Modeling Anomalous Electron Transport in a Fluid Hall Thruster Code
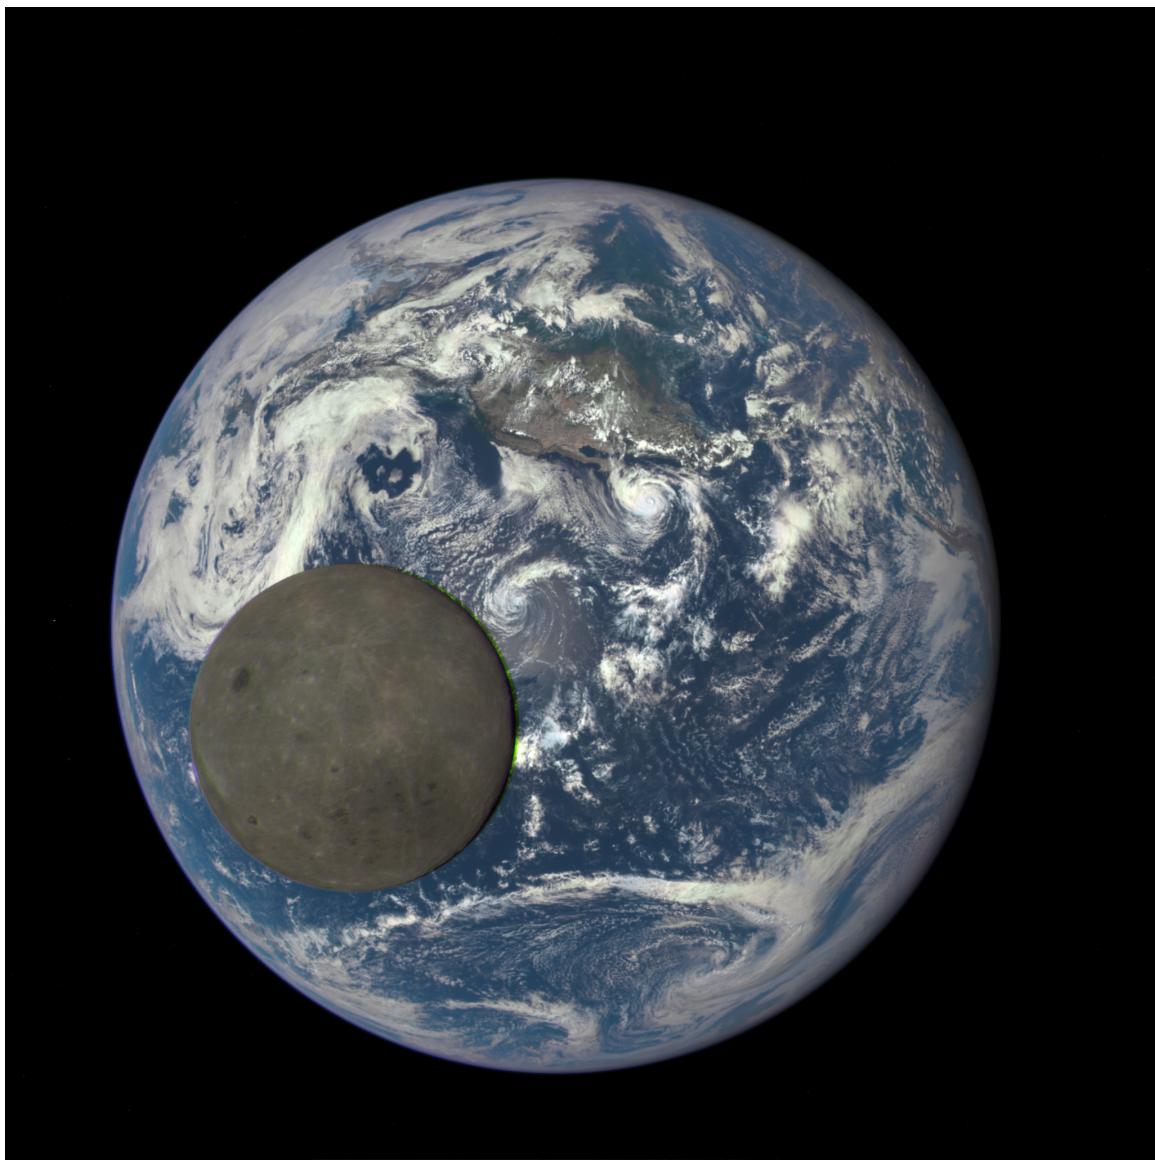
by

Thomas Archer Marks

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering and Scientific Computing)
in the University of Michigan
2023

Doctoral Committee:

        Associate Professor Benjamin A. Jorns, Chair
        Professor Alec D. Gallimore
        Assistant Professor Alex A. Gorodetsky
        Professor Mark J. Kushner
        Dr. Alejandro Lopez Ortega, Jet Propulsion Laboratory

The Moon crossing in front of the earth, as seen from the Deep Space Climate Observatory (DISCOVR) satellite. Photo credit: **NASA/NOAA**

From the rocket we can see the huge sphere of the planet in one or another phase of the Moon. We can see how the sphere rotates, and how within a few hours it shows all its sides successively ... and we shall observe various points on the surface of the Earth for several minutes and from different sides very closely. This picture is so majestic, attractive and infinitely varied that I wish with all my soul that you and I could see it.

Konstantin Tsiolkovsky (1911)

Thomas Archer Marks

marksta@umich.edu

ORCID iD: 0000-0003-3614-6127

ACKNOWLEDGEMENTS

It is often said that it takes a village to raise a child, but the same is true for Ph.D students. I could not have completed this dissertation without the support of so many wonderful people. I would first like to thank my advisor, Professor Benjamin Jorns. While I didn't start out as a member of your lab group, I am grateful to you for taking me in and accepting me into the fold as the only computationalist among experimentalists. Under your tutelage, I have grown substantially as a researcher and a science communicator. Thanks for your patience and understanding as I very slowly learned about plasma waves, and I'm sorry that when you ducked your head into my office and asked "any breakthroughs?" the answer was usually no. Still, I'm happy with how this project turned out, and I look forward to working with you in the years to come.

Next, I would like to thank my committee–Professor Mark Kushner, Professor Alex Gorodetsky, Provost Alec Gallimore, and Dr. Alejandro Lopez Ortega. Professor Kushner, I know we haven't met much during my Ph.D., but I appreciate you agreeing to serve on my committee, as well as your constant efforts to keep the Michigan Institute of Plasma Science and Engineering (MIPSE) going. I have learned so much attending the various seminars and symposiums during my time here, and I only hope that one day I can return, triumphant, to earn the coveted MIPSE mug. Professor Gorodetsky, I have enjoyed discussing the possibilities for accelerating Hall thruster modeling with you. You have forgotten more about Bayesian statistics than I will ever learn, and I look forward to working with you as a postdoctoral researcher. Dr. Gallimore, I am glad I got to meet you and hear you share some of your stories about the early days of Hall thruster research in the US. Thanks for checking in despite your busy schedule as Dean of the College of Engineering (and now Provost at Duke!). Alejandro, thanks for serving as my mentor when I interned at JPL during the pandemic. The circumstances were less than ideal, but I learned a ton and would not be where I am today without your assistance. I would also like to thank you (and Yiangos) for graciously allowing us to use Hall2De here at Michigan.

My labmates have been a huge source of strength and support to me during my time at Michigan. I would first like to thank everyone in the old NGPDL gang. Mike, Sam, Candice, Kaelan, Mackenzie, Alex, Astrid, and Pawel, thanks for looking after me in my first year

here at Michigan. My memories of after-lunch coffee time in the office are some of the fondest I have from the first year of graduate school, despite the fact that I don't even like coffee. After switching lab groups, I got to know another great group of people at PEPL. Given some of the horror stories you hear about other labs, I genuinely think we have one of the best labs out there.

I'll start with those labmates who have since moved on to better places. Shad, thanks for being a cool guy and coming out to speak German with me. Josh, I loved hearing about your dreams, even if no one else did. Ben (Wachs), I thought you hated me for like two years, but since then I have enjoyed hanging out with you, even if the nights tended to get out of control. Zach, I will always remember your poise and confidence, as well as your dancing ability. There are so many more talented and fun people who have come and gone through this lab—Sarah, Marcel, Angela, Pete, Sophia, Sal, Konrad, Paul and Luca—and I am glad to have known each of you.

Next, to those labmates who are still here (though hopefully not for too much longer in some of these cases). Chris—my eternal conference roommate—thanks for being a solid member of the gasdynamics quals cohort, for introducing me to the Angostura sour, and for jamming with me and letting me play your bass guitar. Tate, thanks for letting Julie and I plant stuff in your garden. I really admire your artistic and design abilities, as well as your competence with all things machining. Collin B. Whittaker, thanks for having us over to play board games so many times, as well as for providing some of the best jams and preserves I've ever had. Hopefully we will one day play Twilight Imperium again. Will, you crack me up. Your relentless positively and supportiveness are a light in the darkness, and you have really grown as a researcher since you started here. Parker (Roberts), I will miss discussing physics with someone as incredible intelligent and yet as friendly and approachable as you. Madison, I have enjoyed seeing you grow as a scientist, and hope you continue to get use out of your Panera Unlimited Sip Club membership. Declan, I hope I've been a little bit useful as I scramble to finish up my dissertation, and I know you'll one day be as confident in writting buggy code as I am. Grace, I'm sad I won't get to see how the RMF develops under your leadership. The bet is still out, so let me know if you hit 3% efficiency. Ari, I know you haven't technically started as a Ph.D. student yet, but you already feel like an integral member of the lab, and I think you're gonna crush it.

Out of all of these people, a special shoutout goes to (Dr.) Leanne Su, who I met on the bus back to the airport on visit day and have remained friends with since. I am glad that we have been able to lean[a] on eachother for support throughout the wild ride that is grad school. We have had so many fun experiences together. We've been labmates for four

---

[a]haha

years, roommates for two years, and friends for many more, I hope we will stay in contact as you move on to the next wild experience in life. An equally special shoutout goes to Leanne's Euchre/life partner, Austin. Thanks for being an amazing friend, even after we were no longer in the same lab. You are so good at making people feel welcome in any situation, and I know I can rely on you for companionship when I'm feeling down. Thanks too for having me to thanksgiving in Iowa (twice!) and making me feel at home with your family.

Next, I want to acknowledge the many good friends I made while living here. To all of the former roommates at Beakes Street—Julia, Jenn, Aagnik, Amanda, Adarsh—thanks for being the best roommates I could ever ask for. If you all weren't so cool, I think we would have killed each other during COVID. Instead, we managed to have fun, even with eight people in the house at one point. To Jana and Parker (Ewen), I have loved getting to know you both. Christina and Josh, I wouldn't have made it through quals without your attention to detail and diligence, and I hope we can visit you in Davis soon. Tom, thanks for being in my wedding party. I'll miss your dry humor and understated wit. Lastly, Kate, whom I met in a kind of unconventional way, thanks for letting us have our wedding reception in your backyard.

Literally and figuratively, I would not be here if it weren't for my parents, Craig and Susan Marks. Thanks for always supporting my love of science, be it dinosaurs, space, or whatever other thing I was interested in at the moment. Dad, thanks for instilling in me a voracious desire to learn new things. Mom, your superpower of identifying obscure actors in literally anything we're watching has always astounded me, and I hope some of it has rubbed off on me.

Lastly, I want to thank my wife, Dr. Julie Rieland. I love you dearly and am so happy to have married you. I have come to rely on you for so much over the last few years. Here's to our next steps together as we move to the DC area and adopt the cats we've been wanting for years.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Hall thrusters are the most widely-used type of in-space electric propulsion device. However, aspects of the physics of their operation remain poorly-understood. Notably, the problem of enhanced cross-field electron transport prevents predictive modeling and simulation of these devices. This "anomalous" electron transport, which is often represented as an effective collisional scattering process likely stems from kinetic instabilities and plasma turbulence. This poses a challenge for incorporating this phenomenon into fluid models suitable for engineering applications.

In this dissertation, several models for the anomalous electron transport in Hall thrusters are reviewed and evaluated, and methods for model validation are assessed. Electron transport models from the scientific literature are reviewed, and several new models are derived. It is first shown that the common practice of evaluating models on calibrated simulation outputs and comparing them to hand-tuned, ad-hoc anomalous collision frequencies does not yield predictive results.

Next, the behavior of four electron transport models from the literature is investigated using a fluid Hall thruster code. The models are calibrated against a baseline experimental condition of a 9-kW-class magnetically-shielded Hall thruster operating at 300 V and 15 A on xenon propellant. The extensibility of the models is then assessed by using this calibrated model to simulate three additional operating conditions—300 V and 30 A, 600 V and 15 A, and 300 V and 15 A operating on krypton propellant. The quality of the model prediction is quantified by comparing the model outputs to experimental measurements of discharge current, thrust, and ion velocity. It is found that while none of the models can predict the ion acceleration characteristics accurately, some compare favorably in terms of the scaling of thrust and discharge current across operating conditions. The limitations of the models are attributed to the coupling between the functional scaling of the closure models with respect to the local plasma properties and the fluid model. The role of the electron energy balance in this coupling is also highlighted.

Finally, several novel models are assessed. These models are derived by altering the underlying assumptions of a previously-discussed first-principles transport model. It is found that some of these models are able to predict thruster performance characteristics—

such as thrust and efficiency—better than the baseline model, while inaccurately reproducing the experimentally-observed spatial variation in ion velocity. Another model is then evaluated which improves upon these results, yielding improved accuracy with respect to these velocity measurements in all but one case. The results of this dissertation are finally discussed in the context of motivating improved closure models of the anomalous electron transport in Hall thrusters.

CHAPTER 1

Introduction

Rocket science has been mythologized all out of proportion to its true difficulty.

John Carmack

1.1 Basic rocket science

In 1904, Konstantin Tsiolkovsky derived the famous *rocket equation*, which relates how much we can change the velocity of a spacecraft to the amount of propellant onboard and the velocity at which it can be ejected. He was not the first nor the last to derive it independently, but it bears his name because he was one of the first people to consider how humankind might one day travel to the stars. In addition to this equation, Tsiolkovsky developed the concept of airlocks, multi-stage rockets, and space elevators. He derived the rocket equation while working as a schoolteacher in Kaluga, a rural town in the Russian Empire, far from any universities or centers of learning. The equation is

$$\Delta v = v_{ex} \ln \frac{m_i}{m_f}.$$ (1.1)

In this expression, Δv is the amount by which we would like to change our velocity, m_i is the initial mass of our spacecraft (i.e. when it is full of propellant), m_f is the final mass of our spacecraft (after it is finished exhausting propellant) and v_{ex} is how fast we eject propellant from our spacecraft. This expression shows that the amount that a craft can change its velocity depends linearly on the exhaust velocity but only logarithmically on the ratio of initial to final mass (i.e. the amount of propellant we carry). This suggests that if we want our rocket to go very fast, and by extension very far, it is more effective to increase our exhaust velocity than to add additional propellant. We can re-arrange this expression to make this conclusion more explicit:

$$\frac{m_i}{m_f} = e^{\Delta v / v_{ex}} \tag{1.2}$$

Here, we can see that the ratio of initial to final mass of our spacecraft depends exponentially on the ratio between the change in velocity we wish to enact and the exhaust velocity. At a fixed exhaust velocity, increasing Δv linearly requires exponentially more propellant, as every additional bit of propellant added also needs to be accelerated, requiring still more propellant, in a cascade known as the *tyranny of the rocket equation*. Thus, if we want to go fast, we want to impart as much momentum as possible to each gram of propellant we exhaust. In spacecraft propulsion, we typically quantify how effectively we achieve this goal via the *specific impulse* (I_{sp}), measured in seconds, which is the effective exhaust velocity divided by the acceleration due to gravity at sea level ($g_0 = 9.81\,\mathrm{m/s}$):

$$I_{sp} = \frac{v_{ex}}{g_0} \tag{1.3}$$

If we want to to reduce the amount of propellant we need to achieve a given Δv, or increase the amount of Δv we can get for a fixed propellant mass, we want an engine with higher specific impulse. For existing propulsion systems, specific impulse ranges from tens of seconds on the low end to several thousand seconds on the high end.

The most basic type of rocket is a *cold gas thruster*, in which we have a pressurized tank of gas which is allowed to exhaust through a valve to produce a small amount of thrust. Since the gas is cold, the particles do not have much kinetic energy and move quite slowly, so these systems can only attain specific impulses around 10 - 100 seconds, depending on the gas used and the pressure in the tank. To further speed up the propellant, we can use a nozzle. Nozzles work by transforming thermal energy (i.e. pressure and heat) into kinetic energy are thus able to increase the speed of the propellant. However, there is not much energy available in a cold gas, so the ability of a nozzle to increase the specific impulse is limited. Despite their poor performance, cold gas thrusters are often employed to propel small satellites due to their simplicity.

1.1.1 Chemical propulsion

If we want our gas to go faster, we can heat the propellant. This gives the particles more thermal energy, which a nozzle can turn into kinetic energy. Traditionally, the easiest way to do this has been to use a chemically-reactive propellant or mix of propellants. If the propellants react exothermically, the breaking and reforming of bonds during the reaction will release heat. In a chemical rocket, we inject our reactive propellants into a combustion chamber, which contains the reaction products as they heat and increase in pressure. This gas then flows into a nozzle, which transforms this heat and pressure into directed kinetic energy. In this way, these rockets can achieve specific impulses between 100 and 500 seconds,[4] depending on the propellant chosen. Lighter molecules move faster at a fixed kinetic energy, so propellants with lower molecular weights can attain higher specific impulses. The F-1 engines which powered the first stage of the Saturn V moon rockets employed a mixture of RP-1 (a type of specially-formulated kerosene) and liquid oxygen and achieved specific impulses of around 250 to 300 seconds.[4] In contrast, the space shuttle main engines used liquid hydrogen, which is very light, and reached specific impulses in excess of 400 seconds. More exotic propellant-oxidizer combinations involving

fluorine have been able to reach 480 seconds,[4] However, such combinations are toxic, corrosive, and highly explosive. They are thus generally avoided.

No matter the propellant, the specific impulse of a chemical rocket is ultimately limited by the amount of energy stored in the chemical bonds of the propellant/oxidizer mixture. By **Equation 1.2**, this limitation on specific impulse means that chemical rockets need a lot of propellant compared to the amount of payload they can carry. The Apollo capsule, service module, and lunar lander only comprised 4% of the total launch mass of the Saturn V rocket. The rest of the mass was propellant, propellant tanks, and engines. Despite this limitation, the power and thrust-to-weight ratio of these rockets means they are the only option for launching payloads from the surface of the earth into space. Once in space, however, the utility of chemical rockets is more limited. If we want to launch a probe to another planet using a chemical rocket, most of the available payload will be propellant for the probe, rather than useful scientific instrumentation. To reduce fuel mass, we usually launch scientific probes on one-way trajectories to other planets (or into interstellar space), and usually usually need to rely on complex slingshot maneuvers to get enough energy to reach the target. Such missions typically take many years to reach their destination, and have limited maneuverability once they arrive. If we want to go faster or execute more complex maneuvers, such as visiting multiple bodies, chemical rockets may not suffice.

1.1.2   Nuclear thermal propulsion

To get around the limitations of chemical propulsion, we might try adding energy to the propellant directly, rather than relying on chemical bonds. One way to do this is with a nuclear reactor. When a radioactive material like uranium or plutonium starts decaying, it emits high energy neutrons which heat everything they encounter. If we flow a propellant (like hydrogen) around this hot slug of radioactive metal to heat it before expanding it through a nozzle, we obtain a *nuclear thermal rocket*. First developed in the 1960s, nuclear thermal rockets can achieve specific impulses of between 800 and 1,200 seconds while

maintaining good thrust levels (although the thrust-to-weight ratio is sometimes lower than chemical systems). Despite their potential usefulness, political challenges and technical issues have to date prevented nuclear thermal rockets from flying.

1.1.3 Electric propulsion

If we are unable to use radioactive decay to heat our propellant, then we can use electricity. Systems which use electricity in some way to energize and accelerate a propellant and produce thrust are called *electric propulsion* (EP) systems. One simple way to do this is to just put a resistive heating element, like the filament of a traditional incandescent lightbulb, in the path of the propellant. The resulting thruster is called a *resistojet* and can achieve specific impulses between 100 and 500 seconds,[5] depending on architecture and propellant choice. *Arcjets* are a related technology, in which a voltage drop is applied across a gap. This creates an electrical arc discharge, which heats the gas before it is accelerated in a nozzle. Arcjets typically reach 500 to 1,000 seconds of specific impulse.[5] Together, resistojets, arcjets, and other similar systems which use electricity primarily to heat the propellant are called *electrothermal* thrusters.

To reach higher specific impulses, we must go beyond simply heating the propellant. *Electrostatic* acceleration systems ionize the propellant gas by some means. Then, the charged particles are accelerated using an electric field to produce thrust. This electric field is created by means of an applied voltage of a few hundred to a few thousand volts. Depending on the choice of propellant and the applied voltage, the specific impulse of these systems can be anywhere from 1,000 seconds to upwards of 10,000 seconds. The most popular electrostatic acceleration systems are *gridded ion thrusters* (GITs) and *Hall effect thrusters* (HETs). Finally, there are devices which employ magnetic fields to confine and/or accelerate an ionized gas. These are called *electromagnetic* propulsion systems, and include technologies such as magnetoplasmadynamic, rotating magnetic field, and magnetic nozzle thrusters.

Due to their high specific impulses, EP systems require less propellant for a given spacecraft mass. This means more of the spacecraft's total mass can be taken up by useful payload instead of fuel or tankage. For this reason, EP systems are common for in-space station-keeping applications, in which a satellite needs to continually correct its orbit over a long period of time so that it stays in place. These same advantages make them equally useful for deep space exploration. Notably, the enabling features of electric propulsion allowed the NASA Dawn mission to visit two asteroid belt objects, Vesta and Ceres, in the same mission.[6] EP facilitated the European Space Agency's small and lightweight SMART-1 moon orbiter,[7] as well as the Hayabusa asteroid sample return missions.[8] In the near future, EP will play an important role as the main propulsion on the upcoming Psyche mission.[9,10] Further down the line, EP is being considered as a key enabling technology for crewed Mars exploration[11] and even for use on missions to the outer solar system,[12] where the limitations of chemical rocket systems are even more restrictive.

1.2   Hall thrusters

Despite the key advantages of electric propulsion, high specific impulse alone is not sufficient. If it were, we could simply attach high-powered flashlights to all of our spacecraft. Since the propellant has the maximum possible exhaust velocity (the speed of light), we could go anywhere we wanted with minimal propellant usage. However, if we did this, we would find we went nowhere fast. We want high specific impulse, so we can go fast eventually, but we also want high thrust, so we can go fast soon. At constant power, these two factors trade off of each other, so that high specific impulse systems have low thrust, and vice-versa. To see this, we can write the expression for *thrust power*, sometimes also called *jet power*:

$$P_T = \frac{1}{2}\dot{m}v_{ex}^2 = \frac{1}{2}T\, I_{sp}\, g_0 \tag{1.4}$$

For a 5 kW system with an exhaust velocity for 20,000 m/s ($I_{sp} \sim$ 2000 s) the maximum thrust force we could possibly get would be 500 mN. This is a bit less than the weight of a typical tennis ball. Keeping the power constant, if we doubled the specific impulse, we would halve the amount of thrust generated. In the limit of maximum possible specific impulse, ignoring relativistic effects, a 5 kW flashlight would produce about 33 micronewtons of force, which is enough force to lift a mass of 0.003 grams in earth's gravity (about the weight of a seasame seed). Compare this to the F-1 engine, which produced 7.7 million Newtons of force at a specific impulse of 263 seconds at sea level, equivalent to 10 GW of thrust power. For chemical rockets, this power comes from combusting the fuel, but for EP systems, we would need to supply this as electrical power. Thus, with extremely-high specific impulse systems, even though we could eventually go very fast, it would take us prohibitively long to accelerate to such speeds, or we would need massive and very efficient power plants. To achieve the same thrust as the F-1 engine with a specific impulse of 2,000 seconds, we would need at least 80 GW of power, assuming 100% efficient conversion of electrical power to thrust power. This is about 50% more than the average power draw of the entirety of my home state of Arizona. In real-world applications, we therefore need to balance thrust, specific impulse, and power to achieve reasonable mission durations on reasonably-sized spacecraft.

Hall thrusters occupy a middle position in this trade space. These devices were first invented in the 1960s by both the U.S. and Soviet Union, and reached maturity in the Soviet Union in the 1960s and 70s,[13,14] as *Stationary Plasma Thrusters* (SPTs). Compared to their main rival, the gridded ion thruster, they have high thrust density.[14] This comes down to the fact that the plasma in a Hall thruster is quasineutral, and thus is not limited by physical laws governing the amount of charge allowed through a certain volume. In contrast, gridded ion thrusters are non-neutral, and their thrust density is limited by the size
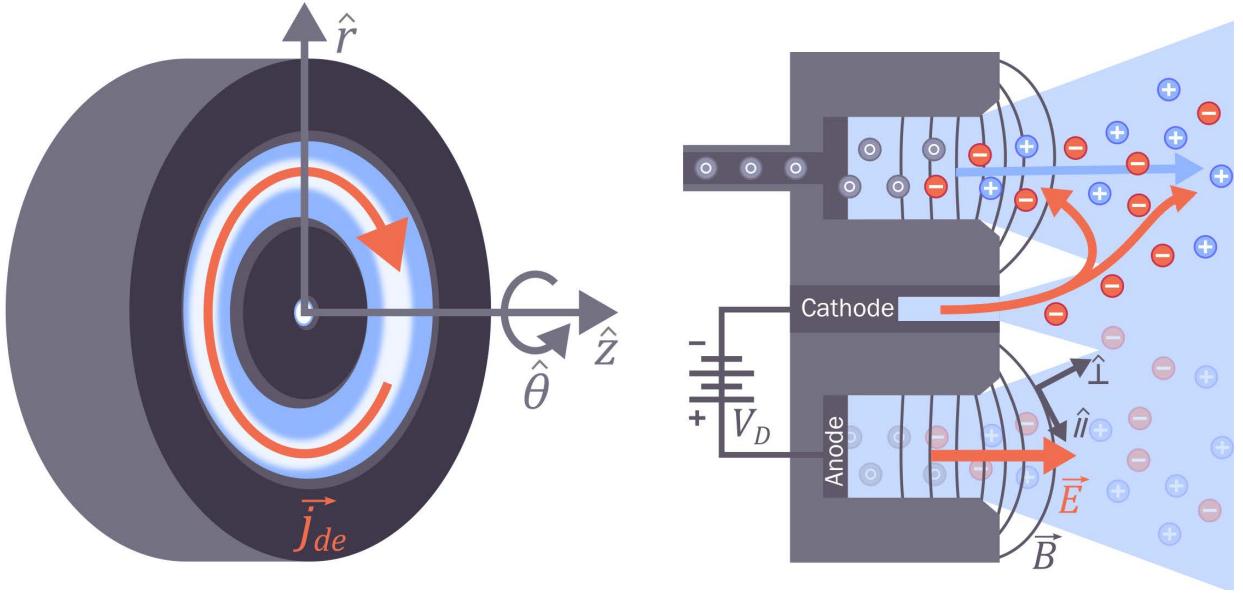
Figure 1.1: Diagram of a Hall thruster, showing particles, forces, and drifts. Ions, neutrals, and electrons are denoted ⊕, ⊚, and ⊖, respectively.
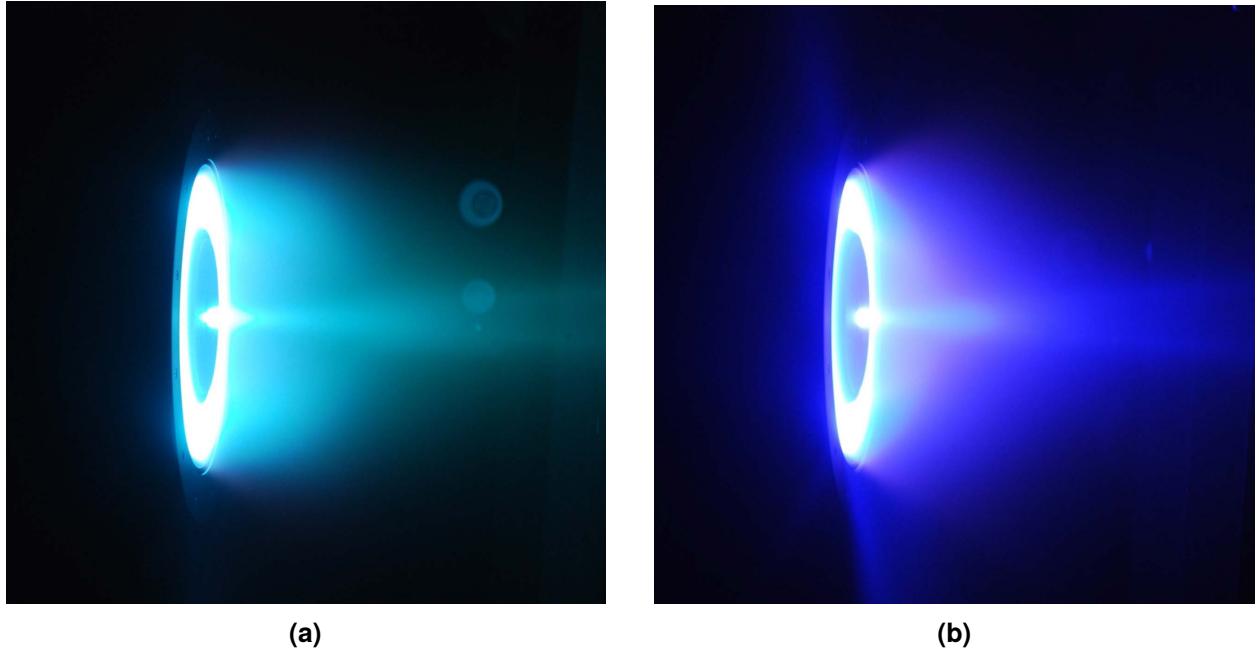
of the openings of the eponymous grid. Despite this, Hall thrusters can achieve specific impulses competitive with those of gridded ion thrusters, between 1500 seconds and 3000 seconds,[14] although even higher specific impulses (up to 8000 s[15]) have been achieved. This means Hall thrusters can be tailored for a wide variety of missions.

Figure. 1.1 depicts a Hall thruster. Hall thrusters are annular (ring-shaped) devices consisting of a discharge chamber, an anode, and a hollow cathode. The right-hand side of the image shows a cross-section of a Hall thruster along the $r - z$ plane. To start a Hall thruster, we apply a discharge voltage (V_D) between the anode and cathode. This creates an electric field (**E**) pointing downstream. Simultaneously, we apply a magnetic field in the radial ($\hat{r}$) direction, perpendicular to the electric field. The hollow cathode produces electrons (labelled ⊖ in the diagram), which follow the electric field toward the anode and into the discharge channel. The crossed electric and magnetic fields cause the electrons to drift in the azimuthal ($\hat{\theta}$) direction. The resulting *Hall* current, from which these devices derive their names, is labelled **j**$_{de}$ in the diagram. Electrons in this drift move very rapidly, with speeds in excess of 100,000 m/s. The magnetic field also acts to halt the electrons

8

on their path from cathode to anode, enabling them to remain in the discharge channel for relatively long periods of time.

At the same time, neutral gas particles (labelled ⊙ in the diagram) are injected at the anode. They drift downstream until they collide with the "buzzsaw" of drifting electrons. These electrons strip one or more of the outer valence electrons from the neutral, leaving behind a positively-charged ion (labelled ⊕). The freed electrons join the Hall drift and in turn ionize even more neutrals, which produces yet more ions until most of the gas is ionized. This cascade is known as *avalanche ionization* and the resulting ionized gas is known as a *plasma*. We will describe plasmas in more detail in **Chapter 2**. The applied electric field accelerates the newly-born ions downstream and eventually out of the device, producing thrust. Enough electrons follow the departing ions so that the plasma remains electrically-neutral. As previously noted, this *quasineutrality*, or the state of having zero net-charge despite the presence of many local positive and negative charges is what enables Hall thrusters to have higher thrust density than competing electric propulsion systems.

Traditionally, Hall thrusters have used xenon as their main propellant.[14] Xenon has three attractive properties. First, it is a relatively massive atom. This means that for a fixed acceleration voltage, we get more thrust per ion, at the cost of reduced specific impulse. It also means that xenon can be stored at high densities, which reduces the amount of spacecraft mass needed for fuel tanks. Second, it is easy to ionize. Since neutral gas that is not ionized is essentially wasted, this property means we need less propellant to produce a given amount of plasma. Lastly, it is a noble gas, and therefore nonreactive. This means it is nice to work with and does not corrode thruster surfaces. However, krypton and argon have recently become popular alternatives. As they are smaller and harder to ionize than xenon, they sacrifice some amount of performance and storage density in exchange for reduced cost, while maintaining the attractive chemical properties of xenon.

(a)   **(b)**

Figure 1.2: The H9 Hall thruster operating on **(a)** xenon and **(b)** krypton at 300 V and 15 A. Photo credit: Leanne Su[1] and the Plasmadynamics and Electric Propulsion Laboratory at the University of Michigan.

In **Figure 1.2** we show photos of a Hall thruster operating on xenon (**Figure 1.2a**) and krypton (**Figure 1.2b**).

1.2.1 Electron transport in Hall thrusters

In the past 30 years, modeling and simulation (M&S) have played an increasingly-important role in the design and qualification of aerospace systems. Just as we can predict the lift, drag, and stall angle of a wing using computational fluid dynamics, we would ideally be able to predict the performance and lifetime of a propulsion system from its geometry and operating conditions alone. However, despite the widespread application of Hall thrusters to a number of domains, there are still aspects of their operation and performance that remain poorly-understood. This lack of understanding precludes use from simulating Hall thrusters in a fully-predictive manner. In the last fifteen years, many interesting new Hall thruster concepts and design modifications have been proposed

and developed. These include magnetically-shielded[16], high-current-density[17,18,19], wall-less,[20] cylindrical,[21] and air-breathing Hall thrusters.[22] Without the ability to simulate these advanced concepts ahead of time and have confidence in the results, we must rely instead solely on experimental testing. Qualifying a Hall thruster for flight requires hundreds of hours of ground testing in a vacuum chamber, which is both time-consuming and expensive.[23] There is thus a clear need for improved M&S capabilities for Hall thrusters. To achieve this, we must tackle some of the unknowns about Hall thruster physics and operation.

Chief among these unknowns is the problem of *anomalous electron transport* in Hall thrusters. Referring again to **Figure 1.1**, we see that in order for electrons to traverse the distance from cathode to anode, they have to cross the applied magnetic field pointing in the radial direction. In the presence of a magnetic field, charged particles will exhibit gyrations about the imposed field lines, in a process known as *Larmor precession*, named after Irish physicist Joseph Larmor. The *Larmor radius* (r_L, **Equation 1.5**) of the resulting orbit scales with the particle's mass, m, and velocity perpendicular to the magnetic field, $v_\perp$, and inversely with the charge of the particle, q, and the strength of the applied magnetic field, B.

$$r_L = v_\perp \frac{m}{|q|B} \tag{1.5}$$

In Hall thrusters, the strength of the magnetic field is tailored such that the Larmor radius of the electrons is much smaller than the size of the device. In contrast, the Larmor radius of the ions, which are thousands of times more massive than the electrons, is much larger than the device scale. The implication of this fact is that the electrons are "trapped" in Larmor procession and unable to migrate across the magnetic field lines, while the ions are free to cross the field lines and exit the device. In principle the electrons thus will remain in the discharge channel, whirling around in the Hall drift, free to ionize neutral
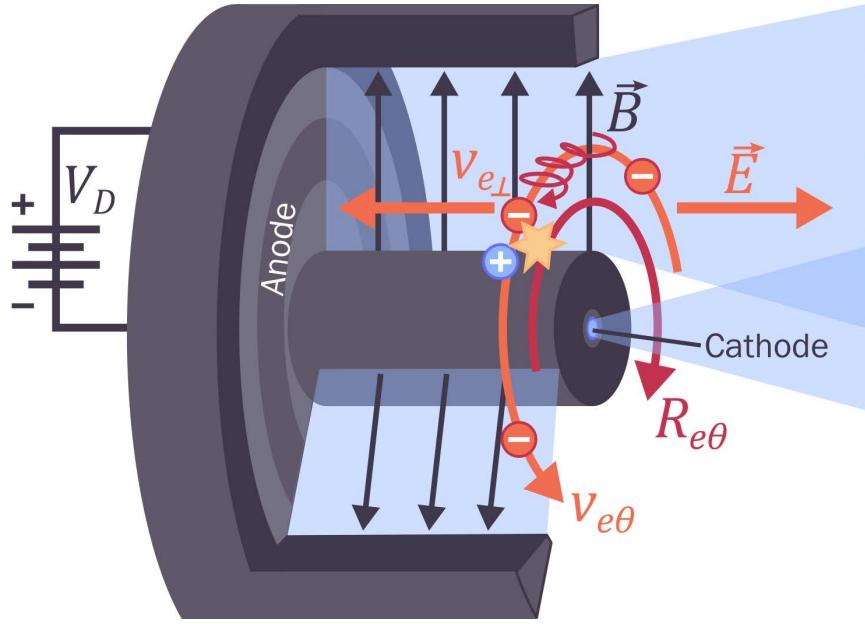
Figure 1.3: How collisions create cross-field transport in a Hall thruster

atoms with minimal additional input power. However, some electrons must eventually leak to the anode to complete the circuit. Classically, this occurs via collisions.

In **Figure 1.3**, we show how electrons are able to cross the magnetic field lines in the presence of collisions. The electrons, which are simultaneously engaged in Larmor procession about the magnetic field lines and an $E \times B$ drift around the device with velocity $v_{e\theta}$, collide with neutrals and ions. This creates a drag force, $R_{e\theta}$ on the electrons in the azimuthal direction. For reasons described mathematically in **Chapter 3**, this induces the electrons to drift across the magnetic field lines in the opposite direction of the applied electric field with velocity $v_{e\perp}$. The resulting cross-field electron current scales with the collision frequency, ν_e, of electrons with heavier species. However, even accounting for all known collision types, the predicted electron current is orders of magnitude lower[13,14] than the electron current measured in Hall thruster experiments. To create this enhanced "anomalous" transport, there must be some additional drag-like force in the azimuthal direction which is much stronger than the collisional drag.

The cause of this enhanced transport has been a focus of Hall thruster research since the very early days of their development. Initially, two theories dominated–*Bohm-diffusion* and *near-wall conductivity*. Bohm diffusion is a type of enhanced transport first observed by David Bohm in 1949 and later found to exist in a wide variety of magnetic confinement fusion devices through the late 1960s. The cause of this enhanced diffusion is plasma turbulence. Specifically, microscopic fluctuations in the ion density induce fluctuations in the electrostatic potential, which then enhance electron transport.[24] In Bohm diffusion, the cross-field electron mobility (μ, a measure of how readily electrons can move in response to an applied electric field) is expressed as:

$$\mu = \frac{1}{16B} \qquad (1.6)$$

The association of Bohm diffusion with Hall thrusters is 60 years old. Alexey Morozov (one of the original Hall thruster designers) recounts that the opponents of the SPT development program within the Soviet Union said that the device would not possibly work, as Bohm diffusion would thwart any attempt to confine the electrons in the discharge channel.[13] While the first SPTs did indeed see increased electron mobility over what was expected from collisions alone, the magnitude of this enhanced transport was lower the value predicted by Bohm diffusion. This lead to the development of the theory of *near-wall conductivity*.[25,13,26] When high-energy electrons collide with a surface, such as the discharge channel walls of a Hall thruster, they may induce that surface to emit additional electrons. This *secondary electron emission* can then enhance the electron current in the device channel. Additionally, any collisions or *backscattering* by electrons with the thruster walls can produce a drag force in the same way that collisions with other particles can. As previously described, this drag force may then enhance transport across the magnetic field lines. Lastly, if the wall material is sufficiently conductive, electrons may *short-circuit* across the magnetic field lines by travelling through the channel wall rather then through the plasma. These three related effects are summarized in **Figure 1.4**.
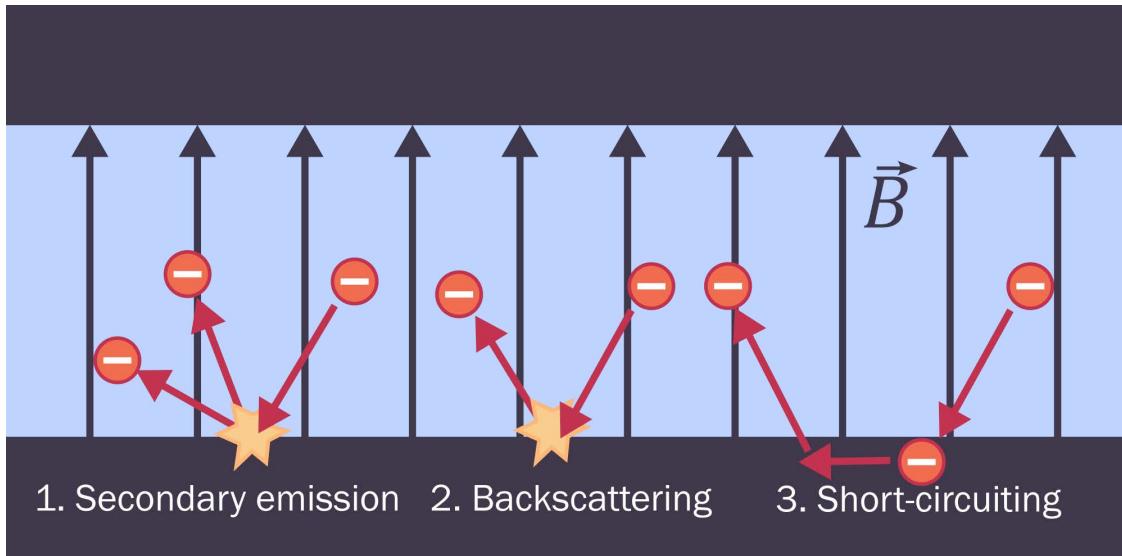
Figure 1.4: Processes inducing near-wall conductivity in Hall thrusters

All of these effects have been observed experimentally in Hall thrusters.[13,26,14] and do serve to enhance the movement of electrons across the applied magnetic field. However, near-wall conductivity alone is not sufficient to describe all of Hall thruster electron transport. In some Hall thrusters, particularly magnetically-shielded and wall-less Hall thrusters, the flux of electrons to the walls is dramatically reduced.[16] If near-wall conductivity was solely responsible for electron transport, then these devices should have significantly-lower electron currents than conventionally-designed Hall thrusters.[14] However, this is not the case.[27,28] Additionally, measurements and simulations of the anomalous electron transport have found that it is much higher *outside* of the discharge channel than inside of it.[29,30,31] Taken together, these indicate that there must therefore be some additional process driving the enhanced cross-field electron transport that is not tied to electron interactions with the walls.

In the last fifteen years, a consensus has emerged that this additional process is plasma turbulence, albeit turbulence which does not follow the Bohm scaling[14]. Soviet physicist Yuri Esipchuk was one of the first researchers to measure high-frequency oscillations in Hall thrusters. In 1976, he identified the cause of these oscillations as

electron drift waves.[32] These waves derive energy from the Hall drift and propagate in the azimuthal direction. By removing energy from the azimuthal drift, they produce an effective drag force on the electrons, which enhances cross-field transport in the same manner as collisional drag. Since the 2000s, experimental,[33,34,35,36,37] theoretical,[38] and computational[39] evidence of the presence of these waves in Hall thrusters and their effect on electron transport has only grown.

Despite this consensus that turbulence originating from electron drift instabilities governs anomalous electron transport in Hall thrusters, we are still unable to model its effect on new Hall thruster designs in a fully-predictive manner. Recent high-fidelity particle-in-cell (PIC) simulations[39] have given us unprecedented insight into the growth and scaling of this transport. In the coming decades, this type of direct numerical simulation may be invaluable in improving our physical knowledge of anomalous transport. However, such simulations currently take months to run and thousands of CPU cores. This makes them prohibitively-expensive for engineering design applications. We would prefer to use simulations which treat the electrons as a continuous fluid instead of particles. These models run dramatically faster—on the order of hours to a few days—but sacrifice the ability to self-consistently resolve the plasma fluctuations which give rise to anomalous transport.

