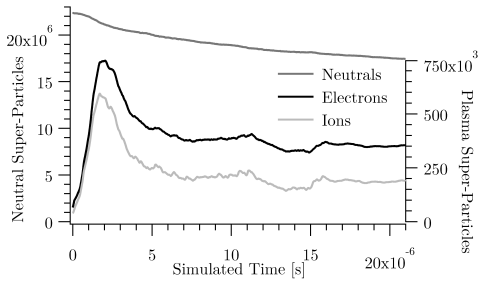

(a) Super-particle counts.

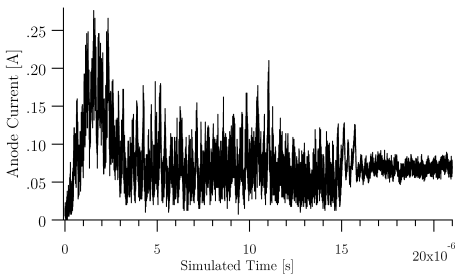

(b) Anode current.

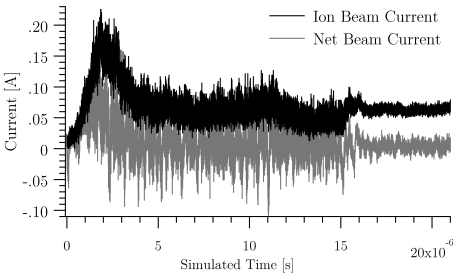

(c) Ion(black) and net beam(gray) currents. The beam current is subjected to an additional boxcar average encompassing 1000 of the pre-averaged data points.

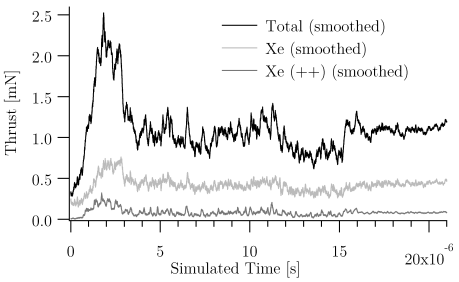

(d) Thrust components, smoothed using an additional 50-point boxcar average.

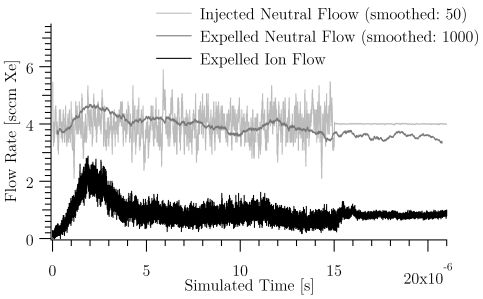

(e) Propellent mass-flow rates. "Smoothed" flows were additionally processed with a boxcar average using the number of specified points.

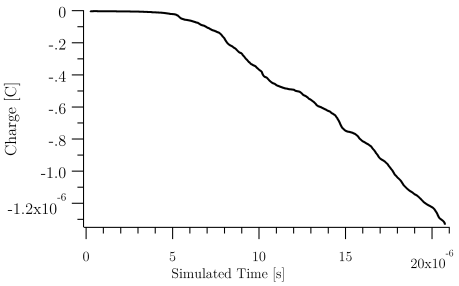

(f) Charge collected by "floating body" electrode, held at 75 $[V]$.

Figure 5-30: Diagnostic data from S_4, where each quantity is time-averaged within groups of 1000 iterations. The simulated data shown here spans an interval $\Delta t^* = 1.15\tau_r$ after completing 21 million iterations.

5.5 Near-Wall Potential Structures: Valleys & Sheaths

A simulated potential distribution within the DCFT chamber is provided in Fig. 5-31(a), showing depressions in potential within the cusps. These features are expected because of stronger electron magnetization, and match trends obtained from flush wall mounted probes [44]. Additionally, traces of potential extending into the bulk are given in Figs.5-31(c)–5-31(e), extracted along the thick black lines shown in Fig. 5-31(a). These plots show that sheaths have formed self-consistently within each cusp.

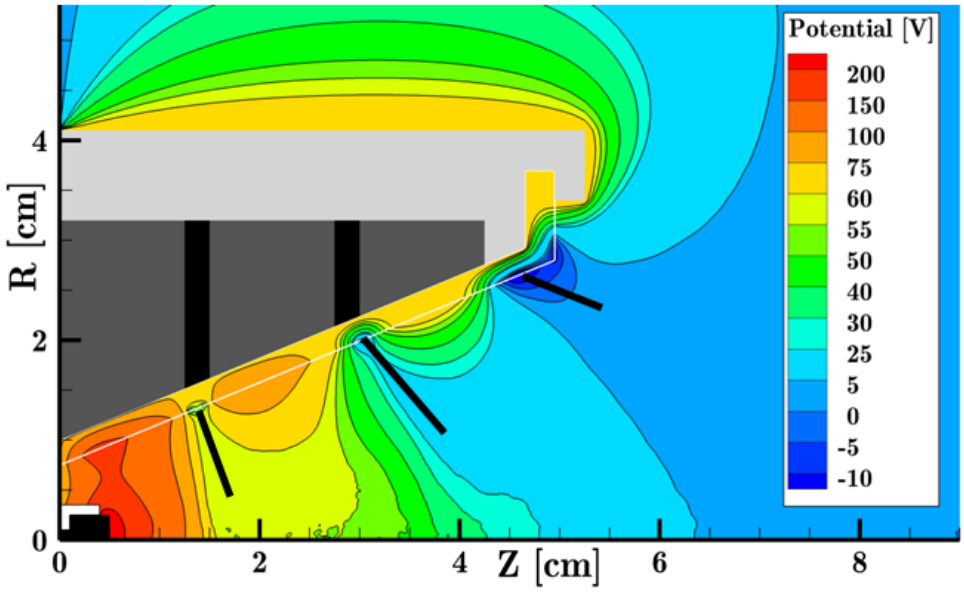

(a) Distribution of electrostatic potential. The insulator wall is shown as a thin white line in this figure.

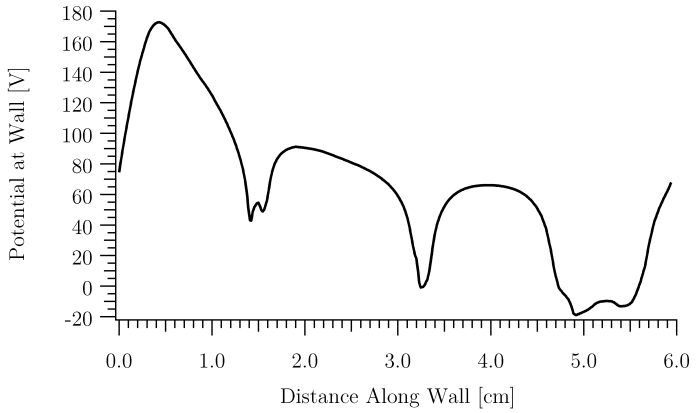

(b) Potential along insulator wall from $(z, r) = (0, .75)$ to $(3.4, 4.95)[cm]$, extracted from data shown in Fig. 5-31(a).

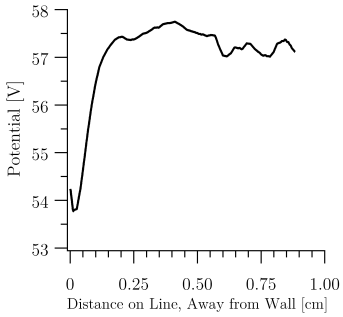

(c) Cusp 1

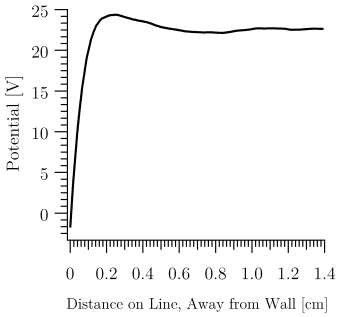

(d) Cusp 2

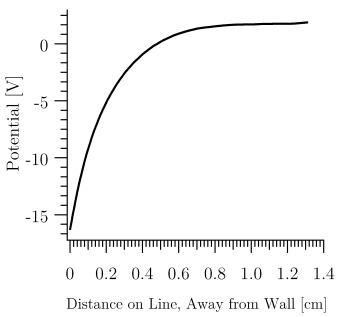

(e) Cusp 3

Figure 5-31: Selected potential distribution showing typical distributions in cusps and along the insulator wall.

Chapter 6

Adaptability to New Designs

In addition to the large modifications for PTpic, described in Chapters 3 and 4, a great deal of effort was put into making PTpic more flexible in terms of how inputs are handled. Many hundreds of error checks, messages and codes were added to PTpic, so that when unanticipated inputs are provided, the user can be notified in such a way that he or she will hopefully be able to identify the offending section of code without having to sift through the several tens of thousands of lines of `C`-code that currently comprise PTpic.

I expect that, in many cases, only two lines of code in PTpic should require modification when simulating a new thruster[a]: the lines in `pic.h` telling PTpic the dimensions of the structured mesh. The macros that store this information are `N_Z` and `N_R`, though more appropriate names would be `N_K` and `N_J`, respectively. After changing `N_Z` and `N_R`, or making any other source code edits, the PTpic executable must be recompiled in order for the program to run according to the provided updates. When the mesh size is altered, modification to `N_Z` and `N_R` are required because many of the necessary variables stored on the mesh are statically allocated, and their size in memory is determined by the product of `N_Z` and `N_R`. After compiling for a certain mesh size, it is a good idea to include the values of `N_Z` and `N_R` in the executable name — `PTpic_100x100` for example, corresponding to the case where `N_Z` and `N_R` are both specified as 100.

Recently, some variables added to PTpic have been dynamically allocated, especially those related to loading the magnetic field from file and storing time-averaged data. This provides

[a]The input parameters will require extensive modifications, however.

the flexibility to free the memory occupied by these variables when they are no longer needed, or to avoid allocating data at all, for instance when time-averaging is not used. A large amount of code editing would be required, but dynamically allocating all variables would allow PTpic to run simulations using different mesh sizes without needing to recompile on a given machine. Dynamically allocating variables will also prevent the code from crashing without useful error messages in the event that the amount of static memory required by PTpic exceeds a given systems capabilities.

After the mesh, magnetic-field, boundary data, and other inputs are provided — the file names for the respective data types are provided by the user in the file `MasterPTpic.in` — simulations can begin. The requirements of the input files can be discerned from reading the source code, but the easiest way is to adapt input files from the DCFT, or other preexisting files, to a new thruster. In order to facilitate this process, I have heavily commented the DCFT input files.

Despite all the work that has gone into making PTpic more adaptable, there are certainly limitations to what can be changed without having to modify the source code of PTpic. Therefore, the author has started recording release notes describing *any* source code modifications that are implement. Each time edits are made, a detailed description should be (and has been) provided in the release notes. At this point, a new version designation is assigned to the new code, and its entire contents are copied to a separate folder. This allows tracking revisions, and reverting to previous versions if desired. Currently, all input files, source code, release notes, and the `Makefile` (this is what organizes how the PTpic source code is compiled) are kept in a single folder. From here, individual executables are compiled and distributed to different simulation directories, where input and output data files from each simulation are read and generated, respectively.

Examples of changes to PTpic that would require modifications to the source code are now listed to provide a sense for the limitations of PTpic as a result of this thesis work:

(a) Injecting neutrals at any boundary other than the one defined by $k = 0$. This is typically the left-hand simulation boundary.

(b) Injecting electrons at any boundaries other than the ones defined by $j = N_R - 1$ or $k = N_Z - 1$. These are typically the “top” and ”right-hand” simulation boundaries.

(c) Using more than one floating body — a body modeled as a perfect conductor whose change in potential is determined by the flux of charged particles to its surface. The floating body may also be converted to an electrode at a fixed potential. In that case, only one electrode could be used, or a single group of electrodes at a common potential.

(d) Using any propellent other than xenon. Changing to a different propellent would require re-defining constants, as well as providing different cross section data for the model. Currently, the cross section data are specified in header files.

(e) Changing which collisions are included in the model. In some cases, there are macros that control which collisions are included, but PTpic has not been tested with any configuration of these macros other than the present.

(f) Changing the secondary electron yield model, or the energy distribution of secondary electrons emitted from insulator surfaces.

(g) Changing how velocity distributions are generated. This is a lacking part of PTpic because it uses a 2D method to generate 3D velocity distributions. This causes the average energy of the resulting distribution to differ from what would be expected from the specified temperature. It also causes the injected distribution to differ from 3D maxwellian.

(h) The re-addition of functions used to calculate and store data related to velocity or energy distribution functions.

(i) Using two different types of dielectrics in the model. Multiple dielectric bodies or surfaces are supported, but they are currently required to have the same relative permittivity.

(j) Likely, many other changes the author has not anticipated up to this point.

Despite the litany of aspects that must remain fixed, by adding many checks for unanticipated options — typically these errors are triggered if PTpic enters an `else` statement — the author has hopefully been able to make modifications to the source code as "painless" as possible for future users.

To test the adaptability of PTpic, a new design was simulated. This design was provided by T. Matlock in the form of a Maxwell SV magnetostatic model, which provided the geometry and magnetic field of the thruster. From this information, high-resolution magnetic field data and a mesh of the domain were developed, as well as information about internal and external boundaries. These data are summarized in Fig. 6-1(a), showing the thruster geometry downstream of the $z = 0[cm]$ plane. The anode extends from $r = 0[cm]$ to $r = 0.6[cm]$, the black region is filled by magnets and iron pole pieces, and the white region represents a BN insulator. This thruster has been given the designation of 3CFT-10, because it is a cusped-field thruster with three ring-cusps and a $10[mm]$ inner diameter.