To close this gap, we require a model for how anomalous transport affects the plasma which does not require us to resolve the microscopic length- and time-scales over which the instabilities grow and propagate. Such models are common in the theory of classical fluid turbulence,[40] and find wide application in commercial, industrial, and research settings. Thanks to such models, we can more-or-less accurately predict the drag across an entire airplane, the flow of water through a rough pipe, or the movement of air over the dimpled surface of a golf ball. While there are always improvements to be made, the development of reliable approximations for the effect of turbulence on mean flow properties (i.e. density, velocity, temperature) has made modeling and simulation an inextricable and

15

indispensable step in the design, testing, and qualification of aerospace and consumer devices.

We would like to find turbulence models that do the same for Hall thruster simulations that they have for conventional fluid simulations. Indeed, many people [41,42,2] have proposed such models, but they have not been well-tested or extensively evaluated in the publicly-available literature. This raises several questions. Once we have a model in hand, how should we evaluate it? To what data should we compare it? If the model has parameters, how should we calibrate them?

1.3   Objectives of this dissertation

In this dissertation, I will try to answer some of the above questions. Unfortunately, I was unable to completely and forever resolve the problem of anomalous transport in Hall thrusters. I leave that problem to the younger generation. However, I made several interesting discoveries and developed some useful methods that should hopefully make the solution of this longstanding problem a little bit easier. I hope to advance the state of the art in developing and testing models of anomalous electron transport in fluid models of Hall thruster plasmas. As part of that goal, I will

1. **Review previously-proposed models for anomalous transport from literature,**

2. **Develop of new methods for calibration of transport models against data,**

3. **Demonstrate that empirically-derived transport profiles alone are not suitable for model calibration,**

4. **Evaluate the predictive performance of several models in a fluid Hall thruster code, and**

5. **Evaluate novel models of anomalous transport and propose several more.**

1.4 Organization

The rest of the dissertation is structured in the following way.

In **Chapter** 2, I overview the basic plasma physics that underpin the dissertation. I describe what plasmas are and how they can be described both statistically as collections of individual particles, and as fluids. I then give the equations of motion for a fluid plasma, and show how these can applied to understanding how electrons move in a Hall thruster. I then show how anomalous electron transport can be incorporated into a fluid description of electron motion.

In **Chapter** 3, I will discuss the types of models of transport in which we are interested. I will then review what models of anomalous electron transport have been proposed to date in the scientific literature, before finally discussing how these models can be altered create new models which may be worth investigating.

Next, in **Chapter** 4, I describe the simulation code, Hall2De, that we employ for the bulk of the work in this dissertation. This includes a description of history of the code, the physics it models, as well as the numerical procedure by which its governing equations are solved. I discuss how I implemented anomalous transport models into the code and describe what other model parameters are needed to run a simulation. I then show the results of grid convergence and sensitivity studies which I performed determine the effect of parameters which are not related to the anomalous electron transport. Lastly, I discuss the thruster being simulated in the remainder of the dissertation, the data to which we compare our simulation results, and the metrics by which that comparison is made.

Then, in **Chapter** 5, I discuss the common practice of fitting empirical anomalous transport profiles to experimental data. I describe my efforts to evaluate models by comparing them to these empirical profiles before demonstrating that this process is flawed and does not work. I then discuss the relevance of these empirical transport profiles to anomalous transport modeling in general.

After that, in **Chapter 6**, I use Hall2De to evaluate four of the models introduced in Chapter 3. I test the models' ability to predict the performance and plasma properties of a magnetically-shielded laboratory Hall thruster. I examine the ability of the models' fit coefficients to generalize beyond their calibration condition and accurately scale across discharge voltages, discharge currents and propellants. I then discuss these results in the context of future transport modeling efforts.

Next, in **Chapter 7**, we test a few models not found in the literature. These were created by tweaking some of the assumptions underlying the models in Chapter 3.

Finally, in **Chapter 8**, I conclude by summarizing my research. I list the major contributions of my thesis work and provide some hopeful avenues for future exploration.

After all of that, I describe in **Appendix A** the development of an open source one-dimensional Hall thruster code, and give some thoughts on the usefulness of such a model in future anomalous transport modeling efforts.

CHAPTER 2

Hall Thruster Plasma Physics

> The electron: may it never be of any use to anybody!
>
> _____
>
> Joseph John Thomson, discoverer of the electron

In the preceding chapter, I said that the goal of this thesis was to to advance the state of the art in developing and testing models of anomalous electron transport in fluid models of Hall thruster plasmas. So far, I have explained what Hall thrusters are, and described (in brief) the problem of anomalous electron transport, and why we would like to develop models for this crucial and poorly-understood process. In this chapter, I will tackle the next part of the sentence, by explaining what a plasma is and describing how electrons in a plasma may be modelled as a fluid. I will then discuss how anomalous transport can be incorporated into a fluid model of a plasma.

2.1 States of matter

All matter is made of particles called atoms. Each atom is the in turn composed of a *nucleus* (itself made of neutrons and positively-charged protons) and a cloud of negatively-charged electrons. These electrons are arranged in a matryoshka-doll like series of shells known as *orbitals*. The electrons in the inner-most orbital are bound tightly to the nucleus,

while those in the outermost orbital are held only tenuously by the positively-charged core. Electrons in this outermost shell are known as *valence* electrons, and interactions between valence electrons of different atoms are responsible for virtually all chemistry. Electrons in an atom typically occupy the orbital with the lowest available energy. An atom in which all electrons exist at the lowest-possible energy level are said to be in the *ground state*. Atoms are most stable when their valence shells are full, a condition which is only true for the noble gases: helium, neon, argon, krypton, xenon, and radon. For other types of atoms to fill their valence shells, they need to exchange or share their valence electrons with neighboring atoms.

Consider salt, or sodium chloride (NaCl). Sodium (Na) has only one electron in its valence shell. The easiest way for it to have a full valence shell would be to lose this solitary electron, exposing the full shell one level below. Chlorine has seven valence electrons, one less than a full set of eight for its outermost orbital. Because both of these atoms are so close to having a full set of valence electrons, they bond with each other readily. Indeed, both sodium and chlorine are highly reactive with almost anything, and do not exist in nature in their atomic forms. A bond like the one formed between sodium in chlorine is known as an *ionic* bond (**Figure 2.1**), because when sodium gives its one valence electron to chlorine, it becomes a positively-charged *ion*, while chlorine becomes a negatively-charged ion. An ion is any atom which has a net charge, positive or negative, in contrast to the normal state of affairs where the number of electrons exactly matches the number of protons. Once both atoms become ions, they are bound tightly together by the electrostatic force of attraction between the positive and negative charge.

Other types of bonds exist, such as covalent bonds. In these bonds, electrons are shared, rather than exchanged, between atoms. There are also metallic bonds, in which many metal atoms share all of their valence electrons in one giant electron cloud which extends across the whole substance. This accounts for the high conductivity of metals, as electrons are very free to move from one side of a chunk of a metal to another.
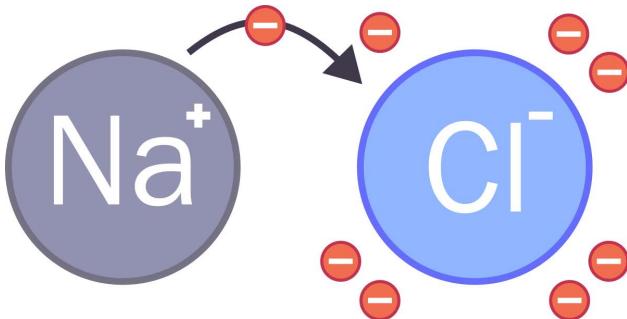
Figure 2.1: Example of an ionic bond between sodium (Na) and chlorine (Cl).

Chemical bonds are the strongest type of force that exists between atoms. These forces are strong enough that the atoms combine into a new unit (known as a molecule), which exhibits completely different properties than its constituent atoms. Ionic bonds and metallic bonds are typically stronger than covalent bonds. One level down in strength are the *interatomic* or *intermolecular* forces. As the name implies, these are forces which attract or repel nearby atoms to or from one another, but which are insufficiently-strong to join the atoms permanently into a molecule.

Intermolecular forces are responsible for what we know as states of matter, and explain why, for example, ice melts and then vaporizes as its temperature increases. The *temperature* of a substance is a statistical property which describes the average kinetic energy of a collection of atoms or molecules. A single molecule cannot have a temperature, only an individual mass, charge, and velocity. In a collection of molecules at a low temperature, the relatively-weak intermolecular forces are strong enough to overcome the low kinetic energies of the particles. The molecules are then able to clump together into a single hunk of substance known as a *solid*. As the temperature increases, the average energy of the particles becomes high enough that they are no longer bound as tightly to one another by the intermolecular forces, but not so high that they can completely escape each others' influence. This state of matter is known as a *liquid*. Finally, as the temperature continues to increase, the kinetic energies of the particles become high enough that the
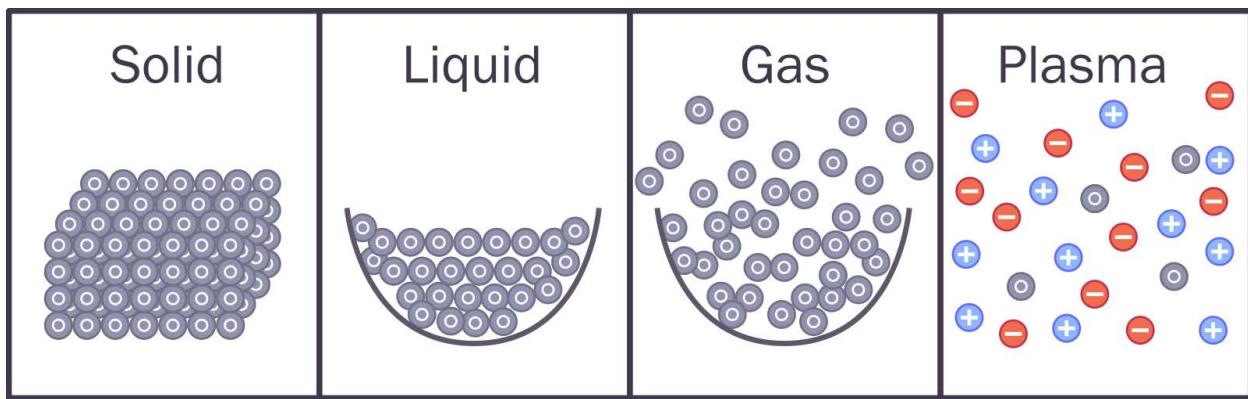
Figure 2.2: The four major states of matter: solid, liquid, gas, plasma.

interatomic forces become inadequate to hold them together, and the liquid breaks apart into a collection of freely-moving atoms, or a *gas*.

What happens if we increase the temperature further? First, the molecules continue to gain kinetic energy, i.e. speed up. Energy can also be deposited into the gas in other ways. The molecules may exhibit rotation or vibration. These energy modes become more important at higher temperatures. Additionally, energy may be deposited directly into electrons. These electrons gain kinetic energy which can allow them to reach an energy level above the ground state. This process of electronic *excitation* increases in importance as the temperature increases. Finally, there comes a point at which the electron which is most loosely-bound to the atom gains enough energy that the attraction of the nucleus can no longer retain it. That electron leaves the atom and begins to float freely among the other atoms in the gas. A positively-charged ion is left behind.

If this process continues, a large fraction of the atoms of the gas will become ions. A substance in this state is known as a *plasma*, the so-called "fourth state of matter". We summarize the four main[a] states of matter in **Figure 2.2**. Plasmas exhibit many features common to gases. For instance, they will diffuse to fill the container they are in. Additionally, they are *compressible*, meaning the volume occupied by of a fixed number of

---

[a]Other more exotic states of matter, such as the Bose-Einstein condensate, exist but are not germane to the present work.

atoms can change if the pressure changes. This is not the case of something like water, which is highly *incompressible* and cannot be easily squeezed into a smaller volume. Despite the similarities with a standard neutral gas, plasmas exhibit a wide variety of emergent properties which make them both interesting to study and useful for a number of engineering purposes.

2.2   Elementary plasma physics

The most visible property of a plasma which makes it different from a normal gas is that plasmas typically glow. The color of this glow depends on which atoms are present, as well as the density and temperature of the plasma. As discussed previously, electrons in a hot gas can be excited. These excited states are typically unstable, meaning that is very easy for an electron to *spontaneously de-excite*. When this happens, the electron falls either back down to ground or to some intermediate excited state. The difference in energy between the excited state and the new state is then released in the form of light, with a wavelength inversely proportional to the energy released electron was. Electrons which fall further release more energy, and the wavelength of the light is shorter (i.e. more blue). In plasmas, since there is enough energy to ionize the atoms, there is also enough energy for a large degree of electronic excitation to occur. As a result, plasmas glow **(Figure 2.3)**.
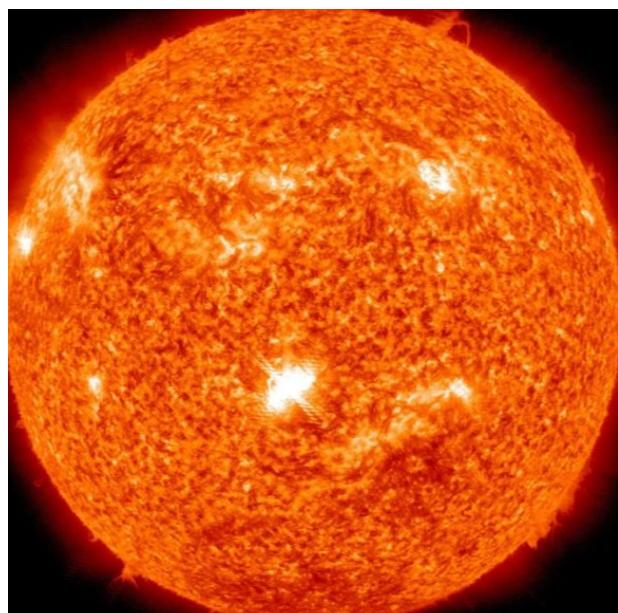
This glow is one of the most useful properties of a plasma. A "neon" light is just a tube filled with a (typically noble) gas at low pressure. When the light is turned on, a voltage is applied across the tube. This creates an electric field which helps accelerate the small number of free electrons present in the gas to high velocities. When these electrons hit the electrons in the outer valence shell of a neutral atom, they can sometimes knock the valence electrons free, ionizing the atom. This newly-freed electron can then be accelerated and ionize another atom. This produces a chain reaction, leading to a a large fraction of the gas becoming ionized. This avalanche ionization, as discussed in

23
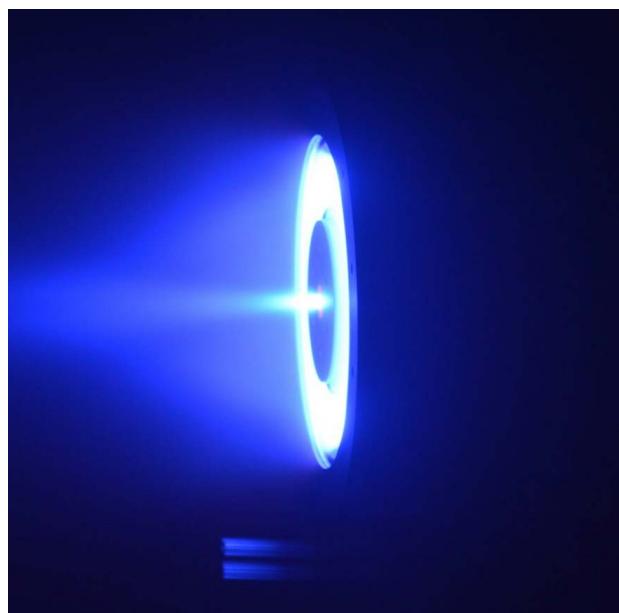
(a)   **(b)**

(c)   **(d)**

Figure 2.3: Some examples of plasmas glowing. (a) A neon sign (b) a compact fluorescent light (c) the sun (d) a Hall thruster. All images either licensed under Creative Commons Share-A-Like License or copyright Plasmadynamics and Electric Propulsion Laboratory.
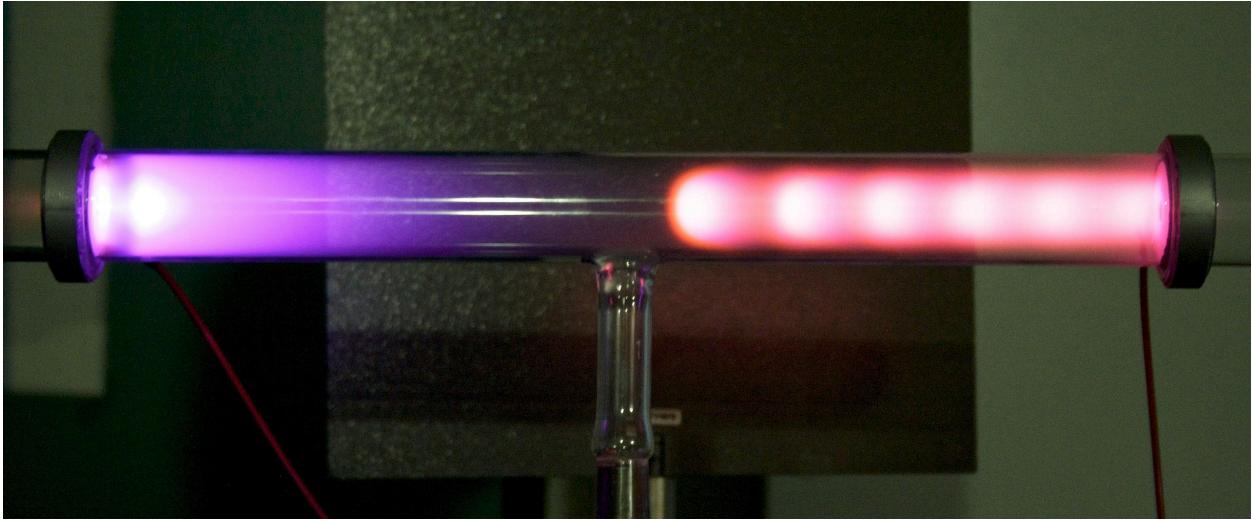
Figure 2.4: A typical DC glow discharge. From the wikipedia page on glow discharges.

the introduction, is the same process that creates ionization in Hall thrusters. The type

plasma found in a neon light is known as a *DC glow discharge* **(Figure 2.4)** and is one

of the simplest plasma devices. Another commonly-encountered type of glow discharge

is in a fluorescent lightbulb, which works in basically the same way as a neon light, but

includes a phosphor coating on the tube which absorbs the light emitted by the plasma

and re-emits it as broad-spectrum white light.

DC (direct current) discharges are plasmas created by the application of a constant

voltage across two conducting surfaces, called *electrodes*. This difference in voltage,

or *electrostatic potential*, creates an electric field pointing from the *anode* (the electrode

at high voltage) to the *cathode* (the electrode at low voltage). This contrasts with AC

(alternating current) discharges, in which the voltage varies at some constant frequency,

just like the AC power coming out of a home power outlet. If the frequency of the applied

voltage/current oscillations is in the radio range (3 kHz to 300 GHz), then we have a

radiofrequency (RF) plasma. The plasma does not need to be ionized by the direct

application of a voltage across electrodes. Electromagnetic waves, whether microwave

or radio, can accomplish the same goal. There is a huge variety of ways in which plasmas

may be generated, but as Hall thrusters are DC discharges, we will not discuss these methods here.

The second important property of a plasma is *electrical conductivity*. Typical gases are not conductive, which is why high-voltage power lines do not typically shock passersby. However, plasmas have lots of electrons which are not bound to atoms and are thus free to move around, making them very conductive. This means plasmas can carry currents, as well as generate and respond to electric and magnetic fields. This property gives plasmas their third important property: *collective behavior*. This term refers to interactions which take place on a larger length scale and involve more atoms than the typical collisional processes which govern neutral gases. In addition to responding to external fields, plasmas can generate and respond to their own fields. This allows them to self-organize, produce large-scale structures, and undergo a rich variety of waves and instabilities. As one example, the plasma in the outermost layer of the sun can form massive arches or filaments larger than the size of the earth **(Figure 2.5)**, with the structure of the arch created and supported by magnetic fields induced by the plasma.

The simplest type of collective behavior is simple plasma oscillation. If we displace displace some of the electrons in a plasma away from the ions, such that get a region of net negative charge and a region of net positive charge are produced, the plasma will respond collectively in a way in such a way as to restore charge neutrality. In this case, the displaced electrons will be attracted back toward the region of positive charge, overshoot, and then reverse course. The electrons will thus oscillate in the same manner as a spring stretched or compressed past its neutral point and then release. The frequency that this oscillation takes place at is the *plasma frequency*, ω_{pe}, and is given by:

$$\omega_{pe} = \sqrt{\frac{q_e^2 n_e}{m_e \epsilon_0}} \quad \text{radians per second (rad/s)} \tag{2.1}$$

In this expression, q_e is the fundamental charge, or 1.6×10^{-16} Coulombs (C), n_e is the *number density* of electrons, which refers to the number of atoms per unit volume ($1/m^3$),
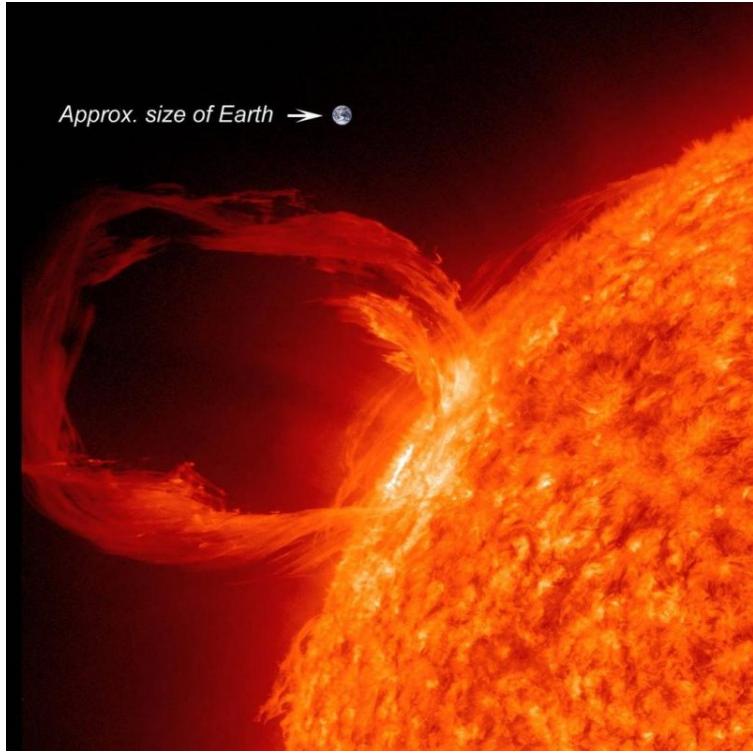
Figure 2.5: A solar filament. Image copyright NASA.

m_e is the mass of an electron, or 9.1×10^{-31} kilograms (kg), and ϵ_0 is a fundamental constant known as the *permittivity of free space*, 8.854×10^{-12} Farads per meter (F/m).

Another form of collective interaction that plasmas experience is *Debye shielding*. Despite being composed of many individual positive and negative charges, plasmas are overall electrically-neutral, meaning that the number of negative charges is equal to the number of positive charges. This condition is called *quasineutrality*. If a small local charge imbalance develops, the charges will re-orient in order to screen out that charge imbalance from being seen by particles far away. The maximum length scale over which a plasma can support small charge imbalances is known as the *Debye length*, λ_D, and is given by:

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{q_e^2 n_e}} \quad \text{meters (m)} \tag{2.2}$$

In the above, k_B is the Boltzmann constant, 1.38×10^{-23} Joules per Kelvin (J/K), which relates the temperature of a substance to the average kinetic energy of the particles, and

T_e is the temperature of the electrons, which is typically on the order of 10,000 - 1,000,000 K for the "low-temperature plasmas" that we are interested in. Typically, we use units of electron-Volts to describe the temperature of a plasma. One electron-Volt is the amount of energy gained by a charged particle with charge q_e accelerated down a one-volt potential drop, and corresponds a temperature of around 11,600 K.

In addition to these longer-range interactions, particles in a plasma can undergo collisions, just as particles in a gas do. The *collision frequency*, generally denoted by the greek letter nu (ν), is the frequency at which a single particle will collide with other particles. However, in addition to collisions between electrons and ions with neutral atoms, there are also collisions between the charged species. These electron-ion (ν_{ei}), electron-electron (ν_{ee}), and ion-ion collisions (ν_{ii}) take place over a larger length scale than typical neutral-neutral collisions as the electrostatic interaction between charged particles means the particles "see" one another from further away than in a neutral gas.

With the Debye length and plasma frequency defined, we can introduce the three basic conditions for a gas to be a plasma. Mathematically, these are:

1. $\lambda_{de} \ll L$

This condition states that the Debye length must be significantly smaller than the characteristic length scale (L) of the system we care about. If this is not true, then we can have charge imbalances on the scale of our system of interest, violating the principle of quasineutrality and the system will not behave as a plasma.

2. $\frac{4\pi}{3}\lambda_{de}^3 n_e \gg 1$

This condition states that the number of particles in a sphere with a radius equal to the Debye length is much greater than one. If this is not true, then there will not be enough particles free to redistribute around to shield any charge imbalances. This will thwart Debye screening and thus quasineutrality.
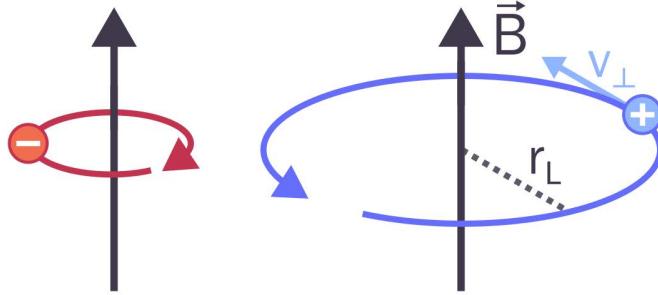
Figure 2.6: Gyration of charged particles around magnetic field lines. Due to their opposite charges, ions and electrons orbit in opposite directions.

3. $\omega_{pe} \gg \nu$

This condition states that the frequency of plasma oscillations is much larger than the collision frequency between particles. If this is not true, then collisions will randomize the velocities of the constituent particles in a way that disrupts collective (long-range) behavior.

Before we move onto developing the equations of motion for a fluid plasma, we briefly touch on the motion of a plasma in a magnetic field. As introduced in the previous chapter, charged particles orbit magnetic fields with a frequency and radius which depend on the magnetic field strength and the particle velocity **(Figure 2.6)**. The *cyclotron frequency*, ω_c is the frequency of this gyration, and is given by

$$\omega_c = \frac{q|\mathbf{B}|}{m} \quad \text{rad/s} \tag{2.3}$$

This is also often referred to as the *gyrofrequency*. In this expression, q is the charge of the particle in Coulombs, $|\mathbf{B}|$ is the magnitude of the magnetic field vector, in Tesla, and m is the mass of the particle in kilograms. The *gyroradius*, also known as the *Larmor radius* (r_L), is then

$$r_L = \frac{v_\perp}{\omega_c} \quad \text{m,} \tag{2.4}$$

where $v_\perp$ is the velocity of the particle perpendicular to the direction of the applied magnetic field.

Due to their larger masses, ions typically experience gyro-orbits with much lower frequencies and larger radii, while electrons circle the field lines in rapid, tight orbits. If the cyclotron frequency of a species is much larger than the frequency at which it experiences collisions, or if the Larmor radius is significantly smaller than the size of the system, then that species is said to be *magnetized*. This means the particle is able to complete many orbits before it experiences a collision or leaves the system. On the other hand, if a species collides much more often then it orbits, or if the Larmor radius is larger than the system size, then it is unlikely that a particle of that species will complete even a single orbit. In this case, that species is said to be *unmagnetized*. Due to the difference in Larmor radii and cyclotron frequencies, it is common in many plasma devices for the ions to be unmagnetized while the electrons are strongly magnetized. We can quantify the degree of magnetization of a plasma using the *Hall parameter*, Ω, which is just the ratio of the cyclotron frequency to the collision frequency:

$$\Omega_s = \frac{\omega_{cs}}{\nu_s}. \tag{2.5}$$

Here, s refers to the type—or *species*—of particle under consideration. In the simplest case, s could be i or e, denoting ions or electrons, respectively. In more complicated situations, we might have multiple types of ions, either of different elements or chemical species, different charge states, or both. If $\Omega_s \gg 1$, then we consider species s to be magnetized. Otherwise, s is unmagnetized.

The propensity of particles in a plasma to gyrate about the magnetic field lines means that the motion of magnetized species perpendicular to the magnetic field lines is inhibited. If the particle is magnetized, any velocity perpendicular to that field, $v_\perp$, will in effect be curved into a circle around the field line. This means that magnetic fields are often used to

30
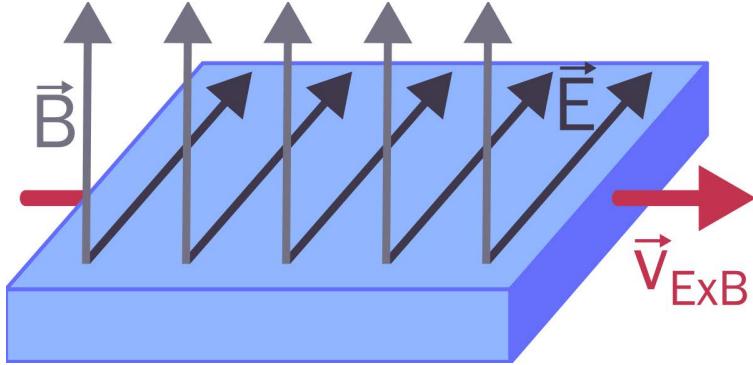
Figure 2.7: Illustration of the $E \times B$ drift in a slab of plasma.

confine and shape plasmas. Indeed, in the absence of collisions and applied electric fields, a single particle should not be able to diffuse across the magnetic field at all. Nonetheless, they do, and the ability of plasmas to diffuse across magnetic fields despite this strong confining effect is a challenge for engineers in virtually every field of plasma science.

If an electric field is applied perpendicular to the magnetic field, charged particles will begin to move or *drift* perpendicular to both of these fields. This phenomenon is known as the $E \times B$ ("E-cross-B") drift **(Figure 2.7)**. The drift velocity is given by cross product of the electric and magnetic field vectors, divided by the square of the magnitude of the magnetic field:

$$\mathbf{v}_{E \times B} = \frac{\mathbf{E} \times \mathbf{B}}{|\mathbf{B}|^2} \quad \text{m/s} \tag{2.6}$$

Note that the direction of this drift does not depend on the sign of the charge of the particles, so electrons and ions will drift in the same direction as each other. Conversely, applying a current across a magnetic field will produce a voltage drop (and thus an electric field) perpendicular to both the applied magnetic field and the the direction of the net current. This phenomenon is known as the *Hall effect*.

With the basics of plasmas out of the way, we can discuss the ways in which we model and simulate plasma devices like Hall thrusters. We can then discuss how, despite the

presence of gyro-orbits, particles can and do diffuse across magnetic field lines in plasmas, and how instabilities and waves can enhance this diffusion.

2.3   Kinetic description of a plasma

The most complete description of a plasma would be given by considering the mutual electrostatic forces between every single pair of particles in the entire system. Even though plasmas, especially the low-temperature plasmas we are interested in, are typically diffuse, there are still an enormous number of particles in even a very small volume of gas. In Hall thruster plasmas, we typically encounter densities on the order of 10^{16} - 10^{18} particles per cubic meter (m^{-3}). For reference, one billion is 10^9, one trillion is 10^{12}, and one quadrillion is 10^{15}. One quintillion (10^{18}) is about one tenth of the number of grains of sand on all of the beaches on Earth combined. Modeling all of these interactions together would be absolutely infeasible.

To get around this, we could instead consider a statistical, or *kinetic* description of the plasma. We do not generally care to know what each individual particle in a system is doing. Instead, we care about the consequences of the summed motion of all of these countless microsopic particles on the macroscopic behavior of the system. To this end, instead of considering the trajectories what the velocity of a single particle is, we might instead consider the *probability* of finding a particle at position **x** and velocity **v** is. If we evaluated this probability across all particles in the system, we would obtain a *probability distribution function* (PDF). If we query this probability distribution function at just one location **x** and ask only what the probability of finding a particle with velocity vector **v** at that location is, then we obtain a *velocity distribution function* (VDF), which we usually denote $f(\mathbf{v})$. This is useful because the VDF characterizes most of the behavior we care about in a plasma while allowing us to discard all of the irrelevant microscopic detail. In **Figure 2.8**, we give an example of a (one-dimensional) VDF. This VDF has two peaks,
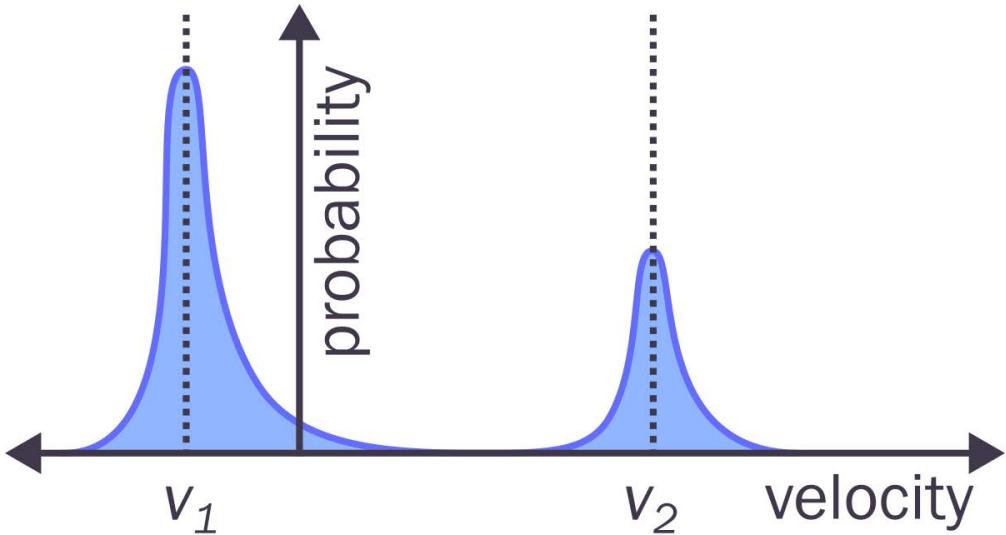
Figure 2.8: Example of a bimodal (two-peaked) velocity distribution function in one dimension.

one at v_1 and one at v_2, indicating that the are many particles with velocities near v_1, fewer with velocities close to v_2, and very few particles with other velocities. In general, a VDF may have arbitrarily many peaks and and arbitrarily-complex shape.

The *Boltzmann equation* describes how a velocity distribution function evolves over time and space in response to external forces and collisions between particles.[b] Considering only electrostatic body forces (i.e. neglecting things like gravity), this equation is

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} + \frac{q}{m} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f}{\partial \mathbf{v}} = \hat{C}(f). \tag{2.7}$$

The first term in Eq. 2.7 is time rate of change of the distribution function. The second term represents the convection of the distribution function with the particles. The third term represents the change in the distribution function due to body forces (in this case only electromagnetic forces). The term on the righthand side of the equation, $\hat{C}(f)$, is the *collision operator*. It is an integral expression for the effect of collisions on the distribution

---

[b]The Boltzmann equation is an averaged description of the *Klimontovich equation*, which is an exact description of the motion of all particles in a system. It is obtained by averaging over small volumes of physical and velocity space.

function. The simplest collision operator is the Bhatnagar-Gross-Krook (BGK) operator,[43] given by

$$\hat{C}(f) = -\nu(f - f_0). \tag{2.8}$$

Here, as described in the preceding section, ν is the characteristic collision frequency, and f_0 is some *equilibrium distribution function*. Physically, this operator encodes the fact that collisions between particles serve to spread out energy among the particles and relax their distribution toward some equilibrium, at which point the distribution function ceases to change further. In order to solve the Boltzmann equation, we need to consider not only the usual three dimensions of space and one of time, but also three dimensions of *velocity space*. This is because in the kinetic description of a gas or plasma, velocity is a coordinate, not a variable. We are interested in the probability of finding a particle *at* some velocity. In its full form, describing a gas or plasma kinetically requires the solution of a seven-dimensional partial differential equation. This is far better than considering every particle individually, but still quite computationally expensive. Additionally, while simple collision operators like the BGK operator suffice for highly collisional plasmas close to equilibrium, they are insufficient in other cases, so a detailed treatment of the collision operator is often required. One example of a more complex collision operator is the *Landau* collision operator[44], which describes the effects of coulombic interactions between charged particles. Directly modeling these operators might turn the relatively tractable seven-dimensional differential equation into a less-tractable seven-dimensional *integrodifferential* equation. Direct solution of the Boltzmann equation is thus typically avoided unless kinetic (i.e. nonequilibrium) effects are so important that they cannot be easily approximated by lower-order methods.

Some tricks exists which allow us to solve the Boltzmann equation approximately while not paying the full price of the solution of a seven-dimensional integrodifferential equation. *Particle-in-cell* (PIC) methods track the motion of *macroparticles*, each of which comprise

millions of real particles. Instead of considering the direct forces of interaction between each particle pair, we instead solve for the global electromagnetic field on some sort of computational grid and move the macroparticles in response to that field. This technique is both elegant and powerful. It allows us to resolve the shape of the distribution function at the cost of losing resolution. Despite these drawbacks, the particle-in-cell method is widely-used throughout computational plasma physics.

Even so, there are limitations to PIC which make other approaches more viable in certain situations. First, the results change depending on the number of particles employed, which manifests as statistical noise that further limits the resolution of the method. Additionally, there are fundamental restrictions on the allowable timestep and grid resolution. If the timestep is larger than the inverse of the plasma frequency ω_p^{-1} or the grid size is larger than the Debye length λ_D, the PIC method becomes unstable or exhibits unphysical behavior, such as excess heating of the particles. In Hall thrusters, the electron plasma frequency is around a picosecond (10^{-12} seconds) and the Debye length is around a micrometer (10^{-6} meters). This means that it becomes prohibitively expensive in most cases to simulate a Hall thruster, which operates on timescales greater than one millesecond (10^{-3} seconds) and length scales of up to a meter, using full particle-in-cell methods in which both electrons and ions treated as particles.

Despite these challenges, recent full-PIC simulations of Hall thrusters have probed the fundamental physics of their operation and the nature and scaling of the instabilities that drive transport.[39] However, such simulations still take a month or more to run, using thousands of CPU cores. Additionally, they are so far limited to small thrusters operating over short timescales. While there may exist techniques to speed up these methods, the cost of full-PIC simulations of Hall thrusters makes them impractical for engineering design applications. For this purpose, we need to reduce the level of fidelity once more and model the electrons as a fluid.

2.4 Fluid description of a plasma

As we have discussed, an individual particle only has a location, a mass, a charge, and a velocity, but a collection particles might have a *temperature* or a *density*. These so-called *fluid* quantities emerge only as statistical properties of large collections of particles. We can express these properties by taking *moments* of the velocity distribution function. The *n*-th moment of a probability distribution function is defined as

$$n\text{-th moment} = \int_{-\infty}^{\infty} v^n \, f(v) \, dv. \tag{2.9}$$

We first consider how we can get the number density, *n* from the velocity distribution function. Consider two systems of particles, in which all properties except density are the same, where we define system 1 to have a higher density than system 2 **(Figure 2.9a)**. In this case, the probability of finding a particle with velocity *v* in system 1 is always higher than the probability of finding a particle in system 2 with the same velocity, simply because there are more particles in system 1. Thus, the *magnitude* of the distribution function contains important information about the number of particles contained within a volume of space, irrespective of the velocity. In fact, the density is the zero-th moment of the velocity distribution function.

$$n \equiv \int_{-\infty}^{\infty} f(\mathbf{v}) d\mathbf{v} \tag{2.10}$$

In one dimension, this corresponds to the area under curve that defines the VDF, while in three dimensions this is a four-dimensional volume. Another property we might care about the the *average velocity* of a collection of particles. Considering 3-dimensional space with x-, y-, and x-coordinates, we can write the average velocity in the x-direction, u_x, as the first moment of the *normalized* velocity distribution function $g(\mathbf{v})$:
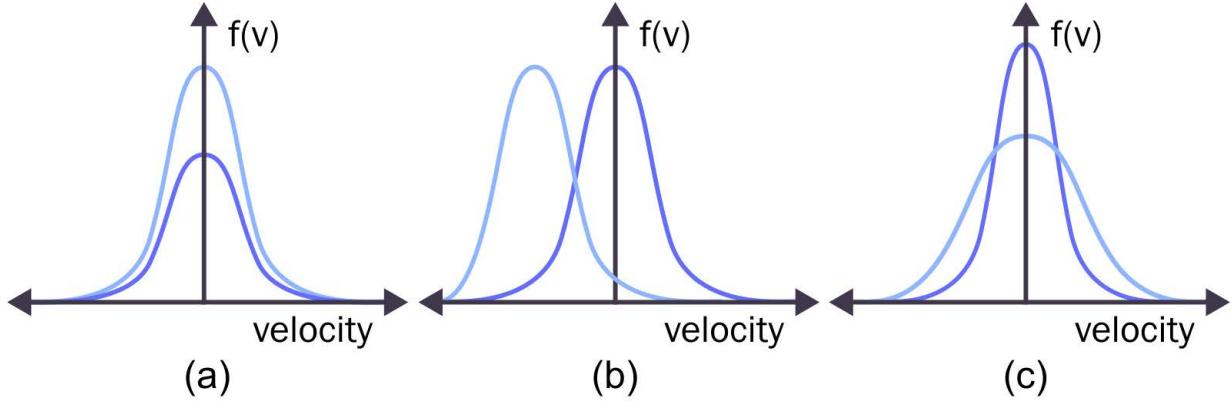
Figure 2.9: Examples of velocity distribution functions with different (a) zero-th moments (i.e. density), (b) first moments (i.e. mean velocity), (c) second moments (i.e. temperature).

$$u_x \equiv \int_{-\infty}^{\infty} v_x g(\mathbf{v}) d\mathbf{v} \tag{2.11}$$

The normalized VDF is just the regular VDF divided by the number density, $g(\mathbf{v}) \equiv f(\mathbf{v})/n$ such that now the "area" under the curve is one. We can write equivalent expressions to Eq. 2.11 to get the y- and z-components of the mean velocity. Intuitively, the mean velocity is the point around which the velocity distribution function is centered. In **Figure 2.9b**, we show two distributions which are identical except for their mean velocity, showing how the location of the mean changes with changing velocity.

Next, we can write the mean particle kinetic energy ϵ in terms of the second moment of the normalized VDF:

$$\epsilon \equiv m \int_{-\infty}^{\infty} \mathbf{v}^2 g(\mathbf{v}) d\mathbf{v}. \tag{2.12}$$

For monatomic species like xenon atoms or electrons, which have no vibrational or rotational energy modes, the total energy is just the sum of the directed and random (i.e. thermal) kinetic energies, or

$$\epsilon = \frac{3}{2} k_B T + \frac{1}{2} m |\mathbf{u}|^2. \tag{2.13}$$

The temperature of a gas corresponds to the width of its velocity distribution function. In **Figure 2.9c** we show two distribution functions which only differ in their second moment, i.e. temperature. At low temperatures, all particles have velocities very close to the mean speed of the population, while at high temperatures particles may have speeds which differ substantially from the mean speed. Consider moving an ice cube with some velocity. All of the particles move with the center of mass of the ice cube. Compare this to piping steam through a tube with the same velocity. While the center of mass velocity might be the same, the water molecules in the steam can spread out ahead of and behind the center of mass due to their differing velocities.

We can continue taking moments of the distribution function past this point. The third moment encodes information about the *skewness* of the distribution, i.e. how asymmetric about the mean the distribution is, while the fourth moment corresponds to the *kurtosis*, which measures how "fat" the tails of the distribution function are compared to the peak. The physical interpretation of these higher moments becomes more difficult, and we typically only employ the zero-th, first, and second moments.

Just as we can take the moments of the distribution function, we can also take moments of the entire Boltzmann equation. If we do this for the equilibrium distribution function, we obtain the fluid equations of motion. This equilibrium is described by the Maxwell-Boltzmann distribution **(Figure 2.10)**, which is given by:

$$f(\mathbf{v}) = n \left(\frac{m}{2\pi k_B T_e} \right)^{\frac{3}{2}} \exp \left(-\frac{m}{2} \frac{(\mathbf{v} - \mathbf{u})^2}{k_B T_e} \right) \tag{2.14}$$

This distribution, commonly referred to as a *Maxwellian* distribution, is a normal distribution (i.e. a bell curve), whose center is given by the mean velocity **u** and whose variance is $k_B T/m$. Using a process called *Chapman-Enskog expansion*[45] (the details of which are out of the scope of this thesis), we can take moments of the Boltzmann equation, assuming only small perturbations from the Maxwellian distribution. In doing this, we lose detailed information about the velocity distribution function but are rewarded with
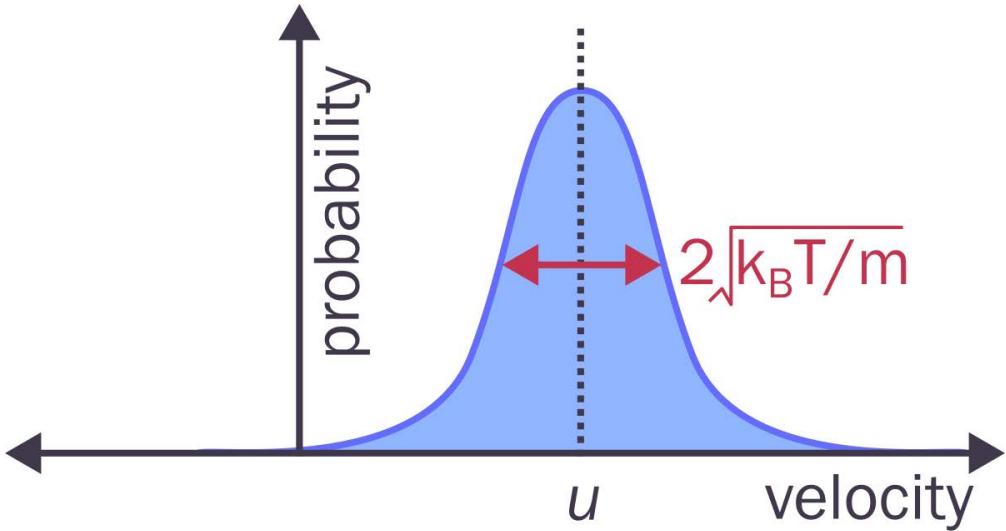
Figure 2.10: An Maxwell-Boltzmann distribution. The standard deviation is given by $\sqrt{k_B T/m}$.

increased computational tractability. For a neutral fluid, we typically employ the BGK collision operator, while for a plasma, we use the Landau operator. For a species s, where s could be electrons, neutrals, ions, or other constituents of a plasma, we can first take the zero-th moment of the Boltzmann equation and arrive at the *continuity equation*, which is

$$\frac{\partial n_s}{\partial t} + \nabla \cdot (n_s \mathbf{u}_s) = \dot{n}_s. \tag{2.15}$$

The continuity equation describes the conservation of mass within a small volume of space. The first term on the left side is the change in density of the fluid over time in that volume. The second term represents the net flux of fluid into or out of that volume, and the last term, $\dot{n}_s$ denotes the creation or destruction of species s due to ionization, recombination, or chemical reactions. Taken together, this equation states that the rate of change of the amount of a substance in a volume is equal to the amount that enters or leaves it, plus the amount that is added due to various reaction processes.

We can next take the first moment of the Boltzmann equation to obtain the *momentum equation*. This is a vector equation describing the change in momentum in each of the coordinate dimensions (i.e. x, y, and z). It is given by

$$\frac{\partial n_s \mathbf{u}_s}{\partial t} + \nabla \cdot (n_s \mathbf{u}_s \mathbf{u}_s + p_s/m) = \frac{q_s n}{m_s} (\mathbf{E} + \mathbf{u}_s \times \mathbf{B}) - \sum_{s'} n_{s'} \nu_{s,s'} (\mathbf{u_s} - \mathbf{u_{s'}}). \tag{2.16}$$

Just as the continuity equation describes the conservation of mass, the above equation(s) describe the conservation of momentum of a fluid. The first two terms denote the change in the amount of momentum in a volume over time and the flux of momentum into and out of the volume. Note that the second term contains a contribution from the gradient of the fluid pressure, $p_s = n_s k_B T_s$. Physically, this stems from the fact that a difference in pressure across a region of space will impart a net force, and thus momentum, to the fluid. The first term on the right-hand side is the *Lorentz force*, which represents the force on a charged fluid due to electromagnetic effects, while the last term denotes the the transfer of momentum due to collisions between species s and all other species (denoted s'). In the above equation, we have neglected the effects of viscosity on the fluid, as plasmas are typically quite diffuse and thus not strongly impacted by viscosity. In classical fluid mechanics, a variant of the above set of equations (including viscosity but not the Lorentz force or collisional terms) form the famous *Navier Stokes Equations*.

Finally, taking the second moment of the Boltzmann equation, we get the *energy equation*, given by

$$\frac{\partial n_s \epsilon_s}{\partial t} + \nabla \cdot \left(\frac{5}{3} n_s \epsilon_s \mathbf{u_s} + \mathbf{q}_s \right) = \mathbf{j}_s \cdot \mathbf{E} + S_{coll} \tag{2.17}$$

In **Equation 2.17**, $\mathbf{q}_s \equiv \kappa \nabla T_s$ is the conductive heat flux vector, where κ is the thermal conductivity tensor, $\mathbf{j}_s \equiv q_s n_s \mathbf{u}_s$ is the current density vector of species s and S_{coll} is the rate

of energy loss due to collisions with other species, as well as to ionization and electronic excitation. The $\mathbf{j} \cdot \mathbf{E}$ term is known as the *Joule heating* or *Ohmic heating* term, and captures the phenomenon that sending a current through a material will produce heat. The amount of heat produced per unit time scales with the amount of current passed through the plasma and the magnitude of the voltage difference (and thus electric field) across the plasma, giving the familiar $P = IV$ scaling. As Hall thrusters feature a strong electric field combined with a net cross-field electron current, this term is responsible for the vast majority of the electron heating in these devices.

In writing the above equations, we have implicitly assumed that the distribution function has the same width in each coordinate direction. As a result, we can consider only a single temperature which applies to all three spatial dimensions.[c] More generally, the temperature might be different in different directions, such as along versus across the magnetic field. A property which varies depending on direction is called *anisotropic*. In saying the temperature does not do this, we are assuming that the temperature, and therefore the pressure, are *isotropic*.

2.5   Hall thruster physics

With all of the fundamentals out of the way, we can return to Hall thrusters. Like the glow discharges discussed in **Section 2.2**, and as illustrated in **Figure 2.11**, Hall thrusters are DC discharges, in which a voltage is applied between an anode and a cathode to produce a plasma. The application of a magnetic field in the radial direction, perpendicular to the applied electric field, causes the particles to undergo the $E \times B$ drift around the discharge channel. Normally, this drift induces the electrons and ions to move with the same speed in the same direction, and thus no net current is produced. However, the strength of the magnetic field is tuned so that the electrons are confined in gyro-orbits around the magnetic

---

[c]Different species of particle may still have different temperatures. In Hall thrusters, the electrons are typically much hotter an the ions.
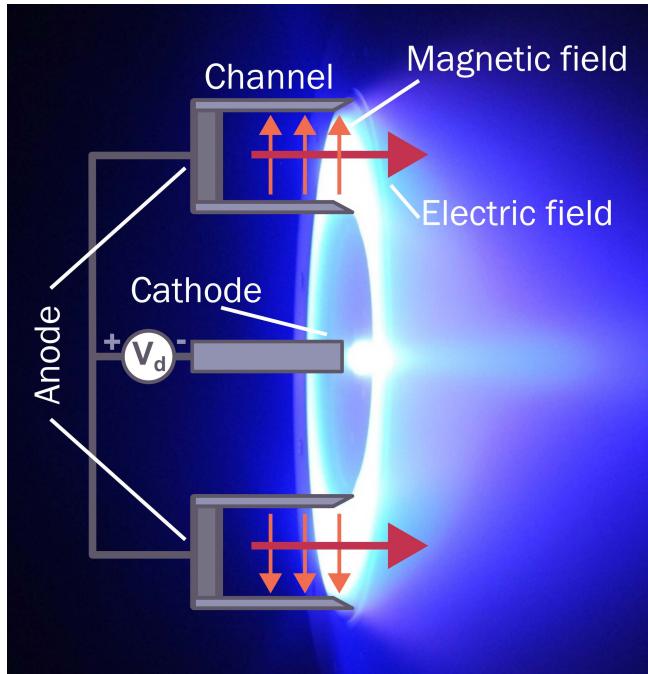
Figure 2.11: Diagram of the operation Hall thruster overlaid on a photograph of the H9 Hall thruster operating at 300 V and 15 A on krypton. Photo credit Leanne Su and the Plasmadynamics and Electric Propulsion Laboratory.

field lines, while the ions are unmagnetized and free to follow the electric field out of the thruster. Put another way, the Larmor radii of the electrons are on the order of a millimeter, while those of the ions are on the order of tens of centimeters to a meter. The ions thus remain in the region of high electric and magnetic fields for a very short period of time and do not undergo the $E \times B$ drift with the electrons. A net current around the device channel, called a Hall current, is thus produced.

To model the physics of Hall thrusters, we would like to use a fluid description for the electrons. When is such a description valid? As we saw in the preceding section, when modeling a plasma as a fluid, we are assuming that the particle velocity distribution function is Maxwellian, or at least very close to Maxwellian. As the Maxwellian distribution function is the equilibrium state of the Boltzmann equation, using a fluid description of the plasma is the same thing as assuming that the plasma is in *local thermodynamic equilibrium*. For this condition to be true, we require that particles collide with each other frequently enough

42

that any perturbations from the equilibrium distribution function are rapidly smoothed out. To illustrate how collisions can return a distribution function to equilibrium, consider a game of pool or billiards **(Figure 2.12)**. Before the cue ball is hit, the balls are in equilibrium, i.e. not moving at all. This means they have a mean velocity of zero, as well a temperature of zero since all of the balls have the same speed. At the moment we strike the cue ball, the velocity distribution function becomes non-Maxwellian. Most of the balls (particles) have zero speed, while the cue ball has a significant amount of speed. This state of affairs continues until the moment cue ball collides with another ball. At that instant, the cue ball transfers some of its energy from into the ball it collides with. Both of these balls then go on to collide with other balls, transferring more and more of their energy into the other balls, and so on. In the absence of friction with the pool table and the walls, the pool balls would quickly assume an equilibrium state in which all of them are moving a small amount.

The consequence of this phenomenon is that we are only justified in treating a system as a fluid if there are sufficient collisions that the distribution function is near equilibrium. To quantify this, we introduce the *mean free path* (λ_{mfp}), which is the average distance a particle travels before colliding with another particle. The mean free path is defined as

$$\lambda_{mfp} = \frac{1}{\sqrt{2}\sigma n} \quad \text{m,} \tag{2.18}$$

where n is the particle number density and σ is the average collisional *cross-section*, which is the effective cross-sectional "area" of particles in a system, measured in square meters. Physically, this expression shows that if the particles are larger, or the density is higher, then collisions are more likely and individual particles travel a shorter distance before colliding with their neighbors. To assess whether the mean free path for a given system is "short enough" to treat the system as a fluid, we employ the *Knudsen number* (*Kn*), defined as
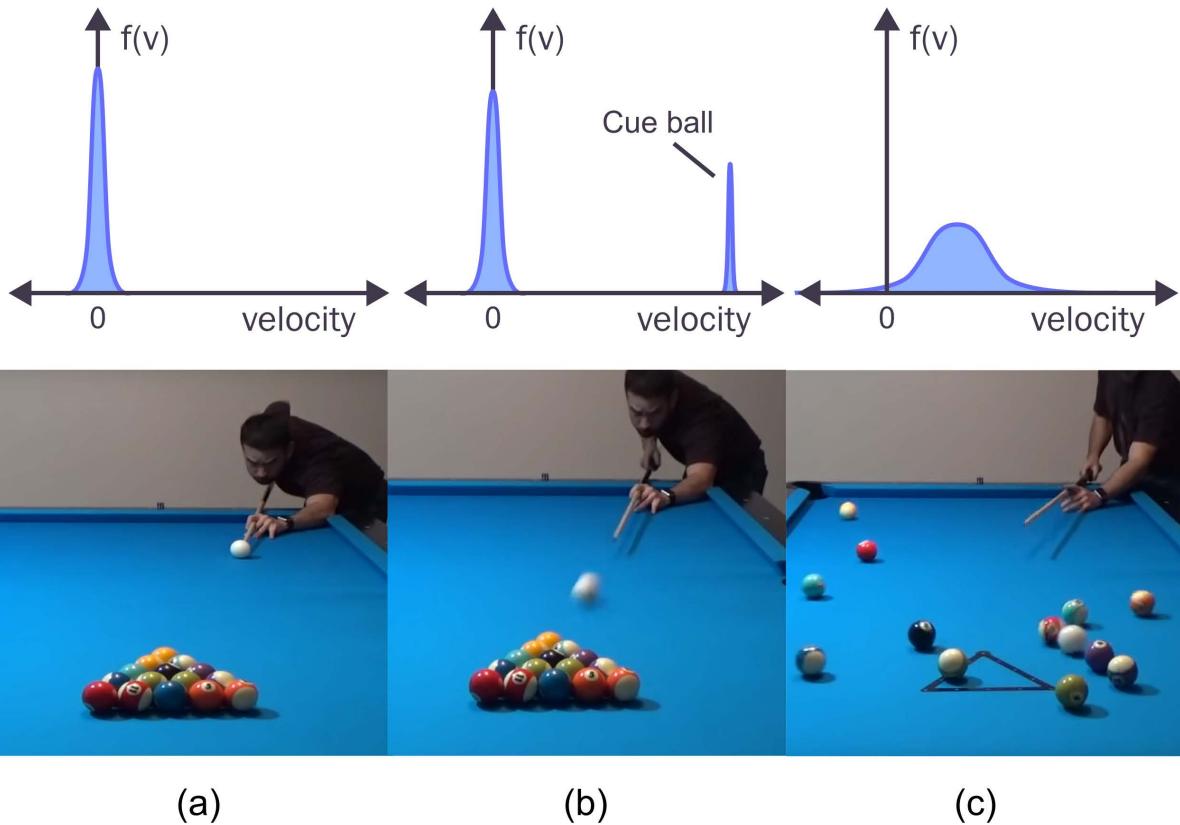
Figure 2.12: Collisional equilibration of a non-Maxwellian velocity distribution function, as illustrated by a break in a game of pool or billiards. In (a), all particles are in equilibrium with nearly zero velocity. In (b), the cue ball has a much higher velocity than the other balls, so the distribution function is non-equilibrium. In (c), collisions between the balls have randomized their velocities and restored equilibrium to the VDF. Photo credit: Lil' Chris on Youtube.

$$Kn = \frac{\lambda_{mfp}}{L}, \tag{2.19}$$

where L is the characteristic "size" of the system. If $Kn < 0.01$, then particles undergo $O(100)$ collisions before transiting the length of the system, and we are justified in treating that system as a fluid. This regime is called *continuum flow*. On the other hand, if $Kn > 10$, then the particle may not collide with anything at all before leaving, and we would need a more detailed kinetic description of our system. This regime is called *free molecular flow*. Between these extremes, we have *transitional flow*, where special techniques may be needed.

Where do Hall thrusters fit on this spectrum? There are three species we would need to consider: neutral atoms, ions, and electrons. Typical neutral densities in a Hall thruster are around 10^{18} m^{-3}, and the diameter[d] of a xenon atom is about 396 picometers, or 3.96×10^{-10} meters. The cross section can then be computed as $\pi d^2/4$. This gives a cross section of around 1.5×10^{-19} square meters. Plugging these numbers into **Equation 2.18**, we get a mean free path for neutral atoms of around four meters. As the neutrals do not circuit the channel, we take the relevant characteristic "size" to be the length of the discharge channel, which is on the order of 1 to 10 centimeters. This gives us a Knudsen number is between 40 and 400. Thus, the neutral atoms in a Hall thruster are typically free molecular and cannot be treated as a fluid.

What about ions? Since ions are positively-charged, they will repel each other before they can collide in the conventional sense. However, these electrostatic repulsion events end up working in much the same way as classical hard-sphere collisions. The effective collision diameter of a population of ions with temperature T_i (in electron-Volts) is

$$d_i = \frac{Zq_e}{4\pi\epsilon_0 T_i}, \tag{2.20}$$

---

[d]Atoms are not really billiard ball-like hard spheres with well-defined diameters, but it suffices for our purposes to treat them as such.

where Z is the ion charge state (i.e. $Z = 1$ if the ion is singly-charged), q_e is the fundamental charge, and ϵ_0 is the permittivity of free space. For ions with a temperature of 500 Kelvins, this gives us a diameter of around 33 nanometers. The cross section is then 9×10^{-16} square meters, giving a Knudsen number between 0.01 and 0.1, which is on the lower end of the transitional flow regime. If the ions are hotter (say 5000 K) the Knudsen number may grow to between 0.8 and 8, which is at the upper end of this range. Other changes in parameters, such as increased density, could push the ions firmly into the fluid regime, while still others could tip them over the line into free molecular flow. Experimental measurements of Hall thruster ion dynamics suggest they are largely Maxwellian,[46] except near the location of the peak electric field. Here, in the aptly-named *acceleration zone*, the ions experience the majority of their acceleration over a very short length (around 10 mm). Taking this as our characteristic scale, the ions in this location would be fully free-molecular. Owing to the transitional nature of the ion dynamics in Hall thrusters, many Hall thruster codes treat the ions kinetically[47], while others[16] treat them as a fluid. The code we employ for this thesis, Hall2De, is one of the latter codes, and treats the ions as a continuum fluid rather than a free molecular flow. We describe Hall2De in more detail in **Chapter 4**.

Following a similar procedure, the mean free path for electrons can be calculated as somewhere between 0.5 and 2 meters, assuming a temperature of between 10 and 30 eV. This would seem to put them firmly in the free molecular category. However, as stated earlier, electrons circle the discharge channel azimuthally as part of the Hall current. The distance they traverse from when they enter the drift to when finally reach the anode can be calculated from their axial drift speed, their azimuthal drift speed, and the channel length. Assuming an axial drift speed of 10,000 meters per second, the electrons will take between 1 and 10 microseconds to traverse the channel. Assuming they circle the discharge channel at around 10^6 meters per second, this means they travel between 1 and 10 meters in that time. This puts their Knudsen number between 0.05 and 2, which

is in the middle of the transitional regime. Thus, while electrons are not immune from non-equilibrium or kinetic effects, it may be reasonable to treat them as a fluid.

Given the previously-discussed challenges with treating electrons kinetically, this is a fortunate result. However, if we were to apply the fluid momentum equation **(Equation 2.16)** directly, we would encounter some difficulties. To solve such an equation computationally, we typically *discretize* it, meaning that we break the region of time and space we want to simulate into a bunch of small chunks. Consider solving the momentum equation on a domain which extends from $x = 0$ to $x = L$. We would break this domain into *cells* with size Δx. Starting from some time t_0, we would then advance the solution in time by *timesteps* of size Δt until we reach our desired final time T. This makes it possible to solve coupled systems of partial differential equations in a reasonable amount of time, but introduces additional difficulties. These include issues of numerical *consistency* (by discretizing the equations, are we sure we are still solving the same problem?), *accuracy* (how small do we need to make Δx and Δt in order to achieve a desired level of accuracy?) and *stability* (how big can we make Δt without the solution blowing up?). While the former two are worthwhile challenges in their own right, it is last condition—numerical stability— which will cause us problems in this instance. The famous Courant Friedrichs Lewy (CFL) condition states that in order for equations of the type we want to solve to be stable, the timestep must obey the following:

$$a\frac{\Delta t}{\Delta x} \lesssim 1 \qquad (2.21)$$

Here, a is the characteristic speed at which information propagates in the flow, and the quantity $a\Delta t/\Delta x$ is the *CFL* or *Courant* number. In a fluid, information can either move with the bulk flow velocity u or as a pressure wave with velocity equal to the speed of sound a. The maximum speed of information propagation is thus (roughly) $u + a$. We can apply this condition to ions and electrons individually to determine the largest timestep we would be allowed to take in our simulation. We assume that $u \sim 20{,}000$ and $\Delta x \sim 1$ mm, while

the ion sound speed is typically around 1,000 or 2,000 m/s, depending on the temperature. This gives us an allowable timestep of approximately 5×10^{-8} seconds. While the electron speed in the axial direction is roughly the same as the ion speed, the electron sound speed is closer to 2,000,000 m/s. This stems from the fact that electrons are very light, so a small amount of energy is able to accelerate them to very high speeds. The maximum allowable electron timestep is then closer to 5×10^{-10} seconds, and can be lower. This means that we need to run our simulation for 100 times longer than if we just considered ions. While some authors do make this sacrifice in the name of increased accuracy,[48] the majority of Hall thruster simulations make additional assumptions to allow us to use timesteps closer to the ion timestep.

Since electrons move so quickly, we could assume that (in the presence of some change in local electrical or plasma properties) the electrons are able to rearrange themselves and reach steady state before the much heavier[e] ions have a chance to notice the change. We can then set terms relating to the time variation of the electrons to zero and consider only their steady-state behavior. This is called the *inertia-less* assumption, as it implies that —due to their low mass— the electrons' inertia has little relevance to the rest of the flow, especially compared to larger factors like the electric and magnetic fields, as well as changes in plasma density and temperature. With this assumption, the electron momentum equation **(Equation 2.16)** reduces to

$$-\frac{q_e n_e}{m_e} \left(\mathbf{E} + \mathbf{u}_e \times \mathbf{B}\right) - \frac{\nabla p_e}{m_e} - \sum_{s'} n_{s'} \nu_{e,s'} (\mathbf{u_e} - \mathbf{u_{s'}}) = 0. \tag{2.22}$$

For simplicity, we can then make the assumption that, since the electrons (e) are much faster than neutrals (n) or ions (i), we can assume $(\mathbf{u}_e - u_n) \approx (\mathbf{u}_e - u_i) \approx \mathbf{u}_e$. We note that this assumption is not critical to making the electron motion computationally tractable, but it simplifies the algebra. With that caveat, the electron momentum equation now becomes:

---

[e]The ratio between the mass of a xenon atom and the mass of an electron is about 240,000.

$$q_e n_e \left(\mathbf{E} + \mathbf{u}_e \times \mathbf{B} \right) + \nabla p_e + m_e n_e \mathbf{u}_e \nu_e = 0, \tag{2.23}$$

where $\nu_e = \nu_{ei} + \nu_{en}$ is the total electron collision frequency, including collisions with both ions and neutral atoms. By eliminating the time-dependent terms, this equation has become *algebraic*, meaning we can use it to solve for electron velocity at a given time, provided we know the electric and magnetic fields, electron density, and electron pressure. As this equation has no timestep constraints, we can use the ion timestep in our simulations rather than the electron timestep, and get simulations which converge in a timely manner. **Equation 2.23** is often called the *generalized Ohm's law*, for reasons we will discuss in a moment.

At no point in this discussion have we mentioned the *anomalous* electron transport. Experiments show that electrons move across the magnetic field lines much faster than **Equation 2.23** predicts. This means that at some point on the road from the Boltzmann equation **(Equation 2.7)** to the simplified algebraic description above, we neglected some critical physics which gives rise to enhanced cross-field electron transport. Perhaps the electron distribution function is not truly Maxwellian, or maybe the electron inertia enhances transport. To get a simulation of a Hall thruster which self-consistently incorporates this effect would require us to increase the fidelity of the model, at the cost of increased computational time. If we want to retain the time benefits of the fluid approach, we could instead include a *model* of the anomalous transport into our simplified description. Ideally, such a model would capture the correct scaling of the transport with other local fluid properties such as density and temperature, while eliding lower-level details which do not impact the behavior of the bulk plasma.

2.6 Anomalous electron transport in a fluid plasma

Before we address how anomalous processes enhance the motion of electrons across the magnetic field lines, we first discuss how classical effects, like collisions, affect this motion. We first decompose the vector form of the generalized Ohm's law **(Equation 2.23)** into its vector components. We choose a coordinate system aligned with the thruster magnetic field, allowing us to write equations for the electron motion in the direction along (parallel, $\parallel$ to) and across (perpendicular, $\perp$, to) the applied magnetic field, as well as around the device channel in the azimuthal (θ) direction. After such a decomposition, the electron momentum balance becomes

$$q_e n_e E_\parallel + \nabla_\parallel p_e + m_e n_e u_{e,\parallel} \nu_e = 0 \tag{2.24}$$

$$q_e n_e \left(E_\perp + u_{e\theta} B\right) + \nabla_\perp p_e + m_e n_e u_{e,\perp} \nu_e = 0 \tag{2.25}$$

$$q_e n_e \left(E_\theta - u_{e\perp} B\right) + \nabla_\theta p_e + m_e n_e u_{e,\theta} \nu_e = 0, \tag{2.26}$$

where we have made use of the fact that in a field-aligned coordinate system, $\mathbf{B} = B\ \hat{\parallel}$. We next employ the assumption of *axisymmetry*. Under this assumption, properties do not vary in the θ-direction, so all gradients or derivatives in this direction go to zero. This eliminates not only the pressure gradient in **Equation 2.26**, but also the electric field, as it is defined as the gradient of the electrostatic potential. The azimuthal component of the electron momentum balance thus reduces to:

$$u_{e,\theta} = \Omega_e u_{e\perp} \tag{2.27}$$

This gives us an explicit expression for the electron azimuthal velocity in terms of the perpendicular velocity and the electron Hall parameter, $\Omega_e = m_e \nu_e / q_e B$. We can then substitute **Equation. 2.26** back into **Equation 2.25** to find an expression for the electron

velocity in the cross-field/perpendicular direction. After some algebra, we arrive at:

$$u_{e,\perp} = -\frac{1}{1 + \Omega_e^2} \frac{q_e}{m_e \nu_e} \left(E_\perp + \frac{\nabla_\perp p_e}{q_e n_e} \right) \tag{2.28}$$

Multiplying through by $-q_e n_e$, we can get an expression for the cross-field electron current density, $j_{e,\perp} = -q_e n_e u_{e,\perp}$. We can then follow a similar procedure for the parallel component of the equation to arrive at a set of equations for the electron current density vector in terms of the other local plasma properties.

$$j_{e,\parallel} = \frac{q_e^2 n_e}{m_e \nu_e} \left(E_\parallel + \frac{\nabla_\parallel p_e}{q_e n_e} \right) \tag{2.29}$$

$$j_{e,\perp} = \frac{1}{1 + \Omega_e^2} \frac{q_e^2 n_e}{m_e \nu_e} \left(E_\perp + \frac{\nabla_\perp p_e}{q_e n_e} \right) \tag{2.30}$$

$$j_{e,\theta} = \Omega_e j_{e,\perp} \tag{2.31}$$

With the above form of the expressions, we can see why Equation 2.23 is called the *generalized Ohm's law*. The classical Ohm's law for circuits is given by $I = V/R$, where I is the current, V is the voltage, and R is the resistance. Above, we have an expression for current density, which is proportional to a difference in voltage (the electric field) and inversely proportional to the plasma resistivity, η, which is defined as $\eta = m_e \nu_e / q_e^2 n_e$. We can see that from **Equation 2.30** that if the Hall parameter is large, then the resistivity in the perpendicular direction is much larger than that in the parallel direction. This captures our intuition that the magnetic field should act to impede or resist electron motion.

From the above equations, we can see that increasing the electric field or electron pressure gradient should increase the electron current density, both along and across field lines. Less intuitive is the effect of collisions. From the expression for resistivity, it seems like collisions should increase the resistance and thus reduce the electron current density. Recall, however, that the Hall parameter Ω_e depends inversely on the electron collision frequency. We can more clearly elucidate the effect of the collisions on the cross-

field transport by invoking one final assumption. If the electrons are magnetized, then the electron collision frequency is much less than the electron cyclotron frequency. Recall, too, that this implies that the electrons are able to complete many orbits around the magnetic field lines before they encounter a particle to collide with. This is equivalent to stating that $\Omega_e \gg 1$. With this assumption, we can reduce our components of Ohm's law one final time:

$$j_{e,\parallel} = \frac{q_e^2 n_e}{m_e \nu_e} \left(E_\parallel + \frac{\nabla_\parallel p_e}{q_e n_e} \right) \tag{2.32}$$

$$j_{e,\perp} = \frac{q_e n_e \nu_e}{\omega_{ce} B} \left(E_\perp + \frac{\nabla_\perp p_e}{q_e n_e} \right) \tag{2.33}$$

$$j_{e,\theta} = \Omega_e j_{e,\perp} \tag{2.34}$$

Now, we can see that in the direction aligned with the magnetic field (the $\parallel$-direction), collisions increase the resistance of the plasma, and decrease the electron current density. However, across the magnetic field lines, collisions *decrease* the plasma resistivity and increase the electron current. This stems from the fact that since collisions randomize particle velocities, they work against the confining effect of the magnetic fields. Incidentally, the quantity $\nu_e / \omega_{ce} B$ is known as the cross-field *electron mobility*, denoted $\mu_{e,\perp}$. We are now ready to include anomalous electron transport in this framework. If electrons are to be moved in some manner across the magnetic field lines, there must be an additional force in the system which we have not accounted for. The nature and scaling of this force is unknown. We can incorporate this anomalous force into the vector Ohm's law by introducing an additional term, $\mathbf{F}_{an}$, so

$$q_e n_e \left(\mathbf{E} + \mathbf{u}_e \times \mathbf{B} \right) + \nabla p_e + m_e n_e \mathbf{u}_e \nu_e - \mathbf{F}_{an} = 0. \tag{2.35}$$

In the absence of any first-principles knowledge of how this force should act, we can make some guesses. First, the force should not violate the laws of physics. In our system,

this means that the force should not be able to accelerate the electrons to arbitrarily high velocities. Hall thrusters do not emit large amount of x-rays[f], so we know that the electrons are not moving close to the speed of light. Typically, forces in nature work against motion. Therefore, we might assume this "anomalous" force which enhances electron mobility across the magnetic field lines works like a drag force. After all, we just saw that collisional drag between electrons and the heavier species (i.e. neutrals and ions) enhances, rather than impedes, electron motion across the magnetic field lines. If we make this assumption, than $\mathbf{F}_{an}$ becomes

$$\mathbf{F}_{an} = -m_e n_e \mathbf{u}_e \nu_{an}. \tag{2.36}$$

This functional form states that the anomalous force should oppose the electron velocity vector, $\mathbf{u}_e$. Additionally, as this force term scales with the plasma density, n_e, we should get a smaller anomalous drag if we have less plasma. Just as in Eq. 2.23, the collisional drag also scales with the effective electron collision frequency, ν_{an}, which we call the "anomalous" collision frequency. This does not imply that the anomalous effects result from collisions, only that we are representing their contribution to the overall plasma solution the form of a collision frequency.

Applying this choice of scaling is convenient for modeling purposes, as we can still directly apply **Equations 2.32-2.27** to describe the motion of electrons across the magnetic field lines. We only need to include the anomalous collision frequency in ν_e, such that

$$\nu_e = \nu_{e,classical} + \nu_{an}. \tag{2.37}$$

---

[f]Electrons deflected by a magnetic field will emit photons with energy that scales with their velocity. For electrons that are near the speed of light, this is called *synchrotron radiation* and typically ranges between high ultraviolet and gamma in wavelength.

Here, $\nu_{e,classical}$ is the sum of the electron-neutral and electron-ion collision frequencies.[g] While convenient, positing the existence of the anomalous collision frequency leaves us with a problem. Namely, we have no clue how it should scale. Mathematically speaking, we now have one more variable than we have equations. We need to introduce an additional equation which describes the functional dependence of ν_{an} on the rest of the plasma in order to *close* the system of equations. We call this state of affairs a *closure problem*. A model for ν_{an} which depends on the other plasma properties would then be termed a *closure model*. The goal of this thesis is thus to develop and test closure models for the anomalous electron collision frequency in Hall thrusters.

This situation is analogous to the problem of turbulence modeling in classical fluids, famously one of the great unsolved problems in physics. In these systems, it is often assumed that the effect of small-scale turbulent fluctuations manifests as an effective "eddy" viscosity, which acts on top of the normal fluid viscosity. The goal of certain types[h] of turbulence modeling is then to find a closure model for the eddy viscosity in terms of the bulk fluid properties, such as density, temperature, and velocity.

Turbulence models have not solved forever the problem of turbulence, but the efforts toward this goal have made it possible to apply predictive fluid simulations to a wide variety of fluid flow problems, such as that over an aircraft wing, through a nozzle, or in a pipe. This has in turn made simulation an integral part of the design process across all fields of fluid mechanics.

---

[g]Electron-electron collisions do not alter the bulk momentum of the electron fluid and thus do not appear in this expression.

[h]Such models are called, fittingly, *eddy viscosity models*, and the assumption of an effective eddy viscosity is called the *Boussinesq approximation*. Other forms of turbulence modeling include *large eddy simulation* (LES) in which the larger parts of the turbulence are directly simulated, while a *subgrid-scale model* is needed to describe the smaller scales, and direct numerical simulation (DNS), in which all scales of the turbulence are resolved. The analogy of DNS in Hall thruster modeling would be full particle-in-cell simulation of both electrons and ions.

2.7 Goals of Hall thruster modeling and simulation

Given their long heritage and history of successful flights, one might wonder whether modeling and simulation need to play a significant role in the Hall thruster design process. Simulation can alleviate several challenges faced during the testing and qualification of Hall thrusters. The first of these is thruster lifetime. The energetic plasmas present in Hall thrusters discharges will gradually erode spacecraft surfaces.[49] After enough erosion, the thruster may cease to function. As these thrusters can run for potentially tens of thousands of hours,[50] it can become extremely expensive to demonstrate thruster life in a ground facility. By predicting the erosion rate of thruster surfaces, models can accelerate this process.

Second, it is known that Hall thrusters perform differently in a vacuum chamber on the ground than in space.[51,52,46] Accurately characterizing how these "facility effects" alter a thruster's performance during the transition from ground testing to flight is a major goal of electric propulsion research. The role of facility effects grows as the thruster power increases, and it may not be possible to test high power thrusters like those we might use to send astronauts to Mars on the ground. Better modeling of the physical phenomena underlying the ground to flight transition can enable higher confidence in the performance of a system on orbit.

Lastly, predictive modeling and simulation tools would enable more rapid design of Hall thrusters to meet new operating regimes. These include thrusters that operate at high power,[53] low power, and using alternative propellants. Ideally, given a geometry, operating condition, and magnetic field configuration, we could predict the performance (i.e. thrust, discharge current, and efficiency) and impact of the plasma on surrounding spacecraft surfaces. Currently, however, evaluating the performance of a thruster after a design change requires a thruster to be built and tested, which can become expensive. Predictive simulation could accelerate this process by allowing new thruster designs to be

evaluated on the computer before they are built. This would enable thruster designs that are better-optimized for a given operating regime.

The first two challenges do not necessarily require the resolution of the problem of anomalous electron transport. Given data from a thruster, we can calibrate simulations to match experiment by adjusting the spatial variation of the anomalous collision frequency. We can then use these calibrated simulations to assess the erosion rate of different thruster surfaces, or determine the flux of plasma to spacecraft surfaces.[6] However, this process only works if there is data available against which we can calibrate our simulations. This requires the thruster to be built, which limits the use of this type of empirically-inferred anomalous transport model in the design of novel Hall thrusters. Therefore, while predictive anomalous transport models cannot resolve all challenges inherent to Hall thruster modeling, they are the single biggest piece of missing physics which limit the use of simulations in thruster design and qualification.

In the next chapter, we explore how we can derive and test models of the effective anomalous collision frequency. We examine models that have been proposed to date in the scientific literature, and show how, by tweaking certain assumptions, a single model can be expanded into an entire family of models.

CHAPTER 3

Models of Anomalous Transport

> Basic research is what I'm doing when I don't know what I am doing.

> Wernher von Braun

In the last chapter, we gave an overview of the assumptions underlying a fluid description of a plasma. We then showed how we can account for anomalous transport in a fluid framework by including an effective anomalous collision frequency, ν_{an}. We discussed how this leads to a problem of closure, which must be resolved by relating ν_{an} to the other variables in the system. In this chapter, I introduce some of the various attempts to solve this problem found in the literature, and propose some new models of our own. In each case, I discuss the assumptions which each model relies upon, and illustrate the impacts of tweaking those assumptions.

3.1 Terminology

Before we discuss what models have been proposed for the anomalous collision frequency, I briefly introduce some terminology. In classical turbulence, models are typically classified by the number of differential equations needed to describe the effect of the turbulence on the mean flow. Thus, a *zero-equation* model (also-known as an

algebraic model) is one which has zero differential equations. Instead, the unknown quantity is an explicit or implicit function of the mean flow properties. A one-equation model, then, is a model in which a differential transport equation must be solved before the turbulent effects on the mean flow can be determined. This terminology continues to two-equation models and further. This work primarily focuses on zero-equation models of the anomalous electron collision frequency, but I will introduce one or two multi-equation models as well.

In addition to the number of equations, models may be characterized by the number of fit coefficients they require. All of the models examined in this work feature one or more parameter which must be adjusted to fit data. This is also a feature of classical turbulence models. These numbers typically represent the impact of higher-order physical phenomena which (hopefully) do not vary much with the parameters under study. As a result, these phenomena may be lumped together under constants, which are adjusted to fit available data. Ideally, these coefficients would be general and thus not need to be adjusted when changing thrusters or operating regimes. In practice, this is not usually the case. In general, the more coefficients that a model has, the better chance it has of fitting data, since the modeler has more ability to tweak and tune the coefficients to match a given situation. However, this usually comes at the cost of generalizablity, as it is easy to over-fit the coefficients to a particular dataset. In our case, this means that a model might only work for a single thruster or operating condition, and thus not be truly predictive. Additionally, the more coefficients a model has, the more expensive it is to calibrate these coefficients. Consider a model with a single coefficient. One easy way to tune this coefficient might be to run ten simulations, each with different values of this coefficient, and then pick the value which gave the best fit with experiment. With two coefficients, we would need to run 100 simulations to get the same resolution. With three coefficients, this balloons to 1000, and so on. A ten-coefficient model would require us to run ten billion simulations if we apply the same strategy that we used to tune the coefficients for the one-coefficient model. This

is one manifestation of the famous *curse of dimensionality*. For *d* coefficients, we have a
$d - dimensional$ parameter space to explore, and the cost scales with N^d, where *N* is the
number of grid points in each dimension. For this reason, we primarily focus on models
with no more than four coefficients. Later in this thesis, I will discuss a few approaches
that we tried to accelerate this calibration procedure.

Finally, I will often make a distinction between "empirical" and "self-consistent" models
of the anomalous collision frequency. This is more a difference in degree than kind.
A self-consistent model is one which is able to reproduce the effect of the anomalous
transport by its functional dependencies alone, and which is allowed to change as the
plasma changes. An empirical model, by contrast, may get the correct degree of transport
but must be tuned to match experiment. It may also not depend to the same degree on
local plasma properties. All of the models examined in this work are in some sense semi-
empirical, as they feature fit coefficients which have to be tuned to data. However, as
they also depend on the plasma, they offer a better chance of being predictive than if they
were fully empirical and had no basis in physics. Before discussing these "self-consistent"
models, we first describe the typical empirical methods by which anomalous transport has
traditionally been incorporated into simulations of Hall thrusters.
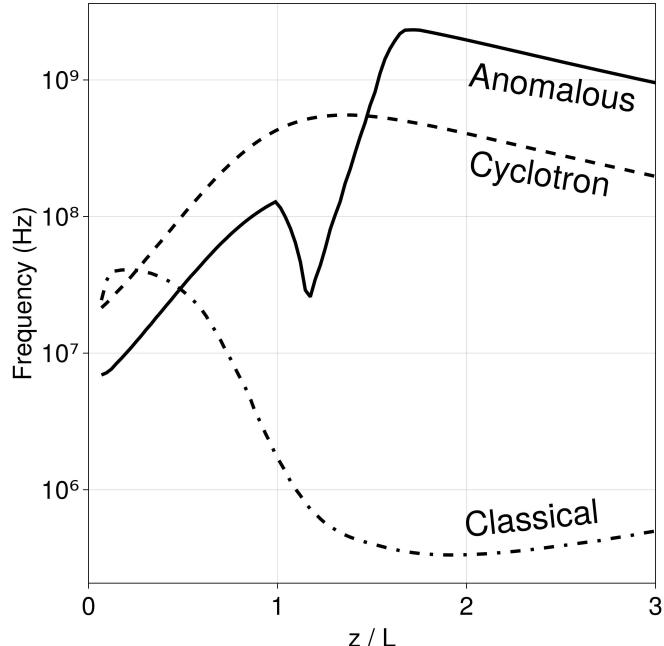
3.2   Empirical models

The most widely-used method of empirically specifying the anomalous collision frequency
is to assume that it is constant in time and with respect to the plasma properties but
may vary spatially. The spatial variation can in turn be adjusted until key aspects of
the simulation predictions, e.g. the acceleration of ions and discharge current, match
experiment. This has resulted in a canonical "shape" for ν_{an}. **Figure 3.1a** depicts an
example of this typical shape in a Hall thruster plotted as a function of distance from the
anode along channel centerline. We generated this result from numerical simulations of a
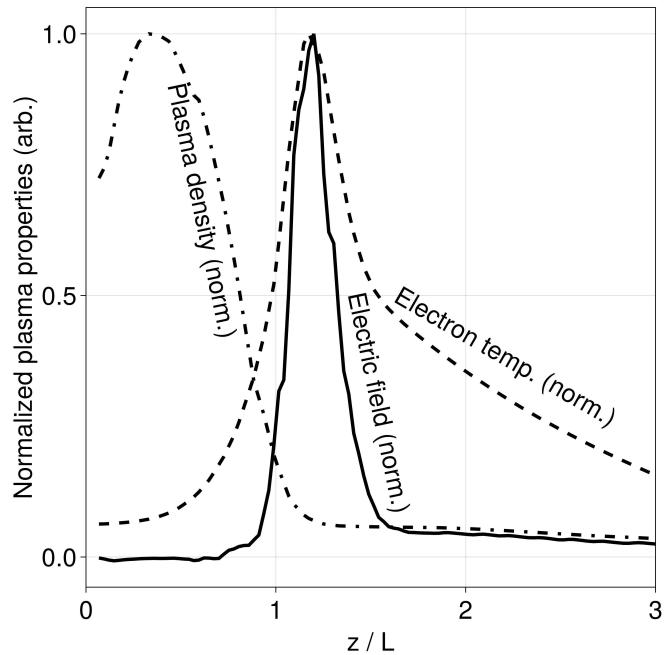
magnetically shielded Hall thruster operating with xenon at 300 V and 15 A. This simulation is described in more detail in **Chapter 5.** We also show in this result the classical collision frequency, ν_c, and electron cyclotron frequency $f_{ce} = \omega_{ce}/2\pi$ Hz. **Figure 3.1b** shows key plasma properties like the electric field strength, electron temperature, and plasma from this simulation.

The properties of the anomalous profile can be understood in the context of the trends in the plasma properties and background magnetic field. In the near-anode region ($z/L \approx 0$), the neutral density and thus the electron-neutral collision frequency are high, so $\nu_c \sim \nu_{an}$. Inside the discharge channel, the anomalous collision frequency ranges from around the same magnitude to an order of magnitude greater than the classical frequency. However, near the channel exit plane, in the region of peak magnetic field strength (and thus maximum cyclotron frequency), the collision frequency decreases by an order of magnitude. The corresponding increase in electron resistivity in this location results in strong electric fields and enhanced Ohmic heating. This is reflected by the profiles exhibited in **Figure 3.1b**. Without this transport barrier, the predicted gradients in plasma properties from simulation would be much more shallow than observed experimentally. Downstream of the exit plane, the anomalous collision frequency increases by up to two orders of magnitude to reach a maximum value in the near-plume region. It then declines with the magnetic field strength with increasing distance from the thruster exit plane.

Adjusting the spatial variation of ν_{an} to reflect these typical features has allowed fluid-based Hall thruster simulations to match experiment to a high degree of fidelity. In later parts of this work, we will refer to these tuned anomalous collision frequency curves as "empirical profiles," as they represent an empirically-inferred estimate of what the anomalous collision frequency needs to be in order to yield a simulation that matches data. Calibrated simulation results obtained via the use of such empirical profiles have been extensively leveraged for the design and qualification of Hall thrusters.[54,55] However, despite the success of these empirical profiles and the commonalities in their shapes,

(a)

(b)

Figure 3.1: (a) Example of an empirical anomalous collision frequency profile for a Hall thruster operating at 4.5 kW, compared to the classical and electron cyclotron frequencies. (b) Normalized plasma properties from a simulation run with the empirical profile in (a). Distance (x-axis) is referenced from the downstream surface of the anode, measured along the channel centerline, and normalized by the channel length (L).