The design was simulated for several million iterations on a uniform, square, orthogonal mesh with nodes spaced by $0.1[mm]$ in both directions. A mass factor of 1000 was used, and the smaller domain size and greater mesh resolution allowed PTpic to use a permittivity factor of 10 in this case. The simulations were not continued past several million iterations to free space on the SPL computer cluster. A convergent simulation was not attempted because there would be no experimental data to compare with, and there are no current plans to build this thruster. Nonetheless, a very small sample of results is provided in Fig. 6-1(b), showing the potential distribution in this imaginary thruster. Here, unlike in the DCFT, elevated potentials extend to the boundary of the domain. This simulation was not allowed to approach anything resembling a fully converged solution. However, I was able to proceed from the point of a design provided in Maxwell SV, to simulating the thruster, in less than a day. The most time-consuming task was calculating the magnetic field data and converting it to the needed format. For a more complex geometry, generating the mesh would have required more time, but using the macro commands available in Tecplot makes this step far less complex than it used to be.

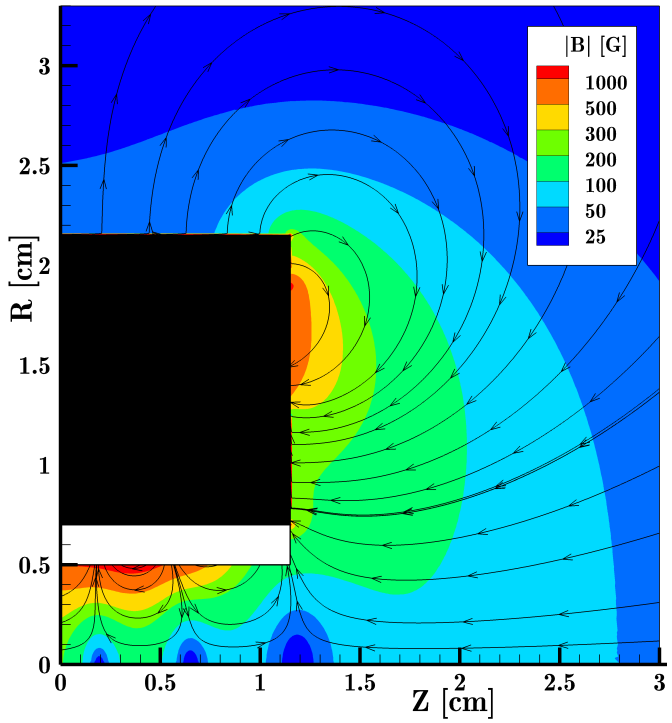

(a) 3CFT-10: Magnetic field strength with selected field lines.

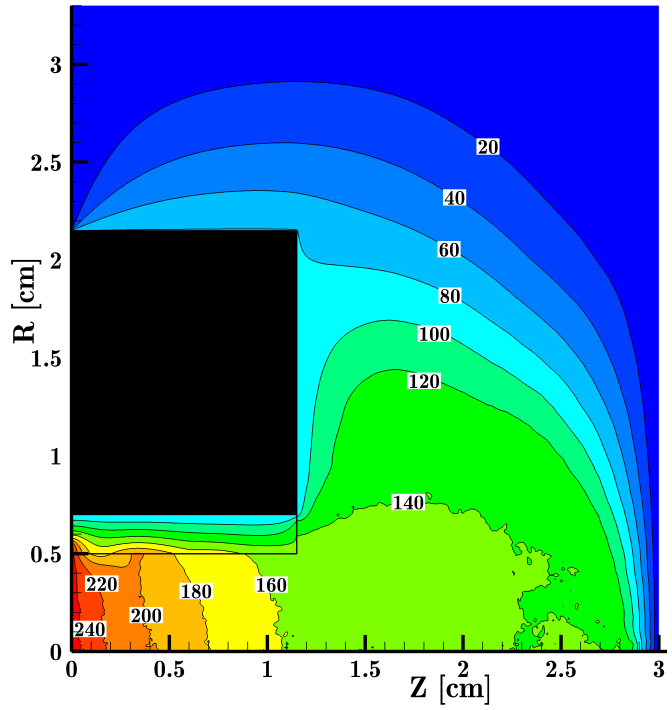

(b) 3CFT-10: Electrostatic potential, given in volts.

Figure 6-1: Geometry, magnetic field, and a sample of simulated potential data from a demonstration simulation using a design concept designated here as 3CFT-10.

Chapter 7

Conclusions

The plasma thruster particle-in-cell (PTpic) simulator has been developed as an extension of the program written by J. Szabo [97], with contributions from J. Fox [111,112] and others [110,116]. The impetus for this began as part of an effort to understand and evaluate the divergent cusped-field thruster as it was initially designed, constructed, and tested by D. Courtney [27]. Opportunities to improve SPL-PIC were identified during the author's first attempts to simulate the DCFT [114]. This ultimately led to the modifications discussed in Chapters 3 and 4, which were subsequently used to generate the results presented in Chapter 5. Concurrently, measurements presented in Chapter 2 of this document — and elsewhere by T. Matlock [32,35,44], N. MacDonald [37], and others [36,159] — have provided the data necessary to evaluate the accuracy of PTpic. These results have also guided the selection of simulated thruster operating conditions, and helped identify which outputs from the model were most appropriate to present and analyze. In the remainder of this chapter, the main conclusions from this work are summarized and discussed. Recommended applications and improvements for PTpic, as well as further studies of cusped-field thrusters, are also given.

7.1 Simulations of DCFT Operation in the Low & High-Current Modes

Attaining a converged simulation of a robust DCFT discharge was a necessary and important milestone for PTpic in the course of this work. In previous efforts, this was only possible using

an overly large permittivity factor γ and reduced magnetic field strengths[114]. The first set of the converged simulations for this work, summarized in Sec. 5.4, employed a mass factor $r = 5000$. Progressing from this vital starting point, the mass factor was reduced to $r = 1000$ for simulations discussed at length in Secs. 5.1–5.3. The lower limit on r imposed here represents a compromise between tractability and realism.

The steady simulation discussed in Sec. 5.1 is representative of the LC-mode of operation. This simulation utilizes an enlarged floating body capacitance C_f in comparison to the value calculated numerically in the absence of plasma. Increasing C_f has the effect of lowering the magnitude of high-frequency floating body potential oscillations, as shown in Fig. 5-27(f), and is justifiable on physical grounds based on the expected influence of the plasma on C_f — see Sec. 4.1.2. Overall, thrust and specific impulse are underpredicted by 50% and anode efficiency is predicted to within 10%. The suspected cause underpinning these discrepancies is the misplacement of a large potential drop $\Delta\phi$ between the anode and first cusp. Emissive probe measurements[32,35,36] and ion velocimetry[37] indicate that the electric fields responsible for the primary acceleration of ions are located near the exit separatrix in cusped-field thrusters. As a consequence of $\Delta\phi$ being so far upstream, where the cross sectional area is smallest and neutral density greatest, the effects of charge-exchange collisions are exaggerated in these simulations. This should naturally lead to an emphasis on correcting this result in future simulations of cusped-field thrusters.

Despite the suspected misplacement of the largest potential drop, PTpic is able to reproduce the hollow conical plume structure of the DCFT. The divergence angle of the hollow jet is approximately 10° less than has been experimentally observed[29,35,132]. Based on data time-averaged for 5×10^5 iterations, the model predicts that the hollow beam has its origins in C2. Here, a portion of the low energy ion beam upstream of C2 intersects a region of elevated axial and radial electric fields, sending the ions away from the centerline and out of the thruster. So, if the simulated main potential drop could be moved to the more likely position (C3), the divergence angle prediction would probably improve. Interestingly, variations in the electric field component perpendicular to the magnetic field ($E_\perp$) are similar in each ring cusp, as shown in Fig. 7-1. The magnitude of $E_\perp$ is peaked at the walls, and again closer to the null point along each separatrix. It is unclear to the author whether or not this pattern would remain

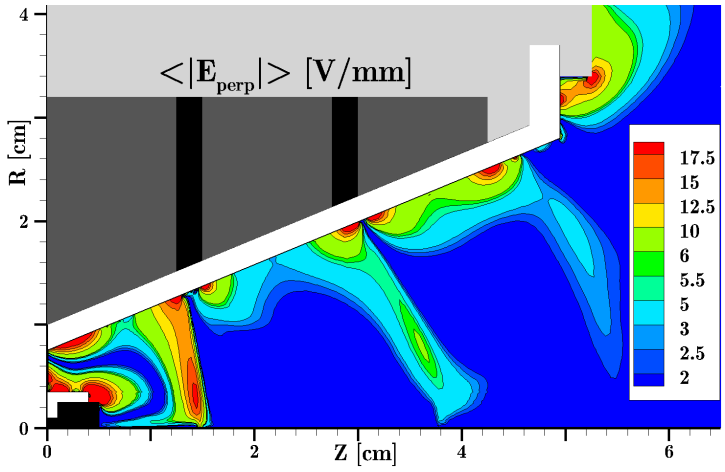

Figure 7-1: Component of the time-averaged electric field perpendicular to the local magnetic field, for the steady simulation discussed in Sec. 5.1. Note that the scale is piecewise linear in order to show similarities in the structure of the perpendicular field in different locations.

if the strongest electric fields were located in C3, but the fact that this pattern is present in each cusp suggests that it might.

To represent the high-current mode of operation, a single period of time was chosen from the oscillatory simulation discussed in Sec. 5.2. This was necessary because this simulation decayed to a more stable discharge later on, but not before completing several cycles resembling what would be expected of a limit cycle. Again, the location of the main potential drop contradicts experimental evidence, and predicted performance metrics are lower than observed. An excess of ion-neutral charge-exchange interactions near the first ring cusp is once more expected to contribute significantly to the lower simulated performance.

Continuing with the list of similar outcomes compared to the steady simulation, a conical beam of fast ions is observed in oscillatory simulations. The hollow distribution of axial ion speed is not present throughout the examined period, but only while the anode current is depressed or rising up toward the broad peak. This trend matches time-correlated ion velocimetry[37] at the only location downstream of the exit separatrix where data is available. Considerably less agreement is found between data recorded upstream of the C3 separatrix, however.

7.2 Erosion Simulations: Comparisons to Measurements & Across Modes

Oscillatory and steady simulations were repeated in order to evaluate PTpic's predictions of erosion rate distributions in the high and low-current modes. The anode voltage and propellent flow rates are identical in these simulations: 250 $[V]$ and 6 $[sccmXe]$, respectively. The only measurements available for comparison were obtained in HC-mode with an anode voltage of 300 $[V]$ and a flow rate of 8.5 $[sccmXe]$. The simulations representing the HC-mode successfully reproduce concentrated erosion profiles in each cusp. One of the two oscillatory simulations correctly locates the peak erosion rate in the second cusp. Between cusps, erosion rates are in agreement amongst measurements and all three simulated data sets. The most notable discrepancy can be found down stream of where the C3 separatrix intersects the wall, where the predicted rates fail to increase toward the exit plane. Agreement in this location could reasonably be expected to improve if the primary acceleration region were properly located there.

Speculation seems unavoidable when addressing how these predictions might change with a corrective shift of the largest potential drop into the third ring cusp. In this case, the magnitude of erosion caused by charge-exchange neutrals upstream of C3 should decrease because ions with energies above the sputtering threshold would not be as abundant in these regions. This statement applies equally to the steady and oscillatory results. Of course, if charge-exchange effects are significantly diminished, more ions may have access to sheath entrances in each ring cusp. In this scenario, at least some erosion would be redistributed from the inter-cusp regions into each cusp because ions are well confined electrostatically along the separatrices — unlike neutrals. The net result of this tradeoff will require using PTpic to predict erosion rates in a simulation with a properly located acceleration region.

At this point, the best indicators of DCFT erosion rates in low-current mode are those presented in Sec. 5.3.2. Comparing the peak *simulated* erosion rates shows that PTpic predicts a 270% longer lifetime in LC-mode than in HC-mode with similar operational parameters. Given that the model has not yet properly located the main potential drop between the anode cathode, such a quantitative prediction should not be interpreted as anything more than a targeted hypothesis. However, it does underscore the magnitude of the differences predicted to

exist by PTpic, in terms of plasma-wall interactions, between the low and high-current modes in the DCFT. The interpretation of these simulations is straightforward: in high-current mode, as ions evacuate the thruster chamber and the anode current approaches its minimum value in the cycle, the plasma shifts away from the axis of symmetry as the hollow divergent ion beam is formed. In the process, ion current densities in cusps are elevated well above the average values present in low-current mode. This causes large amounts of erosion to occur over a fraction of each breathing cycle, and consequently the peak high-current mode erosion rate exceeds the highest rates present in low-current mode. Evaluating the accuracy of this prediction will require improvements in future simulations, and ultimately, measurements of erosion in low-current mode.

7.3 The Current State of PTpic

Numerous changes were made in transitioning from the program inherited by the author in 2007 from J. Fox into what is now referred to as PTpic. Rather than listing all the changes that were made, many of the most important enabling and limiting aspects of PTpic will now be summarized. Some of these capabilities existed in SPL-PIC, but upgrades have been attempted for each functionality listed here.

(a) The adaptability of PTpic is discussed and demonstrated in Chapter 6, but it is sufficient to note here that the model can easily be applied to new geometries, magnetic fields, and other settings. This includes inputs for Hall thrusters. PTpic is in no way limited to modeling the DCFT, if other applications are desired.

(b) Save-states can be written to file at a user-specified frequency to allow assessment of which states to keep or discard at a later time. The simulation can be restarted from any one of these save-states later on.

(c) The author has also written a post processing program, dubbed "ppPTpic" that is used to time average information written into the large diagnostic data files PTpic creates. These files are updated at a user-specified frequency from a memory buffer, and contain important data like anode current, thrust, particle counts, etc. from *every* completed iteration. The time-averaging makes the data easier to interpret and reduces the sizes of files that are

most often used to generate plots. This post processing tool is also useful for combining diagnostic data from multiple runs into a single file.

(d) The calculation of time-averaged particle moments, and other data, can be activated or deactivated based on user inputs. The averaging interval is also provided by the user. The moments calculated, whether instantaneous or averaged, include density, bulk velocity, the full pressure tensor, the resulting temperature, and the heat-flux vector. The kinetic definitions of these moments are listed in Appendix C.

(e) In addition to calculating diagnostic, electric field, and all species' moments, two dimensional distributions of particle-particle and particle-boundary interactions (excluding erosion rates) are automatically calculated and written to file when time-averaging is active.

(f) The user can specify whether erosion data should be calculated for a given simulation based on switches in the input files. Erosion measurements can span multiple simulations, and accumulated erosion data can be reset whenever specified by the user.

(g) Calculating time-averaged moments slows down the execution of PTpic. The amount varies depending on the situation, but for the simulations run for this work, time averaging increased execution times by at least a factor of two. This scales poorly for the same reasons that calculating particle moments and implementing collisions currently scales poorly.

(h) Presently, the number of processes[a] assigned to run PTpic *cannot* be changed from one simulation to the next.

(i) For problems of the sizes discussed here, the use of additional cores has been observed to *decrease* execution times when using 32 or more. To address this, the manner in which particles are distributed in memory amongst each process must be changed.

7.4 Suggestions for Continuing Work

An examination of higher frequency components of DCFT anode current oscillations, specifically in the $10 - 100$ $[kHz]$ and $100 - 500$ $[kHz]$ ranges is motivated by T. Matlock's finding that

[a]The intended distinction between the terms "process", and "processor" is described at the beginning of Sec. 3.3.

oscillations in the range of several tens of kilohertz coincide with mode transitions[35], and by predictions of higher frequency oscillation components in this work — specifically in the range of $200-250$ $[kHz]$. To complement the testing that will soon begin on a new cusped-field design — intended to test ideas regarding how to reduce beam divergence — methods for improving the utilization efficiencies of these thrusters should also be tested. One example would involve placing rings of solid material in between cusps to block neutrals from passing along the walls in these regions. This may increase utilization efficiency in the DCFT, or other cusped-field thrusters, because the simulated distribution of ionization lies primarily on-axis, as shown in Fig. 5-6(b). The use of a variable cross-sectional area has recently been tested in Hall thrusters, and the authors claim improvements in utilization efficiency of about 10% at matched operating conditions[161].

7.4.1 PTpic Applications

Simulations of the DCFT should continue in order to understand what parameters in the model should be adjusted to place the potential drop near the C3 separatrix. Simulations of the new MIT cylindrical cusped-field thruster design should proceed as well. In addition to these thrusters, simulating a standard Hall thruster, or a cylindrical Hall thruster, would ease restrictions on required timesteps because of their weaker magnetic fields relative to the DCFT. With simpler magnetic topologies and more data regarding electron transport, adjusting parameters to produce desired potential profiles in a Hall thruster would be a good way to gain familiarity with how to go about tuning PTpic for a cusped-field thruster. Additionally, Hall thruster simulations may be able to examine the influence of discharge dynamics on erosion in these devices as well.

A more basic idea involves benchmarking PTpic against data from experiments without a magnetic field, such as a glow discharge or Franck-Hertz (FH) experiment. If PTpic could reproduce aspects of the FH experiment, the original form of which did not employ a magnetic field, this would more firmly establish confidence in its ability to model inelastic collisions accurately, without having to confront the issue of electron diffusion in magnetic fields. Simple modifications to PTpic that might be required for this include providing inputs for a different working gas (there may not be much data for a FH type experiment using xenon) and ensuring

that PTpic is logically equipped to represent a closed system. Other minor changes related to boundary conditions might also be necessary, depending on the specifics of the experimental setup. After comparing results to one of these simpler experiments, a systematic experimental and computational examination of discharges attributes, as magnetic fields of varied intensities and topologies are applied, would provide an interesting perspective on electron transport in magnetic fields.

7.4.2 Continued Development of PTpic

Reorganizing how particles are distributed amongst processes could provide significant improvements in the scalability and performance of PTpic. Currently particles are randomly assigned to a process. If each process were instead assigned particles based on their physical locations, then tasks such as calculating moments and implementing collisions can be done in parallel with less need for communications operations between distributed memory processes. This is discussed in more detail in Sec. 3.3.1. In the author's opinion, this would require major modifications to the PTpic source code, but has the potential to vastly improve the model's computational efficiency.

Currently, PTpic only takes advantage of distributed memory parallelization through the using of `openMPI` libraries. However, by bringing `openMP` into the fold, PTpic could benefit from the exploitation of shared memory parallelization. This is also discussed in Sec. 4.2.2. In order to gain the most from using `openMP`, it may be necessary to discontinue the use of linked-lists to manage memory for each species' list of particles in PTpic. Implementing this idea will require a lot of alterations to the PTpic source code, and should be considered a major undertaking.

From a structural standpoint, the performance of PTpic might gain considerably from detaching the `MASTER` process from computationally intensive work and altering its responsibilities to exclusively handle certain communication and redistribution operations, as well as writing data to file. However, it would be wise to measure the performance of the PTpic as it is running to make an accurate assessment of whether this would be a useful change to attempt.

Another possibility is to implement a fully self-consistent calculation of the floating body potential, along the lines of the discussion put forward in Chapter 4. However, approximations might still need to be made unless the entire floating body was bounded within the simulation

domain. The topic should be subject to additional feasibility studies before it is considered seriously.

Part II

Appendices

Appendix A

Basic Concepts in Propulsion

Thrust $\underline{F}$ and specific impulse I_{sp} are two important metrics of thruster performance. The thrust is equal to the momentum flux integrated over the exit area of the thruster, or in terms of magnitude, $F = \dot{m}c$, where c is the effective speed of the propellent leaving the thruster. I_{sp} is defined as the magnitude of the provided thrust impulse per unit of propellant consumed. It is therefore a measure of the fuel efficiency of a thruster. As a result, for constant thrust and mass flow rate $\dot{m}$, the I_{sp} can be simplified as indicated in Equation A.1, and is proportional to the effective speed of the propellant. The constant g is the acceleration of gravity at the Earth's surface, and is included in the definition of I_{sp} so that the number is invariant in different unit systems. For example, the I_{sp} of a rocket with an exhaust speed of 1,000 $^{m}/_{s}$ is approximately 100 s. As a reference, the I_{sp} of the space shuttle main engines is approximately 450 s in space, the highest value for any flight qualified chemical rocket engine[162]. Mature electric space propulsion (EP) technologies, such as Hall or ion thrusters, have specific impulses well in excess of 1,000 s[7].

$$I_{sp} \equiv \frac{\int_{t_o}^{t} F d\tau}{g \int_{t_o}^{t} \dot{m} d\tau} \Rightarrow \frac{c}{g} \tag{A.1}$$

For a spacecraft in an idealized environment devoid of external forces, the propulsive demand of a maneuver is described by the integrated magnitude of acceleration provided by the thruster Δv, commonly referred to as the "delta-V". In the real world, the thruster must also contend with external forces such as gravity, which will cancel the thrust force to varying degrees. Therefore, to characterize the total propulsive demand of a mission, the relationship between

Δv and thrust is defined in Equation A.2. The mass of the spacecraft m, will be reduced as fuel is consumed. Again, in the simplified case of constant thrust and fuel consumption rate, Equation A.2 reduces to Equation A.3 using $dm = -\dot{m}d\tau$. The total mass of the rocket before and after the maneuver is given by m_B and m_A, respectively. Equation A.3 is then easily solved for the mass fraction of the propellent that must be consumed in order to complete the mission. Equation A.4 is referred to as the Tsiolkovsky or ideal "rocket equation," where m_p is the mass of propellant consumed.

$$\Delta v \equiv \int_{t_o}^{t} d\tau \frac{|\underline{F}(\tau)|}{m(\tau)} \tag{A.2}$$

$$\Delta v = -\frac{F}{\dot{m}} \int_{m_B}^{m_A} \frac{dm}{m} = -c\ ln\left(\frac{m_A}{m_B}\right) \tag{A.3}$$

$$\frac{m_p}{m_B} = 1 - e^{-\Delta v/c} \tag{A.4}$$

The rocket equation, taken by itself, would seem to indicate that an arbitrarily high I_{sp} is desirable to minimize the propulsion system mass on a satellite for launch. In EP systems, however, limitations on electrical power levels onboard satellites set an upper-bound on optimal propellent speeds in practice. The thrust efficiency of an electric propulsion system is defined as the quotient of the useful kinetic power imparted to the propellant and the electrical power $\mathbb{P}$ consumed in doing so, shown in Equation A.5. From this equation, the amount of power consumed is proportional to the square of the effective exhaust speed and the rate of propellant consumption. As a general rule, the mass of a satellite power supply increases with the amount of power it can deliver [pg. 7 [163]]. Therefore, the integrability of the propulsion subsystem into the satellite as a whole will limit the power available for thruster operation. Equation A.5 also shows that, at constant efficiency, the product of thrust and I_{sp} is proportional to power. This means that a tradeoff exists between operating at high thrust and high specific impulse in power limited EP systems. The greater complexity of an electric propulsion system compared to a chemical one does not justify its use for low I_{sp} applications. Electrothermal thrusters, such as arcjets or resistojets, operate with intermediate values of I_{sp} and thrust [7]. Intrinsically higher I_{sp} and power limitations require that electric thrusters operate with very low levels of thrust compared to chemical rockets. As an example, a 1 $[kW]$ thruster with a 50 % efficiency and an I_{sp} of 2000 $[s]$ has a thrust of 50 $[mN]$. In order to provide the total impulse demanded

of satellite missions, electric thrusters must be capable of firing for many hundreds or even thousands of hours. Recently, a 4.5 $[kW]$ thruster demonstrated over 10400 $[h]$ of operation during ground testing, supplying a total impulse of 8.7 $[MN-s]$ [164]. The projected lifetime of the thruster is estimated as $20-40$ thousand hours.

$$\eta_F = \frac{\frac{1}{2}\dot{m}c^2}{\mathbb{P}} = \frac{gFI_{sp}}{2\mathbb{P}} \tag{A.5}$$

A common choice of propellent in purely electric or electromagnetic propulsion systems is gaseous xenon. Xenon is attractive because of its relatively low ionization energy of 12.1 $[eV]$ and large atomic mass of 131 $[amu]$. A heavy propellant is preferable when operating at a given specific impulse and particle current[a] because the heavier propellent will provide more thrust per particle. Xenon is also stored easily onboard a satellite, is inert, and relatively safe to work with in a laboratory environment. The major drawback with xenon as a propellent is its high cost. This factor, along with a desire to utilize higher specific impulses for deep-space missions, has led to the serious consideration of krypton as an alternative [130,165,166]. Krypton has a larger ionization potential of 14.0 $[eV]$ and a lower atomic mass of 83.8 $[amu]$. Iodine is another propellant of interest. Monatomic iodine has a lower ionization potential (10.5 $[eV]$) than xenon and a similar atomic mass (126.9 $[amu]$).

[a]The particle current is the particle flux integrated over an area.

Appendix B

Basic Concepts & Definitions in Plasma Physics

Electrons and ions have very different masses, where the mass ratio $\frac{m_i}{m_e}$ between a xenon nucleus and an electron is approximately 240,000. Therefore, even with equivalent energies, electrons have much greater speeds than ions, by a factor of $\sqrt{\frac{m_i}{m_e}}$. However, the electron temperature T_e[a] is much larger than ion or neutral temperatures in Hall thrusters. This increases the average speed ratio to $\sqrt{\frac{m_i}{m_e}\frac{T_e}{T_H}}$, where T_H represents the temperature of a heavy species of particles — ions or neutrals.

When a gas is sufficiently ionized, it becomes highly conductive and electrons are able to quickly rearrange themselves to screen out applied electric fields such that the bulk plasma is quasi-neutral. Assuming the gas to be singly ionized leads to a simple, approximate relation between ion and electron number densities: $n_i \approx n_e$. Disturbed from this equilibrium, electrons react on a time scale characterized by the inverse electron plasma frequency $\omega_{p,e}$, given on the left side of Equation B.1. The length scale over which electron thermal energy is able to sustain potential differences on the order of $\frac{T_e}{e}$ is called the electron Debye length $\lambda_{D,e}$, given on the right in Equation B.1. The Debye length is also the approximate distance an electron traveling at the thermal speed can travel in a period set by the inverse electron plasma frequency. Equivalent expressions can be defined for ions, but their greater inertia prevents them from contributing to

[a] It is common practice in plasma physics to denote the product of the Boltzmann constant and a temperature kT as a single factor T.

the screening process as effectively as electrons. Therefore, the subscript identifying electrons is sometimes omitted to simplify notation.