61

they are not extensible. An empirical profile calibrated to work on one thruster will not work with a different thruster or even a different operating condition on the same thruster. For each new thruster or operational state, we instead must infer a new empirical profile from experimental measurements of the thruster.[56,53] This limits our ability to leverage empirical profiles for fully predictive simulations. In the next sections, we will examine different approaches toward this goal. In Chapter 5, we will return to a discussion of these empirical profiles and how they might (or might not) be used to develop more predictive models of the anomalous collision frequency.

3.3   Bohm diffusion

The simplest and most widespread closure model for anomalous transport in Hall thrusters (as well as other magnetized plasma devices) is Bohm diffusion. In this model, the anomalous electron collision frequency is given by

$$\nu_{an} = c_1 \omega_{ce},$$

(3.1)

where c_1 is an adjustable constant. The classical value of c_1 is 1/16, but it is often adjusted to match data. Equation 3.1 was originally an empirical scaling law, but it can be derived by assuming that the plasma in question exhibits fluctuations in electrostatic potential with length-scales comparable to the electron Larmor radius and magnitudes proportional to the electron temperature.[24] The diffusion is thus assumed to result from micro-turbulence. As discussed in the **Chapter 1**, Bohm diffusion alone is inadequate to predict anomalous electron transport in Hall thrusters. This stems from the fact that a single coefficient is inadequate to capture the previously-discussed "canonical" anomalous collision frequency profile. However, due to the simplicity of the Bohm scaling, it is often used as the basis for the spatially-varying empirical anomalous transport profiles discussed in the preceding section. To generate these profiles, the scaling constant between the anomalous collision

frequency and electron cyclotron frequency is adjusted as a function of axial distance from the anode until a satisfactory match with data is obtained. One common "two-zone" approximation is given by[57,58]

$$
\nu_{an}(z) = \begin{cases} \frac{1}{160}\omega_{ce} & z < L_{ch} \\[2mm] \frac{1}{16}\omega_{ce} & z \geq L_{ch}, \end{cases}
\tag{3.2}
$$

where L_{ch} is the length of the discharge channel. This captures the observation that the anomalous transport inside the discharge channel is lower than the traditional Bohm scaling would predict. By including more zones and adjusting the scaling coefficients, good agreement with spatially-resolved plasma measurements can be obtained. However, such profiles are not extensible across thrusters or operating conditions, so Bohm diffusion is not a predictive closure model.

3.4   Models derived by analogy to classical turbulence

As outlined in the preceding chapter, the approach we take toward including anomalous transport derives heritage from the eddy viscosity-type models employed in classical fluid mechanics. It would make sense, then, to attempt to apply the same methods used in that field to obtain models for the anomalous collision frequency. Classical turbulence obeys what is known as an *energy cascade*. This is a relationship between the amount of energy stored in turbulent fluctuations and the size of those fluctuations. The energy for the turbulence in normal flows derives from the bulk velocity of the fluid. Small differences in velocity give rise to shear, which causes the fluid to rotate. This produces large vortices, contain the majority of the turbulent energy. By the same process, the larger vortices decay into smaller vortices with less energy, and so on down the line. At a certain point, the fluctuations get small enough that they are able to be totally damped out or *dissipated* by viscosity. Thus, the energy "cascades" down across length-scales, from bulk flow to

large vortices to small vortices. In 1922, this process was immortalized by Lewis Fry Richardson in the form of a rhyming quatrain,[59] as

> Big whorls have little whorls
>
> Which feed on their velocity,
>
> And little whorls have lesser whorls
>
> And so on to viscosity.

While vorticity does not play as much of a role in Hall thruster plasma turbulence as classical turbulence, other processes may give rise to an energy cascade qualitatively similar to that seen in neutral fluids.

3.4.1   Model of Cappelli et al. (2015)

Making use of classical scaling relationships between mean flow and turbulent properties, Cappelli et al. derived a simple zero-equation closure model for the anomalous collision frequency asssuming the presence of a turbulent energy cascade with a characteristic dissipation rate.[42] In this model, it is assumed that the scattering of electrons by plasma turbulence is the dominant driver of electron thermalization and heating. It is also assumed that the small length scales at which dissipation occurs scale with the electron Larmor radius. Finally, the authors assume that speed at which the turbulent eddies transport is the same as the electron mean speed, $c_e = \sqrt{8k_B T_e / \pi m_e}$. Under these assumptions, the anomalous electron mobility can be written as

$$\mu_{an} = c_1 \sqrt{\frac{|\mathbf{j}_e \cdot \mathbf{E}|}{K n_e e c_e^2 B^3}}. \tag{3.3}$$

Here, c_1 is an empirical scaling constant, which would need to be tuned to match data. This expression for electron mobility can then be related to an effective electron collision frequency using the relationship $\mu \approx \nu / \omega_{ce} B$, giving us

$$\nu_{an} = c_1 \omega_{ce} \sqrt{\frac{|\mathbf{j}_e \cdot \mathbf{E}|}{K n_e e c_e^2 B}}. \tag{3.4}$$

In the paper first describing this model, Cappelli et al. compare the results of a 2D axisymmetric hybrid particle-in-cell (hybrid-PIC) simulation carried out using their model to experimental data from Stanford's Z-70 Hall thruster. They find that the model correctly captures the expected decrease in mobility near the exit plane and produces simulations with velocity profiles similar to those found experimentally via laser-induced fluorescence. However, the researchers did not compare the predicted performance of the simulation to experimental performance measurements in that work.

3.4.2   Jorns two-equation model (2019)

Jorns took similar inspiration from classical turbulence to derive an analogy to the famous $k - \epsilon$ ("k-epsilon") turbulence model for application to Hall thrusters.[60] Here, k refers to the turbulent kinetic energy, and ϵ is the dissipation rate. By describing the transport of these properties, we can learn both where turbulence is, and where it is growing/decaying. The $k - \epsilon$ model is one of the most widely-used eddy viscosity models, and while it is not suitable for all turbulent flows, it provides a reasonable approximation of the turbulence in a variety of situations. The model of Jorns, as in the $k - \epsilon$ model, solves two differential equations and is thus a *two-equation model*. The first equation describes the transport of wave energy (W) as it convects with the plasma. This is analogous to the k in the $k - \epsilon$ model. The second equation describes transport of the growth rate (ω_i), analogous to ϵ.

In deriving this model Jorns notes that, in contrast to classical turbulence, Hall thruster plasmas exhibit an *inverse energy cascade*.[37] This means that the turbulent energy in the plasma transfers primarily from small length scales and high frequencies to longer scales and lower frequencies. Additionally, since the plasma in a Hall thruster is diffuse and largely inviscid, the processes that mediates the energy transport are tied to wave-wave

and wave-particle interactions rather than viscosity and vorticity. Lastly, Jorns assumes that the plasma waves largely follow an ion acoustic-like dispersion relation. We will discuss this assumption in more detail in the next section. Armed with these assumptions and employing a few additional mathematical simplifications, we can write the model as

$$\frac{\partial W}{\partial t} + \mathbf{u}_i \cdot \nabla W = 2W \left(\omega_i - \omega_{loss} - \nabla \cdot \mathbf{u}_i\right) \tag{3.5}$$

$$\frac{\partial \omega_i}{\partial t} + \mathbf{u}_i \cdot \nabla \omega_i = 2\omega_i \left(c_1 M_e \omega_{pi} - \omega_i\right) - c_2 \frac{1}{M_e} \omega_i^2 \tag{3.6}$$

$$\nu_{an} = c_3 \sqrt{\frac{m_i}{m_e}} \frac{1}{M_e} \omega_i \frac{W}{n_e T_e}, \tag{3.7}$$

where the first equation describes the transport of wave energy, W, and the second describes the transport of the growth rate, ω_i, and the final equation relates the anomalous collision frequency to these two transported quantities. Note that since the last equation is algebraic, it does not affect this model's status as a "two-equation model". In the above equations, ω_{loss} represents the rate at which energy is dissipated from the wave. Jorns makes the assumption that this is proportional to the ion-neutral collision frequency, ν_{in}. Additionally, M_e is the electron Mach number, or the ratio between the electron drift speed and electron sound speed. Additionally, there are three fit coefficients, denoted c_1, c_2, and c_3, which must be adjusted to fit data. This model was able to match time-averaged measurements of the anomalous collision frequency from the H9 Hall thruster (described later in this work), using $c_1 = 0.01$, $c_2 = 400$, and $c_3 = 130$. However, Jorns did not implement this model directly into a Hall thruster simulation so its predictive capabilities have not been assessed in the open literature.

3.5 Models based on plasma instabilities

Just as in the model of Jorns, one can derive models of the anomalous collision frequency by making assumptions about the type and scaling of plasma instabilities in the device. As discussed in **Chapter 1** and the preceding section, anomalous transport in Hall thrusters is commonly believed to result from an instability which gives rise to plasma turbulence. This turbulence derives energy from the electron drift, and thus produces an effective drag force on the electrons' azimuthal motion. Per our derivation in the preceding chapter, gives rise to enhanced cross-field electron current. In order to turn this insight into a model, it is important to understand precise manner in which the instability grows, convects, and decays. I begin this section by briefly outlining some of the experimental and computational evidence of instability-induced transport in Hall thrusters. I then discuss two models of how instabilities affect the anomalous transport before finally discussing some of the assumptions of these models and how they may be adjusted to yield new models.

3.5.1 Theory of plasma instabilities

Although we have discussed waves can create anomalous transport, we do not yet know how to turn these kinetic effects into a useful closure model for the effective anomalous collision frequency. To accomplish this, we will derive some basic plasma wave theory, and show how wave theory can be used to derive a closure model for anomalous electron transport.

The primary objective of plasma wave theory is to describe the *spectrum* of the waves in a given plasma. Much like how sunlight contains light of many frequencies, plasma turbulence may take place over a wide range of frequencies. By determining which frequencies contain more energy than others, and how energy is transferred between frequencies, we can gain insight into which physical processes give rise to the turbulence, and hopefully how to model it. Like light, too, the frequency, ω (in radians per second) of

a plasma wave is related to its wavelength, λ (in meters). For an electromagnetic wave travelling in a vacuum (like a light ray), this relation is:

$$\lambda = \frac{2\pi c}{\omega} \tag{3.8}$$

Here, c is the speed of light in a vacuum, 299,792,458 meters per second. This relationship shows us that higher frequency light waves have shorter wavelengths. Instead of the wavelength, we may consider the *wavenumber*, k, defined to be equal to $2\pi/\lambda$. The wavenumber represents the "spatial frequency" of a wave. We can write **Equation 3.8** in terms of the wavenumber as

$$\frac{\omega}{k} = c \tag{3.9}$$

A relationship such as this one between the frequency of a wave, ω and its wavenumber, k, is known as a *dispersion relation*. This is the basic object of consideration when studying plasma waves. While the scaling relationship between the frequency and wavenumber for a light wave is a simple constant factor, dispersion relations for waves in plasmas may get very complex indeed. While light rays are electromagnetic, meaning both the electric and magnetic fields fluctuate, we consider only electrostatic waves in this work. This is because the strong applied magnetic field in Hall thrusters inhibits the propagation of many electromagnetic waves.

3.5.2   Electrostatic dispersion relation for a plasma

We now turn to deriving a few basic example dispersion relations common to the types of plasma we are interested in. As in **Section. 2.4**, our analysis begins with the Boltzmann equation, which describes the evolution of a velocity distribution function $f(\mathbf{x}, \mathbf{v}, t)$ of particles with charge q and mass m over physical ($\mathbf{x}$) and velocity ($\mathbf{v}$) space, subject to an

electric field **E**, a magnetic field **B**, and collision forces represented by an integral operator $\hat{C}$:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} + \frac{q}{m} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f}{\partial \mathbf{v}} = \hat{C}(f) \tag{3.10}$$

To analyze how plasma waves propagate and determine how they might affect electron transport, we employ *quasi-linear* wave theory, which relies on the following assumptions

1. The effect of the wave on the distribution function is small, so we can decompose the distribution function into an equilibrium part and a fluctuating part ($f \approx f_0 + f'$).

2. Since the fluctuations are small, products of two fluctuating quantities are even smaller, so we neglect them.

3. The wave is electrostatic, so the electric field is a fluctuating quantity, while the magnetic field is constant.

4. The equilibrium electric field is zero.

5. Collisions are approximated using the Bhatnagar-Gross-Krook (BGK) collision operator, so $\hat{C}(f) \approx \nu(f_0 - f)$, where ν is a collision frequency and f_0 is the equilibrium distribution function. Physically, this represents the tendency of collisions to move a particle distribution toward equilibrium.

6. Perturbations are modelled as plane waves, such that for a fluctuating quantity Q' that we are interested in, $Q' \approx \hat{Q} \exp(i(\mathbf{k} \cdot \mathbf{x} - i\omega t)$, where $\hat{Q}$ is the (in general complex) amplitude of the wave, ω is the frequency of the wave, and **k** is the wave-vector.

Subject to these assumptions, the Boltzmann equation for a set of arbitrary particles becomes

$$\left[\frac{\partial f_0}{\partial t} + \mathbf{v} \cdot \frac{\partial f_0}{\partial \mathbf{x}} + \frac{q}{m}(\mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_0}{\partial \mathbf{v}}\right] + \frac{\partial f'}{\partial t} + \mathbf{v} \cdot \frac{\partial f'}{\partial \mathbf{x}} + \frac{q}{m}\mathbf{E} \cdot \frac{\partial (f_0 + f')}{\partial \mathbf{v}} = -\nu f'$$

The bracketed term in the above expression is just the left-hand side of **Equation 3.10** with zero electric field. Since the collision operator is zero for an equilibrium plasma, the entire term vanishes. Additionally, we invoke the second and fourth assumptions to neglect the term $\frac{q}{m}\mathbf{E} \cdot \frac{\partial f'}{\partial \mathbf{v}}$. With these terms removed, we are left with the perturbed Boltzmann equation:

$$\frac{\partial f'}{\partial t} + \mathbf{v} \cdot \frac{\partial f'}{\partial \mathbf{x}} + \frac{q}{m}\mathbf{E} \cdot \frac{\partial f_0}{\partial \mathbf{v}} = -\nu f' \tag{3.11}$$

To continue, we invoke assumption 3 and replace all quantities with plane waves. We also replace the electric field with the potential by noting that $E = -\partial\phi/\partial\mathbf{x} = -i\mathbf{k}\phi$.

$$\frac{\partial}{\partial t}\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right] + \mathbf{v} \cdot \frac{\partial}{\partial \mathbf{x}}\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right] + \frac{q}{m}\left[\hat{\mathbf{E}}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right] \cdot \frac{\partial f_0}{\partial \mathbf{v}}$$
$$= -\nu\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right]$$

$$-i\omega\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right] + i(\mathbf{v}\cdot\mathbf{k})\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right] - i\frac{q}{m}\mathbf{k}\phi\cdot\frac{\partial f_0}{\partial \mathbf{v}}$$
$$= -\nu\left[\hat{f}\exp(i(\mathbf{k}\cdot\mathbf{x} - i\omega t)\right]$$

$$-i\omega f' + i\mathbf{k}\cdot\left[\mathbf{v}f' - \frac{q}{m}\phi\frac{\partial f_0}{\partial v}\right] = -\nu f' \tag{3.12}$$

We can then solve for f':

$$f' = -\frac{\mathbf{k}\cdot\frac{q}{m}\phi\frac{\partial f_0}{\partial \mathbf{v}}}{\omega - \mathbf{k}\cdot\mathbf{v} + i\nu} \tag{3.13}$$

Up until this point, we have been dealing with a distribution function composed of arbitrary particles. To get more specific, let us suppose we have two species—negatively-charged electrons and positively-charged ions. Quantities relating to electrons will have the subscript e and those relating to ions will have the subscript i. Let us further assume that since ions are much more massive than electrons, collisions between the two species will mostly alter electron trajectories, leaving ions unchanged. The governing equations of our system then become:

$$f'_e = -\frac{\mathbf{k} \cdot \frac{q_e}{m_e} \phi \frac{\partial f_{e,0}}{\partial \mathbf{v}}}{\omega - \mathbf{k} \cdot \mathbf{v} + i\nu_e} \tag{3.14}$$

$$f'_i = -\frac{\mathbf{k} \cdot \frac{q_i}{m_i} \phi \frac{\partial f_{i,0}}{\partial \mathbf{v}}}{\omega - \mathbf{k} \cdot \mathbf{v}} \tag{3.15}$$

Additionally, the ions and electrons are coupled through the electrostatic potential via Gauss's law:

$$-\nabla^2 \phi = \mathbf{k}^2 \phi = \frac{1}{\epsilon_0} (q_i n_i + q_e n_e) = \frac{1}{\epsilon_0} \left[q_i(n_{i,0} + n'_i) + q_e(n_{e,0} + n'_e) \right] \tag{3.16}$$

For notational simplicity, we denote $k \equiv |\mathbf{k}|$. We assume the plasma is quasi-neutral at equilibrium, so $n_{i,0} = n_{e,0}$. We can express the perturbed densities in terms of the perturbed distribution functions to get

$$k^2 \phi = \frac{1}{\epsilon_0} \int_{\mathbf{v}} \left(q_i f'_i + q_e f'_e \right) d\mathbf{v} \tag{3.17}$$

Finally, we substitute Eqs. 3.14 and 3.15 into the above relation to arrive at a general dispersion relation for electrostatic waves propagating in a two-component warm plasma:

$$\epsilon(\omega, \mathbf{k}) = 1 + \frac{1}{\epsilon_0 k^2} \int_{\mathbf{v}} \left[\frac{\mathbf{k} \cdot \frac{q_i^2}{m_i} \frac{\partial f_{i,0}}{\partial \mathbf{v}}}{\omega - \mathbf{k} \cdot \mathbf{v}} + \frac{\mathbf{k} \cdot \frac{q_e^2}{m_e} \frac{\partial f_{e,0}}{\partial \mathbf{v}}}{\omega - \mathbf{k} \cdot \mathbf{v} + i\nu_e} \right] d\mathbf{v} = 0 \tag{3.18}$$

This dispersion relationship is the starting point for deriving expressions for the effects of the instabilities which drive electron transport in Hall thrusters. Next we will illustrate how

Equation 3.18 is used by deriving reduced dispersion relations for some simple plasma waves.

3.5.2.1 Simple plasma oscillation

This is the simplest plasma wave. We assume the ions are unperturbed, and thus neglect the ion contribution to **Equation 3.18**. We also assume that the electrons are cold, so $f_{e,0}(\mathbf{x}, \mathbf{v}) = n_e \delta(\mathbf{v} - \mathbf{v}_0)$, where δ signifies the Dirac Delta function[a] and $\mathbf{v}_0$ is the equilibrium velocity of the electrons, which we set to zero without loss of generality. Finally, we consider one-dimensional propagation, reducing $\mathbf{k}$ to a scalar. This gives us the following dispersion relation:

$$\epsilon(\omega, k) = 1 + \frac{1}{k} \frac{q_e^2 n_e}{m_e \epsilon_0} \int_\mathbf{v} \frac{\frac{\partial}{\partial \mathbf{v}} \delta(\mathbf{v})}{\omega - k\mathbf{v} + i\nu_e} d\mathbf{v} = 0$$

Integrating by parts, we find that

$$1 - \frac{q_e^2 n_e}{m_e \epsilon_0} \int_\mathbf{v} \frac{\delta(\mathbf{v})}{(\omega - k\mathbf{v} - i\nu_e)^2} d\mathbf{v} = 1 - \frac{q_e^2 n_e}{m_e \epsilon_0} \frac{1}{(\omega + i\nu_e)^2} = 0,$$

and thus,

$$\omega = \sqrt{\frac{q_e^2 n_e}{m_e \epsilon_0}} - i\nu_e \equiv \omega_{pe} - i\nu_e. \tag{3.19}$$

This dispersion relation captures the fact that a mono-energetic population of charged particles will oscillate in an fluctuating electic field with a characteristic frequency ω_{pe}. If there are collisions, they act to damp the oscillation by adding a negative imaginary component to the frequency.

---

[a]This is a distribution with a "width" of zero, so all particles have the same velocity. It is defined such that the area under this distribution is one, so as to preserve the properties we expect from a normalized distribution function.

72

3.5.2.2   Warm plasma waves and Landau damping

We can move up a level of complexity by allowing the electrons be warm, i.e. $T_e > 0$. The dispersion relation for warm electrons in one dimension (where v is along the direction of propagation of the wave) is

$$\epsilon(\omega, k) = 1 + \frac{1}{k}\frac{q_e^2}{m_e\epsilon_0}\int_v \frac{\partial f_{e,0}/\partial v\, dv}{\omega - kv + i\nu_e} = 0,\tag{3.20}$$

where the electrons have a one-dimensional Maxwellian distribution along the direction of propagation:

$$f_{e,0}(\mathbf{v})dv = \frac{n_e}{\sqrt{\pi}v_{th,e}}\exp\left[-\frac{v^2}{v_{th,e}^2}\right]dv\tag{3.21}$$

$$\frac{\partial f_{e,0}}{\partial \mathbf{v}}dv = -2\sqrt{\pi}n_e\left(\frac{1}{\sqrt{\pi}v_{th,e}}\right)^2\frac{v}{v_{th,e}}\exp\left[-\frac{v^2}{v_{th,e}^2}\right]dv\tag{3.22}$$

In the above, we have introduced the *thermal speed* of a species s, defined as $v_{th,s} = \sqrt{2k_B T_s/m_s}$. The integral in Eq. 3.20 has a singularity on the complex plane where $\omega - k\mathbf{v} + i\nu_e = 0$. To handle this, we make use of the *Sokhotski–Plemelj theorem*,[61] which states that for some value x,

$$\lim_{\epsilon\to 0}\frac{1}{x+\epsilon} = \mathcal{P}\frac{1}{x} - i\pi\delta(x)\tag{3.23}$$

where $\mathcal{P}$ denotes the Cauchy principle value. The integral then becomes

$$\int_v \frac{\partial f_{e,0}/\partial v\, dv}{\omega - kv + i\nu_e} = \int_v \frac{\partial f_{e,0}}{\partial v}\left[\mathcal{P}\frac{1}{\omega - kv + i\nu_e} - i\pi\delta(\omega - kv + i\nu_e)\right]dv$$

$$= \mathcal{P}\int_v \frac{\partial f_{e,0}/\partial v\, dv}{\omega - kv + i\nu_e} - i\pi\left(\frac{\partial f_{e,0}}{\partial v}\right)_{v=(\omega+i\nu_e)/k}\tag{3.24}$$

We can then expand the denominator using a Taylor series. After integration, we obtain our final dispersion relation:

$$\omega \approx \omega_p(1 + 3k^2\lambda_d^2) - i\left(\nu_e + \frac{1}{2}\sqrt{\frac{\pi}{2}}\frac{\omega_p}{(k\lambda_d)^2}\exp\left(-\frac{1}{2k\lambda d^2}\right)\right). \tag{3.25}$$

The real part of the above equation is known as the *Bohm-Gross* dispersion relation. The imaginary part (the part multiplied by *i*) accounts for both collisional damping ($i\nu_e$) and a new term, known as *Landau damping*. Physically, this term means that waves with wavelengths near the Debye length ($k\lambda_d \sim 1$) will be damped.

By relaxing more assumptions (i.e. including ion contributions) and performing similar analyses to the two examples above, we can arrive dispersion relations which more accurately describe the sorts of wave which are thought to exist in Hall thrusters.

3.5.3   Plasma instabilities in Hall thrusters

Since the 1970s,[32] it has been known that Hall thrusters play host to azimuthally-propagating waves with frequencies in the megahertz (MHz) range. Since 2000, significant experimental,[33,34,36,62,37] computational,[63,39] and theoretical[38] effort has gone into understanding the dispersion of these instabilities. This work has produced a consensus that these azimuthal drift waves are a leading contributor to electron transport, and has shed light on the precise nature and scaling of these waves. It is believed that the anomalous electron transport results from the *electron cyclotron drift instability* (ECDI). This instability derives its energy from, and is driven unstable by, the *E × B* drift. The energy enters the turbulence as fluctuations at small wavelengths and high frequencies, and transfers to longer wavelengths via wave-wave interactions in an inverse energy cascade.[37] The dissipation of the wave energy at the longer length scales may primarily result from convection, as the energy is transported downstream and out of the Hall thruster. Other loss mechanisms, such as ion-neutral collisions and ion Landau damping may also play a role, though their importance is debated.[2,37]

The dispersion of this instability is characterized by sharp *cyclotron resonances*. These are peaks in the frequency and growth rate of the instability which occur for wavenumbers near cyclotron harmonics. Mathematically, these resonances occur when

$$k_\theta v_{de} \approx j\omega_{ce}, \quad j = 1, 2, 3, ... \tag{3.26}$$

where v_{de} is the drift speed of the electrons. While the instability primarily propagates in the azimuthal direction, i.e. $\mathbf{k} \approx k_\theta \hat{\theta}$, it also has a small component in the radial direction, parallel to the applied magnetic field. In **Figure 3.2**, we show qualitatively how these resonances respond to an increased radial wavenumber. As this component of the wave-vector grows larger, the cyclotron resonances seem to "smooth out", and the dispersion relation approaches that of the modified ion acoustic instability.[38] This theoretical result has been borne out in numerical simulations, with particle-in-cell simulations often displaying ion-acoustic like wave dispersions.[63,39] Experimental evidence is less conclusive, however, and probe-based measurements of this instability[37] have shown the presence of strong cyclotron resonances. Ultimately, the dispersion of this instability will differ based on local plasma properties and thruster geometry, so whether or not strong resonances are observed may be a function of the thruster and operating condition, as well as where in the discharge (and at what frequencies and wavenumbers) the waves are measured.

3.5.4 Model of Lafleur et al. (2016)

Applying insights from particle-in-cell simulations of the ECDI, Lafleur et al. derived a zero-equation closure model based on the modified ion acoustic dispersion relation. This dispersion relation can be derived by following a procedure similar to the one we used to derive the for the warm plasma dispersion relation, but including the contributions of ions and allowing both ions and electron to drift relative to each other with characteristic velocities $\mathbf{u}_i$ and $\mathbf{u}_e$, respectively. This dispersion relation can be written as
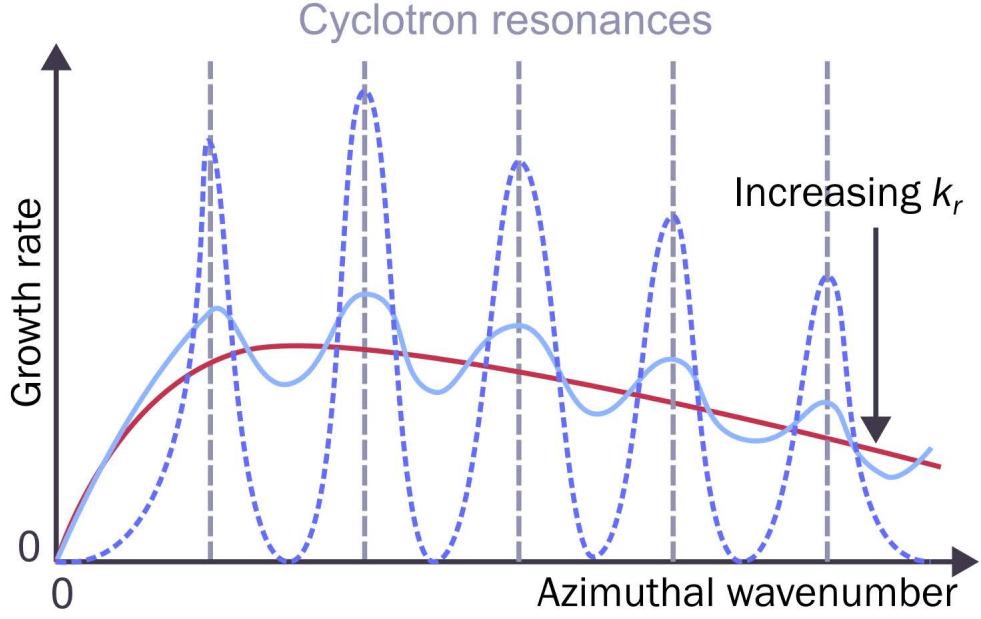
Figure 3.2: Qualitative evolution of the growth rate of the electron cyclotron drift instability with increasing radial wavenumber. The ion acoustic limit is shown in red.

$$\omega_r = \mathbf{k} \cdot \mathbf{u}_i \pm \frac{kc_s}{\sqrt{1 + k^2\lambda_{de}^2}} \tag{3.27}$$

$$\omega_i = \pm\sqrt{\frac{\pi m_e}{8m_i}} \frac{\mathbf{k} \cdot \mathbf{u}_e}{\left(1 + k^2\lambda_{de}^2\right)^{3/2}}, \tag{3.28}$$

where ω_r and ω_i are the real and imaginary parts, respectively, of the frequency. The imaginary part of the frequency is known as the the *growth rate*, often denoted γ. In obtaining the above dispersion, Lafleur et al. have neglected both collisional damping as well as ion Landau damping. They next make use of a conservation equation for the wave energy density, W

$$\frac{\partial W}{\partial t} + \nabla \cdot \left(\mathbf{v}_g W\right) = 2\gamma W. \tag{3.29}$$

At steady state, the time-dependent term vanishes. If the group velocity, $\mathbf{v}_g$, of the wave is equal to the ion drift speed, $\mathbf{u}_i$, then the growth rate can be written as

$$\gamma = \frac{\nabla \cdot (\mathbf{u}_i W)}{2W}. \tag{3.30}$$

Physically, this reflects an assumption that the wave is saturated everywhere, so the growth rate is equal to the loss due to convection, and that convection is the dominant loss mechanism. Next, the authors argue that the waves saturate due to ion trapping. If this is the case, then the wave energy density should be proportional to the electron thermal energy density. Using this assumption, and after some additional manipulation, an expression for the anomalous collision frequency can be derived. This expression is

$$\nu_{an} = c_1 \frac{q_e \nabla \cdot (\mathbf{u}_i n_e k_B T_e)}{m_e n_e c_s v_{de}}, \tag{3.31}$$

where $v_{de} = |\mathbf{E}|/|\mathbf{B}|$ is the electron $E \times B$ drift speed. Lafleur et al. show that, under appropriate scaling of the fit coefficient c_1, this model agrees well with empirical anomalous collision frequency profiles such as those discussed earlier in this chapter. However, they do not implement it directly into a Hall thruster model, so its behavior when treated as a true self-consistent model was not yet known.

3.5.5 Model of Katz et al. (2016)

The model of Katz et al.[64] uses the same theoretical starting point of the model of Lafleur et al., but considers ion Landau damping, rather than convection, to be the dominant loss mechanism. Additionally, Katz et al. consider the effects of different energies in ion populations. They consider a low-energy ion population that is not subject to strong Landau damping, and a high energy ion population which is. Ions born at a potential of less than 100 V are assigned to the low-energy population, while those born at higher potentials are assigned to the high-energy population. The total anomalous collision frequency is then taken as the sum of the anomalous collision frequencies originating in each of these populations. Mathematically, this is written as

$$\nu_{an,1} = \omega_{ce} M_e \left(\frac{T_{i,1}}{T_e} \right)^{3/2} \exp \left(\frac{T_e}{3T_{i,1}} \right) \tag{3.32}$$

$$\nu_{an,2} = \omega_{pe} \frac{n_2^2}{n^2} \tag{3.33}$$

$$\nu_{an} = \nu_{an,1} + \nu_{an,2} \tag{3.34}$$

where the subscripts 1 and 2 denote ions belonging to the high and low energy populations, respectively. The authors evaluate their model using a one-dimensional fluid Hall thruster code. Instead of implementing the model directly into the code, they instead evaluate **Equations 3.32 - 3.34** on the output of a simulation which used a calibrated empirical anomalous collision frequency. The result of this procedure was used as a new empirical collision frequency profile. The process was repeated until convergence. The model produced features very similar to the calibrated profile, including low anomalous transport near the exit plane and high transport downstream. However, it was not implemented into the code in a truly self-consistent manner.

3.5.6   Effect of modifying model assumptions

Both the model of Lafleur et al. and that of Katz et al. made use of the assumption that the instability leading to anomalous electron transport is governed by the modified ion acoustic dispersion relation, and that the turbulence is *marginally stable*.[65] This latter assumption means that at all spatial locations, the growth rate of the wave, γ is equal to the damping or *loss* rate, γ_{loss}, so the net growth rate is zero. The consequence of this assumption is that we can derive expressions for the positive part of the wave growth rate by considering only the mechanisms that cause the wave to lose energy. Combining this with an assumed scaling of the wave energy density, we can derive an expression for the anomalous electron

collision frequency, assuming the waves follow an ion-acoustic-like dispersion. For some choice of W and γ, the anomalous collision frequency can be written as

$$\nu_{an} \propto \frac{2\gamma W}{m_e n_e c_s v_{de}}.$$ (3.35)

Physically, the effective collision frequency scales with the rate at which energy is taken from the plasma, γW, which matches our intuition that the wave acts as an effective drag on the electron drift. Given different choices for the scaling of the wave energy density and growth rate, different transport models can be obtained. For instance, if the wave energy density saturates at the electron kinetic—rather than thermal—energy density, but still assumed that convection was the dominant loss mechanism, then we could derive a Lafleur-like model of the form

$$\nu_{an} \propto \frac{\nabla \cdot \left(\mathbf{u}_i n_e |\mathbf{u}_e|^2 \right)}{n_e c_s v_{de}}$$ (3.36)

Alternatively, we could include multiple loss mechanisms (such as ion Landau damping and collisions), to arrive at a model like

$$\nu_{an} \propto \frac{q_e \left(\nabla \cdot (\mathbf{u}_i W) + W(\gamma_{loss,collision} + \gamma_{loss,Landau}) \right)}{m_e n_e c_s v_{de}}$$ (3.37)

In principle, each of the terms in the numerator could have a different scaling constant, i.e. c_1, c_2, c_3. One could also assume a different group velocity, so the waves no longer convect with the ions. Lastly, we could assume a alternate form of the convection term. Depending on how one derives the conservation equation for wave energy, one of three forms is possible:

1.

$$\frac{\partial W}{\partial t} + \nabla \cdot (\mathbf{v}_g W) = 2\gamma W$$ (3.38)

2.

$$\frac{\partial W}{\partial t} + \mathbf{v}_g \cdot \nabla W = 2\gamma W$$ (3.39)

3.

$$\frac{\partial W}{\partial t} + n_e \mathbf{v}_g \cdot \nabla \left(\frac{W}{n_e} \right) = 2\gamma W \tag{3.40}$$

Thus, by altering some of the key assumptions underlying the models of Lafleur et al. and Katz et al., we have generate a novel family of possible anomalous transport models in a common framework. We will return to evaluating a few models of this type in **Chapter 7**.

3.6 Shear-based models

In many magnetized plasma devices, it is common to adopt models for the anomalous collision frequency in which the transport is essentially Bohm-like, or follows some other well-known scaling, but is reduced in areas of high electron shear.[66] The *electron shear rate*, *s*, is given by

$$s = \nabla v_{de}, \tag{3.41}$$

where v_{de} is defined, as in the Lafleur et al model., as the $E \times B$ drift speed. Scharfe et al.[41] proposed a semi-empirical shear-suppressed transport model with the following functional form:

$$\nu_{an} = \begin{cases} \frac{c_1 \omega_{ce}}{1 + (c_2 s)^{c_3}} & z < L_{trans} \\ c_1 \omega_{ce} & z \geq L_{trans}. \end{cases} \tag{3.42}$$

In this expression, c_1, c_2, and c_3 are adjustable scaling constants and L_{trans} is the axial distance at which the effect of electron shear suppression is assumed to disappear. This is typically taken to be either the discharge channel location or the location of peak magnetic field. This model has the advantage of reducing the electron transport near regions of high electron velocity. As discussed in preceding sections, this is a common feature of empirical anomalous collision frequency profiles and is necessary to capture the strong electric field and steep ion acceleration profiles near the channel exit plane.

In contrast to the others in this work, this model has been implemented into Hall thruster simulations and evaluated against different operating conditions and propellants. Cha et al simulated the Stanford Hall thruster on a few different propellants using this model and found to perform reasonably well,[67] although the coefficients are not transportable across operating conditions. Despite this heritage, this model has a few disadvantages. First, the cutoff in shear suppression at an arbitrary axial location is ad-hoc and unphysical. Second, the reliance on Bohm diffusion means that the coefficient c_1 will not be applicable for all thrusters, as empirical Bohm-like transport profiles are not generally extensible. Finally, the coefficient c_2, which generally has a magnitude on the order of $10^{-7}s^{-1}$ is not well-motivated physically.

To resolve the first objection, we might simply try eliminating the positional dependence. Second, as the functional form of this model is somewhat ad-hoc to begin with, we might try replacing the Bohm collision frequency with a more physically-relevant collision frequency, such as one of the variants of the Lafleur et al. / Katz et al. model described previously. Augmenting these models with shear suppression could help steepen the gradients in the acceleration region and produce ion velocity profiles more in line with experiment. To address the third objection, we turn to previous shear-suppressed turbulence models from the literature, such as those originating in the fusion community. In such models, the shear-induced anomalous collision frequency is given by[66]

$$\nu_{an} = \frac{c_1 \omega_{ce}}{1 + (s/\gamma_{max})^2} \tag{3.43}$$

In this expression, c_1 is as defined in the model of Scharfe et al., and γ_{max} is the growth rate of the dominant instability in the device. Thus, c_2 in the model of Scharfe et al. likely represents an effective growth rate. If we also make the assumption that $c_3 = 2$, as Cha et al. and Scharfe et al. do, and apply a non-Bohm base anomalous collision frequency, $\nu_{an,base}$, we can derive a more general shear-suppressed transport model for Hall thrusters,

$$\nu_{an} = \frac{\nu_{an,base}}{1 + \left(\frac{s}{\gamma_{max}}\right)^2} \tag{3.44}$$

Employing the assumption of marginal stability as in the previous section, we could then assume that γ_{max} is equal to losses resulting from convection, collisions, ion Landau damping, or other processes. Both the base collision frequency, $\nu_{an,base}$, and the growth rate, γ_{max} will likely feature fit coefficients which must be tuned to match data.

3.6.1 Chodura model

While not explicitly a shear-suppressed transport model, the Chodura model[68] shares many features with the model of Scharfe et al. Namely, this model posits that the anomalous collision frequency is equal to the ion plasma frequency, but is reduced where the electron drift speed is high. Mathematically, this can be expressed as

$$\nu_{an} = c_1 \omega_{pi} \left(1 - \exp\left(-\frac{v_{de}}{c_2 c_s} \right) \right), \tag{3.45}$$

where c_1 and c_2 are fit coefficients. In contrast to many of the previously-discussed models, the Chodura model was not derived to explain Hall thruster anomalous transport. Instead, this semi-empirical model describes the effect of microturbulence in a few select types of plasma device, such as theta-pinch and field-reversed-configuration devices.[69,70,71] More recently, Simmonds et al.[19] proposed that a Chodura-like transport model might apply to Hall thrusters. They investigated the effect of such a scaling on the theoretical thrust density limit for Hall thrusters. However, the model has not yet been implemented self-consistently in a Hall thruster code.

3.7 Data-driven models

The last type of model we investigate comes from a recent investigation by Jorns.[3] In this work, a training dataset of empirically-derived anomalous transport profiles such as those described in Section 3.2 was combined with the outputs of simulations performed using those profiles. This dataset included four Hall thrusters at nine distinct operating conditions. Jorns applied symbolic regression to attempt to determine how the empirically-inferred anomalous collision frequencies scale with changes in plasma properties. One of the more promising of these models was given by

$$\nu_{an} = C_1 \omega_{ce} \frac{|\mathbf{u}_i|}{C_2 c_s + V_{de}}. \tag{3.46}$$

The above expression implies the anomalous collision frequency should exhibit a Bohm-like functional dependence, scaled by the ratio between the ion velocity and the electron $E \times B$ drift speed. While no physical explanation for such a scaling was readily apparent, this model was able to predict the empirical profile of a thrusters not present in the training dataset. However, this model was not implemented directly into a thruster code, so its true predictive capabilities were not assessed in a self-consistent manner.

3.8 Conclusion

Over the course of this thesis work, we tested hundreds of models for the anomalous collision frequency. In this chapter, we have discussed a few of the more promising models. We have not described all of the models we tried, nor will we investigate in detail all of the models introduced in this chapter. By listing and describing these closure models, however, we hope that we have given the reader a good overview of the attempts that have been made to develop a fluid closure model for the anomalous collision frequency. In addition, we have demonstrated how new models may be generated by adjusting some

of the assumptions underlying these models. As more experimental data about plasma turbulence in Hall thrusters arrives, the validity—or lack thereof—of these assumptions may become more clear, and the models may be fine-tuned to better match data.

As each of these models has one or more fit coefficients, it is important to find out how to best fit them to data. In order to calibrate these coefficients, we need to define what "fit" means. It turns out that this is a more complicated question than it seems. We will return to this question in Chapter 5, in which we investigate the data-driven approach proposed by Jorns[3] and determine the relevance and efficacy of comparing self-consistent transport models to empirical transport models. Before we do that, we turn in the next chapter to discussing the simulations we perform in the remainder of the thesis, as well as the thruster we simulate.

CHAPTER 4

Thruster Simulations in Hall2De

> "A theory is something nobody believes, except the person who
> made it. An experiment is something everybody believes, except the
> person who made it."

Albert Einstein

In this chapter, we introduce Hall2De, the Hall thruster code we use for all of our simulations. We then describe the thruster and operating conditions being simulated and the data against which we compare our simulation results. Finally, we overview the metrics by which we compare our simulation results to the data.

4.1   Modeling and simulation of Hall thrusters

As we discussed in **Chapter 2**, full particle-in-cell methods (PIC), in which electrons and ions are treated as particles, are infeasible for engineering simulations of Hall thrusters. Instead, the electrons need to be treated as a fluid. Ions and neutrals may still be treated as particles. Simulations in which electrons are modelled as a fluid but ions and neutrals are modelled as particles are called *hybrid-PIC* simulations.

Additionally, while the fluid approximation for electrons enables us to simulate larger devices and longer timescales than we can using full PIC, it can still be very expensive to perform three-dimensional Hall thruster simulations. To reduce simulation times, a common approach is to use the assumption of *axisymmetry*. This means that the gradients in the azimuthal direction are zero, and we can therefore reduce our simulations to two-dimensions (axial and radial). In making this assumption, we lose the ability to model differences in plasma properties in the azimuthal dimension (such as rotating spokes), but gain a large reduction in computational time. For most engineering applications, then, 2-D axisymmetric hybrid-PIC or fluid Hall thruster codes have become the standard Hall thruster simulation tool.

In **Table 4.1**, we list published two-dimensional axisymmetric Hall thruster codes similar to Hall2De. Of these codes, only Hall2De and HYPHEN simulate the electron dynamics in a fully two-dimensional manner. The other hybrid-PIC codes employ a quasi-one dimensional approach to the electron dynamics, in which the magnetic field lines are assumed to be isothermal and follow the Boltzmann relation, so the electrostatic potential, ϕ, can be determined by

$$\phi = \phi_0 + \frac{k_B T_e}{q_e} \log \left(\frac{n_e}{n_{e,0}} \right) , \tag{4.1}$$

where ϕ_0 and $n_{e,0}$ are reference potential and density values. Hall2De is able to recover the generally isothermal nature of Hall thruster magnetic field lines[54] but does not presuppose it, allowing for deviations from isothermality to occur when necessary.

Name	Type	Authors	Date	Reference
HPHall	Hybrid-PIC	J. M. Fife, M. Martinez Sanchez	1998	47
HPHall-2	Hybrid-PIC	F. I. Parra, J. M. Fife, E. Ahedo, M. Martinez Sanchez	2006	72
HALLIS	Hybrid-PIC	J. P. Boeuf, L. Garrigues, G. J. M. Haagelar	2002	73
Hall2De	Multi-fluid/Hybrid-PIC	I. Mikellides, I. Katz, A. Lopez Ortega	2012	16
HYPHEN	Hybrid-PIC	A. Dominguez Vazquez, D. Perez Grande, J. Zhou, E. Ahedo	2019	74

Table 4.1: List of published 2-D axisymmetric Hall thruster codes

4.2 Hall2De

All simulations in this work were conducted using Hall2De,[16] a multi-fluid/particle-in-cell thruster code developed at the Jet Propulsion Laboratory by Ioannis Mikellides, Ira Katz, and Alejandro Lopez Ortega. In this work, we use only the multi-fluid version of the code, which models both electrons and ions as fluids, while neutrals are treated as a free molecular flow.