In the presence of a uniform magnetic field, charged particles will complete circular orbits at an angular speed referred to as the cyclotron, Larmor, or gyro frequency $\omega_{c,s}$, defined on the left in Equation B.2. The radius of the circle is called the Larmor radius $r_{L,s}$, as defined in the middle of Equation B.2, where $v_{s,\perp}$ is the speed of the charged particle perpendicular to the local $\underline{B}$-field. The electron cyclotron frequency is much larger than the corresponding ion frequency, meaning that $r_{L,e} \ll r_{L,i}$ for equivalent $\underline{B}$-fields. The ratio of the cyclotron frequency to the collision frequency is called the Hall parameter Ω_s, defined on the right in Equation B.2. A group of particles with $\Omega_s \gg 1$ and a Larmor radius much less than the dimensions of the thruster is said to be magnetized — meaning that the particle trajectories are dominated by effects related to the direction of the magnetic field, known as guiding-center drifts.

$$\omega_{p,e} = \sqrt{\frac{e^2 n_e}{\epsilon_o m_e}} \quad , \quad \lambda_{D,e} = \sqrt{\frac{\epsilon_o T_e}{e^2 n_e}} \tag{B.1}$$

$$\omega_{c,s} = \frac{eB}{m_s} \quad , \quad r_{L,s} = \frac{v_{s,\perp}}{\omega_{c,s}} \quad , \quad \Omega_s \equiv \frac{\omega_{c,s}}{\nu_s} \tag{B.2}$$

In Hall thrusters, where maximum field strengths are typically a few hundred Gauss or less, electrons are strongly magnetized while ions are not because the thruster dimensions are smaller than $r_{L,i}$. In the DCFT, the magnetic field over a significant volume of the chamber is between 1-4 kG[b], weakening to below 500 G near the exit plane. Even with a peak magnetic field 25 times stronger than those typically used in Hall thrusters, only a small fraction of the low-energy ion population can be magnetized. Figure B-1 compares values of $r_{L,i}$ to the minimum and maximum radii of the DCFT chamber for magnetic fields ranging from 0.024-0.4 T [c]. The field along the centerline of the thruster, see Figure 1-2, does not exceed 2 kG until less than 1 cm from the downstream face of the anode — meaning that $r_{L,i}$ may still be comparable to, or larger than, a particle's distance from the wall in high-field regions. All of this ignores completely Ω_i, which must also be much larger than unity for ions to be considered magnetized. Even if they are not tightly bound to field lines, ions are coupled strongly to the highly magnetized electrons through the electric field, however.

[b] For comparison, the geomagnetic field strength at the surface of the Earth ranges from 0.3-0.6 G [pg. 56] [167]
[c] 1 T = 10^4 G

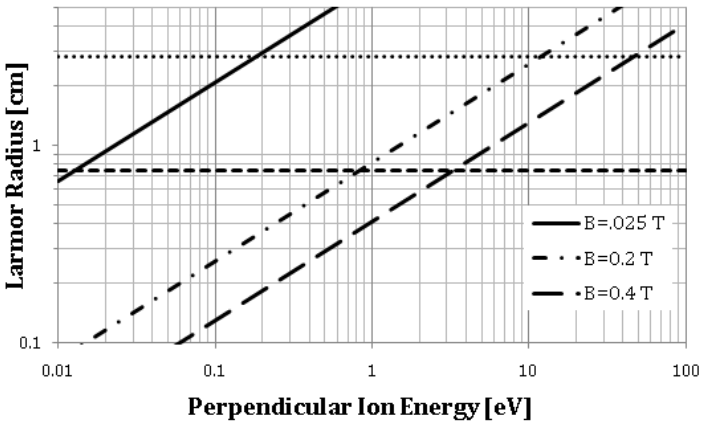

Figure B-1: Larmor radii of Xe^+ for a range energies due to velocities perpendicular to a varied magnetic field strength. The topmost, dotted horizontal line corresponds to the radius of the original DCFT design at the exit plane (2.8 cm), while the other horizontal line denotes the minimum radius of the insulating cone (0.75 cm).

An important parameter for ions is the Bohm speed, given in Equation B.3. When entering regions of strong electric fields, the ion speed at the "entrance" must be at least equal to the Bohm speed in order for the potential to vary smoothly from the quasi-neutral plasma bulk to regions where space-charge effects are influential.

$$V_B = \sqrt{\frac{T_e}{m_i}} \tag{B.3}$$

Appendix C

Fluid Models: Approximations for Closure of the Moment Equations

The Boltzmann equation for each species (s) is shown in Equation C.1. The single species distribution f_s is a function of position $\underline{x}$, microscopic velocity $\underline{w}$, and time t, representing the phase-space density of particles at position $\underline{x}$ contained in an infinitesimal volume in velocity-space about $\underline{w}$. This form of the Boltzmann equation is valid for point particles subject to the full Lorentz force because the electric force is independent of $\underline{w}$ and the magnetic force will be perpendicular to $\underline{w}$ [pp. 327-328 [71]]. The subscript indicating the gradient with respect to the vector coordinate $\underline{x}$ is dropped in subsequent equations, such that the gradient operator is assumed to be with respect to spatial coordinates unless indicated. This equation is only uncoupled in situations where one species is present, or only self-interactions are important. When interactions with different species are important (such as ionization, dissimilar particle scattering, excitation, etc.), they must be accounted for on the right-hand side of the Boltzmann equation through modification of the collision term $\mathbb{C}_s$, as in Equation C.2.

$$\frac{\partial}{\partial t}\left(f_s\right) + \underline{w}_s \cdot \nabla_{\underline{x}}\left(f_s\right) + \left[\frac{\underline{F}_s}{m_s}\right] \cdot \nabla_{\underline{w}_s}\left(f_s\right) = \left[\frac{\partial}{\partial t}\left(f_s\right)\right]_{coll} \equiv \mathbb{C}_s \tag{C.1}$$

$$\mathbb{C}_s \equiv \mathbb{C}_{ss} + \sum_r \mathbb{C}_{sr} \tag{C.2}$$

Multiplying the Boltzmann equation by a function of velocity $\phi\left(\underline{w}_s\right)$ and integrating over velocity space provides a generalized moment relationship, as shown in Equation C.4. ϕ can be

either a scalar or a vector, but is neither a function of $\underline{x}$ nor t. Brackets in equations represent quantities averaged over all of velocity space, as in Equation C.3. Again, Equation C.4 is not valid for moments of functions that depend on position or time, such as $\phi = \underline{w} - \underline{u}\left(\underline{x}, t\right)$, where $\underline{u} \equiv \langle \underline{w} \rangle$.

$$n_s \langle \phi_s \rangle \equiv \int \phi\left(\underline{w}_s\right) f_s\left(\underline{x}, \underline{w}_s, t\right) \; d\underline{w}_s \tag{C.3}$$

$$\frac{\partial}{\partial t}\left(n_s \langle \phi_s \rangle\right) + \nabla \cdot \left[n_s \langle \phi_s \underline{w}_s \rangle\right] - n_s \left\langle \left(\frac{\underline{\underline{F}}_s}{m_s} \cdot \nabla_{\underline{w}_s}\right) \phi_s \right\rangle = \int \phi_s \mathbb{C}_s \; d\underline{w}_s \equiv \Delta\left[\phi_s\right] \tag{C.4}$$

The conservation equations — also referred to as "moment" or "transfer" equations — for more commonly measured quantities such as density, momentum, and energy can be found by taking three important moments, shown in Equations C.12, C.13, and C.14. The average, or bulk, velocity $\underline{u}_s$ is defined as $\langle \underline{w}_s \rangle$, and the random velocity is locally defined as the difference of the microscopic and average velocities, given in Equation C.5. Note that the average value of the random velocity is zero, by definition. The pressure tensor, or momentum flux tensor, $\underline{\underline{P}}_s$ for a given species is given by Equation C.6. The species temperature, heat flux and current density are defined in Equations C.7-C.9. The definitions in Equations C.6-C.9 are valid for arbitrary distribution functions. Also general are the definitions of the scalar pressure p_s ($p_s \neq \underline{\underline{P}}_s$) and viscous-stress tensor $\underline{\underline{\tau}}_s$ in Equations C.10 and C.11. These are the common definitions of terms found in the literature, though analogous variations may be presented when they serve

to simplify notation. In the following equations, $\underline{\underline{I}}$ is used to denote the unit tensor.

$$\underline{c}_s \equiv \underline{w}_s - \langle \underline{w}_s \rangle = \underline{w}_s - \underline{u}_s \tag{C.5}$$

$$\underline{\underline{P_s}} \equiv m_s n_s \langle \underline{c}_s \; \underline{c}_s \rangle \tag{C.6}$$

$$\frac{3}{2} k T_s \equiv \frac{1}{2} m_s \langle c_s^2 \rangle = \frac{1}{2} m_s \langle w_s^2 \rangle - \frac{1}{2} m_s u_s^2 \tag{C.7}$$

$$\underline{q}_s \equiv \frac{1}{2} n_s m_s \langle c_s^2 \underline{c}_s \rangle \tag{C.8}$$

$$\underline{j}_s \equiv q_s n_s \underline{u}_s \tag{C.9}$$

$$p_s \equiv \frac{tr\left(\underline{\underline{P_s}}\right)}{3} = m_s n_s \frac{\langle c_s^2 \rangle}{3} = n_s k T_s \tag{C.10}$$

$$\underline{\underline{\tau_s}} \equiv p_s \underline{\underline{I}} - \underline{\underline{P_s}} \tag{C.11}$$

$$\frac{\partial}{\partial t} n_s + \nabla \cdot (n_s \underline{u}_s) \equiv \Delta [1]_s \tag{C.12}$$

$$\frac{\partial}{\partial t} (m_s n_s \underline{u}_s) + \nabla \cdot (m_s n_s \underline{u}_s \underline{u}_s) + \nabla \cdot \underline{\underline{P_s}} - q_s n_s \underline{E} - \underline{j}_s \times \underline{B} \equiv \Delta [m_s \underline{w}_s] \tag{C.13}$$

$$\frac{\partial}{\partial t} \left[\frac{3}{2} n_s k T_s + \frac{1}{2} m_s n_s u_s^2 \right] + \nabla \cdot \left[n_s \underline{u}_s \left(\frac{3}{2} k T_s + \frac{1}{2} m_s u_s^2 \right) + \underline{u}_s \cdot \underline{\underline{P_s}} + \underline{q}_s \right] - \underline{j}_s \cdot \underline{E} = \Delta \left[\frac{1}{2} m_s w_s^2 \right] \tag{C.14}$$

The momentum equation (C.13) can be simplified by evaluating the partial derivative with respect to time and the divergence of the directed momentum flux tensor $m_s n_s \underline{u}_s \underline{u}_s$, so that the left hand side (LHS) of the continuity equation (C.12) appears within a term on the LHS of the momentum equation. A simplified energy equation, given in Equation C.19, is derived in an identical way by expanding the energy flux terms $n_s \underline{u}_s \left(\frac{3}{2} k T_s + \frac{1}{2} m_s u_s^2 \right)$ in C.14. Further simplification comes from expanding the divergence of the volumetric rate of work done by pressure $\nabla \cdot \left(\underline{u}_s \cdot \underline{\underline{P_s}} \right)$, according to Equation C.16. The term with a dot product between the species bulk velocity and divergence of the pressure tensor can be replaced after taking the dot product of $\underline{u}_s$ with the simplified momentum equation shown in Equation C.18. The total derivative operator $\frac{d}{dt_s}$ is defined in Equation C.17.

Up to this point, the representation of the collision operator has been as abstract as possible.

This is because a closed set of equations can never be obtained by repeatedly taking moments of the Boltzmann equation. Closure of the moment equations is a critical step in the development of any model, and a variety of methods are used. Two common examples are closure by an *ansatz* (which is just another word for a guess), or by expansions of the distribution function in orders of deviation from a specified zeroth order distribution f_0. A common choice of f_0 is a Maxwellian distribution f_M, given in Equation C.15 for a three dimensional velocity space. However, f_0 may also be chosen as a non-Maxwellian isotropic function of velocity[pp. 351-353, 375 [71]], or even a non-isotropic distribution with multiple temperatures [pp. 204, 292 [168]]. The Chapman-Enskog and Burnett expansions are both based on the choice $f_0 \equiv f_M$ — the Chapman-Enskog set of equations[a] being valid to first order in Kn while the Burnett equations are accurate up second order [67]. Further discussions of fluid model closure are given in the literature, with the review by Robson, et al. [169] being particularly pertinent to this discussion. However, the discussion falls short of suggesting a method of closure when inelastic interactions are important, such as ionization. Interestingly, an ansatz meant to include ionization effects has been utilized in a 1-D fluid model of the Franck-Hertz experiment [80], reproducing effects that were previously inaccessible to fluid models.

$$f_M = n_s \left(\frac{m_s}{2\pi k T_s} \right)^{3/2} \exp \left[\frac{-m_s}{2kT_s} \left(\underline{w}_s - \underline{u}_s \right)^2 \right] \tag{C.15}$$

$$\nabla \cdot \left(\underline{u}_s \cdot \underline{\underline{P_s}} \right) = \underline{\underline{P_s}} : \nabla \underline{u}_s + \underline{u}_s \cdot \left(\nabla \cdot \underline{\underline{P_s}} \right) \tag{C.16}$$

$$\frac{d}{dt_s} \equiv \frac{\partial}{\partial t} + \underline{u}_s \cdot \nabla \tag{C.17}$$

$$n_s m_s \frac{d}{dt_s} \underline{u}_s + \nabla \cdot \underline{\underline{P_s}} - q_s n_s \left(\underline{E} + \underline{u}_s \times \underline{B} \right) = \Delta \left[m_s \underline{w}_s \right] - m_s \underline{u}_s \Delta \left[1 \right]_s \tag{C.18}$$

$$n_s \frac{d}{dt_s} \left(\frac{3}{2} k T_s \right) + \nabla \cdot \underline{q}_s + \underline{\underline{P_s}} : \nabla \underline{u}_s = \Delta \left[\frac{1}{2} m_s w_s^2 \right] - \underline{u}_s \cdot \Delta \left[m_s \underline{w}_s \right] - \Delta \left[1 \right]_s \left(\frac{3}{2} k T_s - \frac{1}{2} m_s u_s^2 \right) \tag{C.19}$$

Still, the validity of a closure depends on the size of the error incurred by use of an ansatz

[a] Equivalent to the Navier-Stokes equations

or truncating the expansion of f_s. For Knudsen numbers that do not satisfy $Kn < 1$, systems of equations derived from expansions in terms of $Kn \ll 1$ will not be accurate. Based on arguments by Struchtrup (2005) [pp. 93-95, 169-170 [69]], for methods not based entirely on the smallness of Kn, convergence in highly rarefied situations is slow, requiring the consideration of hundreds of moment equations. Also, these higher order moment methods require specifying boundary conditions of physically ambiguous quantities, as well as restrictive assumptions about molecular interactions — the author has only come across examples in the literature where Maxwellian interaction potentials and the Bhatnager-Gross-Krook model [170] have been used. No complications from the inclusion of the full Lorentz force have been considered in these higher order methods either, to the author's knowledge. Presumably, the computational effort required to solve the high-order moment equations lies somewhere between the Navier-Stokes equations and the full Boltzmann equation.

Appendix D

Relevant Cross Section Data & Estimates

This section includes a discussion of the cross sections used to describe various collisional processes in the MIT PIC plasma thruster model. The source of each cross section is cited, and in some cases the accuracy of the data is discussed. Outlines of the methods by which certain interactions are represented in the model are also mentioned. To be clear, the assumed bulk plasma properties near the anode and at the exit plane are repeated in Table D.1. The cross section estimates given in this section are summarized in Table 1.3, located in Section 1.2.7.

Table D.1: Estimated electron, ion, and neutral temperatures and densities at two locations in the DCF thruster chamber.

		Anode	Exit
T_e	[eV]	5	25
n_e	[cm^{-3}]	$5{\cdot}10^{12}$	$5{\cdot}10^{11}$
n_n	[cm^{-3}]	10^{14}	10^{12}
T_i	[eV]	1	1
T_n	[K]	500	500

Electron-Neutral (e-n) Interactions

The model includes representations of elastic scattering, excitation and ionization from (e-n) collisions. The cross section data, along with the polynomial fits used in the code, are given by Szabo [pp. 94-99[97]]. The relative speed between electrons and neutrals is dominated by the electron motion, because $m_e \gg m_H$ and $T_e \gg T_n$. For the DCFT, the total (e-n) scattering cross-section is estimated to be between $2.5 - 4 \times 10^{-15} cm^2$. The contributions to this cross section from electron impact-ionization and excitation are $0 - .5 \times 10^{-15} cm^2$ each, over the considered range of parameters. The ionization cross section accounts for about half the total cross section at a relative energy of 100 eV. Neutral excitation includes only a single level, 8.32 eV above the ground state. Excited neutral particles are assumed to radiate immediately, so no collisions involving pre-excited neutrals are accounted for.

Ion-Neutral (i-n) Interactions

Elastic scattering and resonant charge-exchange collisions between ions and neutrals were both included in Szabo's code. The charge-exchange cross section Q_{in}^{cex} is estimated using Equation D.1[171], where g_{in} has the units $^{cm}/_{s}$[a]. For estimation purposes, g_{in} is taken as the ion Bohm speed corresponding to the estimated values of T_e, which is much larger than the neutral speed. The resulting values for Q_{in}^{cex} are about $8 \times 10^{-15} cm^2$. Equation D.1 is based on the work of Sakabe and Izawa, and is utilized in the MIT PIC model. Rapp and Francis[172] provide an alternate set of charge-exchange cross sections, with values smaller than provided in Equation D.1 by $1 - 2 \times 10^{-15} cm^2$. Different charge-exchange cross sections have also been used by other authors[173].

$$Q_{in}^{cex} = 1.1916 \left[1.81 \times 10^{-14} - 2.12 \times 10^{-15} \log_{10}(g_{in})\right] \ [cm^2] \tag{D.1}$$

[a]The units of g_{ij} are $^{m}/_{s}$ in all other equations given here, except Equation D.2, where the units of all terms are in CGS units

$$Q_{in}^{scat} = 2.21\pi \left(\frac{\alpha e^2}{\mu g_{in}^2} \right)^{1/2} cm^2 \, [CGS \ units] \tag{D.2}$$

$$\begin{aligned} Q_{in}^{scat} &= \frac{2.21}{4\epsilon_o} \left(\frac{\alpha e^2}{\mu g_{in}^2} \right)^{1/2} m^2 \\ &= \frac{6.42 \times 10^{-12}}{g_{in}} cm^2 \end{aligned} \tag{D.3}$$

The cross section for (i-n) elastic scattering is given in Equations D.2 [174] and D.3. The electric dipole polarizability α of xenon is $4.500 \times 10^{-40} Fm^2$, or $4.044 \times 10^{-24} cm^3$ in CGS units [175]. The (i-n) scattering cross section is based on the theoretical arguments by A. Dalgarno, et al. and numerical calculations by Langevin and Hasse [174]. The cross section was incorrectly converted to SI units by Oh [101], overestimating this particular cross section by 30%. The incorrect value was carried over by Szabo, but was corrected for PTpic. However, experimental comparisons to this model were not found in the literature and this error likely had negligible impacts on previous uses of the MIT PIC model. The value shown on the second line of Equation D.3 is the same as that given by Boyd and Dressler in a Hall thruster plume model [173]. Estimates of Q_{in}^{scat} for the DCFT are betwen $1.5 - 3.4 \times 10^{-15} cm^2$, which is of the same order as the charge exchange cross section. Ions created via electron impact ionization of neutral xenon atoms are assigned a velocity equal to the bulk neutral velocity in that cell. Cross sections for these slower ions will be larger than the estimates made here, but these particles will quickly gain energy in the electric field. Ion excitation is not currently included in the model.

Coulomb Interactions

Typically, cross sections are evaluated by integrating the differential cross section over all possible impact parameters $b \in [0, \infty)$, or scattering angles $\chi \in [0, \pi]$, where the relationship between b and χ depends on the interaction potential. However, for charged particle encounters, or Coulomb collisions, the integrals diverge due to the long range nature of the interaction potential $\left(\propto \frac{1}{r}\right)$ between particles. Selecting a maximum impact parameter b_{max}, beyond which it is assumed no interaction occurs, allows the integrals to have finite values. The evaluated cross section must be insensitive to the chosen cutoff in order to be meaningful.

With that in mind, we look at two cross sections for binary collisions between charged

particles. The procedure for calculating these cross sections is presented in many texts related to the physics of plasmas[71,176,177]. Evaluating the integral scattering cross section, which is just the cross section for any amount of scattering, it is found that $Q = \pi b_{90}^2 \Lambda = \pi b_{max}^2$, where b_{90} is the impact parameter corresponding to $\chi = 90^o$, and $\Lambda \equiv \frac{b_{max}}{b_{90}}$. Equation D.4 gives the definition of b_{90}, where the numerator contains the charges of the affected particles, μ is the reduced mass of the pair, and the meaning of g remains as the magnitude of the relative velocity of the two particles. This is expected, because selecting a cutoff is identical to specifying the area over which any interaction can occur. A typical value for b_{90} in Hall and cusped-field thrusters is $50 - 100\ pm$, and it can be an order of magnitude lower for collisions between high energy ions. Evidently, the integral cross section is extremely sensitive to the chosen value of b_{max}.

The cross section for momentum transfer Q^m is different in that it weights impact parameters that cause severe deflections much more strongly than weak deflections. The dependence of Q^m on the relative energy of two charged particles, as well as the cutoff b_{max} is given in Equation D.5. Here, the dependence on b_{max} is much weaker, through the $\ln \Lambda$ term, called the "Coulomb logarithm". A natural choice for b_{max} is the electron Debye length λ_D, which represents the length scale over which electrons can shield time-independent electric fields. As a result, particle pairs with impact parameters greater than λ_D are essentially incapable of interacting due to the collective rearrangement of the surrounding electrons. In reality, the cutoff is not so clear cut, but the dependence of Q^m on $\ln \Lambda$ means that the choice of b_{max} need only be of the correct order. For a typical Hall thruster Debye length of $5 - 50\ \mu m$, we see that $\Lambda \gg 1$ and $\ln \Lambda \approx 10 - 15$. Calculations with different parameters show that the Coulomb logarithm does not vary strongly over a range of plasma temperatures and densities spanning many orders of magnitude[pg. 59 [71]]. Physically, the Coulomb logarithm represents the approximate factor by which contributions to Q^m from small angle collisions $b > 2b_{90}$ are greater than contributions for collisions with large scattering angles $b < 2b_{90}$ [pp. 108-111 [129]].

$$b_{90} \equiv \frac{|q_1 q_2|}{4\pi\epsilon_o \mu g^2} \tag{D.4}$$

$$Q^m = 4\pi b_{90}^2 \ln \sqrt{1 + \Lambda^2} \Rightarrow 4\pi b_{90}^2 \ln \Lambda, \text{ for } \Lambda \gg 1 \tag{D.5}$$

These results assume that charged particles interact on a one-to-one, or binary, basis. However, the number of plasma particles within the shielding distance is given by the plasma parameter $N_D = \frac{4\pi}{3} n_e \lambda_D^3$, with a typical value of $10 - 300 \times 10^3$ for plasmas considered here. This is discussed further at the end of Section 1.2.7.

Coulomb Cross Section Estimates

The MFP for Coulomb interactions will be estimated based on the Coulomb momentum transfer cross section Q_{jk}^m, whereas the previous estimates were made using using scattering cross sections. This is a reasonable approach because Q_{jk}^m represents the cross section for the loss of directed momentum, which is analogous to the total cross sections for scattering, excitation, ionization, and charge-exchange because the interactions involve greater deflections of the involved particles. The Coulomb cross sections between electrons, electrons and ions, and ions can all be calculated from Equation D.5, where the differences arise from the relative energy dependence of b_{90}. Notice also that the cross section depends on the ambient electron temperature and density, though this effect is weak as mentioned previously. Continuing to use the average thermal speed for electrons and the Bohm speed for ions, estimates of the Coulomb momentum transfer cross sections are given in Table D.2. The sensitivity of b_{90} to the relative speed is evident in the wide range of values estimated. For electron interactions, the range is the result of estimations using a low electron temperature of 5 eV, and a larger value of 25 eV, which effects the average relative speed much more strongly than the value of $\ln \Lambda$. $Q_{ee}^m = 4Q_{ei}^m$ because the reduced mass is half as big for (e-e) as in (e-i) interactions. The cross section amongst ions at the Bohm speed is very high, especially at low electron temperatures.

Table D.2: Estimated Coulomb momentum transfer cross sections Q^m for electrons and Xe^+. The last cross section $Q_{ii,HE}^m$ is for (i-i) interactions amongst ions with large relative energies. A relative ion energy of 250 V was chosen for this example.

Cross Section	[cm^2]
Q_{ei}^m	1-18 $\times 10^{-15}$
Q_{ee}^m	5-65 $\times 10^{-15}$
Q_{ii}^m	20-400 $\times 10^{-15}$
$Q_{ii,HE}^m$	.02 $\times 10^{-15}$

Neutral-Neutral Interactions

The (n-n) elastic scattering cross section for xenon atoms is given by Equation D.6. The form and coefficients of the cross section are given by Koura and Matsumoto[178] based on fits to experimental data that are valid for normalized relative energies between 300-15000 K[b]. Oh and Boyd both use the same expression to evaluate this cross section in models of a Hall thruster plume. The estimated value of this cross section within the DCFT is $5.5 \times 10^{-15} cm^2$.

$$Q_{nn}^{scat} = 2.117 \times 10^{-14} g_{nn}^{-.24} cm^2 \tag{D.6}$$

Despite being comparable or larger than the cross sections for neutral encounters with ions and electrons, (n-n) scattering collisions are ignored in the model because of the added complexity that would be necessary to model such collisions accurately. The standard Monte Carlo collision (MCC) procedure does not properly model the effects of collisions between identical particles. A more correct representation of these collisions is available using the direct simulation Monte Carlo (DSMC) technique, whose logical structure is quite distinct and more complex than the MCC method. A more detailed comparison of the MCC and DSMC methods is given in Section 1.2.3. Szabo decided that accounting for the self-equilibration of the neutral population was not worth increasing the complexity of the code, especially with his main objective being the simulation of the discharge plasma.