We first describe the algorithm used to compute the neutral dynamics. The present version of the code uses a line-of-sight view factor algorithm described in **Reference 75**. When the simulation is initialized, view-factors of each thruster surface are computed for all points in the domain (i.e. cells, edges, vertices). The *view factor* of a surface at a point is the fraction of one's view that is taken up by that surface if one was to stand at that point. The neutral velocity at a point can then be computed by considering the temperatures and speeds of the neutral populations that enter the domain at each surface, weighted by the view factor of that surface at that point. By considering convection from these emission surfaces, the neutral density is then updated every timestep using the velocity computed from the view factors. Ionization is treated as a sink term which reduces neutral density in locations of high ionization rate.

The version of Hall2De employed in this thesis solves continuity, momentum, and energy equations for several ion species. These species include both ions of different charge states (singly-charged to triply-charged), and up to four populations which are distinguished by the electrostatic potential at their origin and thus by their energies. Axial, radial, and azimuthal ion momentum equations are solved. As the code is axisymmetric, however, no gradients in the azimuthal direction are considered. For each ion species s, the governing equations are

$$\frac{\partial n_{is}}{\partial t} + \nabla \cdot (n_{is}\mathbf{u}_{is}) = \dot{n}_{is} \tag{4.2}$$

$$\frac{\partial n_{is}\mathbf{u}_{is}}{\partial t} + \nabla \cdot (n_{is}\mathbf{u}_{is}\mathbf{u}_{is} + n_s k_B T_{is}/m_i) = Z_s \frac{q_e}{m_i} n_{is} (\mathbf{E} + \mathbf{u}_{is} \times \mathbf{B}) - \mathbf{R}_s \tag{4.3}$$

$$\frac{\partial}{\partial t}\frac{3}{2}(n_{is}k_B T_{is}) + \nabla \cdot \frac{5}{2}(\mathbf{u}_{is}n_{is}k_B T_{is}) = Q_{fric,s} + Q_{an,s}. \tag{4.4}$$

In the above, $\dot{n}_{is}$ is the total rate of change of the density of ion species s due to ionization reactions. This includes both increases in density due to ionization of neutrals or ions of lower charge states, as well as decreases in density due to ionization of species s to higher charge states. This also accounts for the loss of ions to charge exchange collisions with neutrals. Z_s is the charge state (i.e. 1, 2, or 3) of species s, and R_s is the net elastic and inelastic frictional force between ions of species s and all other ion and neutral species. Lastly, $Q_{fric,s}$ and $Q_{an,s}$ are the rates at which the ions of species s heat due to the aforementioned inter-species friction and anomalous effects, respectively.

Next, we can write the electron governing equations, which consist of quasi-neutrality, current conservation, a generalized Ohm's law, and an energy equation. Mathematically, these are given by

$$n_e = \sum_s Z_s n_{i,s} \tag{4.5}$$

$$\nabla \cdot (\mathbf{j}_e + \mathbf{j}_i) = 0 \tag{4.6}$$

$$q_e n_e (\mathbf{E} + \mathbf{u}_e \times \mathbf{B}) + \nabla(n_e k_B T_e) + m_e n_e \sum_s (\mathbf{u}_e - \mathbf{u}_{is}) = 0 \tag{4.7}$$
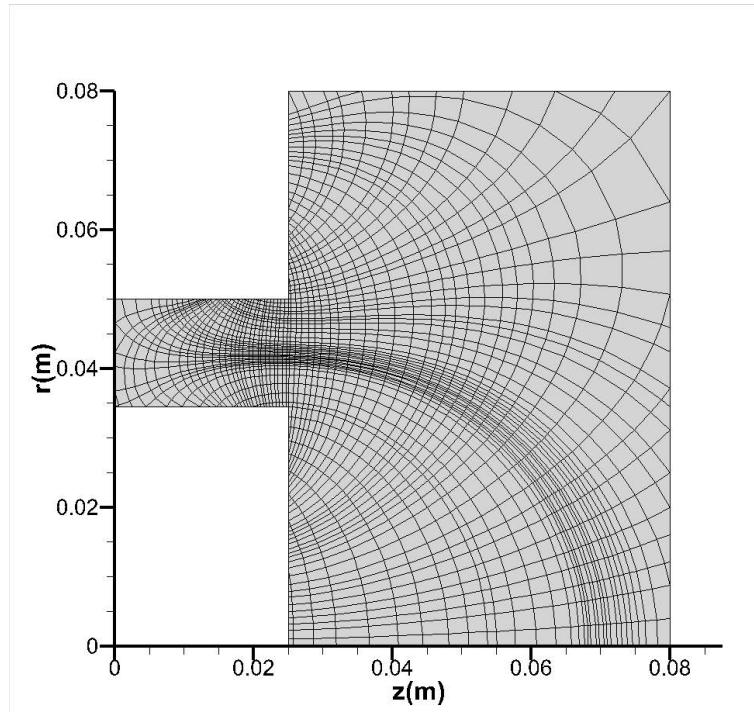
$$\frac{\partial}{\partial t}\frac{3}{2}(n_e k_B T_e) + \nabla \cdot \left(\frac{5}{2}\mathbf{u}_e n_e k_B T_e + \kappa_e \nabla T_e\right) = Q_{fric} - Q_{iz} - Q_{ex}, \tag{4.8}$$

where $\mathbf{j}_e$ and $\mathbf{j}_i$ are the electron and ion current densities, respectively, κ_e is the electron thermal conductivity coefficient, Q_{fric} is the electron heating rate due to classical and

89
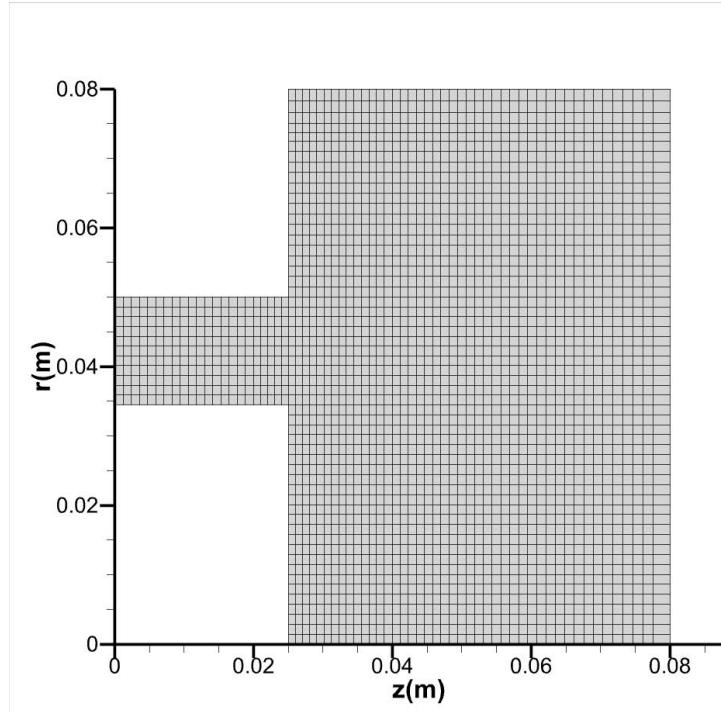
anomalous friction forces (i.e. the Ohmic heating), and Q_{iz} and Q_{ex} represent the reduction in electron thermal energy due to ionization and excitation, respectively. The generalized Ohm's law (Equation 4.7) is then broken into field-aligned, field-perpendicular, and azimuthal components, as outlined in **Chapter 2**. This transforms the above into a set of six equations which together solve for the electron density, n_e, the electric field, **E**, the electron parallel, perpendicular, and azimuthal current densities, $j_{e,\parallel}$, $j_{e,\perp}$, and $j_{e,\theta}$, and the electron temperature, T_e.

In order to solve these equations numerically, they are discretized using the first-order finite volume method on appropriate computational meshes. As electrons are strongly-magnetized in Hall thrusters, Hall2De solves for their behavior using a mesh aligned with the applied magnetic field. This magnetic field-aligned mesh (MFAM) reduces numerical diffusion and preserves the largely equipotential and isothermal nature of the magnetic field lines while still solving a two-dimensional electron energy equations. However, such field-aligned meshes inherently have small, high-aspect-ratio and high-skewness cells near boundaries and in regions where the magnetic field is convergent. Small timesteps are thus required to accurately and stably solve the governing equations in these regions. In order to increase the allowable timestep and improve numerical stability, Hall2De provides the option to solve the equations of motion for the unmagnetized ions and neutrals on a rectilinear grid with more uniform cell sizes. Similar approaches have previously been applied to hybrid particle-in-cell (hybrid-PIC) Hall thruster simulations.[47,76] As the geometry and magnetic field configuration of the H9 Hall thruster are export-controlled, we cannot show the simulated mesh in this work. To illustrate the difference in grids, we show instead in **Figure 4.1a** a field-aligned mesh for the SPT-100 Hall thruster. In contrast, we show a rectilinear mesh constructed for the same thruster in **Figure 4.1b**.

We employed a MFAM with 3925 cells for the electrons and a rectilinear grid with 3955 cells for the ions. The resolution of each of these grids was selected following a grid convergence study described in the next section. To increase the allowable timesteps, the

Figure 4.1: Example (a) field-aligned and (b) rectilinear meshes for an SPT-100 simulation in Hall2De.

ion equations of motion, as well as the electron temperature equation, are discretized semi-implicitly in time. **Figure 4.2** depicts the domain of the simulation with labelled boundaries. The domain extended eight channel lengths downstream from the thruster anode and the same distance radially outward from the channel centerline.

4.2.1   Boundary conditions

In order to close our system of PDEs, we are required to specify properties at the simulation boundaries. These boundaries are labelled in **Figure 4.2**. We can treat solid walls, such as the discharge channel or pole covers, as either insulating or conducting. If insulating, then no net current is allowed to enter the surface, and $\mathbf{j}_e \cdot \hat{n}$ is set equal to $\mathbf{j}_i \cdot \hat{n}$, where $\hat{n}$ is the outward-facing wall normal vector. If conducting, then a sheath boundary condition is used to determine the amount of current collected at the boundary, following the model of Hobbs and Wesson.[77,5] A similar sheath boundary condition is also used at the anode, with the electrostatic potential at the anode surface set to the sum of the applied discharge voltage and the computed sheath potential. If the discharge channel walls are set to conducting, they are also set to anode potential. We compute the ion current at solid boundaries by assuming a pre-sheath potential drop equal to the electron temperature. This leads to the ions impacting the walls at the Bohm speed, $u_{is,wall} \cdot \hat{n} = \sqrt{Zq_e k_B T_e / m_i}$. Ions that hit the walls recombine and are re-emitted as neutrals. At outflow boundaries, we assume the ions can leave freely, so we apply Neumann (zero-gradient) boundary conditions for the ion density, momentum, and temperature. To ensure no net current leaves the device, we also set the electron current equal to the ion current at such boundaries. The electron temperature and density are fixed to a user-specified value at the cathode orifice. The ion velocity at this location is assumed to be a fraction of the Bohm speed, while the electron current is computed using Ohm's law. Finally, along the symmetry axis, all properties are set to have zero gradient. This captures the fact that there should not be a net flux of any property across this axis, or else it would not symmetric.
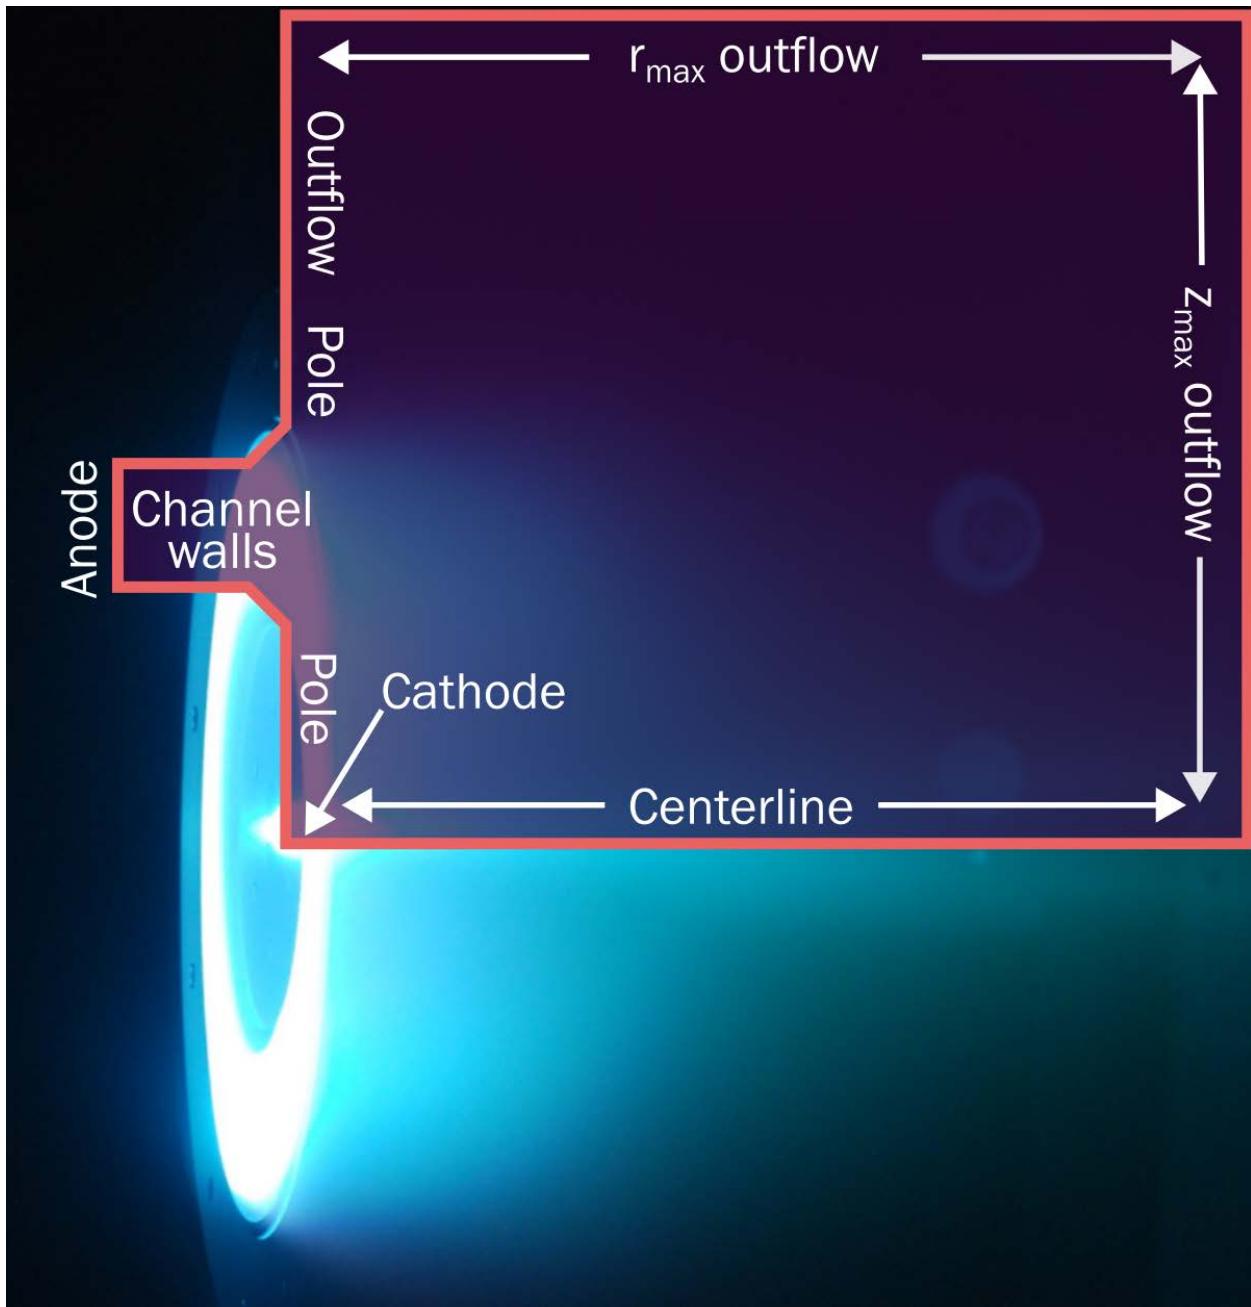
Figure 4.2: Simulation domain overlaid on the H9 Hall thruster, with labelled boundary conditions
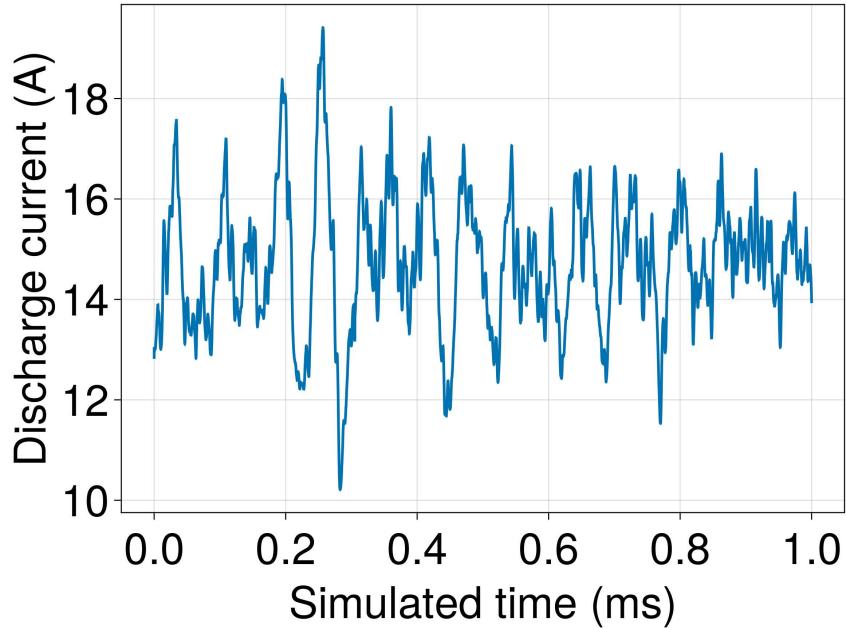
Figure 4.3: Example of a "converged" discharge current trace for a simulation at 300 V and 15 A over 1 ms of simulated time.

4.2.2 Simulation setup

For the majority of this work, we simulated three ion charge states and two ion populations. The first ion population comprised ions originating in the main beam, and the second consisted of ions emitted from the cathode or born in the cathode plume. This yielded a total of six ion species.

As Hall thrusters are inherently oscillatory devices,[14] our simulations did not converge to a steady state. Instead, after an initial transient, we have found the simulated discharge current will eventually exhibit a quasi-perodic oscillation with constant frequency and amplitude (**Figure 4.3**), with no change in the time-averaged properties. We ran each of our simulations for two milliseconds of simulated time, which was sufficiently long for the simulation to converge to this quasi-stationary state in all cases.

4.2.3 Parameter sensitivity

In addition to the anomalous collision frequency of the main Hall thruster beam, Hall2De requires the user to specify several other parameters related to the plasma and simulation numerics. In this section, we discuss our choices for these parameters, and assess the sensitivity of the simulation to some of these numbers.

While some of these parameters do not have a large effect on the Hall thruster plasma, or can be directly measured, others are more uncertain. Chief among these are the cathode plasma properties. **Figure 4.5a** depicts a centrally-mounted hollow cathode with its associated magnetic field and field-aligned coordinate system. Electrons emitted from the cathode stream outward along the magnetic field lines until they eventually encounter the main Hall thruster ion beam. Some of these electrons remain with the beam to keep it electrically neutral as it leaves the device, while the remaining others then stream toward the anode across the field lines. While Hall2De simulates the cathode plume, it does not simulate the interior of the cathode. As a result, the properties at the cathode orifice, namely electron temperature and ionization fraction, must be specified. Following standalone experimental measurements of a cathode similar to the one used in the H9,[78] we fixed the electron temperature and ionization fraction at 3 eV and 5%, respectively. In **Figure 4.4**, we show how the thrust and discharge current of an example simulation vary as these parameters are altered. We find that these global performance metrics are largely insensitive to the cathode orifice electron temperature and ionization fraction, especially near our chosen values. The one exception to this is that the simulation thrust increases by 3% as the electron temperature at the cathode orifice increases from 3 to 5 eV.

An additional uncertain parameter relates to the question of anomalous electron transport in the cathode plume. As with the thruster channel, experimental and computational evidence have shown that anomalous transport exists in the plumes of hollow cathodes like those used in Hall thrusters.[79,80] There may be two primary types of anomalous electron transport in the cathode plume.[80] The first is ion-acoustic in nature
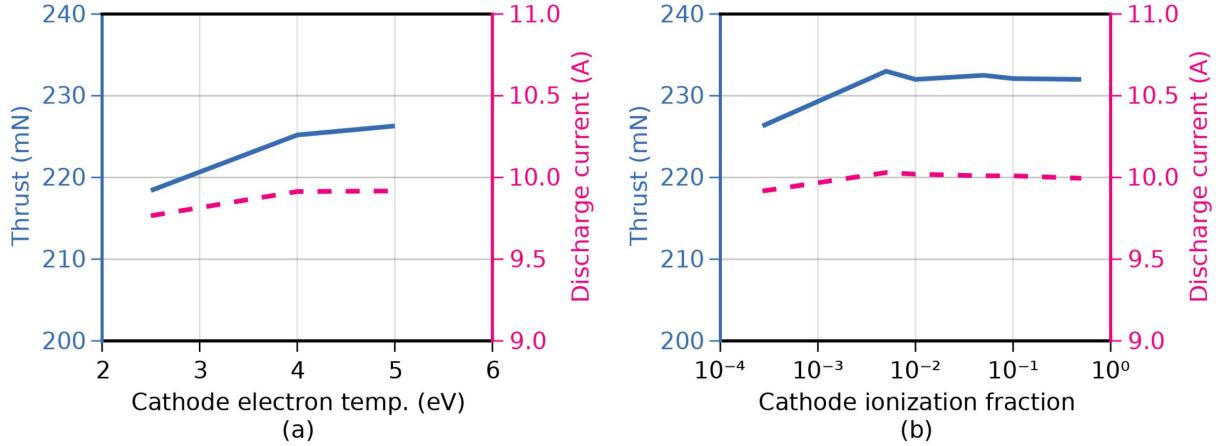
Figure 4.4: Performance quantities versus changing cathode electron temperature (a) and ionization fraction (b).

and acts primarily along the magnetic field lines. The other results from azimuthally-propagating anti-drift waves[81] and yields cross-field electron transport. In this work, we consider the effect of the parallel transport only. In thruster models, the effect of this anomalous collision frequency on the electron dynamics can be represented with a generalized Ohm's law in the same manner as the anomalous transport in the rest of the thruster, albeit in the field-parallel rather than field-parallel direction. Following the derivation in Chapter 2, we obtain the following equation:

$$\frac{m_e \nu_e}{e^2 n_e} j_{e\parallel} = E_\parallel + \frac{\nabla_\parallel p_e}{e n_e}. \tag{4.9}$$

Here, we see that the field-aligned potential and pressure gradients $E_\parallel$ and $\nabla_\parallel p_e$ scale with the anomalous collision frequency for a fixed density n_e and parallel electron current $j_{e\parallel}$. We can see from this equation introducing an additional anomalous collision frequency in the cathode plume can lead to steeper gradients in both pressure and electric field. This is consistent with experimental measurements. Indeed, it has been shown that incorporating anomalous transport in the cathode plume can significantly increase the agreement between Hall thruster simulations and experiments.[82] As is the case with the Hall thruster channel dynamics, however, introducing this additional collision frequency
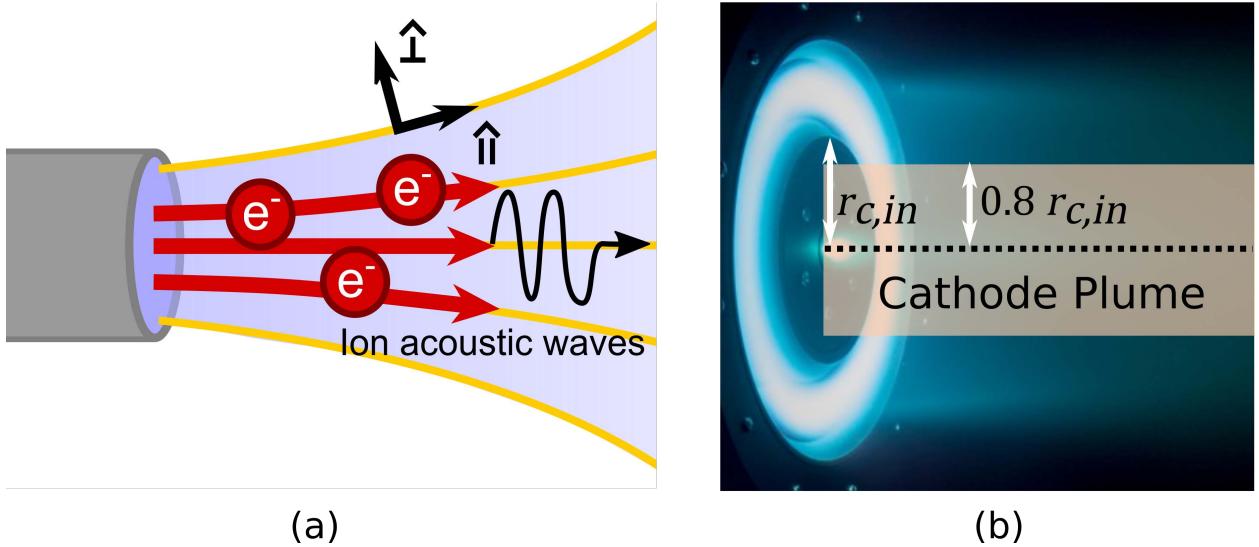
Figure 4.5: (a) Magnetic field lines (yellow) and magnetically-aligned coordinate system of a centrally-mounted hollow cathode. Ion acoustic waves propagate primarily in the field-aligned direction. (b) The region of the plume in which the cathode anomalous collision frequency model is applied. Here, $r_{c,in}$ is the inner-channel radius.

again opens the governing fluid equations. We therefore need to select a closure model for the cathode anomalous collision frequency.

Following the work of Mikellides et al[79] and Jorns et al,[80] we adopt a model for the transport based on the Sagdeev model[83]:

$$\nu_{an,cathode} = \alpha M_e \omega_{pi} \frac{T_e}{T_i}. \tag{4.10}$$

Here, M_e is the electron Mach number, ω_{pi} is the ion plasma frequency, and α is a proportionality constant of order 0.01[84,85]. In this model, the spectrum of the ion acoustic turbulence is saturated by nonlinear ion Laundau damping. The wave grows linearly with the electron drift speed, leading to a dependence on the electron Mach number. When applying this model to the cathode plume, we typically assume that the ion acoustic turbulence is able to heat the ions, so the ion temperature, T_i, self-consistently depends on the turbulent collision frequency.
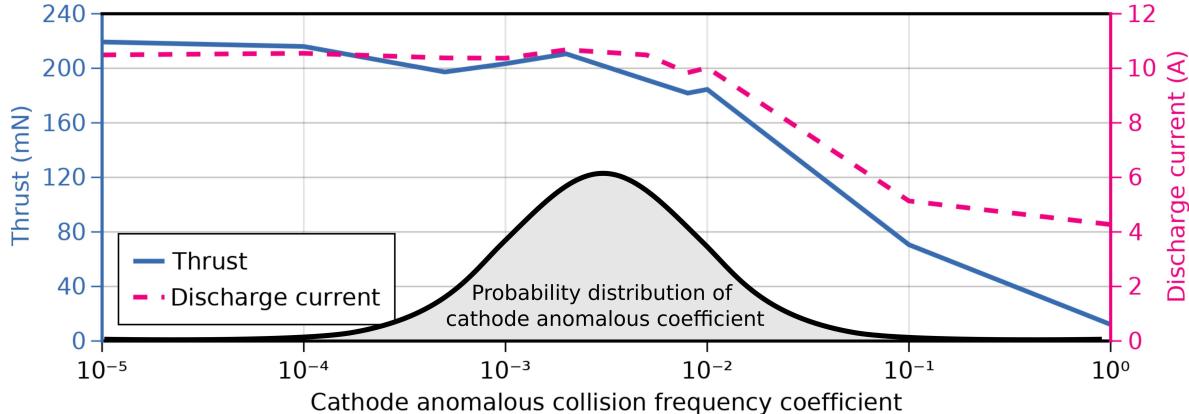
Figure 4.6: Performance quantities versus cathode anomalous collision frequency.

While it has been shown that this model may not be able to accurately capture oscillations in the cathode plume[86], it has been used with great success to improve the fit between experiment and simulation in hollow cathodes[79] and Hall thrusters.[82] In terms of actual implementation for this closure, unlike the case of the wave-driven effects in the Hall thruster channel, the IAT will not persist far downstream of the cathode. To represent this in practice, we therefore confine the region where this closure is adopted. We show this region in **Figure 4.5b**. Outside of this region, the cathode contribution to the anomalous cathode collision frequency is assumed to be negligible. In **Figure 4.6**, we show how the thrust and discharge current vary with increasing values of the cathode anomalous collision frequency. In some of our initial work (described in the next chapter), we treated this value as uncertain and sampled from a probability distribution function, also shown on **Figure 4.6**. However, the global performance metrics were largely insensitive to the cathode anomalous collision frequency coefficient up to a value of 0.01, so in subsequent work we used a fixed value of 0.005 for this coefficient.

Consistent with the experimental operating conditions described in the next section, we set the cathode flow rate to be 7% of the anode flow rate. In addition, we fixed both the neutral gas temperature and wall temperature to 500 K, which are consistent with

Parameter	Value
Maximum charge state	3+
Number of fluids	2 (beam and cathode)
Number of cells (MFAM)	3925
Number of cells (rectilinear grid)	3955
Cathode flow fraction	7%
Cathode ionization fraction	5%
Cathode electron temperature	3 eV
Cathode plume anomalous collision frequency coefficient	0.005
Neutral temperature	500 K
Wall temperature	500 K

Table 4.2: Numerical parameters employed in this work

typical operating temperatures of the Hall thruster being studied. Finally, in Table 4.2, we summarize the simulation parameters employed throughout this thesis.

4.2.4 Grid convergence

In order to ensure that the field-aligned mesh had sufficient resolution, we carried out a grid convergence study. We tracked how the thrust and integrated velocity residual (a measure of how the ion velocity as a function of space differs from experiment, defined in more detail in the final section of this chapter) varied as we increased the number of cells in the computational domain. The results of this study are presented in **Table 4.3**. Additionally, we plot in **Figure 4.7** the fractional error in these performance metrics against the grid resolution. We find that the error decreases with increasing resolution. At the penultimate resolution of 3438 cells, all quantities differ by less than 5% from those at the finest studied resolution of 3925 cells. We note that the convergence is not monotonic—increasing the number of cells does not always decrease the error. This is likely due to the manual field-aligned mesh generation procedure, which makes it difficult to ensure that the resolution smoothly increases across the whole domain as the number of cells is increased. Despite these caveats, our chosen resolution of 3925 cells for the field-aligned mesh appears to be sufficiently converged.

Number of cells	T	I_d	IVR
949	244.2 mN	13.57 A	0.200
1423	254.1 mN	16.49 A	0.118
1949	226.6 mN	15.70 A	0.057
2176	243.1 mN	14.64 A	0.101
2750	232.9 mN	16.44 A	0.060
3438	249.2 mN	15.95 A	0.073
3925	258.3 mN	15.20 A	0.072

Table 4.3: Thrust (T), discharge current (I_d), and integrated velocity error for different grid resolutions.
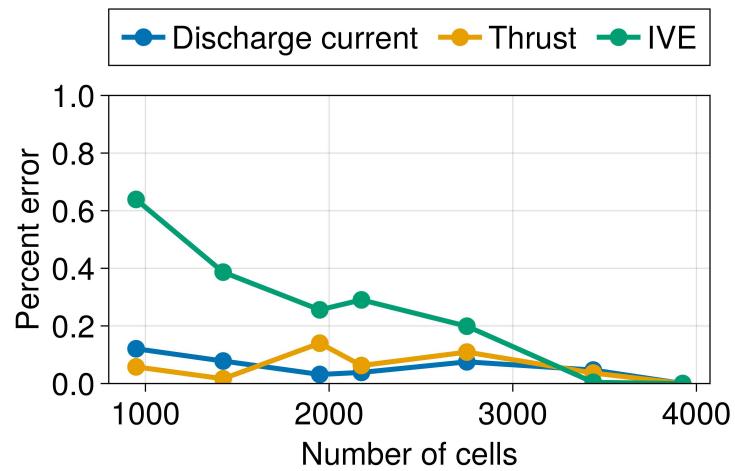
Figure 4.7: Variation of error in quantities of interest with increasing grid resolution.

All simulations in this this thesis were carried out on the Great Lakes supercomputing cluster at the University of Michigan. This enabled us to run many simulations in parallel. As Hall2De is designed to be run on workstation-class computers, each instance of the code used eight CPU cores and a single compute node. The time required to complete a simulation varied with the grid resolution, simulation duration, and chosen timestep, but typical values ranged between 12 and 24 hours.

4.3 Thruster

In this work, we primarily simulate the H9 Hall thruster. This is a 9-kW class laboratory Hall thruster developed in a collaboration between the University of Michigan, the Air Force Research Laboratory, and the Jet Propulsion Laboratory.[87,28] In **Figure 1.2a**, we show the H9 operating at a discharge voltage of 300 V and a discharge current of 15 A on its design propellant of xenon. In the years since its development, the H9 has subsequently been tested extensively using krypton[1] **(Figure 1.2b)**, as well as at discharge current densities well outside of its design envelope (up to 150 A).[18] In this work, however, we confine our investigations to discharge currents between 15 and 30 A, and discharge voltages between 300 and 600 V, due to the lack of spatially-resolved data at these extreme conditions.

Throughout this thesis, we compare our simulation results to experimental data obtained from this thruster. These data include global performance metrics—such as the thrust, discharge current, and specific impulse—as well as the spatially-resolved mean ion velocity along channel centerline. These latter data were obtained via the method of laser-induced fluorescence (LIF).[88] This type of local experimental plasma data is preferred for model validation as it is collected non-invasively. In contrast, probe-based methods to perform local measurements have shown in some cases to perturb the plasma properties[89]. We use the experimental thrust and discharge current data taken by Su and Jorns in 2021,[1] in combination with ion velocity profiles measured by the same authors in

2022.[53] All experimental data were collected in the Alec D. Gallimore Large Vacuum Test Facility at the University of Michigan.

4.4   Metrics for comparison to experiment

In order to quantitatively compare the performance of our simulations to experimental data, we consider five metrics. The first of these metrics is the thrust, which is the amount of propulsive force generated by the device. To compute the thrust from the simulation results, we integrate the flux of axial momentum over the outflow boundaries of the simulation:

$$T = \iint_{\text{outflow}} \sum_{f=1}^{2} \sum_{j=0}^{3} M \, n_{f,j} u_{z,f,j} (\vec{u}_{f,j} \cdot \hat{n}) \, dS. \tag{4.11}$$

In the above expression, $\vec{u}_{f,j}$ is the velocity vector of particles belonging to fluid f with charge j, $u_{z,j,f}$ is the component of that vector parallel to the $\hat{z}$ axis, $n_{f,j}$ is the number density of fluid f with charge j, dS is the differential surface area, $\hat{n}$ is the surface normal vector, and M is the mass of a xenon atom.

The second metric we consider is the discharge current I_d, which is the current carried by ions and electrons from anode to cathode. We compute I_d from our simulations by integrating the sum of the ion and electron currents over the anode boundary surface:

$$I_d = \iint_{\text{anode}} \left[\left(\sum_{f=1}^{2} \sum_{j=1}^{3} j \, q \, n_{i,f,j} \vec{u}_{i,f,j} \right) + \vec{j}_e \right] \cdot \hat{n} \, dS. \tag{4.12}$$

Here, q is the fundamental charge and $\vec{j}_e$ is the electron current density vector. The third metric for comparison is the anode efficiency η_a, which measures the fraction of the discharge power converted into useful thruster power. It is defined as

$$\eta_a = \frac{1}{2} \frac{T^2}{\dot{m}_a V_d I_d}, \tag{4.13}$$

where $\dot{m}_a$ is the mass flow rate injected through the anode and V_d is the discharge voltage. The fourth metric we employ is the *integrated velocity residual* (IVR). We also refer to this as the *integrated velocity error* (IVE) in various places in this dissertation. This measures how well the simulation predicts the ion velocity along the discharge channel centerline. We define this as

$$\text{IVR} = \sqrt{\frac{\int_{z_0}^{z_N} (u_{i,1}(z) - u_{i,2}(z))^2 dz}{\int_{z_0}^{z_N} u_{i,1}^2(z) dz}}. \tag{4.14}$$

In the above, z_0 and z_N are the axial locations of the first and last ion velocity data-points, respectively, and $u_{i,1}$ and $u_{i,2}$ are the axial components of the ion velocity obtained from two different measurements or simulations. Higher values of the IVR correspond to worse agreement between simulation and experiment. For example, if the ion velocity curves differ by 10% on average, then the IVR should be approximately 0.1.

The final metric we use in this work is the *integrated anomalous collision frequency residual* (IAR). This is a measure of how well the anomalous collision frequency profiles from two simulations agree with each other. This has the same functional form as the integrated velocity error:

$$\text{IAR} = \sqrt{\frac{\int_{z_0}^{z_N} (\nu_{AN,1}(z) - \nu_{AN,2}(z))^2 dz}{\int_{z_0}^{z_N} \nu_{AN,1}^2(z) dz}}. \tag{4.15}$$

We note that all of the metrics we consider here are time-averaged, and we do not consider time-dependent metrics like the oscillation frequency and amplitude in the present dissertation. The oscillation characteristics are sensitive to changes in simulation parameters not related to the anomalous collision frequency, such as the temperature and velocity of the injected neutrals. This means that they may be sufficiently diagnostic to use in evaluating closure models. Additionally, as we will see, the present state of anomalous transport modeling in Hall thrusters is such that matching these time-averaged metrics is very difficult, so it may be premature to consider time-dependent properties.

4.5 Summary

In this chapter, we described the code we use for the simulations in the remainder of the thesis, as well as the thruster being simulated and the data to which we compare our simulation results. We then introduced the metrics that allow us to evaluate the efficacy of our simulations. In the next chapter, we will discuss the first application of these tools in this thesis.

CHAPTER 5

Model Calibration Against Empirical Transport Profiles

> Engineering is too important to wait for science.

————————————————————————————

Benoît Mandelbrot

In **Chapter 3**, we described the common practice of tuning the spatial variation of the anomalous collision frequency to match experimental data. We also introduced the data-driven approach to anomalous transport modeling, which relies on these "empirical" anomalous collision frequency profiles as a "ground truth" from which physical scaling laws can be inferred. However, as these data-driven models had not yet been implemented directly into a Hall thruster code, it was still unclear how predictive they were. In this chapter, we address this question in two parts. First, we implement one of these data-driven models into Hall2De in order to investigate its performance. We run hundreds of simulations of the H9 at multiple operating conditions. Each of these simulations uses a set of model coefficients sampled from a distribution calibrated to match empirical profiles from thrusters not including the H9. This gives us probabilistic predictions of thruster performance using the data-driven model. Second, we examine the broader question of how valid it is to tune a model's coefficients to match empirical transport profiles in the

first place. For a single operating condition, we calibrate an empirical anomalous collision frequency profile for the H9 thruster. We then tune the model coefficients of the data-driven model and the model of Lafleur et al. to match the empirical profile. After implementing the models into Hall2De, we investigate how well the initial calibrated results correspond to the outputs of our self-consistent simulations. Finally, we discuss our results in the context of future closure modeling efforts.

5.1 Probabalistic evaluation of a data-driven transport model in Hall2De

This work was previously presented at the 2021 AIAA Propulsion and Energy Forum as **Reference 90**.

As discussed in **Chapter 3**, Jorns[3] applied symbolic regression to a dataset of ten empirical anomalous collision profiles inferred from experimental measurements to propose algebraic closure models for this coefficient. This effort produced several models capable of predicting the anomalous collision frequency profile of the H9 operating at 300 V and 15 A with greater accuracy than previously-proposed first principles models. These first principles models included some of the models we introduced in **Chapter 3**, namely those of Cappelli et al., Lafleur et al., and Scharfe et al. Notably, the H9 was not in the training dataset, so this result held promise for the extensibility of the model's fit coefficients beyond its training dataset. However, the models were not self-consistently implemented into a Hall thruster simulation. Instead, they were computed on the time-averaged plasma properties output by calibrated Hall thruster simulations, such as those shown in **Figure 3.1b**. As a result, it was unclear whether these predictions would yield converged simulations with thruster performance and plasma properties similar to experiment.

To assess this, we expand the implementation of one of these data-driven closure models to be applied locally everywhere in the domain, making for a more self-consistent test. We also expand the dataset for comparison by simulating five operating conditions on the H9 magnetically-shielded Hall thruster. We perform a statistical characterization of the model's performance by running $O(100)$ simulations per condition, each simulation using a coefficient set drawn from an inferred distribution. Lastly, we treat the scaling coefficient of the cathode anomalous collision frequency as uncertain, and sample from a distribution of this coefficient in concert with the other sampled model fit coefficients. This process yields probabilistic predictions of thrust, discharge current, anode efficiency, and centerline plasma properties at each condition, which are then compared to experiment.

5.1.1 Model calibration

While the data-driven model we presented in **Chapter 3** featured only two fit coefficients, the original version had three. This version of the model is given by

$$\nu_{an} = \omega_{ce} \left(c_0 + \frac{c_1 |\mathbf{u}_i|}{c_2 c_s + v_{de}} \right).$$
(5.1)

Jorns found that $c_0 = -3.37 \times 10^{-2}$, $c_1 = 2.39$, and $c_2 = 3.32$ gave the best fit to the empirical anomalous collision frequency profiles in the training dataset, but in this work we attempt to calibrate these coefficients in a probabilistic manner.

Both the data-driven model and the cathode anomalous collision frequency model have free parameters that must be calibrated against data. However, there is inherent uncertainty in the coefficients stemming primarily from uncertainty in the proposed model. We want to capture this rigorously and in turn be able to quantify how it impacts certainty in the integrated model predictions. To this end, we formulate our model calibration process as a Bayesian inference problem. In this approach, we represent the model coefficients, c_0, c_1, etc. probabilistically, i.e. as distribution functions. The nature of these distribution

functions can be inferred by comparing the model predictions to data. The algorithm for this is based Bayes' theorem **(Equation 5.2)**, where we obtain an expression for the probability P of that model parameters $\Theta = \{c_0, c_1, c_2\}$ being correct given our model f that we have observed the data d:

$$P(\Theta \mid d, f) = \frac{P(d \mid \Theta)P(\Theta \mid f)}{P(d)}. \tag{5.2}$$

Here $P(\Theta \mid d)$ is the *posterior probability*, or the probability of the model parameters being correct after seeing the data, given the model; $P(d \mid \Theta, f)$ is the *likelihood*, or the probability of seeing the data assuming the model parameters are correct, given the model; $P(\Theta)$ is the *prior probability*, or the probability of the model parameters being correct before seeing the data, and $P(d)$ is the *evidence*, or the probability of seeing the data.

5.1.1.1 Closure model for cross-field transport

In calibrating the our data-driven model, we take d to refer to the set of nine piecewise-linear calibrated multi-zone Bohm anomalous collision frequency profiles from validated simulations described in **Reference 3**. This data spans four thrusters at multiple operating conditions, with powers ranging from 1 to 6 kW.