[b]The relative energy is calculated as ϵ/k, where $\epsilon \equiv \frac{1}{2}\mu g_{nn}^2$. In these expressions, μ is the reduced mass of the interacting particles, and k is the Boltzmann constant. This explains why the normalized energy units are expressed in terms of the Kelvin scale.

References

[1] Gorshkov, O., Belikov, M., Muravlev, V., Shagayda, A., Shanbhogue, K., Premkumar, S., Nandalan, V., and Jayaraman, M., "The GEOSAT Electrical Propulsion Subsystem Based on the KM-45 HET," *Acta Astronautica*, Vol. 63, No. 1-4, 2008, pp. 367 – 378.

[2] Belikov, M. B., Gorshkov, O. A., Dyshlyuk, E. N., Lovtzov, A. S., and Shagayda, A. A., "Development of Low-Power Hall Thruster with Lifetime up to 3000 Hours," 30^{th} *International Electric Propulsion Conference*, Florence, IT, September 2007, IEPC-2007-129.

[3] Hargus, Jr., W. and Nakles, M. R., "Ion Velocity Measurements Within the Acceleration Channel of a Low-Power Hall Thruster," *Plasma Science, IEEE Transactions on*, Vol. 36, No. 5, October 2008, pp. 1989 –1997.

[4] Saverdi, M., Signori, M., and Biagioni, L., "Experimental Characterization of the HT-100 Hall Thruster in Twin Engine Cluster Configuration," 30^{th} *International Electric Propulsion Conference*, Florence, IT, September 2007, IEPC-2007-320.

[5] Loyan, A. V. and Maksymenko, T. A., "Performance Investigation of SPT-20M Low Power Hall Effect Thruster," 30^{th} *International Electric Propulsion Conference*, Florence, Italy, September 2007, IEPC-2007-100.

[6] Kozubskii, K., Murashko, V., Rylov, Y., Trifonov, Y., Khodnenko, V., Kim, V., Popov, G., and Obukhov, V., "Stationary Plasma Thrusters Operate in Space," *Plasma Physics Reports*, Vol. 29, 2003, pp. 251–266.

[7] Martínez-Sánchez, M. and Pollard, J. E., "Spacecraft Electric Propulsion - An Overview," *Journal of Propulsion and Power*, Vol. 14, No. 5, 1998, pp. 688–699.

[8] Hofer, R. R., Mikellides, I. G., Katz, I., and Goebel, D. M., "BPT-4000 Hall Thruster Discharge Chamber Erosion Model Comparison with Qualification Life Test Data," 30^{th} *International Electric Propulsion Conference*, Florence, IT, September 2007, IEPC-2007-267.

[9] Raitses, Y. and Fisch, N. J., "Parametric Investigations of a Nonconventional Hall Thruster," *Physics of Plasmas*, May 2001.

[10] Diamant, K. D., Pollard, J. E., Raitses, Y., and Fisch, N. J., "Ion Velocity Measurements Within the Acceleration Channel of a Low-Power Hall Thruster," *Plasma Science, IEEE Transactions on*, Vol. 38, No. 4, April 2010, pp. 1052–1057.

[11] Kaufman, H. R., Robinson, R. S., and Seddon, R. I., "End-Hall Ion Source," *Journal of Vacuum Science and Technology A*, Vol. 5, No. 4, 1987, pp. 2081–2084.

[12] Polzin, K. A., Raitses, Y., Gayoso, J. C., and Fisch, N. J., "Comparisons in Performance of Electromagnet and Permanent-Magnet Cylindrical Hall-Effect Thrusters," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-6695.

[13] Arkhipov, B., Bober, A., Day, M., Gnizdor, R., Kozubsky, K., and Maslennikov, N., "Extending the Range of SPT Operation - Development Status of 300 and 4500 W Thruster," 32^{nd} *Joint Propulsion Conference & Exhibit*, Lake Buena Vista, FL, July 1996, AIAA-1996-2708.

[14] Belikov, M. B., Gorshkov, O. A., Muravlev, V. A., Rizakhanov, R. N., Shagayda, A. A., and Vasin, A. I., "Qualification of Hall Thruster for Small Spacecraft Applications," *The Tenth International Workshop on Combustion and Propulsion*, Lerici, IT, September 2003, paper 19.

[15] Cheng, S. Y. and Martínez-Sánchez, M., "Hybrid Particle-in-Cell Erosion Modeling of Two Hall Thrusters," *Journal of Propulsion and Power*, Vol. 24, No. 5, 2008, pp. 987–998.

[16] Misuri, T., Tavoni, S., and Andrenucci, M., "Monte Carlo Erosion Model for HT-100 Lifetime Prediction," *ESA/3AF Space Propulsion Conference*, San Sebastian, ES, May 2010.

[17] Dignani, D., Ducci, C., Cifali, G., Rossetti, P., and Andrenucci, M., "HT-100 Hall Thruster Characterization Tests Results," 32^{nd} *International Electric Propulsion Conference*, Wiesbaden, DE, September 2011, IEPC-2011-191.

[18] Rossetti, P. and Andrenucci, M., "HT-100 Development Status," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-126.

[19] Kim, V., Kozlov, V., Lazurenko, A., Popov, G., Skrylnikov, A., Clauss, C., Day, M., and Sancovic, J., "Development and Characterization of Small SPT," 34^{th} *Joint Propulsion Conference & Exhibit*, Cleveland, OH, July 1998, AIAA-1998-3335.

[20] Jacobson, D. T. and Jankovsky, R. S., "Test Results of a 200 W Class Hall Thruster," 34^{th} *Joint Propulsion Conference & Exhibit*, Cleveland, OH, July 1998, AIAA-1998-3792.

[21] Loyan, A. V., Maksymenko, T. A., Rybalov, O. P., and Podgorny, V. A., "The Investigations of Erosion of Discharge Chamber of Low Power Stationary Plasma Thruster During Long-Continued Lifetime Test," *АВИАЦИОННО-КОСМИЧЕСКАЯ ТЕХНИКА И ТЕХНОПОГИЯ*, Vol. 65, No. 8, 2009, pp. 121–124.

[22] Goebel, D. M. and Katz, I., *Fundamentals of Electric Propulsion*, John Wiley and Sons, Inc., 2008.

[23] Janes, G. S., Dotson, J., and Wilson, T., "Momentum Transfer Through Magnetic Fields," *Proceedings of Third Symposium on Advanced Propulsion Concepts, Vol. 1*, Cincinnati, OH, October 1962.

[24] Kornfeld, G., Koch, N., and Harmann, H.-P., "Physics and Evolution of HEMP-Thrusters," 30^{th} *International Electric Propulsion Conference*, Florence, Italy, September 2007, IEPC-2007-108.

[25] Harmann, H.-P., Koch, N., and Kornfeld, G., "Low Complexity and Low Cost Electric Propulsion System for Telecom Satellites Based on HEMP Thruster Assembly," 30^{th} *International Electric Propulsion Conference*, Florence, IT, September 2007, IEPC-2007-114.

[26] Kornfeld, G., Harmann, H.-P., and Koch, N., "Status and Limited Life Test Results of the Cylindrical HEMP 3050 Thruster," 41^{st} *Joint Propulsion Conference & Exhibit*, Tucson, Arizona, July 2005, AIAA-2005-4223.

[27] Courtney, D. G., "Development and Characterization of a Diverging Cusped Field Thruster and a Lanthanum Hexaboride Hollow Cathode," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2008.

[28] Daspit, R., Lozano, P., and Martínez-Sánchez, M., "Characterization and Optimization of a Diverging Cusped Field Thruster with a Calibrated Horizontal Accelerometer," 48^{th} *Joint Propulsion Conference & Exhibit*, San Diego, CA, July 2011, AIAA-2011-6069.

[29] Gildea, S., Martínez-Sánchez, M., Nakles, M., and Hargus, Jr., W. A., "Experimentally Characterizing the Plume of a Divergent Cusped-Field Thruster," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-259.

[30] Gildea, S. R., Matlock, T. S., Lozano, P., and Martínez-Sánchez, M., "Low-Frequency Oscillations in the Diverging Cusped-Field Thruster," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-7014.

[31] T. Matlock, a. a., "Spectroscopic and Electrostatic Investigation of the Diverging Cusped-Field Thruster," 45^{th} *Joint Propulsion Conference & Exhibit*, Denver, Co, August 2009, AIAA-2009-4813.

[32] Matlock, T., Gildea, S., Hu, F., Becker, N., Lozano, P., and Martínez-Sánchez, M., "Magnetic Field Effects on the Plume of a Diverging Cusped-Field Thruster," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-7104.

[33] Matlock, T. S., Hu, F., and Martínez-Sánchez, M., "Controlling Plume Divergence in a Cusped-Field Thruster," 32^{nd} *International Electric Propulsion Conference*, Wiesbaden, DE, September 2011, IEPC-2011-178.

[34] Gildea, S. and Martínez-Sánchez, M., "Improvements in Divergent Cusped-Field Thruster Modeling," *ESA/A3F Space Propulsion Conference*, San Sebastián, ES, May 2010.

[35] Matlock, T. S., "An Exploration of Prominent Cusped-Field Thruster Phenomena: The Hollow Conical Plume and Anode Current Bifurcation," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, September 2012.

[36] Young, C., Smith, A., and Cappelli, M., "Preliminary Characterization of a Diverging Cusped Field (DCF) Thruster," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-166.

[37] MacDonald, N. A., "Laser Induced Fluorescence Characterization of Cusped Field Plasma Thrusters," Ph.D. Thesis, Stanford University, Standford, CA, June 2012.

[38] Cedolin, R. J., Hargus, Jr., W. A., Storm, P. V., Hanson, R. K., and Cappelli, M. A., "Laser-Induced Fluorescence Study of a Xenon Hall Thruster," *Applied Physics B*, Vol. 65, 1997, pp. 459–469.

[39] MacDonald, N. A., Cappelli, M. A., and Hargus, Jr., W. A., "Ion Velocity Distribution in a Low-Power Cylindrical Hall Thruster," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-7103.

[40] Williams, Jr., G. J., Smith, T. B., Gulczinski, III, F. S., Beal, B. E., Gallimore, A. D., and Drake, R. P., "Laser Induced Fluorescence Measurement of Ion Velocities in the Plume of a Hall Effect Thruster," 35^{th} *Joint Propulsion Conference & Exhibit*, June, address = Los Angeles, CA, note = AIAA-1999-2424 1999.

[41] Hargus, Jr. , W. A. and Cappelli, M. A., "Laser-induced Fluorescence Measurements of Velocity Within a Hall Discharge," *Journal of Applied Physics B*, Vol. 72, 2001, pp. 961–969.

[42] Haas, J. M., "Low-perturbation Interrogation of the Internal and Near-Field Plasma Structure of a Hall Thruster Using a High-speed Probe Positioning System," Ph.D. Thesis, Univeristy of Michigan, Ann Arbor, MI, 2001.

[43] MacDonald, N. A., Cappelli, M. A., Gildea, S. R., Martínez-Sánchez, M., and Hargus, Jr., W. A., "Laser-Induced Fluorescence Velocity Measurements of a Diverging Cusped Field Thruster," *Journal of Physics D: Applied Physics*, Vol. 44, 2011, pp. 295203.

[44] Matlock, T., Lozano, P., and Martínez-Sánchez, M., "Discharge Chamber Wall Flux Measurements in a Diverging Cusped-Field Thruster," *ESA/A3F Space Propulsion Conference*, San Sebastián, ES, May 2010.

[45] Gildea, S. R., Matlock, T. S., Martínez-Sánchez, M., and Hargus, Jr., W. A., "Erosion Measurements in a Low-Power Cusped-Field Plasma Thruster," *Journal of Propulsion & Power*, 2012, accepted for publication.

[46] King, L. B. and Gallimore, A. D., "Mass Spectral Measurements in the Plume of an SPT-100 Hall Thruster," *Journal of Propulsion and Power*, Vol. 16, No. 6, 2000, pp. 1086–1092.

[47] Thomas, C. A., G. N., and Cappelli, M. A., "Nonintrusive characterization of the azimuthal drift current in a coaxial ExB discharge plasma," *Physical Review E*, Vol. 74, 2006, pp. 056402.

[48] Gascon, N., Dudeck, M., and Barral, S., "Wall Material Effects in Stationary Plasma Thrusters. I. Parametric Studies of an SPT-100," *Physics of Plasmas*, Vol. 10, No. 10, 2003, pp. 4123–4136.

[49] Simon, A., "Ambipolar Diffusion in a Magnetic Field," *Phys. Rev.*, Vol. 98, No. 2, Apr 1955, pp. 317–318.

[50] Yim, J. T., "Computational Modeling of Hall Thruster Channel Wall Erosion," Ph.D. Thesis, The University of Michigan, Ann Arbor, MI, 2008.

[51] Cho, S., , Komurasaki, K., and Arakawa, Y., "Lifetime Simulation of a SPT-Type Hall Thruster by Using a 2D Fully Kinetic PIC Model," 48^{th} *Joint Propulsion Conference & Exhibit*, Atlanta, GA, August 2012, AIAA-2012-4016.

[52] Jameson, K. K., "Investigation of Hollow Cathode Effects on Total Thruster Efficiency in a 6 kW Hall Thruster," Ph.D. Thesis, Univeristy of California Los Angeles, Los Angeles, CA, 2008.