The prior probability, $P(\Theta|f)$, encapsulates our knowledge of what the model coefficients should be prior to the inference procedure. We assume that the prior distributions are Gaussians (i.e. normal distributions). This has the advantage of allowing the parameters to take any value, while biasing them toward regions of higher confidence. In the absence of more detailed information, we assume the prior probability of the coefficient set $\{c_0, c_1, c_2\}$ to be the product of three independent normal distributions centered on previous estimates of these coefficients. We use $c_0 = -0.0337, c_1 = 2.39, c_2 = 3.32$ as the mean of our prior distribution, as these were the coefficients initially described in **Reference 3**. We then

have $c_0 \sim \mathcal{N}(-0.0337, \sigma_{p0})$, $c_1 \sim \mathcal{N}(2.39, \sigma_{p1})$, and $c_2 \sim \mathcal{N}(3.32, \sigma_{p2})$. Here, σ_{pi} is the standard deviation of the prior distribution of parameter i. The notation $x \sim \mathcal{N}(\mu, \sigma)$ refers to sampling from a normal distribution with mean μ and standard deviation σ.

The overall prior probability density function for a model with N_p parameters with normal priors is given by

$$P(\theta) = \prod_{i=1}^{N_p} \frac{1}{\sigma_{pi}\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\theta_i - \mu_{pi}}{\sigma_{pi}}\right)^2\right]. \tag{5.3}$$

The prior standard deviation of coefficient i, σ_{pi}, is a hyperparameter, an uncertain number which is not formally part of the parameter set being inferred. Tuning hyperparameters from data is a nuanced problem[91,92,93] and one outside of the scope of this work. Instead, we simply assume $\sigma_p = 3.0$ for all three parameters, which gives the prior distributions a similar variance to the previous work without explicitly constraining them to a fixed range of values.

As with our priors, we assume the likelihood, $P(d|\Theta, f)$ is Gaussian in form. The full expression for our likelihood is:

$$P(d \mid \Theta, f) = \frac{1}{N}\sum_{i=1}^{10}\sum_{j=1}^{N} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\log_{10}\nu_{AN}(T_i, z_j) - \log_{10}f(\Theta, T_i, z_j)}{\sigma(T_i, z_j)}\right)^2\right]. \tag{5.4}$$

For each of the training datasets $\{T_1, T_2, ..., T_{10}\}$ we have N grid points $\{z_1, z_2, ..., z_N\}$ at which various surrogate plasma data, including the calibrated anomalous collision frequency, are known, and at which our model f for the anomalous collision frequency can be evaluated. At each grid point, we assume the data is equal to the prediction plus some Gaussian noise. $\sigma(T, z)$ is the standard deviation of this noise at axial location z on the centerline of thruster T. We infer the base-ten logarithm of the anomalous collision frequency so that we do not bias our predictions toward regions where the anomalous

collision frequency is high. This is important because the anomalous collision frequency is typically low in the ionization and acceleration regions, and the solution is very sensitive to the value of ν_{an} in this region.[94] Taking the logarithm of the anomalous collision frequency in the likelihood evaluation means a factor of two difference when $\nu_{an} \sim 10^6$ Hz is treated the same as when $\nu_{an} \sim 10^9$ Hz, prioritizing a solution which matches the data well at all orders of magnitude.

We adopt a three-zone model for the noise. Mikellides and Lopez Ortega[94] have shown that Hall2De is insensitive to changes of the anomalous collision frequency of up to two orders of magnitude when $z/L \lesssim 0.65$ and up to a factor of two in the plume, when $z/L \gtrsim 1.5$, where L is the thruster channel length. The noise model we use in this work weights these regions accordingly:

$$\sigma(T_i, z_j) = \begin{cases} 3\sigma_N & z/L < 0.65 \\ \sigma_N & 0.65 \leq z/L \leq 1.5 \\ 2\sigma_N & z/L > 1.5 \end{cases} \tag{5.5}$$

Here, the "base noise" σ_N is another hyperparameter, which we set to 2.

The evidence term in **Equation 5.2**, $P(d)$, is typically difficult to evaluate. However, if data d is fixed with respect to the model f and parameters Θ, it acts solely as a normalizing constant on the posterior probability. The Markov chain Monte Carlo methods (MCMC) we employ in this work can sample from un-normalized posterior distributions, so evaluating the evidence is unnecessary.

Armed with these inputs, we now have a prescription for evaluating the probability distribution of coefficients. To visualize these distributions and to obtain sample coefficient sets with which we can run simulations, we employ the Delayed-Rejection Adaptive Metropolis algorithm[95] to sample randomly from the distribution. **Figure 5.1** shows the resulting joint and marginalized probability distributions for the model coefficients from **Equation 5.1** after drawing one million samples. The three histograms on the diagonal
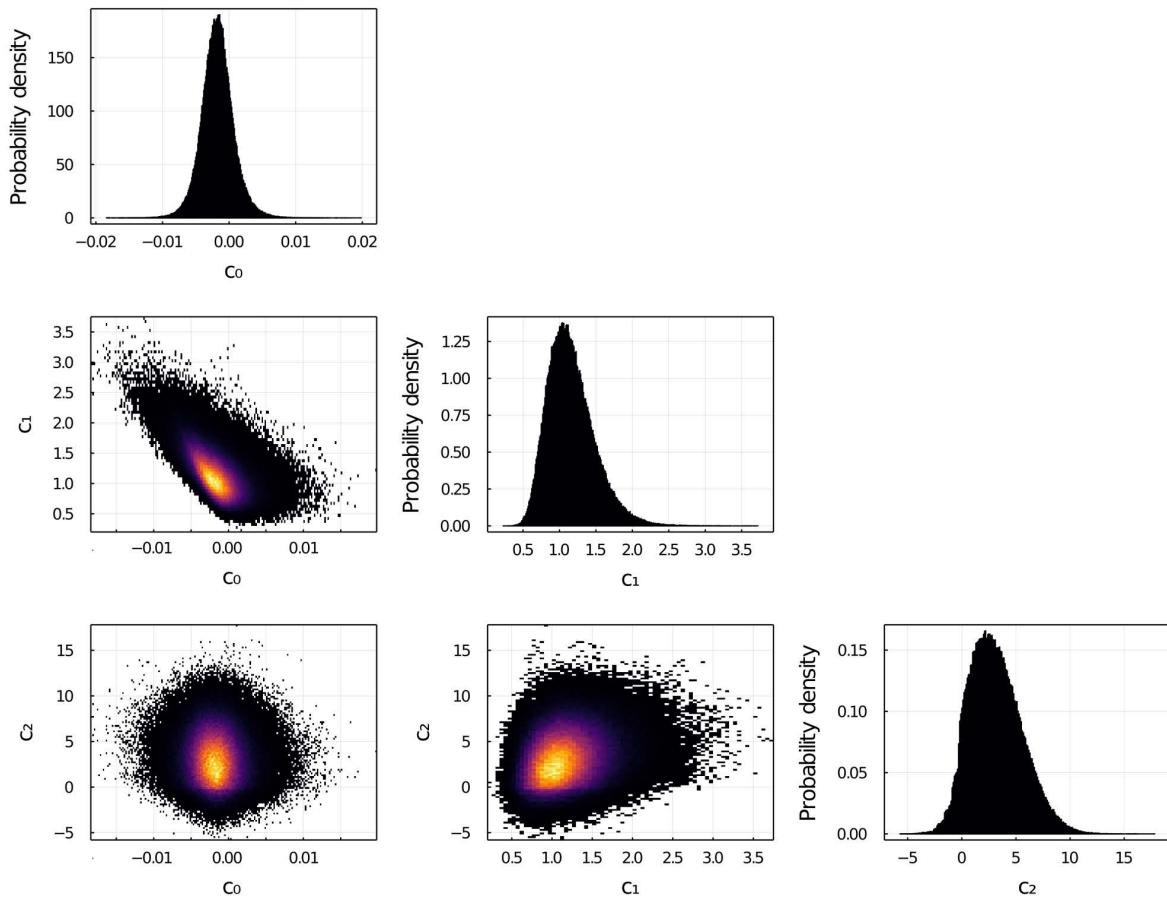
Figure 5.1: MCMC posterior probability distribution function on the model coefficients. One million samples were drawn from the posterior. The acceptance percentage was 40%. Bright regions are regions of high probability (many accepted samples), while dark regions have low probability (few samples).

of **Figure 5.1** are the marginal probability density functions of each of the parameters c_0, c_1, and c_2. These represent the overall probability of seeing each of the the three parameters without reference to the other two parameters. The three plots in the bottom left of **Figure 5.1** represent the joint distributions of the coefficients with respect to each other, and show how, if at all, the coefficients are correlated. We can see that the joint distribution of c_0 and c_1 appears roughly elipsoid, with the long axis pointing down and to the right. This indicates that c_0 and c_1 are negatively correlated, which, with reference to **Equation 5.1**, makes sense, as a smaller (more negative) value of c_0 would require a larger value of c_1 to give a similar result. The other parameters appear uncorrelated, as their joint distributions are roughly circular in shape.

Using our probability distribution of fit coefficients, we can visualize how the variance in c_0, c_1, and c_2 affect the predictions of our model. To this end, we show in **Figure 5.2** the resulting predictions for the centerline anomalous collision frequency for two of the nine training points. To generate these plots, we evaluate the model on the data at each sampled set of fit coefficients. At each grid point, we compute the median predicted value of ν_{an} and a 95% high density credible interval (the smallest interval containing 95% of the predictions). We plot the median prediction of ν_{an}/ω_{ce} in black and the credible interval in blue. The red dashed line represents the empirical anomalous collision frequency profile from the validated simulations that we were calibrating against.

The most probable coefficients according to our procedure are $\{c_0, c_1, c_2\}$ = $\{-0.0015, 1.04, 2.16\}$, which differ somewhat from those obtained in **Reference 3**. This can be attributed to the fact that our updated likelihood de-emphasizes the downstream anomalous collision frequency in favor of a closer match in the acceleration region.

Figure 5.2a depicts the worst fit among the ten produced, while **Figure 5.2b** is among the better fits. We see that both fits agree well with the inferred anomalous collision frequency just downstream of the exit plane ($1 \leq z/L \leq 1.25$) while over-predicting the upstream collision frequency and under-predicting the downstream collision frequency.
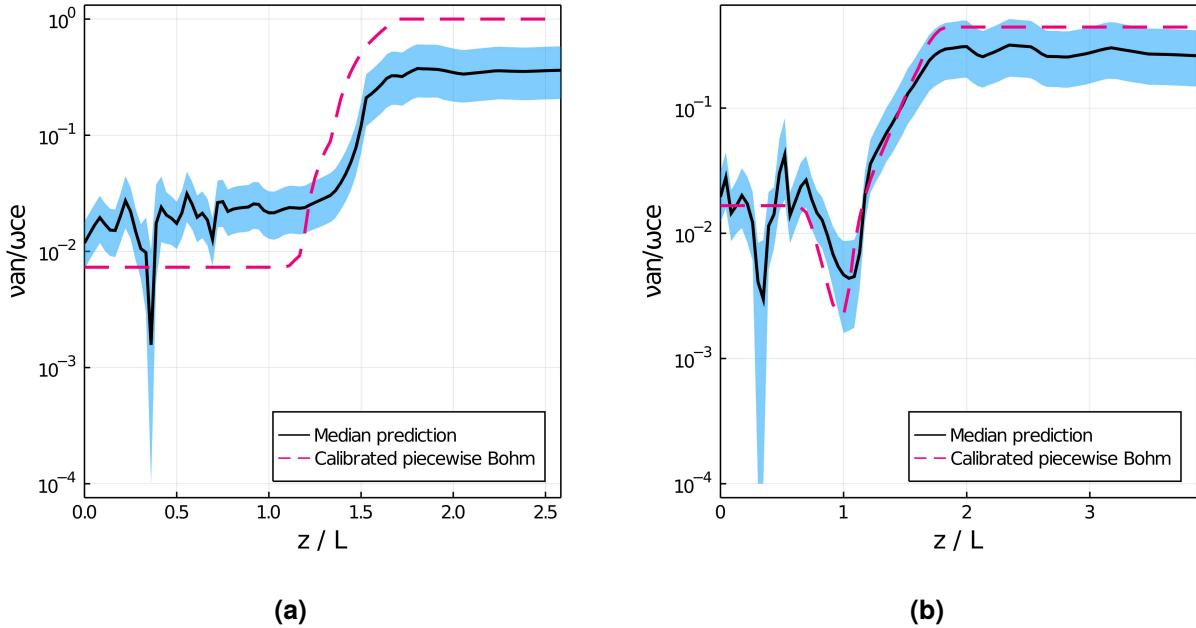
Figure 5.2: Simulated centerline anomalous collision frequency profile as predicted by the data-driven model for thrusters 1 (a) and 5 (b) in the training data. The blue filled area represents 95% credible interval.

Referring to **Figure 5.1**, we see that c_0 is typically very small, c_1 is close to unity, and c_2 is $O(5)$. Since the ion sound speed c_s is typically a few orders of magnitude smaller than the electron drift velocity v_{de}, we can simplify Equation 5.1 to show that our data-driven model scales with:

$$\frac{\nu_{an}}{\omega_{ce}} \approx c_1 \frac{|\mathbf{u}_i|}{v_{de}}.$$

This means that the goodness of fit is largely constrained by the ratio of the ion velocity to the electron drift velocity, and that the shape of the profile is insensitive to the coefficient choice. The coefficient with the largest impact, c_1, only serves to scale the profile up or down. While we get a very good fit considering the simplicity of the model, the only coefficient with the ability to affect the shape of the curve, c_2, has only a minor impact on the solution, so we fit some of the test data points better than others.

5.1.1.2 Closure model for cathode plume

There is significantly less data available for calibrating the cathode anomalous collision frequency profile than for the beam collision frequency. The data we rely on for this work comes from Georgin and Jorns,[86] who compared the predictions of several cathode anomalous collision frequency models to time-resolved measurements. They found a best fit coefficient of α = 0.003 for the Sagdeev closure model which we employ in this work. **(Equation 4.10)**. In the absence of other data to compare against, we assume that $\log_{10}(\alpha)$ is normally distributed about $\log_{10}(0.003)$, with a standard deviation of a factor of three in either direction:

$$\log_{10}(\alpha) \sim \mathcal{N}(\log_{10}(0.003), \log_{10}(3)). \tag{5.6}$$

This distribution, while centered about the best fit point in Ref[86], encapsulates our uncertainty about α in two ways. First, other studies have found a value of $\alpha \sim 0.01$ gives good agreement between Hall thruster simulations and data,[82] and **Equation 5.6** gives significant probability weight in this region. Second, Georgin's data comes from the cathode employed in the H6 thruster,[78] which is a slightly different cathode than the one used in the H9. It is therefore unreasonable to constrain our distribution of α too narrowly around his estimate.

5.1.2 Simulations

For this part of the work, we examined five operating conditions, summarized in **Table 5.1**. Each case used xenon as the propellant. The performance data in this table was collected in 2021 at the University of Michigan.[1] We compare our simulation results to both this data and laser-induced fluorescence measurements of the ion axial velocity in the acceleration region. This data was collected in 2018, also at the University of Michigan. We note that this (unpublished) data was collected on a slightly different thruster configuration (the so-

Case	V_d (V)	I_d (A)	$\dot{m}_a$ (mg/s)	T (mN)	P_b (Torr)	Simulations
1	300	15	14.8	292.9 ± 3.5	4.8×10^{-6}	100
2	300	20	18.5	377.6 ± 3.8	5.8×10^{-6}	98
3	400	15	15.4	350.8 ± 3.4	5.0×10^{-6}	99
4	500	15	16.1	405.1 ± 3.4	5.3×10^{-6}	97
5	600	15	15.4	447.2 ± 3.0	5.4×10^{-6}	100

Table 5.1: Experimental discharge voltage V_d, discharge current I_d, anode mass flow rate $\dot{m}_a$, measured thrust T, and measured background pressure P_b, and number of simulations for all test cases.

called H9v1) than the LIF data we described in Chapter 4, which we use in the remainder of this work and was obtained in 2022.[53]

In this work, we evaluate data-driven model for the anomalous collision frequency on all field-perpendicular edges and in all cells. We specify the model coefficients in the simulation input file. For each operating condition, we sample 100 coefficient sets (c_0, c_1, c_2, α) from the data-driven coefficient distributions in **Figure 5.1** and the cathode anomalous collision frequency distribution **(Equation 5.6)**. We set up a simulation input file for each of the coefficient sets, and then run a simulation using that input. Each simulation uses a pre-computed initial condition corresponding to its operating condition. These were generated by running cold starts at that case conditions flow rate and voltage for 0.5 ms of simulation time. The simulations are then run in parallel on the Great Lakes supercomputer cluster at the University of Michigan. Each simulation uses 8 cores, and was run for 0.6 milliseconds of simulation time, at a timestep of 9 nanoseconds per iteration. Each batch of 100 parallel simulations took on average 15 hours to run. We include two ion fluids (one corresponding to the main beam and the other to the cathode beam) and up to triply-charged ions. In contrast to later work in this dissertation, we do not model the ion temperature or electron azimuthal momentum.
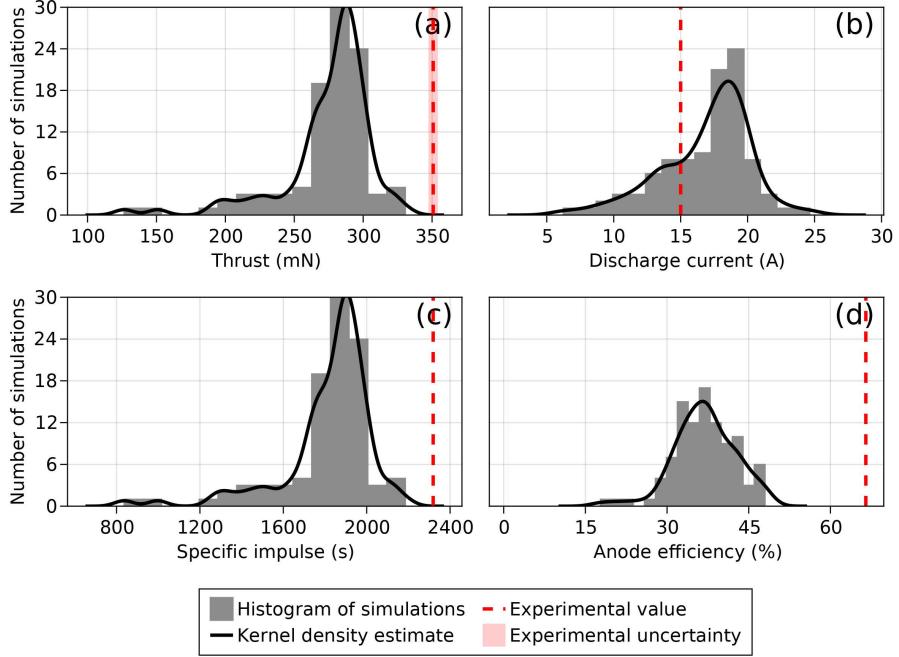
Figure 5.3: Histograms of performance metrics at the 400 V and 15 A condition.

5.1.3 Results

In this section, we first present probabalistic predictions for the performance of the 400 V and 15 A case, before showing how these trend to other operating conditions. In **Figure 5.3**, we present histograms of the predicted performance at the 400V and 15 A condition. The thrust is under-predicted, with a median thrust of 285 mN and a 95% credible interval of 210 - 310 mN, 19% lower than the experimental value of 350 mN **(Table 5.1)**. The discharge current is over-predicted, with a median of 18 A (95% credible interval of 12 - 20 A) compared to the 15 A of the experiment. The low thrust leads to a low specific impulse, with a simulated median I_{sp} of 1900 s (95% credible interval 1500 - 2200 s) compared to 2350 s in the experiment. Finally, we find a median simulated anode efficiency of 39% (95% credible interval of 30-45%) compared to 65% in the experiment. The experimental thrust, efficiency, and specific impulse are not contained within the simulated credible intervals.

5.1.4 Comparison of median performance metrics and experiment

To see how performance trends across voltages, we plot the median predictions and 95% credible intervals against the experimental data for all five conditions in **Figure 5.4**. We see that, like the experiment, the thrust increases linearly with voltage and is underpredicted by 40-60 mN for all conditions except the 600 V and 15 A condition, where the thrust is underpredicted by 120 mN. For 300 V, 400V, and 500 V (at 15 A), the discharge current does not change significantly with voltage, which matches experiment, though the simulated discharge currents are 3 to 4 A higher than experiment. At 600 V and 15 A, the median discharge current drops to a value closer to experiment. Interestingly, at the 300 V and 20 A condition, the median simulated discharge current is predicted nearly exactly, though there is high uncertainty in this value. As I_{sp} scales linearly with thrust, we see the same trends as we see in the **Figure 5.4a**, with a linear increase with voltage for all conditions except 600 V and 15 A. Finally, while the experimental the anode efficiency trends slightly upward at higher voltages, the median simulated anode efficiency trends slightly downward or remains flat.

In addition to accurately capturing the experiment trends with respect to voltage, the data-driven model also seems to predict the correct trends with increasing mass flow rate, though we only have two points for comparison (300 V and 15 A and 300 V and 20 A). The thrust increases with higher flow rate, though the increase in thrust in the simulation between 15 A and 20 A is only 40 mN, compared to a difference of nearly 100 mN in the experiment. As we would expect, the discharge current increases as well with increased flow rate, though again by a smaller magnitude than in the experiment (3 A in the simulation vs 5 A in the experiment for the same change in flow rate). The specific impulse and anode efficiency change little at higher flow rates in both simulation and experiment.

With these figures in hand, we can summarize the data-driven model's predictive performance. It underpredicts thrust by 40-75 mN at most conditions, overpredicts the discharge current by 2-4 A, underpredicts the specific impulse by 300-400s and
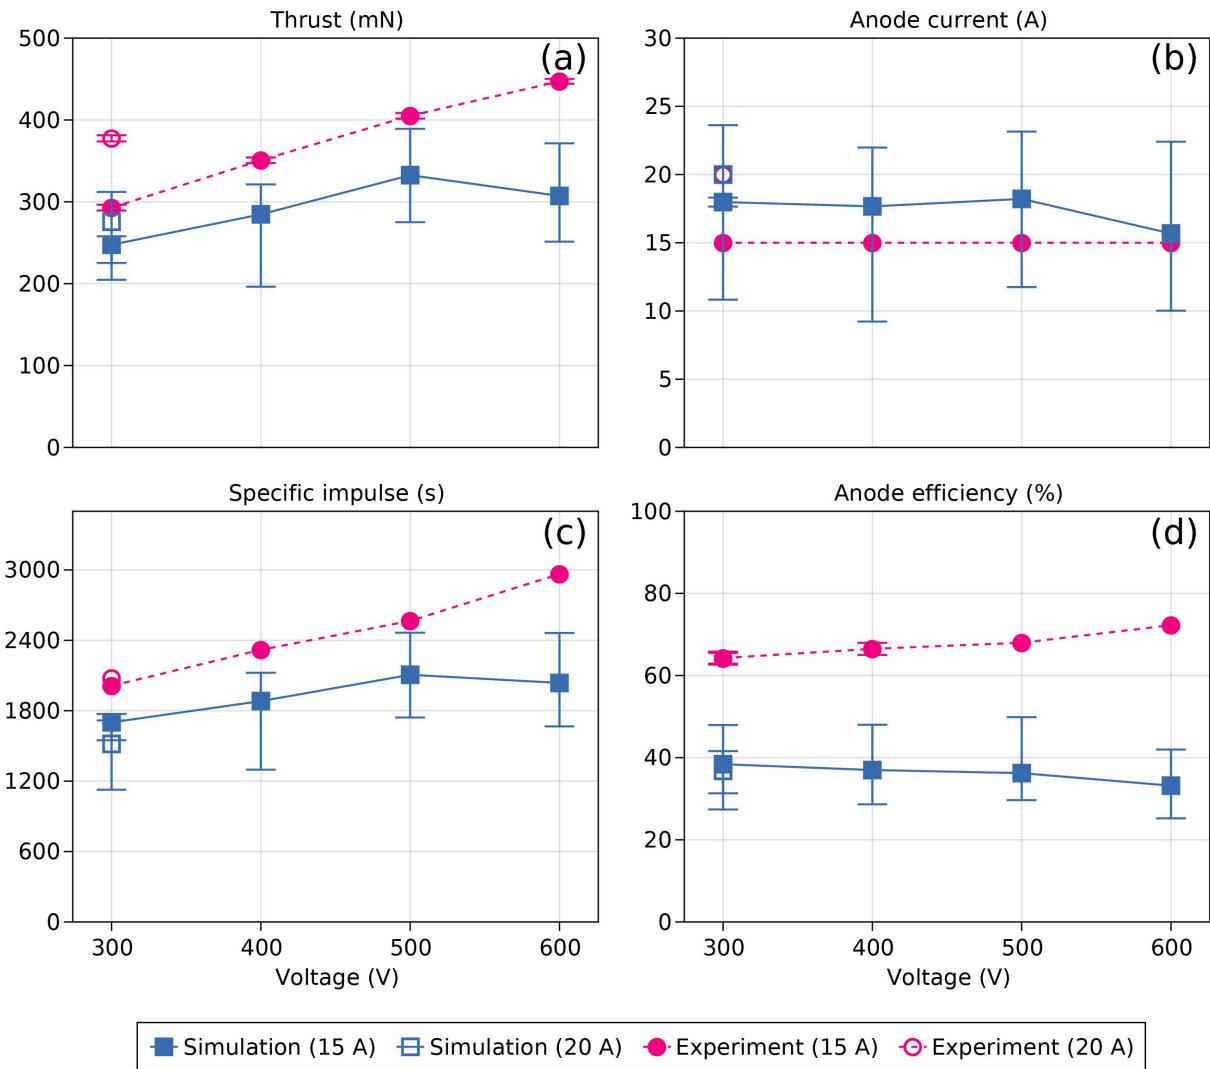
Figure 5.4: Evolution of a) thrust, b) discharge current, c) specific impulse and d) anode efficiency with respect to discharge voltage for both simulation (blue) and experiment (magenta). The markers represent median predicted values and the uncertainty bars represent 95% credible intervals on the predictions.

underpredicts the anode efficiency by less than a factor of two. However, with the exception of the 600 V and 15 A condition, we qualitatively reproduce experimental trends for thrust, specific impulse and discharge current with respect to both voltage and mass flow rate. We can now turn to analyzing what drives this lower performance.

In **Figure 5.5** we present simulated centerline profiles of the axial ion velocity, electric potential, normalized anomalous collision frequency, and electron temperature for the 400 V and 15 A condition. Compared to the laser-induced fluorescence measurements of ion velocity in **Figure 5.5a**, we find that the simulated ion velocity profile is significantly more shallow than the experiment, with an ion acceleration region two thruster-lengths long. This is four times the length of the experimental region. Consistent with this more protracted acceleration, the potential profile in **Figure 5.5b** is also shallow. The anomalous collision frequency profile **(Figure 5.5c)** looks qualitatively similar to those in **Figure 5.2a** and **Figure 5.2b**, but axially-stretched and shifted downstream, with the inflection point occurring at 3 thruster-lengths downstream of the exit plane. Near the exit plane, at $z/L =$ 1, the anomalous collision frequency is between 1/30 and 1 times the cyclotron frequency. This is unusual, as the anomalous collision frequency is typically a factor of two or three lower than the lowest value predicted by our model in this region. Lastly, we see that the electron temperature curve is very wide, nearly three thruster-lengths in extent. It reaches a median peak of 45 eV at an axial location of $z/L = 2$, with a variance of nearly 75% in both directions.

While the variance in the anomalous collision frequency is approximately uniform throughout the domain, the variance in predictions at the 400 V and 15 A condition for ion velocity, plasma potential, and electron temperature reach their maximum in the acceleration region, and decline deeper into the plume. These trends are representative of the results at other operating conditions.
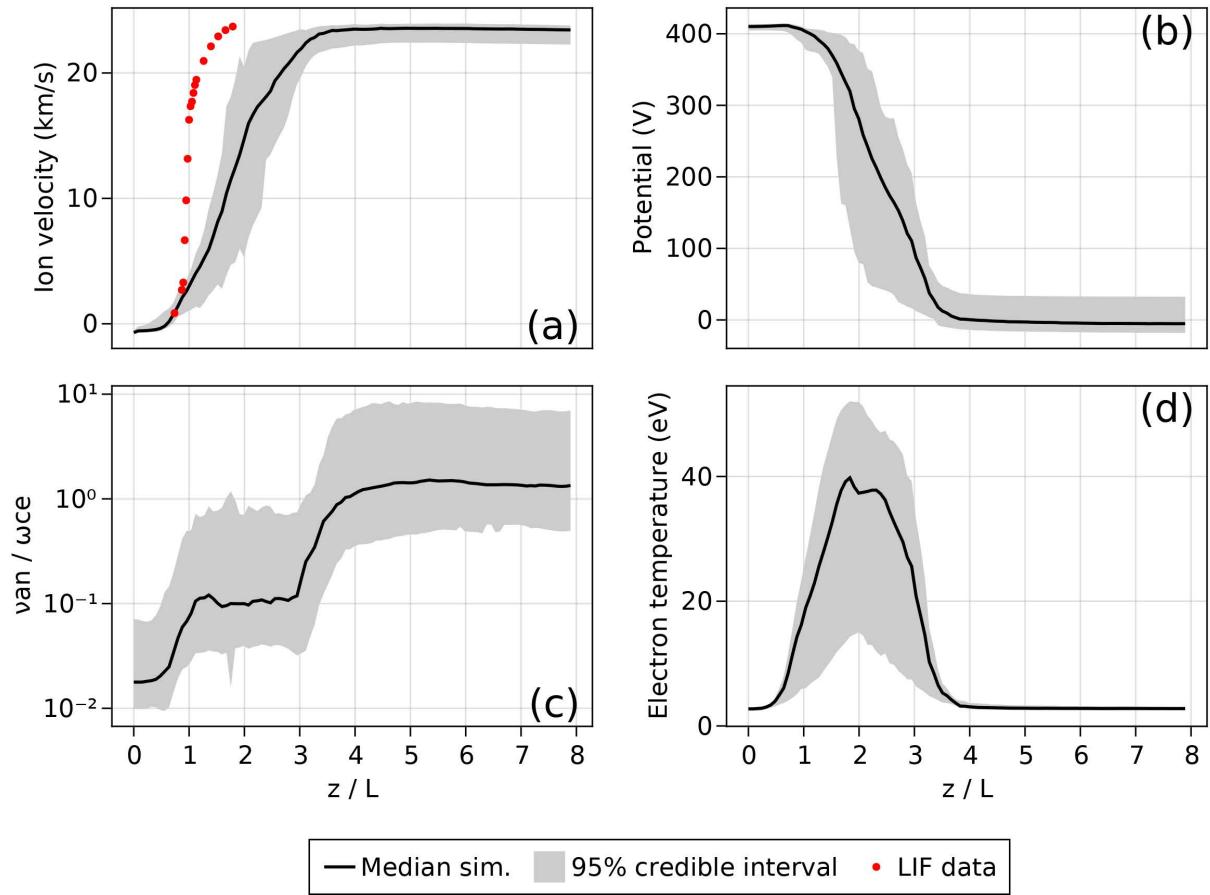
Figure 5.5: Simulated centerline profiles of a) ion velocity, b) potential, c) normalized anomalous collision frequency and d) electron temperature at the 400V and 15 A condition. Ion axial velocity data obtained via laser-induced fluorescence is depicted as red markers in (a).

5.1.5   Discussion and conclusion

As we discussed in the preceding chapters, the presence of low anomalous collision frequency near the thruster exit plane ($z/L \approx 1$) is critical for establishing the correct ion acceleration behavior in Hall thruster simulations. This transport barrier produces a strong electric field and thus steep ion acceleration. As there is no strong transport barrier in this region in our simulations, it is unsurprising that our predicted ion velocity profiles are very shallow. This has knock-on implications for the reduced global performance metrics we observed. The primary effect of this relaxed ion acceleration profile is increased plume divergence, which results in more of the ion momentum being lost in the radial direction. This reduces the overall thrust and efficiency of our simulations compared to experiment.

While the data-driven model was able to correctly predict some trends in thruster performance with changing discharge voltage and current, it was unable to reproduce experimentally-observed trends in axial plasma properties. Given the demonstrated ability of this closure model to replicate empirically-inferred anomalous collision frequency profiles, this is surprising. This result suggests that the ability of transport models to match these empirical profiles may not correlate well with their predictive performance when implemented self-consistently into a Hall thruster code.

If this is true, then the probabilistic calibration procedure implemented in this section of the dissertation may not be applicable. Instead, model calibration would need to be performed by running many Hall2De simulations. This would make it much more time consuming to determine the optimal fit coefficients, and potentially make the type of uncertainty quantification we demonstrated here computationally infeasible.

In the next section, we address these questions by tuning two models to match the H9's empirical profile directly. In doing so, we put aside the question of extensibility for the moment and focus on the whether and to what degree a model's ability to predict an empirically-inferred collision frequency profile translates to good predictive performance.

5.2 Challenges with the use of empirical profiles for closure model calibration and evaluation

This work was previously published in Plasma Sources Science and Technology as **Reference 96***.*

In light of the wide range of proposed closure models models for the anomalous collision frequency, and the results in the preceding section, it is an open question as to how best to compare these models and evaluate their fidelity. It is common to use "empirical" profiles of the anomalous transport that have been calibrated against experimental data. These profiles are usually found by assuming the transport is static in time but varies spatially in the thruster. The profiles are then adjusted until key predictions of the code agree with experimental measurement.[53,16,29,56,94,97] This process yields useful insight into the state of the Hall thruster plasma in regions of the discharge that are not accessible by non-perturbative diagnostics, i.e. within the discharge channel. The resulting empirical profiles in turn are commonly used as references for evaluating the quality of proposed closure models.[2,3,41,42]

Despite this practice, these empirically-determined profiles may not be sufficiently representative of the electron dynamics to be used as reference for benchmarking closures. Experimental measurements, for example, have shown that the transport is time-dependent, which is a departure from the static assumption of these profiles.[31] Additionally, it has been shown that these profiles are non-unique in that the empirical transport may vary by an order of magnitude or more in certain regions of the discharge without significantly impacting the predicted operations of the thruster.[94] This makes the validation of closure models challenging, as two models may agree well in regions of the discharge most impacted by anomalous transport, but diverge in regions less sensitive to the electron dynamics. These previous observations may suggest that closure models that accurately predict empirical profiles may not do so when self-consistently implemented in

a simulation. In light of this possibility—which would significantly impact efforts to develop predictive Hall thruster models—there is an apparent need to evaluate the fidelity of closure models trained on empirical data. The goal of this section is therefore to determine if empirical profiles are appropriate for validating closure models of electron transport in Hall thrusters.

5.2.1 Closure models for anomalous collision frequency

In order to evaluate closure models for the Hall thruster anomalous collision frequency, it is common to compare them to empirically-inferred profiles.[56,41,2] More specifically, one can evaluate the closure model by using the time-averaged outputs from a simulation run with the empirical profile as inputs. The quality of the model is then evaluated based on how well the two curves match. As discussed in **Reference 3**, following this metric, two of the most effective closure models are the first-principles (FP) model from Lafleur et al[2] and the data-driven (DD) model by Jorns.[3] In this study, we use the forms of these models presented in **Chapter 3**.

In **Figure 5.6**, we compare these two closures to the calibrated empirical anomalous collision frequency profile from **Figure 3.1a**. For both models, we have used the time-averaged output from the resulting simulation to determine the plasma property values in Eqs. 3.31 and 3.46. As can be seen, the first-principles model (with c_1 = 0.1245) and the data-driven model (with c_1 = 2.39 and c_2 = 3.32) both capture the major features of the empirically-inferred anomalous collision frequency profile and are able to correctly predict the location of the minimum in this parameter. There are discrepancies between the models both upstream and downstream of the minimum. However, as shown in **Reference 94**, the anomalous collision frequency may vary by up to an order of magnitude in these regions without substantially impacting predictions for performance and plasma profiles.

Despite how well these models are able to predict the empirical anomalous collision frequency profile, it remains to be seen how they might perform when coupled with the rest

123
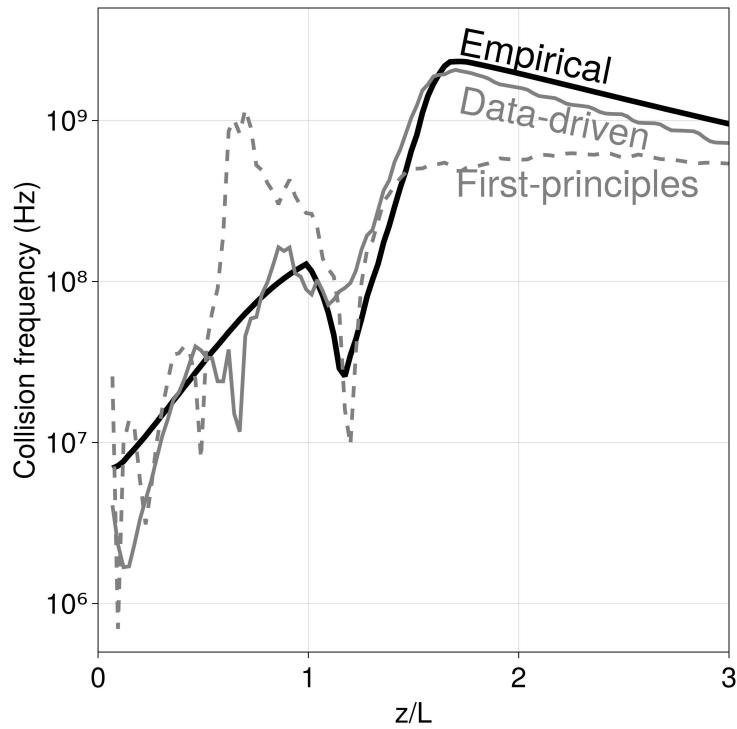
Figure 5.6: Comparison of the first-principles[2] **(Equation 3.31)** and data-driven[3] models **(Equation 3.46)** to the reference empirical profile depicted in **Figure 3.1a**.

of the plasma. This is the most rigorous measure of the predictive power of these models. To this end, we now turn to integrating these models into a full Hall thruster simulation.

In this work, as in the rest of the dissertation, we implemented the models into Hall2De. As in the preceding section, we used a field-aligned mesh for both the ions and electrons. We used a global timestep of 9 ns and ran each simulation for 2 ms of simulated time. As global performance properties and axial plasma properties are largely insensitive to the cathode anomalous collision frequency coefficient, we employ a fixed value of 0.005 for this number. The remaining simulation parameters, simulated thruster, and data for comparison are as described in **Chapter 4**. Additionally, we use a more recent set of LIF data for comparison, as described in **Chapter 4**,

5.3   Results

In this section, we first describe the results of our calibrated reference simulation. We then compare the performance of the closure models against that of the reference simulation.

5.3.1   Reference simulation

To assess the performance of our chosen closure models, we compare them to a calibrated reference simulation of the H9 at 300 V and 15 A. We choose to make comparisons to a reference simulation instead of directly to experiment in order to directly assess the impact of the anomalous transport on the simulation. Since all factors except for the anomalous transport model have been held equal between the reference simulation and the two self-consistent simulations, any deviations are due to the transport model and not due to physics that may differ from experiment. We briefly discuss here how this simulation was generated with Hall2De. In particular, we obtained the empirical anomalous collision frequency profile by varying its spatial dependence iteratively until the discharge current was within 0.5 A of the experimental value (**Table 5.2**) and the

125

Case	T	I_d	η_a	IVR
Experiment	292.9 ± 3.5 mN	15.0 A	64.4 ± 1.5 %	0.0
Ref. sim.	258.3 mN	15.2 A	49.4%	0.0722

Table 5.2: Comparison of reference simulation output to experimental data from Su and Jorns (2021).[1]
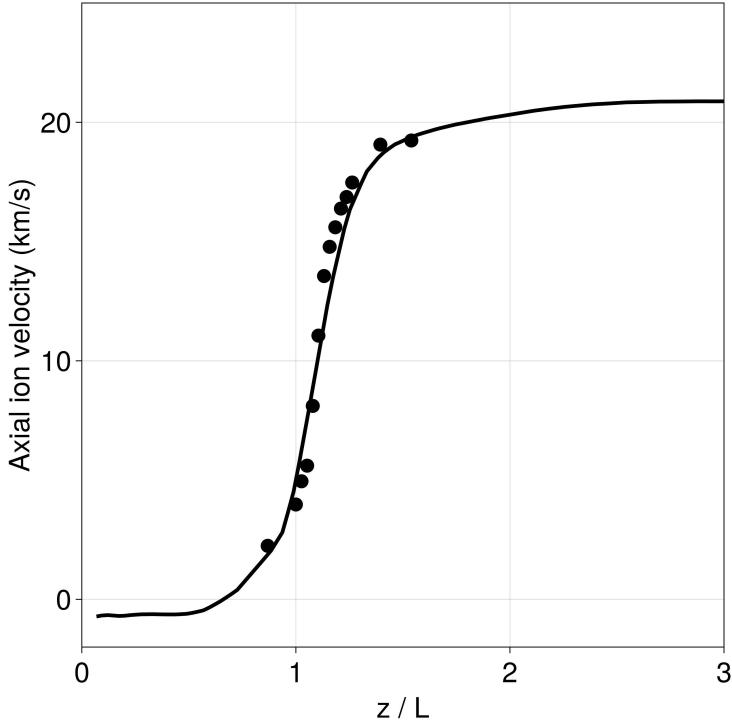
Figure 5.7: Empirical velocity curve (solid line) compared to LIF measurements of the mean ion velocity (markers).

ion velocity along the channel centerline matched LIF measurements to within IVR < 0.1. This approach to determining the empirical profile is common practice for Hall2De simulations.[82,94] Achieving this match in our case required 29 iterations. We show the results of the reference simulation as a solid line in **Figure 5.8b)**. Experimental data is depicted as discrete markers.

We compare in **Table 5.2** performance metrics from our simulations to experimental data. As can be seen, the current and IVR of the reference simulation agree to within 8% of experiment, while the thrust is 12% lower and anode efficiency 20% lower than the

measured values. This is a result of the fact that the calibration was based on matching experimental discharge current and ion velocity profile while we did not consider thrust. Comparable levels of reduced thrust and anode efficiency compared to experiment have been reported in previous Hall2De simulations with empirically-based anomalous collision frequency profiles.[82]

With that said, as we discussed in the beginning of this part of the chapter, despite under-predicting aspects of global performance, empirical profiles calibrated in this way to match ion velocity and current are often treated as surrogate measurements of the true anomalous collision frequency in these devices. Our goal is thus to determine how well closure models calibrated against such profiles perform when implemented directly into a Hall thruster code. We address this comparison in the following section.

5.3.2 Closure models

We consider in this section the results of the first-principles and data-driven when self-consistently implemented in Hall2De. As we discussed in this previous section, we calibrated these closures by evaluating them on the time-averaged plasma parameter outputs of the reference simulation (**Figure 3.1b** and **Figure 5.8b**) and adjusting the model coefficients until we obtained quantitative agreement with the empirical anomalous collision frequency profile (**Figure 3.1a**). The underlying assumption in this case is that a model that matches the empirical profile will yield similar predictions.

In **Table 5.3**, we compare the results of the closure models described in Sec. 5.2.1 to the reference simulation. Note here that in this table, we have computed the integrated velocity residual with respect to the reference simulation rather than the experimental data. As can be seen, when we self-consistently implement these models, neither agree with the reference simulation on all five metrics described in **Chapter 4**. The first-principles (FP) model of Lafleur et al. shows the closest agreement with the reference discharge current and thrust—the thrust is 1 mN lower and discharge current 4.5 A higher than

Case	T	I_d	η_a	IVR	IAR (initial)	IAR
Ref. sim.	258.3 mN	15.2 A	49.4%	0.0	0.0	0.0
FP	257.6 mN	19.7 A	37.9%	0.307	0.653	1.084
DD	333.6 mN	31.1 A	40.3%	0.344	0.309	0.826

Table 5.3: Performance metrics from closure models compared to reference simulation.

those of the reference simulation. On the other hand, the simulated anode efficiency is 11.5% lower than the reference. In contrast, the data-driven model of Jorns (DD) predicts a discharge current twice that of reference simulation and thrust 70 mN higher than the reference simulation, with anode efficiency 9% lower. Both simulations have high integrated velocity errors —0.307 for the FP model and 0.344 for the data-driven model— indicating poor agreement with the reference ion velocity curve. Lastly, in both cases the integrated anomalous collision frequency residuals (IAR) have nearly doubled compared to the non-self-consistently ("initial") implemented values. This suggests that the anomalous collision frequency profiles have significantly diverged from the original empirical profile. In summary, the major implication of the collected results in Table 5.3 is that, despite the fact the closure models agree with the empirical profile, their self-consistent implementation does not yield results that match the empirical profile's outputs.

We can explain the global metrics reported in **Table 5.3** by examining the time-averaged anomalous collision frequency profiles of the two models. To this end, we show in **Figure 5.8a** the resulting anomalous collision frequency profiles varying as a function of position after we have implemented them self-consistently in the code. For comparison, we also show the empirical profile. We emphasize here that these are the self-consistent results for anomalous collision frequency and differ from the spatial profiles shown in **Figure 5.6**. In this earlier case, we simply trained the models against the outputs of the reference simulation with the empirically determined collision frequency profile. As we already had concluded globally from the IARs, the two models have diverged from the empirical profile. The first-principles model still qualitatively matches the empirical profile in some respects – namely, the profile exhibits a minimum collision frequency with the right
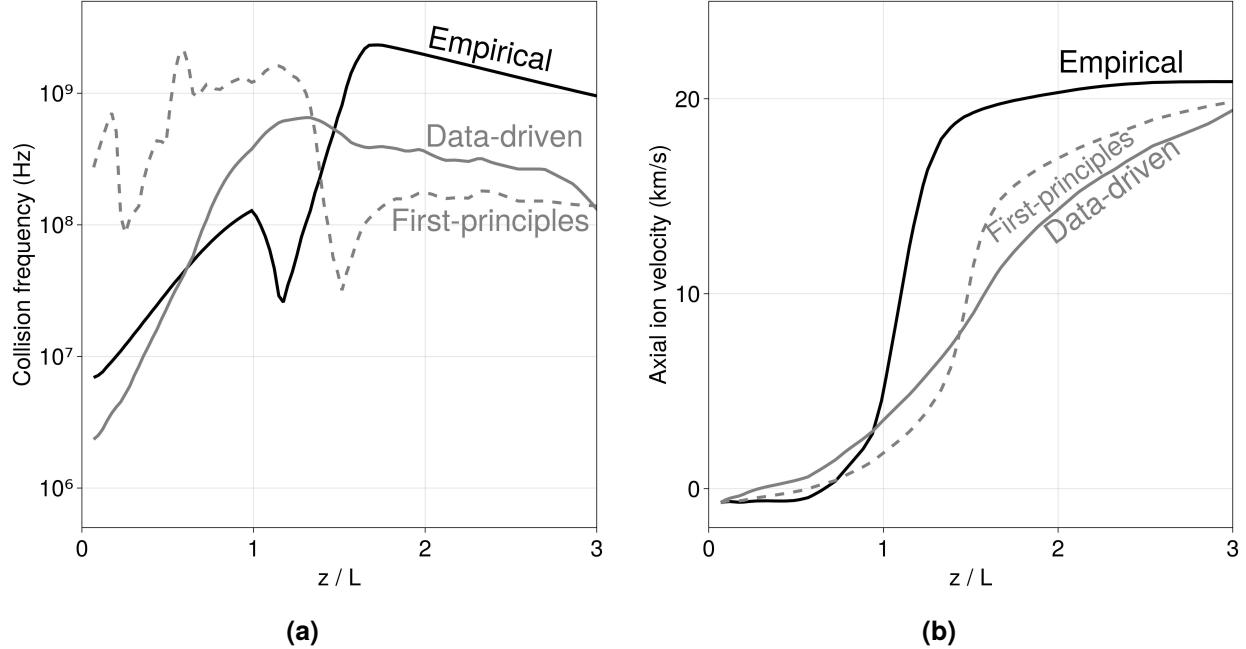
Figure 5.8: (a) Comparison of simulated anomalous collision frequency profiles to the empirical anomalous collision frequency. (b) Comparison of simulated ion velocity curves from each of the results of the reference simulation.

.

order of magnitude – but the location of this minimum is shifted downstream by half of a channel-length from the original value. The minimum collision frequency predicted by the first-principles model is also about 50% higher than the minimum of the empirical profile. This leads to enhanced electron transport in this region over the reference simulation, explaining the increased discharge current and reduced anode efficiency reported in **Table 5.3**. In contrast, the data-driven profile shows an anomalous collision frequency profile which reaches a maximum, rather than a minimum, downstream from the exit plane with a collision frequency almost ten times higher than the empirical profile in this region. As in the first principles model, this has the effect of increasing the current and reducing the efficiency.

We can in part explain the differences in integrated velocity error between the two models by comparing the spatially resolved ion velocity curves to the anomalous collision frequency profiles in **Figure 5.8a**. **In Figure 5.8b**, we show the predicted ion velocity

129

curves from the two models next to that of the reference simulation. Commensurate with its lower IVR, the ion velocity of the first-principles model matches the empirical result more closely than that of the data-driven model. Specifically, the first-principles model captures the steep slope of the empirical ion velocity curve while the data-driven model exhibits gradual ion acceleration. This results from the characteristics of the minimum in anomalous collision frequency exhibited by the first-principles model (**Figure 5.8a**). The point of minimum collision frequency corresponds to the point of maximum electron impedance, which in turn promotes a steep electric field. This field leads to rapid ion acceleration at the location coincident with this minimum. However, as can be seen, this profile is shifted downstream from the empirical profile, thus leading to the delayed ion acceleration profile. The relatively gradual acceleration of the data-driven model can similarly be understood by noting the high collision frequency near the thruster exit plane, which produces low electron impedance, a relaxed electric field, and thus slow ion acceleration.

In summary, we have shown that models trained on time-averaged data to match empirical profiles diverge from simulation results obtained using the original empirical profile, in terms of both global performance metrics and spatially-resolved plasma properties. The differences in these quantities can be understood in terms of the shapes of the anomalous collision frequency profiles produced by the models.

5.3.3 Discussion

In this section, we discuss possible explanations for our results, outline the implications of our findings for developing models of anomalous transport, and provide recommendations for future Hall thruster model calibration efforts.

5.3.3.1   Divergence between self-consistent models and empirical profiles

We consider in the following sections three possible explanations for the divergence between modeling results when we calibrate the closures on an empirical profile versus self-consistent implementation. These include non-linearity of the governing equations, use of time-averaged data for calibration, and non-uniqueness of the empirical profile. As the models were compared to a calibrated reference simulation instead of directly to experimental data, we do not expect that differences between the Hall2De physics model and the true physics governing the thruster operation are able to explain our results.

5.3.3.2   Non-linearity of the governing equations

As the dynamics of Hall thruster plasmas are highly non-linear, small initial variations in initial conditions can compound into larger differences over time. As a result, the small differences between our calibrated closure models and the empirical profile (**Figure 5.6**) may have led to increasingly divergent behavior, resulting in the self-consistent profiles that do not match the initial empirically-informed ones. Moreover, in the absence of any differences between the closure models and the empirical profile (i.e. a perfect match), we suspect plasma oscillations inherent to Hall thrusters may still be sufficient to cause the model to diverge from the empirical profile.

5.3.3.3   Use of time-averaged data for calibration

Another contribution to the diverging profiles may stem from the use of time-averaged plasma properties to calibrate our closure models. We can illustrate the problem with a simple example. As Hall thruster discharges are oscillatory, we can consider two plasma quantities, $A(t)$ and $B(t)$, which oscillate at a frequency ω, where B has a constant phase offset from A of ϕ radians. The time average operator is

$$\langle f(t) \rangle = \lim_{T \to \infty} \frac{1}{T} \int_0^T f(t) dt. \tag{5.7}$$

If $A(t)$ and $B(t)$ are sinusoidal functions with zero mean, we find after integration that

$$\langle AB \rangle = \cos(\phi)/2 \qquad\qquad \langle A \rangle \langle B \rangle = 0,$$

where ϕ denotes the phase between A and B. The implication of this result that the product of the averages can be fundamentally different than the average of the products. Therefore, when we make calculations using time-averaged oscillatory quantities, the result may differ from what we would find if we had performed our computation at all times and then averaged the result. When evaluating models of anomalous transport of time-averaged data, there are thus likely to be averaging artifacts if the model involves products of oscillating quantities.

We can apply this reasoning to illustrate the effect of time averaging on one of our closure models, the data-driven result **(Equation 3.46)**. For this analysis, we neglect the ion sound speed c_s for simplicity to arrive at

$$\nu_{an} = c_1 \omega_{ce} \frac{|\vec{u}_i|}{v_{de}}. \tag{5.8}$$

We furthermore restrict the problem to one dimension along the axial coordinate z and restrict our analysis to locations immediately downstream of the ion stagnation point, so that $\vec{u}_i \approx u_i > 0$ and $\vec{E} \approx E > 0$. We can then collect factors which do not vary in time (e, $|B|$, and m_e) into a constant K. After these simplifications, we have

$$\nu_{an} = K \frac{u_i}{E}. \tag{5.9}$$

We can then assume both the ion velocity and electric field oscillate sinusoidally with a frequency ω and that the electric field oscillates out of phase with the ion velocity with a

phase difference of 90 degrees. This is consistent with previous measurements of the breathing mode .[98] This yields

$$\nu_{an} = K \frac{\langle u_i \rangle + \tilde{u}_i \sin(\omega t)}{\langle E \rangle + \tilde{E} \cos(\omega t)}, \tag{5.10}$$

where variables in brackets ($\langle .. \rangle$) denote mean quantities and those with tildes ($\tilde{..}$) denote amplitudes. This expression is also a periodic quantity with period $2\pi/\omega$. The time-averaged anomalous collision frequency is thus given by

$$\langle \nu_{an} \rangle = K \frac{\omega}{2\pi} \int_0^{2\pi/\omega} \frac{\langle u_i \rangle + \tilde{u}_i \sin(\omega t)}{\langle E \rangle + \tilde{E} \cos(\omega t)} dt. \tag{5.11}$$

This yields a final result that scales as

$$\langle \nu_{an} \rangle = K \frac{\langle u_i \rangle}{\langle E \rangle} + \epsilon(z, \langle E \rangle, \tilde{E}, \langle u_i \rangle, \tilde{u}_i), \tag{5.12}$$

where ϵ is a positive number which is in general a function of space, E, and u_i. This simple illustration demonstrates that we cannot replace the arguments of a closure model with time-averaged values and assume the result will map to the time averaged collision frequency. This may in part explain why we can achieve quantitative agreement with the empirical profiles (that are implicitly time-averaged) while not matching the results with a self-consistent implementation. In practice, if this is a contributing factor, calibration of closure models would require time-dependent measurements.

5.3.3.4 Non-uniqueness of empirical profiles

As a final note, we consider the implications of the fact that the empirical profiles we have used for calibration are non-unique. Mikellides and Lopez-Ortega have demonstrated[94] that stationary empirically-calibrated anomalous collision frequency models can vary significantly in certain aspects with only minor impacts on the simulation output. In

particular, changing the anomalous collision frequency far upstream or far downstream of the thruster exit-plane has a smaller impact on the shape of the potential drop in the channel than changing it in the vicinity of the exit plane. Additionally, altering how the anomalous collision frequency changes away from the thruster centerline changes little about the observable device behavior beyond small changes in the plume divergence angle. Practically, this insensitivity is fortuitous as it allows more latitude in calibrating simulations to match data. However, it also poses significant challenges if we wish to use calibrated empirical profiles to inform our understanding of how anomalous transport actually varies in Hall thrusters. The non-uniqueness suggests there may be no effective "ground truth" for calibrating or validating a closure model in certain regions of the thruster. To this point, given how insensitive the simulation is to the anomalous collision frequency in the region upstream of the exit plane, there may be multiple empirical profiles that give the same results as our reference simulation that we could use for regressing our closure models. These may ultimately have yielded improved predictive capabilities. One possible way to resolve this ambiguity would be to increase the amount of data used for determining the empirical profiles, e.g. additional velocity measurements off centerline and measuring other plasma properties. Alternatively, direct experimental measurements of the anomalous collision frequency would resolve this non-uniqueness.[31]

5.3.3.5 Data-driven generation of empirical profiles

While we have shown the limitations of the self-consistent implementation of closure models calibrated against empirical profiles, there is still a practical application for using these expressions to increase the efficiency of calibrating empirical models against new datasets. Given experimental measurements of the time-averaged plasma properties, we could use a given closure model to estimate what the shape of the empirical anomalous collision frequency profile should be. Indeed, Jorns showed that closure models generated by regressing a dataset of empirical profiles were able to predict the empirical profiles of

134

thrusters and operating conditions not in the training dataset.[3] Reinforcing this point, the fit between the data-driven model and the H9 empirical profile displayed in **Figure 5.6** was obtained with no alteration to the model coefficients in **Reference 3**, despite the fact that the data-driven model was not trained on data from this thruster. Using data-driven methods to speed the rate of empirical model calibration is an active area of research.[99]

5.3.3.6   Implications for future closure model calibration and discovery

Self-consistent closure models for electron transport are critical for enabling predictive simulations. As the community continues to propose new models based on first principles or data-driven methods, it is necessary to develop methods for calibrating and validating them. Our work has shown that using "measurements" of the steady-state electron collision frequency inferred from calibrating against experimental data may not be sufficient for evaluating new models. With this in mind, as we look to future efforts, there are a number of possible strategies that may be employed for generating data sufficient for evaluating and calibrating new closure models.

The most useful data for calibrating transport models would be direct experimental measurements of the anomalous collision frequency. Such measurements have been made in the past,[31,78,100,101] though they are often subject to simplifying assumptions about the plasma state that introduce uncertainty in the values. Furthermore, these measurements often do not extend upstream of the exit plane, which leaves the problem of non-uniqueness stemming from the value of the transport in the near-anode region unresolved. With that said, as new experimental methods for characterizing the electron properties in situ become available, directly measuring the time and spatially-resolved transport may become more tractable.[35]

In complement to developing methods to measure directly the anomalous collision frequency, it may be possible to infer the time-resolved electron dynamics by combining time-resolved data with real-time state estimation techniques[102] with the standard,

empirical calibration approach. This method still relies on inferring the collision frequency indirectly by running a model iteratively and comparing the outputs to data. However, this approach has the advantage of the resulting empirical profile being time-dependent. In light of the discussion in the previous section, this would provide a much higher-fidelity dataset for the calibration of closure models.

As an alternative approach, we could in principle calibrate and evaluate closure models by self-consistently implementing them in Hall thruster codes and iterating until the results match experiment. This has the advantage that it is not necessary to directly measure the collision frequency for comparison. The drawback is that this approach may be more time-intensive than evaluating the model against empirically calibrated or direct measurements. Reduced-order and multi-fidelity modeling may accelerate this process.[103,104]

As a last comment, we remark that while our goal remains to find self-consistent closures for predictive-fluid modeling, empirically-inferred transport profiles remain critical to thruster development and qualification. They have been leveraged, for example, for discovering new designs to that increase thruster lifetime[16] and for evaluating plasma conditions in regions of the discharge which are difficult to measure .[53] With this in mind, we have shown that closure models calibrated against empirical collision frequency datasets may play an important role in accelerating the calibration procedure for new thrusters and operating conditions.

5.4 Summary

In this chapter, we investigated the relationship between static, empirically-inferred profiles of Hall thruster anomalous collision frequency and self-consistent closure models calibrated on these profiles. I first generated probabilistic performance predictions of the H9 Hall thruster using a data-driven model and a calibration procedure based on Bayesian inference. We then self-consistently implemented two closure models into

Hall2De and compared the results to a validated reference simulation. We ultimately found that closure models tuned to match empirical profiles when using time-averaged simulation data diverged from the empirical results when implemented self-consistently into simulations. These results may be attributed to a number of factors such as the non-linearity in the Hall thruster fluid model, artifacts introduced by time-averaging, and the non-uniqueness of empirically-inferred collision frequency profiles. This inability to use steady-state profiles poses a challenge for the development and validation of new closure models for the anomalous transport.

CHAPTER 6

Self-Consistent Implementation of Closure Models in Hall2De

All models are wrong, but some are useful.

George P. Box

In the previous chapter, we showed that empirically-inferred electron transport profiles cannot be relied upon to calibrate or evaluate models of anomalous electron transport. To truly evaluate a model's efficacy, it must be implemented self-consistently into a Hall thruster code. In the cases where the models in **Chapter 3** have been incorporated into fluid models, the predictions were shown to deviate from experimental data, and the lack of a parametric comparison across operating conditions invites questions about model extensiblity. More broadly, there has not yet been published comparison of proposed closures on the same thruster with the same fluid model. As a result, differences between transport models cannot be unambiguously shown to result from the anomalous transport models themselves, instead of differences in simulation parameters. In order to advance the state of the art in predictive engineering simulations of Hall thrusters, there is thus a need to systematically evaluate proposed transport closure models with a common fluid model and to compare their predictions to experimental data.

Model	Expression for ν_{an}	Eqn. #	Typical coefficients
Cappelli et al.	$\nu_{an} = c_1 \omega_{ce} \sqrt{\dfrac{|j_{e\perp} E_\perp|}{e n_e c_e^2 B}}$	(6.1)	$c_1 = 1/\sqrt{2.5} \ (\approx 0.632)$
Lafleur et al.	$\nu_{an} = c_1 \dfrac{|\nabla \cdot (\vec{u}_i n_e T_e)|}{m_e c_s n_e v_{de}}$	(6.2)	$c_1 = 1/4\sqrt{6} \ (\approx 0.102)$
Chodura	$\nu_{an} = c_1 \omega_{pi} \left[1 - \exp\left(-\dfrac{1}{c_2}\dfrac{v_{de}}{c_s}\right)\right]$	(6.3)	$c_1 = 1, \ c_2 = 3$
Data-driven	$\nu_{an} = c_1 \omega_{ce}\left(\dfrac{|\vec{u}_i|}{c_2 c_s + v_{de}}\right)$	(6.4)	$c_1 = 2.39, \ c_2 = 3.32$

Table 6.1: Summary of anomalous transport models investigated in this chapter.

We focus in this chapter on four previously-published attempts at algebraic closure, which we introduced in **Chapter 3** and review in **Section 6.1**. We implement these models into Hall2De and evaluate their predictive performance across four operating conditions of the H9 Hall thruster. Finally, we analyze their performance in the context of the functional dependence of the models on the local plasma properties. We then discuss the implications of our results and provide recommendations and guidelines for future closure modeling efforts.

6.1   Investigated closure models

We list in summary form in **Table 6.1** the four closure models we investigate in this chapter. Each of these was previously described in **Chapter 3**. These four were selected because they have been shown in a select series of cases to yield qualitative agreement with experimental measurements. Practically, each of these models also has only one or two tunable coefficients. This simplifies the process of optimizing each model to find the best agreement with experiment.

The first chosen model is that of Cappelli et al.,[42] which was derived based on an analogy to classical fluid theory. The second is the model of Lafleur et al.[2], which is based on a first-principles derivation of the effect of ion acoustic turbulence on Hall thruster anomalous transport. The third is the Chodura model,[68] which is a semi-empirical model that has found use in field-reversed configuration and theta-pinch plasma devices. We compare these three models to the data-driven model of Jorns,[3] which we investigated in detail in the previous chapter. Due to difficulties countered in model calibration, we do not investigate in this chapter any of the shear-suppressed transport models in the vein of Scharfe et al.[41] With these models described, we turn in the next section to describing the methods employed in this study.

6.2   Simulations and data for comparison

In this chapter, as in the rest of the dissertation, we simulate the H9 Hall effect thruster. We compare our simulations results to data obtained in two 2021 test campaigns by Su and Jorns.[1,17] This data includes performance figures such as thrust and discharge current, as well as spatially-resolved ion velocity measurements obtained via laser-induced fluorescence. To investigate how well the models examined in this work can generalize across operating conditions, we consider the subset of operating conditions from Refs. 1 and 17 summarized in **Table 6.2**. In addition to the nominal 4.5 kW operating condition, where the discharge voltage (V_D) was 300 V, the discharge current (I_D) was 15 A and the propellant was xenon, we also simulate high-voltage (600 V and 15 A), high-current (300 V and 30 A) and krypton (300 V and 15 A) cases. For each of the cases, we include the mass flow rate ($\dot{m}$) and background pressure (P_B) in the facility, which are needed as simulation inputs, as well as the measured thrust (T), to which we compare our simulation results.

Case #	V_D (V)	I_D (A)	Propellant	$\dot{m}$ (mg/s)	P_B (μTorr)	T (mN)
1	300	15	Xe	14.8	4.6	292.9 ± 3.5
2	300	30	Xe	26.0	7.6	539.0 ± 4.0
3	600	15	Xe	15.4	5.4	447.2 ± 3.0
4	300	15	Kr	11.8	4.6	235.8 ± 2.5

Table 6.2: H9 operating conditions simulated in this chapter.

6.2.1 Simulation details

In line with the rest of this disseration, we conducted all of our simulations using Hall2De. Each simulation was run for 2 ms of simulation time with a timestep of between 1 and 2 nanoseconds.

All of the anomalous transport models described in **Section 6.1** were implemented using under-relaxation, such that the anomalous collision frequency at timestep n was given by a weighted average of the instantaneously-computed value and the value computed at timestep $n-1$. This improved stability and reduced iteration-to-iteration numerical oscillations, at the cost of a potential reduction in time-accuracy. Since we were primarily interested in time-averaged simulation properties, we considered this an acceptable compromise in favor of improved numerical convergence.

6.2.2 Metrics

In addition to the those introduced in Chapter 4, we include in this chapter two additional metrics for assessing the discrepancy between our simulations and experiment. The first is the *thrust residual*. This is the normalized residual between the experimental and simulated thrust and is defined as

$$\text{Thrust residual} = |T_{exp} - T_{sim}|/T_{exp}, \tag{6.5}$$

where T_{exp} is the experimental thrust for the current operating condition. The second is the *total residual*. This is a global metric for the validity of the experiment, created by summing the previously-defined integrated velocity residual (IVR) with the thrust residual.

$$\text{Total residual} = \text{IVR} + \text{Thrust residual.} \tag{6.6}$$

Physically, this represents how close a given simulation is to the experimental data, with lower numbers indicating a closer match. It allows for a direct comparison across closure models.

6.3   Model calibration and validation

6.3.1   Empirically-derived reference profiles

As discussed in **Chapter 3**, models of anomalous transport are typically benchmarked against calibrated reference anomalous transport profiles. In order to assess how well the investigated closure models predict the spatial variation of plasma properties, we generated a multi-zone empirical Bohm[94]-like profile for each of the four operating conditions described in **Section 6.2**. These profiles were characterized by four "nodes", (z_i, α_i), with the anomalous collision frequency varying logarithmically between these nodes. Specifically, the anomalous collision frequency was given by

$$\frac{\nu_{an}}{\omega_{ce}} = \begin{cases} \alpha_1 & z < z_1 \\ \exp\left[\log(\alpha_i) + (z - z_i)\frac{\log(\alpha_{i+1}) - \log(\alpha_i)}{z_{i+1} - z_i}\right] & z_i < z \leq z_{i+1} \\ \alpha_4 & z \geq z_4 \end{cases} \tag{6.7}$$

To obtain the final profiles, we iteratively adjusted the values of z_i and α_i until the simulated ion velocity profiles matched LIF data (IVR < 0.1) and the thrust residual was less than

10%. In doing so, we obtained a "ground truth" for how plasma properties such as the anomalous collision frequency and electron temperature varied as a function of axial distance from the anode. We note that this is the same calibrated reference profile for the H9 employed in the previous chapter.

6.3.2 Closure model coefficients

As outlined in **Table 6.1**, all the investigated closure models have either one or two coefficients. For single-coefficient models (**Equations 6.1** and **6.2**), we adjusted c_1 in order to match the experimental discharge current. For the two-coefficient (**Equations 6.3** and **6.4**) models, we first performed a parameter sweep over c_2. For each value of c_2, we then varied c_1 until the experimental discharge current matched within $\pm$ 0.5 A. We then selected the combination of c_1 and c_2 which minimized the total residual (**Equation 6.6**).

To this end, we introduce in this paper a simple on-line calibration procedure for one of the anomalous transport coefficients (in this case, c_1). We employ a proportional-integral (PI) controller built into Hall2De to automatically adjust c_1 as the simulation progresses until the time-averaged discharge current matches experiment. To mitigate the effects of initial transients, we evaluate time-averaged quantities such as the discharge current using an exponential moving average as opposed to a simple average over all times. For the discharge current, this has the form:

$$I_{D,avg}(t_n) = \alpha_n I_D(t_n) + (1 - \alpha_n) I_{D,avg}(t_{n-1}) \tag{6.8}$$

$$\alpha_n = 1 - \exp\left(-\frac{\Delta t_n}{\tau}\right). \tag{6.9}$$

In the above, $I_{d,avg}(t_n)$, α_n, and Δt_n are the average discharge current, exponential smoothing coefficient, and timestep, respectively at iteration n and time t_n and τ is a user-

adjustable time constant. The discrete form of the PI controller is given by

$$\log c_1(t_n) = \log c_1(t_{n-1}) + K_p \left[\left(1 + \frac{\Delta t_n}{T_i} \right) e(t_n) - e(t_{n-1}) \right]. \tag{6.10}$$

Here, K_p is the proportionality constant and T_i is the integral timescale. Following numerical experiments in our open-source one-dimensional Hall thruster code, HallThruster.jl, we set K_p to 0.06 and $T_i = \tau = 0.4$ ms. This choice of τ has the advantage of averaging the discharge current over multiple breathing mode cycles, so the controller does not try to damp out these oscillations. We discuss this 1D code further in **Appendix A**. In **Figure 6.1**, we show the results of one of our numerical experiments on a one-dimensional simulation of the H9 Hall thruster at 300 V and 15 A using a simple two-zone anomalous collision frequency model of the form:

$$\nu_{AN} = \begin{cases} c_1 \omega_{ce}/16 & z < L_{ch} \\ c_1 \omega_{ce}/160 & z \geq L_{ch}. \end{cases} \tag{6.11}$$

In **Figure 6.1a**, we show the results of an initial simulation with $c_1 = 3$. The simulation reaches to a non-oscillatory steady state with an average discharge current of 30 A. In **Figure 6.1b**, we show the results of applying our PI controller to c_1. We let the simulation settle for 0.2 ms before turning on the controller. After 2 ms, the controller drives the simulation into a stable breathing mode oscillation with a mean current of 15 A and a peak to peak amplitude of 32 A. **In Figure 6.1c**, we show how the value of c_1 changes over time in the controlled simulation. We see that c_1 declines rapidly until $t = 1$ ms, after which is stabilizes into an oscillation about its final mean value. In order to demonstrate that this small oscillation does not artificially drive the observed breathing mode oscillations, we perform a final simulation in which we fix the value of c_1 at 0.813, which was mean value that the controller found (red dashed line in **Figure 6.1c**). We find that the uncontrolled
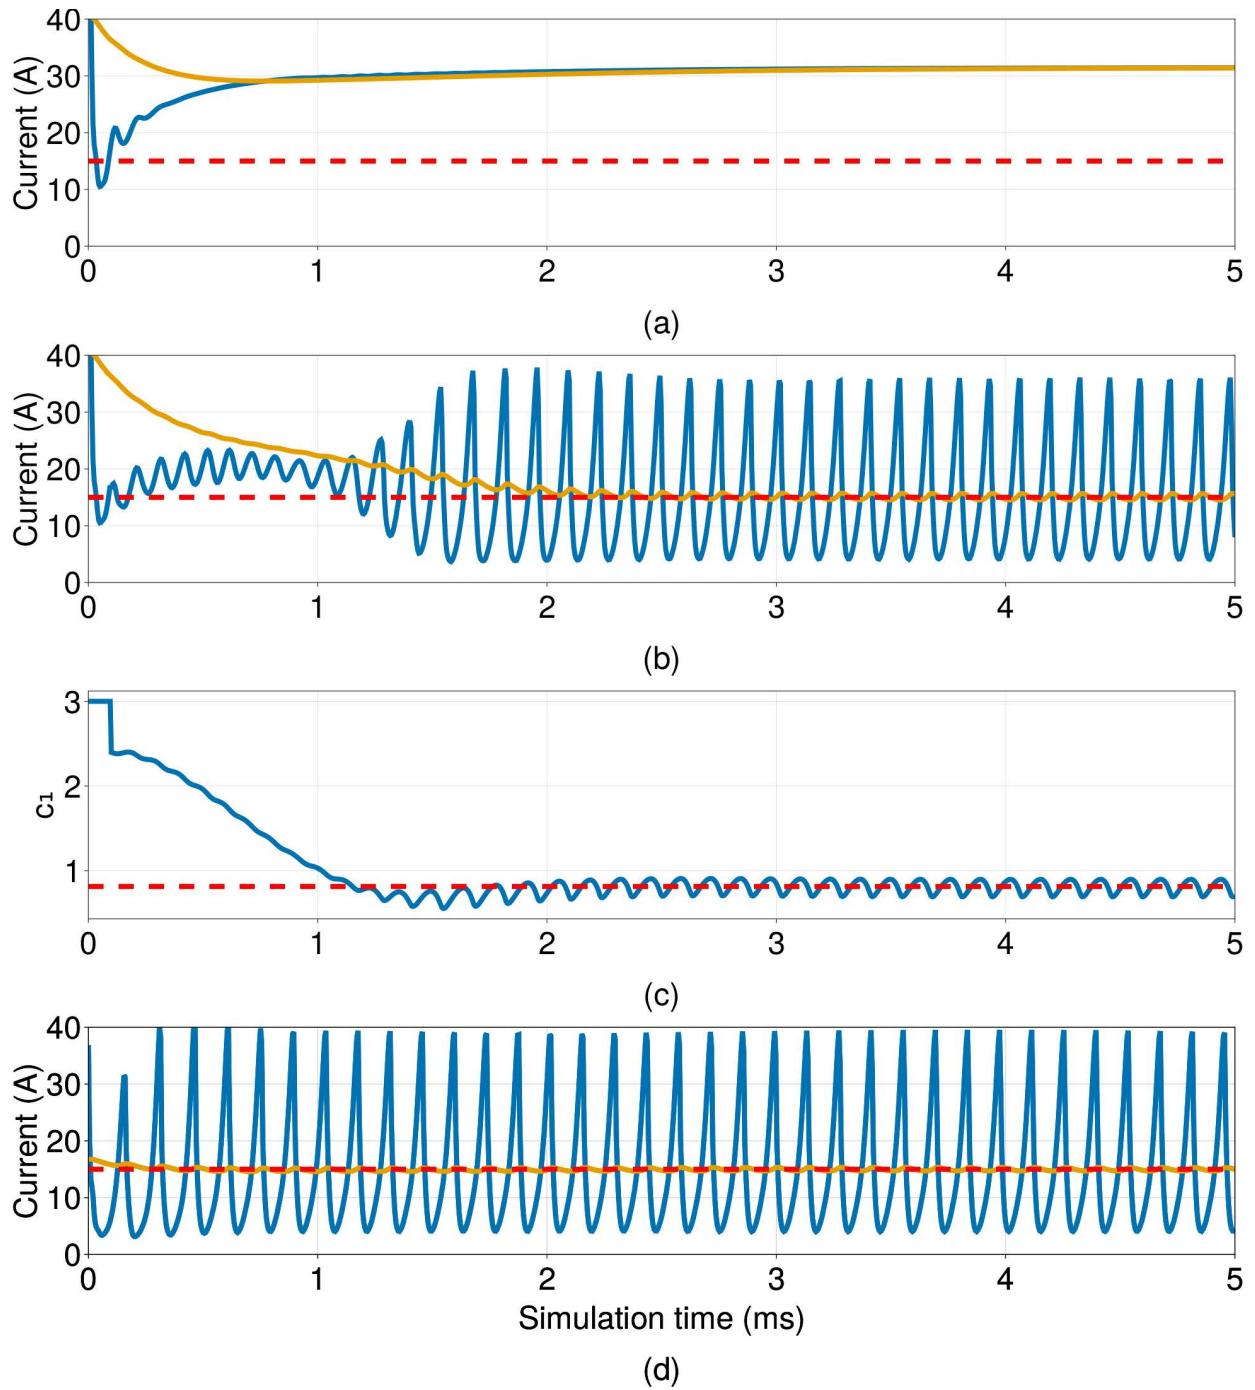
Figure 6.1: Results of test of PI controller in 1D fluid Hall thruster code. (a) Discharge current (blue) and exponential moving average of discharge current (yellow) with $c_1 = 3$ and no control. (b) Discharge current with active PI control. Setpoint of 15 A shown as red dashed line. (c) Time history of c_1 for controlled simulation shown in (b). Final mean value of c_1 shown as red dashed line. (d) Discharge current with no control using mean c_1 found by controller in (c).

	This work		Literature		
Model	c_1	c_2	c_1	c_2	Reference
Cappelli et al.	1.7	-	0.632	-	[42]
Chodura	14.5	500.0	1.0	3.0	[70]
Data-driven	1.295	100.0	2.39	3.32	[3]
Lafleur et al.	0.057	-	0.102	-	[2]

Table 6.3: Best-fit coefficients for the 300 and 15 A Xenon condition, obtained by calibration procedure described in Sec. 6.3, compared to values from literature (c.f. Tab. 6.1).

simulation exhibited the same oscillations as the controlled simulation, albeit with a slightly larger peak to peak amplitude of 37 A.

This proof of concept has shown that the PI controller can successfully tune the magnitude of the anomalous collision frequency to match a target value, even in the case of quite strong breathing mode oscillations of > 200 % peak to peak amplitude. The final value of c_1 may oscillate, but these oscillations have a negligible effect on the breathing mode oscillations exhibited by the plasma. Additionally, once c_1 converges to a stationary value, the PID controller can be turned off and the simulation continued with a static value of c_1.

We only applied this calibration procedure for each closure model to one condition, the 300 V, 15 A operating point. Once we obtained values of c_1 and c_2 for each model on this condition, we evaluated the extensibility of each model by simulating the other three operating conditions using these coefficients. With this approach in mind, we now turn to presenting the results of our study.

6.4 Results

Case	Condition	I_D (Exp., A)	I_D (Cappelli et al., A)	I_D (Chodura, A)	I_D (Data-driven, A)	I_D (Lafleur et al., A)
1	Xe, 300 V, 15 A	15	14.9	15.2	15.0	14.9
2	Xe, 300 V, 30 A	30	25.6	29.6	27.8	26.5
3	Xe, 600 V, 15 A	15	22.9	18.9	19.3	14.5
4	Kr, 300 V, 15 A	15	14.0	13.3	16.3	10.8

Table 6.4: Discharge current (I_D) predicted by each of the models compared to experimental discharge current.

Case	Condition	T (Exp., mN)	T (Cappelli et al., mN)	T (Chodura, mN)	T (Data-driven, mN)	T (Lafleur et al., mN)
1	Xe, 300 V, 15 A	292.9 ± 3.5	270.3	276.2	258.3	215.5
2	Xe, 300 V, 30 A	539.0 ± 4.0	404.8	473.9	440.4	398.2
3	Xe, 600 V, 15 A	447.2 ± 3.0	400.7	403.9	442.3	273.5
4	Kr, 300 V, 15 A	235.8 ± 2.5	220.7	210.7	238.7	149.0

Table 6.5: Thrust (T) predicted by each of the models compared to experimentally-measured thrust.

Case	Condition	IVR (empirical)	IVR (Cappelli et al.)	IVR (Chodura)	IVR (Data-driven)	IVR (Lafleur et al.)
1	Xe, 300 V, 15 A	0.069	0.520	0.400	0.417	0.761
2	Xe, 300 V, 30 A	0.095	0.463	0.339	0.393	0.730
3	Xe, 600 V, 15 A	0.065	0.614	0.583	0.538	0.793
4	Kr, 300 V, 15 A	0.059	0.457	0.272	0.347	0.669

Table 6.6: Integrated velocity residual (IVR) of each of the four models compared to IVR of reference simulation
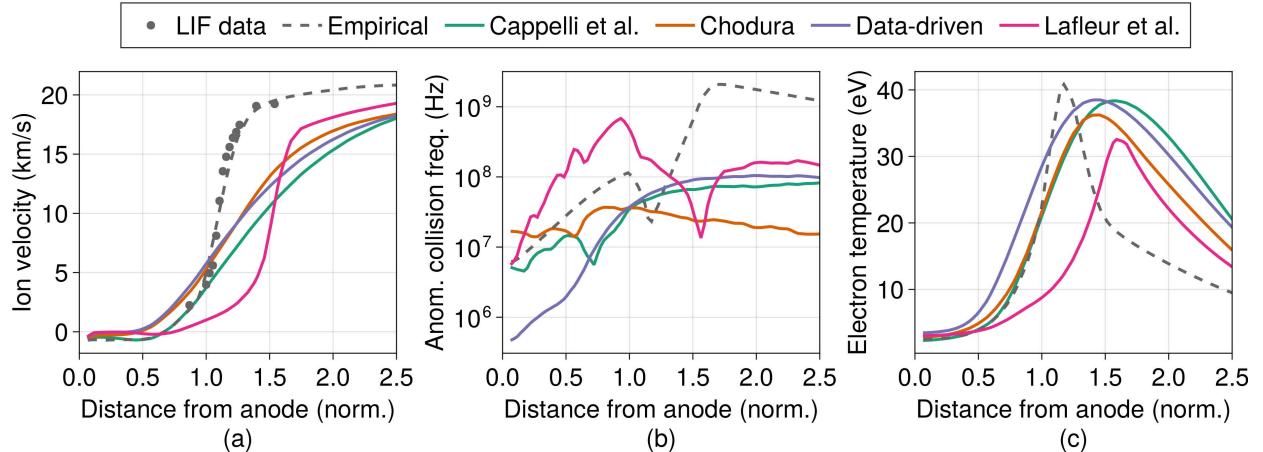
Figure 6.2: Comparison between model predictions, experimental data, and calibrated reference results of the ion velocity (a), anomalous collision frequency (b), and electron temperature (c) extracted along channel centerline for Case 1 (300 V and 15 A using xenon). All distances are normalized by the length of the discharge channel.

In **Table 6.3**, we report the best-fit values of c_1 and c_2 for all four models, as obtained by the calibration procedure discussed in the preceding section. We compare these values to results from the best-fit parameters given in previous works. Notably, our coefficients differ by at least a factor of two in each case from the values reported in literature. This indicates that coefficients inferred for these models may not be extensible across different thrusters and operating conditions, or that the models are not accurately capturing the true electron transport in these devices.

In **Table 6.4,** we present the discharge current predicted by each of models at the four investigated operating conditions using the coefficients displayed in **Table 6.3.** We see that for the calibration case (Case 1, 300 V and 15 A on xenon), the obtained discharge currents are within 0.2 A of the target value of 15 A. This indicates that the PID controller was able to successfully tune the fit coefficients to match the experimental discharge current. In Case 2 (300 V and 30 A on xenon), all four of the models correctly predict that increasing the mass flow rate will increase the discharge current but under-estimate the predicted increase by between 2 and 5 A. Moving to Case 3 (600 V and 15 A on xenon), we find that three of the models (Cappelli et al., Chodura, and Data-driven) output discharge currents

between 4 and 8 A greater than experiment while the model of Lafleur et al. outputs a discharge current of just 0.5 A lower than the experimental value. Lastly, in Case 4 (300 V and 15 A on krypton), both the data-driven model and the model of Cappelli et al. show discharge currents which differ from the measurement by less than than 1.5 A . In contrast, the model of Lafleur et al. shows a 5 A decrease in discharge current when switching from xenon to krypton propellant, while the Chodura model shows a 3 A increase.

Next, in **Table 6.5,** we compare the thrust predicted by the models to the experimental values. In Case 1, all four models under-predict the thrust compared to experiment. The model of Lafleur et al. outputs the lowest thrust (77 mN lower than the experimental value), while the Chodura model predicts a thrust only 7 mN less than experiment. This discrepancy increases in the 30 A case, where the model with thrust closest to experiment (the data-driven model) still predicts thrust 100 mN less than experiment. The agreement with experiment is improved in Case 3, where the data-driven model predicts a thrust value just 7 mN less than the experimental value. In Case 4, the data-driven model again has the best agreement with experiment, this time over-predicting the thrust by just 3 mN. In all four cases, the model of Lafleur et al. has the worst performance of the four models in terms of thrust .

Lastly, in **Table 6.6** we present the integrated velocity residual (IVR) of each of the models at each operating condition. For comparison, we also present the IVR of the empirical reference simulations described in **Section 6.2.** By design, each of the calibrated empirical reference profiles have IVR less than 0.1, indicating a less than 10% disagreement with the experimental data obtained by LIF across the entire LIF measurement domain. In contrast, the best IVR that any of the models achieves is 0.229 (Chodura model, Case 4), but values closer to 0.5 are more typical. As in the thrust results, the model of Lafleur et al. has the worst agreement with the spatially-resolved ion properties as measured by the integrated velocity residual, with IVRs between 0.67 and 0.79.
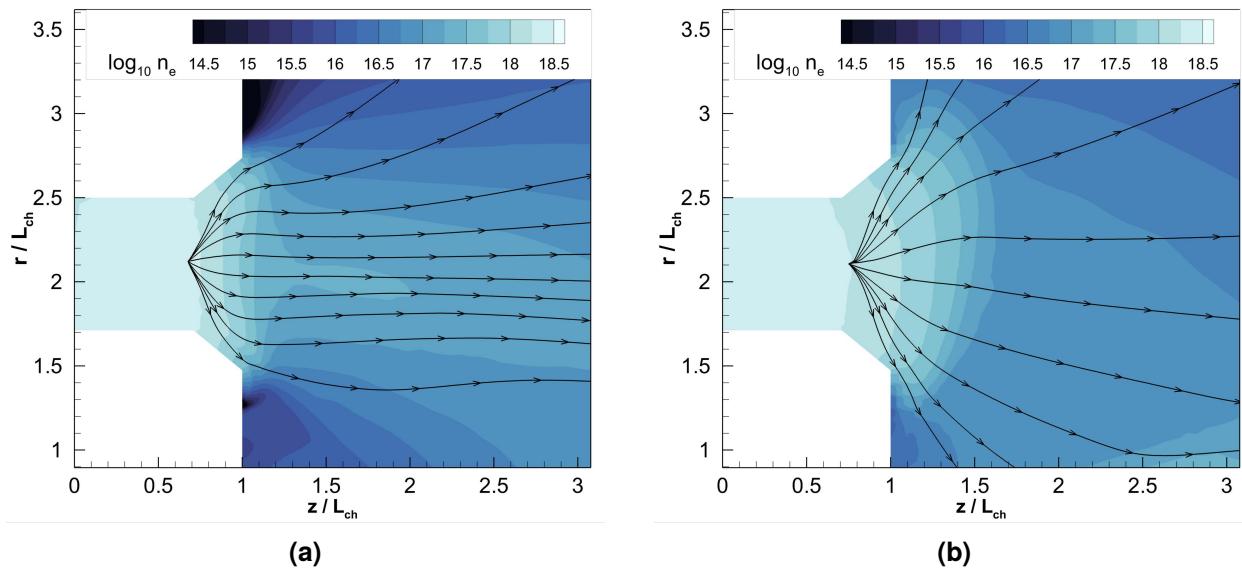
Figure 6.3: Ion current density streamlines overlaid on plasma density contours for (a) the calibrated reference simulation and (b) the simulation using the model of Lafleur et al. Plasma density is measured in units of m^{-3} and distances are normalized by the device channel length.

There are two major implications of these results. First, while some of these models are able to capture thruster performance and scaling to different operating conditions, none is able to capture all of these trends. Second, even when a model is able to correctly predict the thruster performance, it may exhibit poor agreement with other data, namely spatially-resolved measurements of plasma properties. Together, these results indicate that these closure models have poor extensibility across thrusters and operating conditions.

To provide insight into the performance of these metrics for the closure models, we present in **Figure 6.2** the spatially-resolved plasma properties from the 300 V and 15 A xenon case extracted along channel centerline. We note that the results at this operating condition are representative of the other three conditions. In **Figure 6.2a**, we compare the anomalous collision frequency to the empirically-inferred reference profile. We also show the electron cyclotron frequency. The empirical profile exhibits a strong minimum coincident with the location of maximum magnetic field strength (and thus maximum cyclotron frequency). In contrast, the models of Cappelli et al. and Chodura show only

weak minima, which are displaced half of a discharge channel-length upstream from the minimum of the empirical profile. The anomalous collision frequency predicted by data-driven model does not exhibit a minimum, while the model of Lafleur et al. has a steep minimum, albeit shifted half of a channel-length downstream of the thruster exit plane (z/L = 1). We return to a discussion of why these predicted anomalous collision frequency profiles differ from the empirical profiles in **Section 6.5**. Before proceeding with this, however, we examine how these differences in anomalous collision frequency translate to the poor agreement with experiment.