[53] Sommerville, J. D., "Hall-Effect Thruster-Cathode Coupling The Effect of Cathode Position and Magnetic Field Topology," Ph.D. Thesis, Michigan Technical University, Houghton, MI, 2009.

[54] Randolph, T., Kim, V., Kaufman, H., Kozubsky, K., Zhurin, V., and Day, M., "Facility Effects on Stationary Plasma Thruster Testing," 23^{rd} *International Electric Propulsion Conference*, Seattle, WA, September 1993, IEPC-1993-097.

[55] Walker, M., "Effects of Facility Backpressure on the Performance and Plume of a Hall Thruster," Ph.D. Thesis, Univeristy of Michigan, Ann Arbor, MI, 2005.

[56] Fife, J. M., "Hybrid-PIC Modeling and Electrostatic Probe Survey of Hall Thrusters," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, February 1999.

[57] Szabo, J., Martínez-Sánchez, M., and Monheiser, M., "Application of 2-D Hybrid PIC Code to Alternative Hall Thruster Geometries," 34^{th} *Joint Propulsion Conference & Exhibit*, Cleveland, OH, July 1998, AIAA-1998-3795, doi: 10.2514/6.1998-3795.

[58] Scharfe, M. K., Thomas, C. A., Scharfe, D. B., Gascon, N., Capelli, M. A., and Fernandez, E., "Shear-Based Model for Electron Transport in Hybrid Hall Thruster Simulations," *IEEE Transactions on Plasma Science*, Vol. 36, No. 5, October 2008.

[59] Hofer, R. R., Katz, I., Mikellides, I. G., Goebel, D. M., Jameson, K. K., Sullivan, R. M., and Johnson, L. K., "Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations," 44^{th} *Joint Propulsion Conference & Exhibit*, Hartford, CT, July 2008, AIAA-2008-4924.

[60] Koo, J. W. and Boyd, I. D., "Modeling of Anomalous Electron Mobility in Hall Thrusters," *Physics of Plasmas*, Vol. 13, 2006, pp. 033501.

[61] Koizumi, H., Komurasaki, K., and Arakawa, Y., "Numerical Prediction of Wall Erosion on a Hall Thruster," *Vacuum*, Vol. 83, 2009, pp. 67–71.

[62] Parra, F. I., Ahedo, E., Fife, J. M., and Martínez-Sánchez, M., "A Two-Dimensional Hybrid Model of the Hall Thruster Discharge," *Journal of Applied Physics*, Vol. 100, 2006, pp. 023304.

[63] Morozov, A. and Savelyev, V., "Fundamentals of Stationary Plasma Thruster Theory," *Reviews of Plasma Physics*, Vol. 21, No. 5, 2000, pp. 203.

[64] Janes, G. S. and Lowder, R. S., "Anomalous Electron Diffusion and Ion Acceleration in a Low-Density Plasma," *The Physics of Fluids*, Vol. 9, No. 6, June 1966, pp. 1115–1123.

[65] Bishaev, A. M. and Kim, V., "Local plasma properties in a Hall-current accelerator with an extended acceleration zone," *Zhurnal Tekhnicheskoi Fiziki*, Vol. 48, September 1978, pp. 1853–1857.

[66] Meezan, N. B., Hargus, Jr., W. A., and Capelli, M. A., "Anomalous Electron Mobility in a Coaxial Hall Discharge Plasma," *Physical Review E*, Vol. 63, 2001, pp. 026410.

[67] Chapman, S. and Cowling, T. G., *The Mathematical Theory of Non-Uniform Gases, 3^{rd} Edition*, Cambridge University Press, 1970.

[68] Grad, H., "On the Kinetic Theory of Rarefied Gases," *Communications on Pure and Applied Mathematics*, Vol. 2, No. 4, 1949, pp. 331–407.

[69] Struchtrup, H., *Macroscopic Transport Equations for Rarefied Gas Flows*, Springer, 2005.

[70] Shkarofsky, I. P., Johnston, T. W., and Bachynski, M. P., *The Particle Kinetics of Plasmas*, Addison-Wesley Publishing Company, 1966.

[71] Mitchner, M. and Kruger, C. H., *Partially Ionized Gases*, John Wiley and Sons, Inc., 1973.

[72] Chew, G. F., Goldberger, M. L., and Low, F. E., "The Boltzmann Equation and the One-Fluid Hydromagnetic Equations in the Absence of Particle Collisions," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, Vol. 236, No. 1204, 1956, pp. pp. 112–118.

[73] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press Inc., 1970.

[74] Levermore, C. D., Morokoff, W. J., and Nadiga, B. T., "Moment Realizability and the Validity of the Navier-Stokes Equations for Rarefied Gas Dynamics," *Physics of Fluids*, Vol. 10, December 1998, pp. 3214–3226.

[75] Mikellides, I. G., Katz, I., Hofer, R. R., and Goebel, D. M., "Hall-Effect Thruster Simulations with 2-D Electron Transport and Hydrodynamic Ions," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-114.

[76] Katz, I. and Mikellides, I. G., "Neutral Gas Free Molecular Flow Algorithm Including Ionization and Walls for Use in Plasma Simulations," *Journal of Computational Physics*, Vol. 230, 2011, pp. 1454–1464.

[77] Katz, I. and Mikellides, I. G., "A New Algorithm for the Neutral Gas in the Free-Molecule Regimes of Hall and Ion Thrusters," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-095.

[78] Mikellides, I. G., Katz, I., Hofer, R. R., and Goebel, D. M., "Magnetic Shielding of the Acceleration Channel Walls in a Long-Life Hall Thruster," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-6942.

[79] Nobelprize.org, "The Nobel Prize in Physics 1925," http://www.nobelprize.org/nobel_prizes/physics/laureates/1925/, August 9, 2012.

[80] Nicoletopoulos, P. and Robson, R. E., "Periodic Electron Structures in Gases: A Fluid Model of the "Window" Phenomenon," *Physical Review Letters*, Vol. 100, March 2008, pp. 124502.

[81] Hargus, Jr., W. A. and Charles, C. S., "Near Exit Plane Velocity Field of a 200-Watt Hall Thruster," *Journal of Propulsion and Power*, Vol. 24, No. 1, 2008, pp. 127–133.

[82] Hargus, Jr. , W. A. and Nakles, M. R., "Evolution of the Ion Velocity Distribution in the Near Field of the BHT-200-X3 Hall Thruster," 42^{nd} *Joint Propulsion Conference & Exhibit*, Sacramento, CA, July 2006, AIAA-2006-4991.

[83] Huang, W., Drenkow, B., and Gallimore, A. D., "Laser-Induced Fluorescence of Singly-Charged Xenon Inside a 6-kW Hall Thruster," 45^{th} *Joint Propulsion Conference & Exhibit*, Denver, Co, August 2009, AIAA-2009-5355.

[84] Smith, T. B., Herman, D. A., Gallimore, A. D., and Drake, R. P., "Deconvolution of Axial Velocity Distributions from Hall Thruster LIF Spectra," 27^{th} *International Electric Propulsion Conference*, Pasadena, CA, October 2001, IEPC-2001-019.

[85] Surrendra, M. and Dalvie, M., "Moment Analysis of RF Parallel-Plate-Discharge Simulations Using the Particle-in-cell with Monte Carlo Collisions Technique," *Physical Review E*, Vol. 48, No. 5, 1993, pp. 3914–3924.

[86] Robson, R. E., Nicoletopoulos, P., Li, B., and White, R. D., "Kinetic Theoretical and Fluid Modelling of Plasmas and Swarms: the Big Picture," *Plasma Sources Science and Technology*, Vol. 17, 2008, pp. 024020, doi:10.1088/0963-0252/17/2/024020.

[87] Birdsall, C. K. and Langdon, A. B., *Plasma Physics via Computer Simulation*, McGraw-Hill Book Company, 1985.

[88] Hockney, R. W. and Eastwood, J. W., *Computer Simulation Using Particles*, Taylor & Francis Group, 1988.

[89] Serikov, V. V., Kawamoto, S., and Nanbu, K., "Particle-in-Cell Plus Direct Simulation Monte Carlo (PIC-DSMC) Approach for Self-Consistent Plasma-Gas Simulation," *Plasma Science, IEEE Transactions on*, Vol. 27, No. 5, October 1999, pp. 1389–1398.

[90] Bird, G. A., "Approach to Translational Equilibrium in a Rigid Sphere Gas," *Physics of Fluids*, Vol. 6, 1963, pp. 1518–1519.

[91] Nanbu, K., "Probability Theory of Electron-Molecule, Ion-Molecule, Molecule-Molecule, and Coulomb Collisions for Particle Modleing of Materials Processing Plasmas and Gases," *Plasma Science, IEEE Transactions on*, Vol. 28, No. 3, June 2000, pp. 971–990.

[92] Skullerud, H. R., "The Stochastic Computer Simulation of Ion Motion in a Gas Subjected to a Constant Electric Field," *British Journal of Applied Physics (Journal of Physics D), Series 2*, Vol. 1, 1968, pp. 1567–1568.

[93] Babovsky, H. and Illner, R., "A Convergence Proof for Nanbu's Simulation Method for the Full Boltzmann Equation," *SIAM Journal on Numerical Analysis*, Vol. 26, No. 1, February 1989, pp. 45–65.

[94] Birdsall, C. K., "Particle-in-Cell Charged-Particle Simulations, Plus Monte Carlo Collisions with Neutral Atoms, PIC-MCC," *Plasma Science, IEEE Transactions on*, Vol. 19, No. 2, April 1991, pp. 65–85.

[95] Takizuka, T. and Abe, H., "A Binary Collision Model for Plasma Simulation With a Particle Code," *J. Comput. Phys.*, Vol. 25, 1977, pp. 205–219.

[96] Verboncoeur, J. P., "Particle Simulation of Plasmas: Review and Advances," *Plasma Phys. Control. Fusion*, Vol. 47, 2005, pp. A231–A260.

[97] Szabo, J. J., "Fully Kinetic Numerical Modeling of a Plasma Thruster," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, January 2001.

[98] Manheimer, W. M., Lampe, M., and Joyce, G., "Langevin Representation of Coulomb Collisions in PIC Simulations," *J. Comput. Phys.*, Vol. 138, 1997, pp. 563–584.

[99] Montgomery, D. C. and Tidman, D. A., *Plasma Kinetic Theory*, McGraw-Hill Book Company, 1964.

[100] Hirakawa, M. and Arakawa, Y., "Particle Simulation of Plasma Phenomena in Hall Thrusters," 24^{th} *International Electric Propulsion Conference*, Moscow, RU, September 1995, IEPC-1995-164.

[101] Oh, D. Y., "Computational Modeling of Expanding Plasma Plumes in Space Using a PIC-DSMC Algorithm," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, February 1997.

[102] Boyd, I. D., "Review of Hall Thruster Plume Modelling," *Journal of Spacecraft and Rockets*, Vol. 38, No. 3, 2001, pp. 381–387.

[103] Adam, J. C., Héron, A., and Laval, G., "Study of Stationary Plasma Thruster Using Two-Dimensional Fully Kinetic Simulations," *Physics of Plasmas*, Vol. 11, No. 1, January 2004, pp. 295–305.

[104] Fernandez, E., Knoll, A., and Cappelli, M. A., "An Axial-Azimuthal Hybrid Simulation of Coaxial Hall Thrusters," 39^{th} *Joint Propulsion Conference & Exhibit*, Sacramento, CA, July 2006, AIAA-2006-4329.

[105] Taccogna, F., Longo, S., Capitelli, M., and Schneider, R., "Plasma flow in a Hall Thruster," *Physics of Plasmas*, Vol. 12, 2005, pp. 043502.

[106] Watson, V., "Computer Simulation of a Plasma Accelerator," Ph.D. Thesis, Stanford University, Standford, CA, 1969.

[107] Hirakawa, M. and Arakawa, Y., "Plasma Particle Simulation in Cusped Ion Thrusters," 23^{rd} *International Electric Propulsion Conference*, Seattle, WA, September 1993, IEPC-1993-2198.

[108] Taccogna, F., Longo, S., Capitelli, M., and Schneider, R., "Self-similarity in Hall plasma discharges: Applications to particle models," *Physics of Plasmas*, March 2005.

[109] Liu, H., Wu, B., Yu, D., Cao, Y., and Duan, P., "Particle-in-Cell Simulation of a Hall Thruster," *Journal of Physics D: Applied Physics*, Vol. 43, 2010, pp. 165202.

[110] Blateau, V., "PIC Simulation of a Ceramic-Lined Hall-Effect Thruster," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, September 2002.

[111] Fox, J. M., "Parallelization of a Particle-in-Cell Simulation Modeling Hall-Effect Thrusters," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, February 2005.

[112] Fox, J. M., "Advances in Fully-Kinetic PIC Simulations of a Near-Vacuum Hall Thruster and Other Plasma Systems," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, September 2007.

[113] Fox, J. M., "Fully-Kinetic PIC Simulations for Hall-Effect Thrusters," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, September 2007.

[114] Gildea, S., "Fully Kinetic Modeling of a Divergent Cusped Field Thruster," S.M. Thesis, Massachusetts Institute of Technology, Cambridge MA, May 2009.

[115] Hobbs, G. D. and Wesson, J. A., "Heat Flow Through a Langmuir Sheath in the Presence of Electron Emission," *Plasma Physics*, Vol. 9, 1967, pp. 85–87.

[116] Sullivan, K. U., "PIC Simulation of SPT Hall Thrusters: High Power Operation and Wall Effects," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2004.

[117] Matyash, K., Schneider, R., Kalentev, O., Raitses, Y., and Fisch, N. J., "Numerical Investigations of a Cylindrical Hall Thruster," 32^{nd} *International Electric Propulsion Conference*, Wiesbaden, DE, September 2011, IEPC-2011-070.

[118] Irishkov, S. V., Gorshkov, O. A., and Shagayda, A. A., "Fully Kinetic Modeling of Low-Power Hall Thrusters," 29^{th} *International Electric Propulsion Conference*, Princeton, NJ, October 2005, IEPC-2005-035.

[119] Azziz, Y., Warner, N. Z., Martínez-Sánchez, M., and Szabo, J. J., "High Voltage Plume Measurements and Internal Probing of the BHT-1000 Thruster," 40^{th} *Joint Propulsion Conference & Exhibit*, For Lauderdale, FL, July 2004, AIAA-2004-4097.

[120] Szabo, J., Gasdaska, C., Hruby, V., and Robin, M., "Bismuth Hall Effect Thruster Development," 53^{rd} *JANNAF Propulsion Meeting*, Monterey, CA, December 2005.

[121] Szabo, J. J., September 2012, Personal communication to the author.

[122] Szabo, J., Martínez-Sánchez, M., and Batishchev, O., "Fully Kinetic Hall Thruster Modeling," 27^{th} *International Electric Propulsion Conference*, Pasadena, CA, October 2001, IEPC-2001-341.

[123] Warner, N. Z., Szabo, J. J., and Martínez-Sánchez, M., "Characterization of a High Specific Impulse Hall Thruster Using Electrostatic Probes," 28^{th} *International Electric Propulsion Conference*, Toulouse, France, March 2003, IEPC-2003-82.

[124] Szabo, J., Warner, N. Z., and Martínez-Sánchez, M., "Instrumentation and Modeling of a High Isp Hall Thruster," 38^{th} *Joint Propulsion Conference & Exhibit*, Indianapolis, IN, July 2002, AIAA-2004-4248.

[125] Szabo, J. J. and Azziz, Y., "Characterization of a High Specific Impulse Xenon Hall Effect Thruster," 29^{th} *International Electric Propulsion Conference*, Princeton, NJ, November 2005, IEPC-2005-324.

[126] Azziz, Y., "Experimental and Theoretical Characterization of a Hall Thruster Plume," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA,, 2007.

[127] Warner, N. Z., "Theoretical and Experimental Investigation of Hall Thruster Miniaturization," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2007.

[128] Gildea, S. R., Matlock, T. S., Martínez-Sánchez, M., and Hargus, Jr., W. A., "Erosion Measurements in a Diverging Cusped-Field Thruster," 32^{nd} *International Electric Propulsion Conference*, Wiesbaden, DE, September 2011, IEPC-2011-149.

[129] Trubnikov, B. A., "Particle Interactions in a Fully Ionized Plasma," *Reviews of Plasma Physics*, Vol. 1, 1965, pp. 105–204.

[130] Nakles, M., Barry, R., Larson, C. W., and Hargus, Jr., W. A., "A Plume Comparison of Xenon and Krypton Propellant on a 600 W Hall Thruster," 31^{st} *International Electric Propulsion Conference*, Ann Arbor, MI, September 2009, IEPC-2009-115.

[131] Courtney, D. G. and Martínez-Sánchez, M., "Diverging Cusped-Field Hall Thruster (DCHT)," 30^{th} *International Electric Propulsion Conference*, Florence, Italy, September 2007, IEPC-2007-39.

[132] Courtney, D., Lozano, P., and Martínez-Sánchez, M., "Continued Investigation of Diverging Cusped Field Thruster," 44^{th} *Joint Propulsion Conference & Exhibit*, Hartford, CT, July 2008, AIAA-2008-4631.

[133] Barral, S. and Ahedo, E., "Low-frequency model of breathing oscillations in Hall discharges," *Physical Review E*, Vol. 79, 2009.

[134] Hargus, Jr., W. A., Nakles, M. R., Tedrake, R., and Pote, B., "Effect of Anode Current Fluctuations on Ion Energy Distributions within a 600 W Hall Effect Thruster," 44^{th} *Joint Propulsion Conference & Exhibit*, Hartford, CT, July 2008, AIAA-2008-4724.

[135] Brown, D. L., "Investigation of Low Discharge Voltage Hall Thruster Characteristics and Evaluation of Loss Mechanisms," Ph.D. Thesis, University of Michigan, Ann Arbor, MI, 2009.

[136] Brown, S. C., *Basic Data of Plasma Physics*, American Institute of Physics Press, 1994.

[137] Martínez-Sánchez, M. and Ahedo, E., "Magnetic Mirror Effects on a Collisionless Plasma in a Convergent Geometry," *Physics of Plasmas*, Vol. 18, 2011, pp. 033509.

[138] Hut, P., Makino, J., and McMillan, S., "Building a Better Leapfrog," *The Astrophysical Journal*, Vol. 433, 1995, pp. L93–L96.

[139] Haus, H. A. and Melcher, J. R., *Electromagnetic Fields and Energy*, Prentice-Hall, 1989.

[140] Skeel, R. D., "Variable Step Size Destabilizes the Stömer/Leapfrog/Verlet Method," *BIT Numerical Mathematics*, Vol. 33, 1993, pp. 172–175.

[141] Richarson, A. S. and Finn, J. M., "Symplectic Integrators with Adaptive Time Steps," *Plasma Phys. Control. Fusion*, Vol. 54, 2012, pp. 014004.

[142] Yoshida, H., "Symplectic Integrators for Hamiltonian Systems: Basic Theory," *Chaos, Resonance, and Collective Dynamical Phenomena in the Solar System*, edited by S. Ferraz-Mello, Vol. 152 of *IAU Symposium*, 1992, p. 407.

[143] Huang, W. and Leimkuhler, B., "The Adaptive Verlet Method," *SIAM J. Sci. Comput.*, Vol. 18, No. 1, 1997, pp. 239–256.

[144] Fuchs, V. and Gunn, J. P., "On the Integration of Equations of Motion for Particle-in-Cell Codes," *Journal of Computational Physics*, Vol. 214, 2006, pp. 299–315.

[145] Seldner, D. and Westermann, T., "Algorithms for Interpolation and Localization in Irregular 2D Meshes," *Journal of Computational Physics*, Vol. 79, No. 1, 1988, pp. 1–11.

[146] Yalin, A. P., Surla, V., Farnell, C., Butweiller, M., and Williams, J. D., "Sputtering Studies of Multi-Component Materials by Weight Loss and Cavity Ring-Down Spectroscopy," 42^{nd} *Joint Propulsion Conference & Exhibit*, Sacramento, CA, July 2006, AIAA–2006–4338.

[147] Cheng, S. Y., "Modeling of Hall Thruster Lifetime and Erosion Mechanisms," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 2007.

[148] Yamamura, Y., "An Emperical Formula for Angular Dependence of Sputtering Yields," *Radiation Effects*, Vol. 80, 1984, pp. 57–72.

[149] StackExchange, "Problem with NonlinearModelFit," `http://mathematica.stackexchange.com/questions/2309/problem-with-nonlinearmodelfit`, September 14, 2012.

[150] Wolfram Research, Inc., *Mathematica, Version 8.0.4.0*, Champaign, IL, 2011.

[151] Saint-Gobain, "Grade HP Combat® Boron Nitride Solids," `http://www.bn.saint-gobain.com/Combat-Solid-GradeHP.aspx`, September 14, 2012.

[152] Dietrich, S. and Boyd, I. D., "Scalar and Parallel Optimized Implementation of the Direct Simulation Monte Carlo Method," *Journal of Computational Physics*, Vol. 126, No. 0141, 1996, pp. 328–342.

[153] Jackson, J. D., *Classical Electrodynamics, 3rd ed.*, John Wiley & Sons, Inc., 1999.

[154] Whipple, E. C., "Potentials of Surfaces in Space," *Reports on Progress in Physics*, Vol. 44, 1981, pp. 1197.

[155] Hutchinson, I. H., "Ion Collection by a Sphere in a Flowing Plasma: 3. Floating Potential and Drag Force," *Plasma Physics and Controlled Fusion*, Vol. 47, No. 1, 2005, pp. 71–87.

[156] International Business Machines Corperation (IBM), "Watson Sparse Matrix Package (WSMP)," `http://researcher.watson.ibm.com/researcher/view_project.php?id=1426`, October 14, 2012.

[157] Blackford, L. S., Choi, J., Cleary, A., D'Azevedo, E., Demmel, J., Dhillon, I., Dongarra, J., Hammarling, S., Henry, G., Petitet, A., Stanley, K., Walker, D., and Whaley, R. C., *ScaLAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.

[158] Choi, J., Demmel, J., Dhillon, I., Dongarra, J., Ostrouchov, S., Petitet, A., Stanley, K., Walker, D., and Whaley, R., "ScaLAPACK: a Portable Linear Algebra Library for Distributed Memory Computers Design Issues and Performance," *Computer Physics Communications*, Vol. 97, No. 12, 1996, pp. 1 – 15, ¡ce:title¿High-Performance Computing in Science¡/ce:title¿.

[159] Daspit, R. M., "Performance Characterization and Optimization of a Diverging Cusped Field Thruster witha Calibrated Counter-Weighted Millinewton Thrust Stand," S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 2012.

[160] Barral, S., Makowski, K., Peradzyński, Z., and Dudeck, M., "Transit-time Instability in Hall Thrusters," *Physics of Plasmas*, Vol. 12, 2005, pp. 073504.

[161] Ning, Z., Zhang, S., Liu, H., and Yu, D., "Study on Ionization Characteristics of Low Power Hall Thruster with Variable Cross-Section Channel," 32^{nd} *International Electric Propulsion Conference*, Wiesbaden, DE, September 2011, IEPC-2011-286.

[162] Isakowitz, S. J., Hopkins, J. B., and Hopkins, Jr., J. P., *International Reference Guide to Space Launch Systems*, American Institution of Aeronautics & Astronautics, 4th ed., 2004.

[163] Jahn, R. G., *Physics of Electric Propulsion*, Dover Publications, Inc., 2006.

[164] de Grys, K., Mathers, A., Welander, B., and Khayms, V., "Demonstration of 10,400 Hours of Operation on a 4.5 kW Qualification Model Hall Thruster," 46^{th} *Joint Propulsion Conference & Exhibit*, Nashville, TN, July 2010, AIAA-2010-6698.

[165] Peterson, P. Y., Jacobson, D. T., Manzella, D. H., and John, J. W., "The Performance and Wear Characterization of a High-Power High-Isp NASA Hall Thruster," 41^{st} *Joint Propulsion Conference & Exhibit*, Tucson, AZ, July 2005, AIAA-2005-4243.

[166] Linnell, J. A., "An Evaluation of Krypton Propellant in Hall Thrusters," Ph.D. Thesis, University of Michigan, Ann Arbor, MI, 2007.

[167] Hastings, D. and Garrett, H., *Spacecraft-Environment Interactions*, Cambridge University Press, 1996.

[168] Mason, E. A. and McDaniel, E. W., *Transport Properties of Ions in Gases*, John Wiley and Sons, Inc., 1988.

[169] Robson, R. E., White, R. D., and Petrović, Z. L., "*Colloquium*: Physically Based Fluid Modeling of Collisionally Dominated Low-Temperature Plasmas," *Reviews of Modern Physics*, Vol. 77, October 2005, pp. 1303–1320.

[170] Bhatnagar, P. L., Gross, E. P., and Krook, M., "A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems," *Phys. Rev.*, Vol. 94, No. 3, May 1954, pp. 511–525.

[171] Sakabe, S. and Izawa, Y., "Cross Sections for Resonant Charge Transfer between Atoms and their Positive Ions: Collision Velocity $\leq$ 1 a.u." *Atomic Data and Nuclear Data Tables*, Vol. 49, 1991, pp. 257–314.

[172] Rapp, D. and Francis, W. E., "Charge Exchange Between Gaseous Ions and Atoms," *Journal of Chemical Physics*, Vol. 37, No. 11, 1962, pp. 2631–2645.

[173] Boyd, I. D. and Dressler, R. A., "Far Field Modeling of the Plasma Plume of a Hall Thruster," *Journal of Applied Physics*, Vol. 92, No. 4, August 2006, pp. 1764–1774.

[174] Dalgarno, A., McDowell, M., and Williams, A., "The Mobilities of Ions in Unlike Gases," *Philosophical Transactions of the Royal Society of London A*, Vol. 250, 1958, pp. 411–425.

[175] "CRC Handbook of Chemistry and Physics, 91^{st} Edition," CRC Press, 2011, pp. 10-188.

[176] Freidberg, J., *Plasma Physics & Fusion Energy*, Cambridge University Press, 2007.

[177] Bittencourt, J. A., *Fundamentals of Plasma Physics, 3^{rd} Edition*, Springer Science & Bussiness Media, LLC, 2004.

[178] Koura, K. and Matsumoto, H., "Variable Soft Sphere Molecular Model for Air Species," *Physics of Fluids A*, Vol. 4, No. 5, May 1992, pp. 1083–1085.