To this end, in order to explain the variations in the IVR, we show in **Figure 6.2b** the axial ion velocity profiles of each of the four models. We compare these results to the LIF data and the ion velocity from the reference simulation. The trends in velocity profiles largely can be explained by the behavior of the collision frequencies in **Figure 6.2a**. In particular, in **Chapter 3** we argued that the peak electric field in the thruster will be coincident with the minimum in anomalous collision frequency and that the magnitude of the electric field should increase as the minimum steepens spatially. With this in mind, given the shallow mininma in the models of Chodura and Cappelli et al. and the absent minimum in the data-driven model, we would expect these models to produce weaker peak electric fields and thus more gradual ion acceleration than the experimental results. This is indeed the case and can explain the large discrepancy with experiment. The model of Lafleur et al. predicts a steep ion acceleration curve in better qualitative agreement with experiment than the other three models. This stems directly from the steep and spatially-localized minimum in collision frequency. However, since this minimum occurs downstream, so too does the bulk of the ion acceleration, which starts half of a channel-length downstream from the experimental measurements. This consequently results in the largest overall IVR.

We next discuss how variations in the collision frequency profile may contribute to poor predictions in thrust. To this end, we first examine the role of electron temperature. As higher electron temperatures correspond to increased ion production (and therefore

thrust), discrepancies in this quantity may serve to explain under/overpredictions in thrust. In this context, we show in **Figure 6.2c** the electron temperature extracted along the thruster channel centerline.. In terms of peak magnitude in temperature, the data-driven model and the model of Cappelli et al. predict peak electron temperatures of 39 eV, while the Chodura model predicts a peak temperature of 36 eV. The Lafleur et al. model predicts the lowest peak electron temperature, at 33 eV. These results could explain in part why the model of Lafleur et al predicts a lower thrust than the other models. Physically, this may result from the fact that the electron temperature is generally assumed to be isothermal along magnetic field lines in Hall thrusters.[5,16,14] As the magnetic field lines downstream from the thruster exit plane are longer than those closer to the exit plane, the electron thermal energy must therefore be spread out across a larger region of space. A downstream displacement of the anomalous collision frequency profile (as exhibited by the model of Laflur et al. model in **Figure 6.2)** and peak electric field could thus shift the location of maximum electron heating in such a way as to reduce the peak electron temperature.

Another explanation for the low predicted thrust of the model of Lafleur et al. could be high plume divergence compared to the other models. If all of the ionization and ion acceleration happens far outside of the thruster channel, then the magnetic field is less able to shape and focus the plasma beam. Indeed, it has been shown that displacing the ion acceleration zone in Hall thrusters by a few millimeters downstream can result in an increase in plume divergence and a corresponding reduction in thrust efficiency.[1] To investigate this hypothesis in the context of our results, we present in **Figure 6.3** a comparison between streamlines of ion current density in the reference case at 300 V and 15 A to the predictions of the model of Lafleur et al. As expected, the streamlines are more divergent in this model than in the reference simulation. This means more of the ion momentum is directed in the radial direction and thus unable to contribute to thrust.

In summary, we have shown that none of the the investigated models is able to successfully reproduce the experimental measurements, either in terms of global performance metrics or in terms of spatially-resolved plasma properties. We can link this behavior to the spatial variation of the anomalous collision frequency as predicted by the four closure models. We turn in the next section to a discussion of the physical processes governing the anomalous collision frequency profiles.

6.5 Discussion

We discuss in this section each model individually. We focus on how the self-consistent coupling of the closure models with with the physics of a fluid Hall thruster model like Hall2De may have led to the trends seen in the previous section. We then highlight the importance of the electron energy balance in this coupling process. Finally, we discuss the implications of this work for informing future closure modeling efforts.

6.5.1 Cappelli et al.

As discussed in the original paper[42] and in subsequent work[3], the model of Cappelli et al. may be reduced to the following simpler expression if the electron pressure gradient is neglected:

$$\nu_{an} = \omega_{ce} \left(c_1 \frac{v_{de}}{c_e} \right)^2 = \omega_{ce}(c_1 M_e)^2,\tag{6.12}$$

where M_e is the azimuthal electron Mach number. In essence, this model predicts that the anomalous collision frequency should follow a Bohm-like scaling, weighted by the square of the electron Mach number. We expect that this scaling should actually produce a peak in the anomalous collision frequency, rather than a minimum, near the peak magnetic field.

Per our discussion in **Chapter 3**, this would then lead to poor agreement with experimental measurements.

To illustrate this, we show in **Figure 6.4** how the electron Mach number along channel centerline varies with space for the empirical reference simulation for the H9 at 300 V and 15 A. As the magnetic field is a design variable, and the ion velocity and thus electric field are measured to a high degree of accuracy, this can be taken as a relatively accurate picture of what the electron Mach number is in the H9 at this condition. We also show the results of applying the model scaling of **Equation 6.12** to the output of the reference simulation. As expected, the electron Mach number peaks sharply just downstream of the thruster exit plane. If we apply the the model of Cappelli et al. to this distribution of plasma properties, we obtain a maximum in the anomalous collision frequency at the precise location where the empirical profile has a minimum. The high transport in this region would act to reduce the peak electric field, gradually smoothing out the electric field profile and making the ion acceleration profile less steep than experiment. This is what we observed in the preceding section.

6.5.2 Chodura

We can rewrite the Chodura model (**Equation 6.3**) in terms of the electron Mach number, as

$$\nu_{an} = \omega_{pi} \left[1 - \exp\left(-\frac{1}{c_2}\sqrt{\frac{m_i}{m_e}} M_e \right) \right]. \tag{6.13}$$

In this case, we have a base collision frequency (in this case the ion plasma frequency, ω_{pi}) weighted inversely by a factor which depends on the electron Mach number. In contrast with the model of Cappelli et al., as the Mach number increases, the collision frequency decreases. We thus expect intuitively that this model should produce a minimum collision frequency coincident with the peak Mach number (as exhibited by the reference
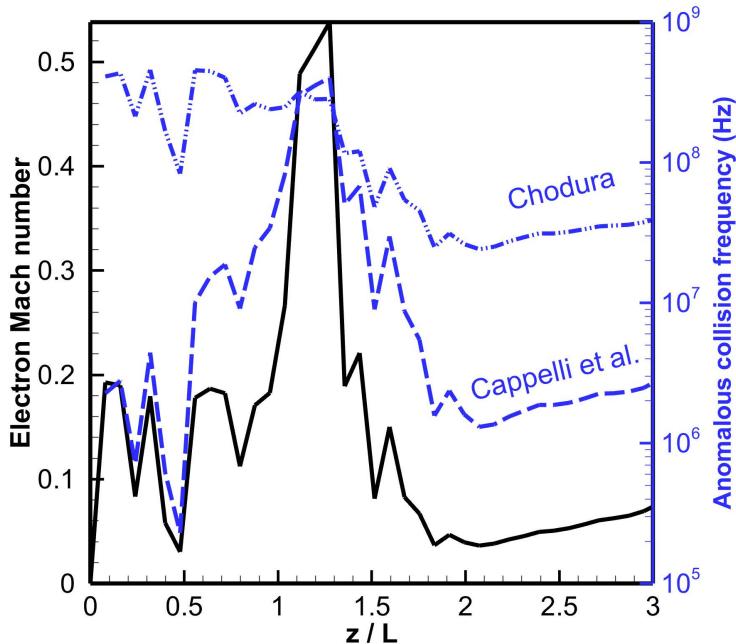
Figure 6.4: Electron azimuthal Mach number (left axis, black solid line) and evaluated anomalous collision frequencies (right axis) of models of Cappelli et al. (blue dashed line) and Chodura (blue dashed-dotted line) along the channel centerline for the calibrated reference simulation at 300 V and 15 A operating on xenon.

simulation). The conflating factor, however, is that the ion plasma frequency scales with the square root of the plasma density, which in turn decays monotonically with distance from the anode. This makes it difficult to simultaneously obtain the required high degree of transport in the near field plume while also having a minimum near the exit plane. These two factors thus compete with each other to produce the collision frequency profile shown in **Figure 6.4** which is relatively flat over large regions of the domain. Without the strong variations in collision frequency needed to localize the peak electric field, the ions experience gradual, largely constant acceleration.

6.5.3 Data-driven

While we can analyze the behavior of the preceding two models by comparing their behavior to the empirical anomalous collision frequency profile and the plasma properties

of the reference simulation, this is not true for the data-driven model, which was developed using exactly these empirical profiles.[3] As shown in the original paper by Jorns[3] and in our recent work[96], the data-driven model is able to predict the empirical collision frequency profile of the H9 operating at 300 V and 15 A on xenon (Case 1) with a high degree of accuracy given only the plasma properties resulting from a simulation run using that empirical profile.

Instead, we consider how the data-driven model couples with the ion dynamics of the Hall thruster simulation. As the model **(Equation 6.4)** depends linearly on the ion velocity, the collision frequency should reduce to the classical collision frequency at the ion stagnation point, where $|\vec{u}_i| = 0$. As shown in **Figure 5.7**, this occurs about 0.7 channel lengths downstream from the anode in the empirical reference simulation. This would recenter the peak electric field (and thus maximum ion acceleration) to this location. The ion stagnation point would then shift further upstream. This process would repeat until the ion stagnation point and minimum collision frequency eventually reach the anode. This explains the nearly monotonic anomalous collision frequency profile seen in **Figure 6.2a**, and thus the gradual ion acceleration exhibited by this model.

One other interesting feature of this model is its ability to predict empirical transport profiles despite its poor performance in self-consistent simulations. We posit that the data-driven model contains a re-arrangement of Ohm's law. If the symbolic regression procedure produced a model which restated Ohm's law, it would be able to predict the anomalous collision frequency from plasma data while appearing to be a unique model of the anomalous transport. When implemented into a simulation however, it would provide no new information to the code about how the anomalous collision frequency should vary with local plasma properties.

To investigate whether the data-driven model contains a restatement of Ohm's law, we invoke the definition of the discharge current, assuming the magnetic field is purely radial and the ions and and electrons have no velocity parallel to that field.

$$I_D = (q_e n_e u_{i,\perp} + j_{e,\perp})A_{ch}$$

$$\implies |u_i| \approx \frac{I_D/A_{ch} - j_{e,\perp}}{q_e n_e}$$

In this expression, A_{ch} is the cross-sectional area of the discharge channel. Next, we note that in Hall thrusters, $c_s \ll v_{de}$. Invoking these two relationships, the data-driven model (Equation 6.4) becomes

$$\frac{\nu_{an}}{\omega_{ce}}v_{de} = c_1\frac{I_D/A_{ch} - j_{e,\perp}}{q_e n_e}$$

Next, we note that $v_{de} \equiv |\vec{E}|/|B|$ and assume that $|\vec{E}| \approx |E_\perp|$ to find:

$$j_{e,\perp} = -\frac{\nu_{an}}{c_1\omega_{ce}B}q_e n_e|E_\perp| + \frac{I_D}{A_{ch}}$$

We can compare this to our earlier definition of Ohm's law (c.f. **Equation 2.33**), neglecting the effect of pressure gradient:

$$j_{e,\perp} = \frac{\nu_e}{\omega_{ce}B}q_e n_e E_\perp.$$

While not exactly identical due to the presence of absolute values and the constant factor of I_D/A_{ch}, it appears that the data-driven model primarily contains information related to Ohm's law, rather than a unique closure model.

6.5.4 Lafleur et al.

While the model of Lafleur et al. performed the worst out of the four models in terms of the global performance metrics, it was the only model able to adequately predict the steepness of the ion acceleration profile and the presence of a strong minimum in the

anomalous collision frequency. However, the location of this minimum was almost 50% further downstream than the location of the minimum in the empirical profile. One possible reason for this behavior can be motivated from the one-dimensional electron internal energy equation (c.f. **Reference 58**):

$$\frac{\partial}{\partial t}\left(\frac{3}{2}n_e T_e\right) + \frac{\partial}{\partial z}\left(\frac{5}{2}u_{e,z}n_e T_e\right) = m_e n_e \nu_e |\vec{u}_e|^2 + u_{e,z}\frac{\partial p_e}{\partial z} - S_{loss} \tag{6.14}$$

In the above, S_{loss} represents energy losses to ionization, excitation, and wall effects. Additionally, we have neglected the heat conduction term. We will discuss the effects of this term in the next section. We make the additional assumption that $|\vec{u}_e| \approx v_{de}$. Substituting in the model of Lafleur et al. (**Equation 6.2**), we obtain

$$\frac{\partial}{\partial t}\left(\frac{3}{2}n_e T_e\right) + \frac{\partial}{\partial z}\left(\frac{5}{2}u_{e,z}n_e T_e\right) - c_1 M_e \sqrt{\frac{m_i}{m_e}}\left|\frac{\partial}{\partial z}u_{i,z}n_e T_e\right|$$
$$= m_e n_e \nu_{e,classical}v_{de}^2 + u_{e,z}\frac{\partial p_e}{\partial z} - S_{loss} \tag{6.15}$$

We note that since $M_e \sim 0.1$, $c_1 \sim 0.1$, $u_i \sim u_e$, and $\sqrt{m_i/m_e} \sim 500$ for xenon, including an anomalous collision frequency of this form introduces an additional convection term to the energy equation with a magnitude comparable to or greater than the original convection term. For the purposes of this discussion, we neglect this term so as to examine the effect of this new convection term in isolation:

$$\frac{\partial}{\partial t}\left(\frac{3}{2}n_e T_e\right) - c_1 M_e \sqrt{\frac{m_i}{m_e}}\left|\frac{\partial}{\partial z}u_{i,z}n_e T_e\right|$$
$$= m_e n_e \nu_{e,classical}v_{de}^2 + u_{e,z}\frac{\partial p_e}{\partial z} - S_{loss} \tag{6.16}$$

The dynamics of this modified energy equation are significantly different from the original equation **(Equation 6.14)**. The most salient difference is that the electron thermal energy now convects with the ion drift speed instead of the axial electron velocity. In a typical one-dimensional Hall thruster discharge, the axial electron velocity is significantly reduced near the location of peak magnetic field and low anomalous transport (c.f. **Equation 2.33**). This, combined with the strong Joule heating at this location, has the effect of localizing the electron temperature profile to this point. In the modified energy equation **(Equation 6.16)**, both of these localizing factors are diminished. As the heating rate now scales with $\nu_{e,classical}$ instead for ν_e, the impact of Joule heating on the energy balance is reduced by at least an order of magnitude. Additionally, the axial ion velocity in a Hall thruster increases monotonically moving downstream of the anode, with no minima which would anchor the new convection term. Taken together, these factors explain how the electron temperature profile (and thus the anomalous collision frequency profile, which depends self-consistently on the electron temperature) are displaced so far downstream in the model of Lafleur et al. compared to the reference simulation. The magnitude of this displacement would then be determined by the balance between the convection term and the other source terms. This suggests that loss terms which are captured under S_{loss} may be more important in determining the one-dimensional structure of the plasma when using anomalous transport models which depend self-consistently on the convection of electron thermal energy.

The other challenge we encountered with the model of Lafleur et al. was its poor performance on Case 4 (krypton at 300 V and 15 A). This suggests that the model may not properly capture the species-dependence of the instability-driven transport.

6.6 Role of electron heat flux

Aside from the discrepancies with experiment already discussed, one feature common to all of the models investigated in this chapter, including the multi-zone Bohm reference profile, are broad electron temperature profiles. In **Figure 6.2c**, we see that the predicted electron temperatures two channel-lengths downstream of the exit plane are on the order of 10 to 20 eV. Experimental measurements of the temperature in this region suggest that these values are too high, with the true temperature likely being between 3 and 7 eV.[31,98] This phenomenon has been observed in other Hall thruster simulation works,[105] and it has been remarked[94] that it is difficult to simultaneously match the experimental electron temperature and ion velocity measurements using an empirically-calibrated multi-zone anomalous collision frequency profile.

One factor which may explain this is the electron heat conduction, which we have neglected to this point. In Hall2De, as in other Hall thruster models,[47] the conductive heat flux is assumed to follow Fourier's law of conduction:

$$\vec{q}_e = \kappa_e \nabla T_e, \tag{6.17}$$

where $\vec{q}_e$ is the electron heat flux vector and κ_e is the anisotropic electron thermal conductivity tensor. The cross-field component of this tensor is given by[106]

$$\kappa_{\perp,e} = 4.7 \frac{\nu_e n_e T_e}{m_e \omega_{ce}^2}. \tag{6.18}$$

This expression is derived under the assumption that the heat conduction is moderated by classical electron-ion collisions, i.e. $\nu_e = \nu_{ei}$. In Hall2De, in order to approximate the effect the anomalous transport on the heat flux, we use the sum of anomalous and classical collision frequencies in this expression, i.e. $\nu_e \approx \nu_{ei} + \nu_{an}$ as we do for the generalized Ohm's law. This is common in other models of low-temperature plasmas in electric propulsion.[47,58] However, the theoretical underpinnings of this substitution are

suspect. While it can be shown that many of the proposed mechanisms that contribute to anomalous electron momentum transport can be approximated using an effective collision frequency in a generalized Ohm's law, there has yet to be a demonstration that the classical Fourier law can be modified in a similar way. Indeed, the Fourier law is valid under the assumption that the collisions serve to relax the electron velocity distribution function back to equilibrium, while the processes that produce the anomalous "collisions" may not have the same effect. For example, if the transport is driven by the growth and convection of plasma waves, the thermal transport may be affected in a more complex manner which cannot be easily represented in a Fourier-like formulation.

The fact that the electron temperature profiles are too wide indicates that incorporating the anomalous collision frequency into **Equation 6.18** may lead to an over-estimate of the true cross-field thermal conductivity. Accurate modeling the effect of anomalous processes on the electron heat flux, as well as the impacts of other non-ideal and anisotropic processes in the electron energy balance,[107] may be important to fully capture the effect of anomalous electron transport on Hall thruster discharge plasmas.

6.7   Implications for closure model development

In our previous work,[96] we showed that the ability of an anomalous collision frequency model to reproduce measured electron transport profiles when evaluated with time-averaged or steady plasma data does not guarantee the closure model will perform well when implemented directly into a Hall thruster simulation. We theorized that the strong coupling between the anomalous transport model and the Hall thruster model can cause the simulation to quickly diverge to a new equilibrium profile. This is supported by the present chapter, in which we found that the behavior of all four models can be explained by considering the feedback between anomalous collision frequency model and fluid Hall

thruster model. With these results in mind, our findings recommend some guidance for refining new closure models of the Hall thruster anomalous electron transport.

First, as suggested by our previous work, it is not sufficient to evaluate a model of the anomalous collision frequency against time-averaged plasma data, whether experimental or numerical. To evaluate the predictive ability of a closure model, it must be implemented directly into a Hall thruster code, and the collision frequency must be allowed to vary self-consistently with the plasma properties it depends on.

Second, as highlighted by the results of the model of Lafleur et al. and our subsequent discussion of the electron heat flux, the electron energy balance appears to play a critical role in the coupling between the functional form of the closure model and the rest of the Hall thruster model. When developing a new closure model for the anomalous collision frequency, it is thus important to consider the effect of the model not just on the momentum balance but also on the energy balance.

Lastly, it is critical that proposed models must be tested against multiple operating conditions, and ideally, multiple thrusters. Extensibility with voltage, current, propellant choice, magnetic field strength, thruster geometry, and background pressure should be considered. Particular emphasis should be placed on matching spatially-resolved plasma properties in addition to targeting the global performance metrics.

6.8   Conclusion

In this chapter, we have assessed the predictive capabilities of closure models for the anomalous electron transport in Hall effect thrusters. We implemented these models self-consistently into Hall2De, a 2D axisymmetric Hall thruster code developed by the Jet Propulsion Laboratory, and compared their performance on four operating conditions of the H9 Hall thruster to experimental data obtained at the University of Michigan. We investigated four models. The first was based on classical turbulence theory, the second

on empirical scaling laws, the third on the scaling of the electron drift instability in Hall thrusters, and the last from data. We calibrated the coefficients of these models to match the experimental discharge current of one of these four conditions, and then applied these calibrated coefficients the other three operating conditions. In this way, we were able to assess the ability of the models to generalize across operating conditions. To our knowledge, this is the first comprehensive comparison of different models of the anomalous electron transport using the same simulation code and operating conditions.

While a few of the models were able to capture some of the scaling trends with discharge current and changes in propellant, most exhibited incorrect scaling of thrust and discharge current with increasing discharge voltage. Additionally, three of the models exhibited relaxed ion acceleration profiles and were unable to capture the strong reduction in anomalous transport coincident with the location of maximum radial magnetic field. In the model based on the electron drift instability, the observed ion acceleration curves were more in line with experimental data, but the location of the acceleration zone was shifted downstream by half of a channel length.

We then explained the performance of each of the models by considering the coupling between the closure models and the simulation dynamics. In all four cases, we showed that the dependence of these models on local plasma properties will drive the converged solution away from the experimental data. Thus, even though these models are predicated on physics-based hypotheses or derived directly from data, their functional scalings are evidently incorrect. Relatedly, one common element to all of the models, even those calibrated to match data, is that they predict electron temperature profiles more relaxed than the experiments suggest. While this can be attributed partially to the form of the collision frequency in each of the models, it may also result from an over-simplification of how anomalous momentum transport impacts anomalous heat conduction. With that said, the insights we have gained from analyzing the limitations of these models suggest some

163

guidelines for future efforts to develop a predictive model of anomalous electron transport in Hall thrusters.

CHAPTER 7

Effect of Changing Assumptions in a First-Principles Model of Anomalous Electron Transport

> Those who are able to adapt and change in accord with the enemy and achieve victory are called divine.
>
> ———————————————————————————
>
> Sun Tzu, The Art of War

So far, we have shown that several of the models for Hall thruster anomalous transport in **Chapter 3** do not seem to work very well when implemented into Hall2De. In this chapter, we investigate the effect of changing some of the assumptions of these models and examining how that affects the model performance. Given its theoretical grounding, we focus on the model of Lafleur et al. **(Equation 3.31)**.

To this end, the first part of this chapter, we examine the effect of changing the wave energy at which instability is assumed to saturate. We also investigate the effect of including loss mechanisms, such as charge exchange collisions and ion Landau damping, on closure model performance. In the second part, we investigate the effect by changing

the form of the convection term in the numerator, which has implications for the role of radial effects on models of the Hall thruster anomalous collision frequency.

7.1   Investigation of a drift-saturated transport model

This work was previously presented at the 2022 AIAA SciTech forum as **Reference 108.**

In Chapter 3, we showed that by relaxing some of the assumptions of the model of Lafleur et al., we can arrive at the following expression for the anomalous collision frequency

$$\nu_{an} = c_1 \frac{\gamma W}{m_e n_e c_s v_{de}} \quad \text{s}^{-1},$$
(7.1)

where c_1 is a unitless constant, γ is the growth rate of the wave, W is the energy density of the wave, $c_s = \sqrt{k_B T_e / m_i}$ is the ion sound speed, $v_{de} = |\mathbf{E} \times \mathbf{B}|/|\mathbf{B}|^2$ is the azimuthal electron $E \times B$ drift speed, m_i is the ion mass, T_e is the electron temperature in K, and Z is the ion charge number.

To relate γ and W to the local plasma properties, we need to make some further assumptions. For this, we employ a conservation equation for wave energy[109]

$$\frac{DW}{Dt} = \frac{\partial W}{\partial t} + (\mathbf{v}_g \cdot \nabla)W = 2W(\gamma - \omega_{loss}).$$
(7.2)

In the above, $\mathbf{v}_g$ is the group velocity of the wave and ω_{loss} is the rate at which the wave loses energy to the plasma via damping. We assume the wave propagates with the ions as they convect downstream, so $\mathbf{v}_g \approx \mathbf{u}_i$, where $\mathbf{u}_i$ is velocity of singly-charged ions. We note that this is a slightly different form for the convection equation than Lafleur et al. employed. Rearranging terms and making the further assumption that $\partial W / \partial t \approx 0$ (i.e., assuming marginal stability), we find that

$$2\gamma W = \mathbf{u}_i \cdot \nabla W + 2\omega_{loss}W. \tag{7.3}$$

We can now substitute the above expression into Eq. 7.1 and fold the factor of two into c_1 to get our closure model master equation

$$\nu_{an} = c_1 \frac{\mathbf{u}_i \cdot \nabla W + 2\omega_{loss}W}{m_e n_e c_s v_{de}} \quad \text{s}^{-1}. \tag{7.4}$$

Armed with this expression, all that we need to obtain closure models for the anomalous collision frequency is to substitute suitable expressions for the damping frequency ω_{loss} and wave energy density W.

7.1.1   Closure models

We consider five models in this section, each of which has been derived from the **Equation 7.4**. In contrast to some previous work, here we assume the wave energy saturates at the electron drift energy density, i.e. $W \propto m_e n_e v_{de}^2$ instead of the electron thermal energy. We then make several different assumptions for the scaling of ω_{loss}. In model A we assume that $\omega_{loss} = 0$, i.e. that there is no collisional or landau damping and that convection is the sole loss mechanism. In models B and C, we assume that the wave can lose energy to ion Landau damping. The ion Landau damping rate for an ion acoustic-like wave with wavenumber k can be approximated as[110]

$$\omega_i = -kc_s\sqrt{\frac{\pi}{8}}\left[\sqrt{\frac{Zm_e}{m_i}} + \left(\frac{ZT_e}{T_i}\right)^{\frac{3}{2}}\exp\left(-\frac{ZT_e}{2T_i} - \frac{3}{2} - 3\frac{T_i}{ZT_e}\right)\right]. \tag{7.5}$$

In the above, we define T_i to be the temperature of singly-charged ions. For our analysis, we assume the majority of the growth occurs at a single wavenumber. In model B, we assume that frequency of maximum growth is the plasma frequency, so that $k \approx$

$1/\sqrt{2}\lambda_{de}$[111]. This gives a damping rate of:

$$\omega_i = -\frac{\sqrt{\pi}\omega_{pe}}{4}\left[\frac{Zm_e}{m_i} + \sqrt{\frac{Zm_e}{m_i}}\left(\frac{ZT_e}{T_i}\right)^{\frac{3}{2}}\exp\left(-\frac{ZT_e}{2T_i} - \frac{3}{2} - 3\frac{T_i}{ZT_e}\right)\right] \qquad (7.6)$$

For Xenon and Krypton, $m_e/m_i = 4.178 \times 10^{-6}$ and 6.546×10^{-6}, respectively so we neglect the first term in the brackets and arrive at our final expression for the ion landau damping rate:

$$\omega_{loss} \propto \omega_{pe}\sqrt{\frac{Zm_e}{m_i}}\left(\frac{ZT_e}{T_i}\right)^{\frac{3}{2}}\exp\left(-\frac{ZT_e}{2T_i} - \frac{3}{2} - 3\frac{T_i}{ZT_e}\right) \qquad (7.7)$$

In model C, we instead assume that the frequency of maximum growth is an electron cyclotron resonance. This gives a loss rate of

$$\omega_{loss} \propto \frac{\omega_{pe}}{M_e}\sqrt{\frac{Zm_e}{m_i}}\left(\frac{ZT_e}{T_i}\right)^{\frac{3}{2}}\exp\left(-\frac{ZT_e}{2T_i} - \frac{3}{2} - 3\frac{T_i}{ZT_e}\right), \qquad (7.8)$$

where $M_e = v_{de}/c_e$ is the electron Mach number and $c_e = \sqrt{T_e/m_e}$ is the electron sound speed. In models D and E, we investigate the effects of classical collisions on the wave energy. In model D we assume that electron-ion collisions damp the wave, with a damping term of

$$\nu_{ei} = \frac{Z^2 e^4 n_i \ln\Lambda}{(4\pi\epsilon_0)^2 m_e^{1/2} T_e^{3/2}} \qquad (7.9)$$

$$\omega_{loss} \propto \nu_{ei}, \qquad (7.10)$$

where $\ln\Lambda$ is the Coulomb logarithm. For model E, we consider ion-neutral charge exchange (CEX) collisions instead of electron-ion collisions. For the purposes of computing the charge exchange collision frequency, both ions and neutrals are assumed

to be cold, with the resulting damping rate having the form

$$\nu_{in,CEX} = n_n u_{in} \sigma_{CEX} \tag{7.11}$$

$$u_{in} = \frac{v_{th,i}}{\sqrt{\pi}} \left[\exp\left(-\tilde{u}_i^2\right) + \frac{\sqrt{\pi}}{2} \left(2\tilde{u}_i + \frac{1}{\tilde{u}_i} \right) \mathrm{erf}(\tilde{u}_i) \right] \tag{7.12}$$

$$\tilde{u}_i = \frac{|\mathbf{u}_i|}{v_{th,i}} \tag{7.13}$$

$$\sigma_{CEX} = 10^{-20} \frac{1}{Z} \left[A - B \log\left(\frac{m_i}{q_e} u_{in}^2 2 \right) \right] \tag{7.14}$$

$$\omega_{loss} \propto \nu_{in,CEX}. \tag{7.15}$$

In the above, n_n is the neutral density, u_{in} is the effective ion-neutral collision velocity, $v_{th,i}$ is the ion thermal speed, equal to $\sqrt{2k_B T_i/m_i}$, $\tilde{u}_i$ is the magnitude of the ion velocity normalized by the thermal speed, σ_{CEX} is charge exchange collision cross section as a function of kinetic energy, and A and B are numerical fit coefficients. For xenon, Miller et al.[112] found $A = 87.3$ and $B = 13.6$, while for krypton, Hause et al.[113] found $A = 80.7$ and $B = 14.7$. We summarize our models in **Table 7.1**.

Model	Loss mechanisms	Expression for ν_{an}
A	Convection	$c_1 \dfrac{\left|\mathbf{u}_i \cdot \nabla(n_e v_{de}^2)\right|}{n_e c_s v_{de}}$
B	Convection, Landau damping at ω_{pe}	$c_1 \left[\dfrac{\left|\mathbf{u}_i \cdot \nabla(n_e v_{de}^2)\right|}{n_e c_s v_{de}} + c_2 \omega_{pe} M_e \left(\dfrac{ZT_e}{T_i}\right)^{3/2} \exp\left(-\dfrac{1}{2}\dfrac{ZT_e}{T_i} - \dfrac{3}{2} - 3\dfrac{T_i}{ZT_e} \right) \right]$
C	Convection, Landau damping at $\frac{\omega_{ce}}{M_e}$	$c_1 \left[\dfrac{\left|\mathbf{u}_i \cdot \nabla(n_e v_{de}^2)\right|}{n_e c_s v_{de}} + c_2 \omega_{ce} \left(\dfrac{ZT_e}{T_i}\right)^{3/2} \exp\left(-\dfrac{1}{2}\dfrac{ZT_e}{T_i} - \dfrac{3}{2} - 3\dfrac{T_i}{ZT_e} \right) \right]$
D	Convection, $e - i$ collisions	$c_1 \left[\dfrac{\left|\mathbf{u}_i \cdot \nabla(n_e v_{de}^2)\right|}{n_e c_s v_{de}} + c_2 \sqrt{\dfrac{m_i}{m_e}} M_e \nu_{ei} \right]$
E	Convection, CEX collisions	$c_1 \left[\dfrac{\left|\mathbf{u}_i \cdot \nabla(n_e v_{de}^2)\right|}{n_e c_s v_{de}} + c_2 \sqrt{\dfrac{m_i}{m_e}} M_e \nu_{in,CEX} \right]$

Table 7.1: Summary of investigated models

7.1.2 Simulations

All simulations in this part of the chapter used the same set-up as that overviewed in **Chapter 4**. The sole exception is the grid resolution. In the present study, we use a MFAM with 2176 cells and a Cartesian grid with 1909 cells. This is not in the full-converged region of the grid-convergence study conducted earlier, but is sufficiently fine so as to allow the resolution of bulk trends. By choosing a coarser grid, we were able to perform a larger number of simulations at reduced computational cost. This allowed finer resolution of performance trends with changing model parameters. The timestep is set globally by the CFL condition,

$$\Delta t = \min_{i \in \text{cells}} \text{CFL}_{\text{max}} \frac{\lambda_{max,i}}{\Delta x_i}, \tag{7.16}$$

where Δt is the timestep, CFL_{max} is a user-configurable maximum allowed CFL number, $\lambda_{max,i}$ is the maximum ion wave speed in cell i, and Δx_i is a measure of size of cell i. In our case, we set Δx_i to the the ratio of cell volume to cell surface area. With the chosen mesh resolution and a maximum CFL number of 0.9, simulation time steps were on the order of 50 nanoseconds.

We also employed the same calibration procedure as in the preceding chapter. In brief, for each of the models, we simulated many values of c_2. At each of these, we employed a PI controller to tune c_1 to 15 A. To calibrate c_2, the damping rate scale coefficient, we perform a parameter sweep in the range of $0.1 - 100$ with 11 simulations distributed logarithmically. For models which do not have a second coefficient, this step is not required. This combination of online and offline calibration methods allows us to optimize our closure model fit coefficients efficiently and effectively using the smallest number of simulations possible.

With our methods defined, we now present the results of the study.

Model	Optimal c_1	Optimal c_2	Optimal η_a	Optimal IVE
A	0.490	N/A	59.4%	0.444
B	0.420	0.1	57.6%	0.443
C	0.470	0.1	59.3%	0.446
D	0.476	0.1	56.6%	0.446
E	0.451	0.25	59.3%	0.445

Table 7.2: Optimal coefficient values and performances for the models investigated in this work.
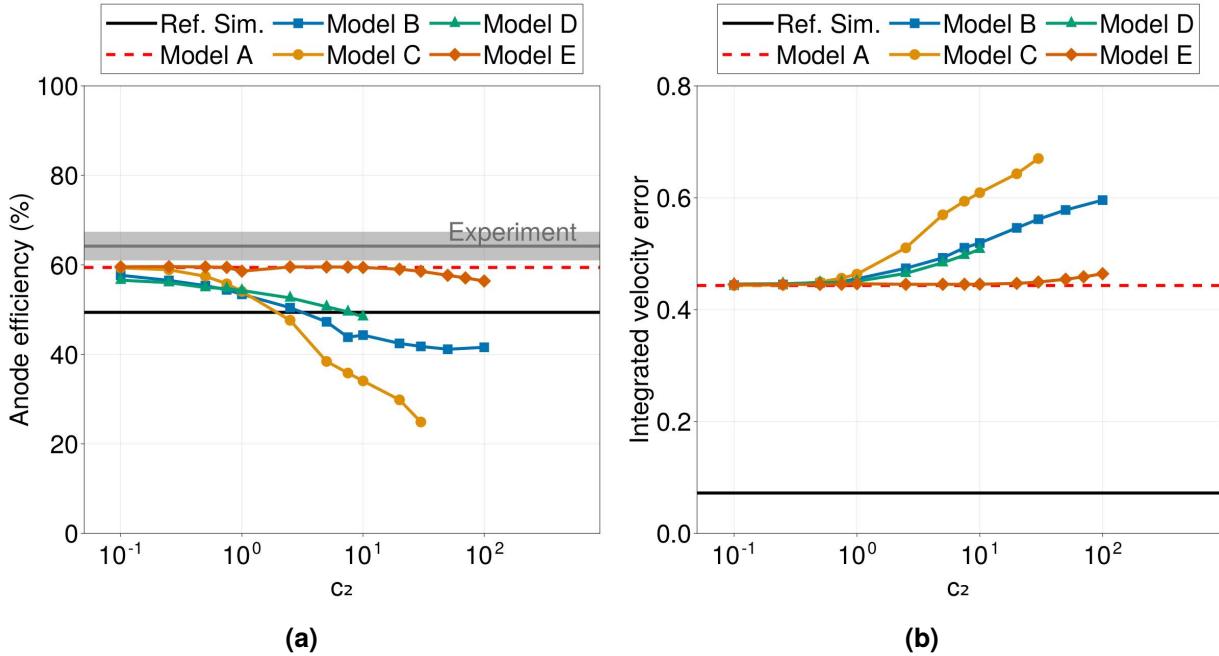
Figure 7.1: (a) Anode efficiency and (b) integrated velocity residual as a function of fit coefficient c_2 for all models. Experimental anode efficiency is indicated as a grey shaded region in (a).

7.1.3 Results

In **Figure 7.1**, we present the results of our study. We were unable to obtain converged solutions for model C with $c_2 > 50$ and model D with $c_2 > 10$. All simulations had a discharge current within 0.5 A of the target value of 15 A. We compare the simulated anode efficiencies **(Figure 7.1a)** and integrated velocity residuals/errors as a function of the second fit coefficient c_2 to experimental values, as well as to a calibrated reference simulation of the H9 at 300 V and 15A.
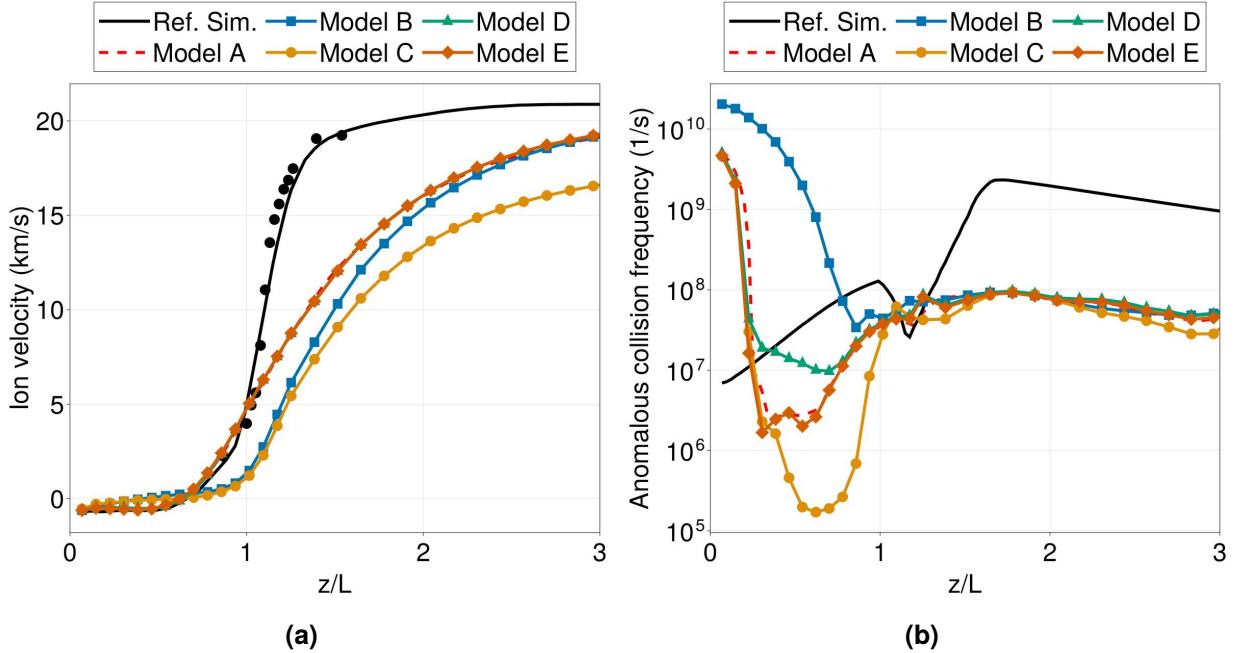
Figure 7.2: (a) Ion velocity and (b) anomalous collision frequency along channel centerline for best-performing cases of each model. LIF data is represented as discrete black circular markers in (a).

First examining the anode efficiencies, all models were able to match the experimental anode efficiency of 64% to within 10% at low values of c_2. The best-performing model was model A (convection with no damping term) with an anode efficiency of 59.4%. In the limit of low c_1, the other models approached result of model A, but none were able to improve upon it. As c_2 increased and thus the damping term became dominant over the convection term in models B - E, the efficiency monotonically decreased. The worst obtained efficiency was 24%, seen in model C (convection with ion landau damping at cyclotron resonances) at high values of c_2. Interestingly, all models were able to improve upon the anode efficiency of the reference simulation at low values of c_2. Calibrated stationary anomalous collision frequency profiles such as the one employed for the reference simulation often have difficulty matching experimental thrust values,[82,53] so this is an encouraging result. However, these empirically-inferred profiles are typically calibrated to match the discharge

current and ion velocity profile, not the thrust, so it is possible that further calibration could rectify this discrepancy.

Next, we examine the integrated velocity residuals. For all models except for model G, this was constant at roughly 0.45 for $0 < c_2 < 1.0$. For $c_2 > 1.0$, the IVR increased for all models in ways that mirrored the trends in efficiency. To better understand why the integrated velocity residuals are so similar across all of the models, we plot in **Figure 7.2** the axial component of the ion velocity **(Figure 7.2a)** and the anomalous collision frequency **(Figure 7.2b)** extracted along the discharge channel centerline for the cases with the highest anode efficiency for each model. These best-fit coefficients are summarized in **Table 7.2**. As expected from the IVRs, the ion velocity profiles for all models are very similar. Models A, D, and E are indistinguishable from each other, while models B and C are somewhat different, with acceleration happening a tenth of a thruster length later than the other simulations. In all cases, the ion acceleration is more relaxed than that of the experimental data and reference simulation.

To understand the structure of the ion acceleration profiles, we now turn to examining the time-averaged anomalous collision frequency along the channel centerline. In all cases, the anomalous collision frequency starts high, then declines to a minimum before recovering somewhat. This qualitatively matches the shape of the empirically-inferred anomalous collision frequency profile used to generate the reference simulation. However, the collision frequency is generally at least an order of magnitude lower than the empirical profile, with the exception of near the anode, where it is two to three orders of magnitude higher. Downstream of the thruster exit plane ($z/L = 1$), all five models attain a nearly constant collision frequency of 10^8 s^{-1}. However, the models differ from each other by several orders of magnitude upstream of this location. Models A and E are indistinguishable from one another, with a minimum occurring at $z/L \approx 0.5$. For model D, the trough in collision frequency is shallower than in models A and E by an order of magnitude. For model B, the collision frequency is much higher upstream than the other

models, with a shallow minimum around z/L = 0.9. As the location of the minimum in collision frequency coincides with the location of the peak electric field strength in Hall thrusters, this feature explains the delayed acceleration profile seen in model B. For model C, the minimum collision frequency is an order of magnitude deeper than models A and E. The location of the minimum in most models is near z/L = 0.5, which is nearly coincident with the ion stagnation point, where u_i = 0, where the convection term should be minimized. Taken with the efficiency and integrated velocity residual data, these results reinforces the idea that the convection term dominates the damping terms, especially downstream of the exit plane.

The high value of the collision frequency near the anode was also observed by Lafleur et al[2] in their analytical model, which shares a similar derivation to ours. It likely stems from the fact that the ion sound speed c_s and electron drift velocity v_{de} are both low near the anode, causing the value of the collision frequency to grow rapidly in this region. To fix this, they cut off the anomalous collision frequency near the anode, with the justification that the wave may not be saturated before the ion stagnation point. With this modification, the magnitude of the collision frequency downstream of the anode minimum would likely need to increase by an order of magnitude in order to match the target discharge current. However, it is not assured that our simulations need to match the empirical reference profile in order to give good results, as these profiles are known in general to be non-unique[94] and may not correspond well with the actual value of the anomalous collision frequency in Hall thrusters. Nevertheless, a different model may be necessary in the near anode region in order to account for differences in the level of wave saturation.

The other major difference between our model predictions and the empirical reference profile is the location of the minimum collision frequency. To better match experiment, it would need to be shifted half a thruster length downstream from its present location. One possible way to achieve this would be if the group velocity of the wave in the axial direction was not $u_{i,z}$ but $u_{i,z} - c_s$. This might shift the minimum to the ion sonic point instead of

the ion stagnation point. In the experimental data, the ion sonic point sits just inside of the channel, near $z/L = 0.9$, while in our simulations, it sits between $z/L = 0.75$ and $z/L = 1.1$. This is still likely too far upstream, but would be an improvement over the present results. More technically, the group velocity of a wave is given by:

$$\mathbf{v}_g = \frac{\partial \omega}{\partial \mathbf{k}}, \tag{7.17}$$

where ω is the real frequency of the wave and $\mathbf{k}$ is the wave-vector of propagation. For an ion acoustic instability with drifting ions with a velocity vector $\mathbf{u}_i$, the real frequency is given by[2,109]

$$\omega = \mathbf{k} \cdot \mathbf{u_i} \pm \frac{kc_s}{\sqrt{1 + k^2\lambda_{de}^2}}. \tag{7.18}$$

In our previous analysis and in experiments, k is primarily in the azimuthal direction, so the sound speed contribution should be very small in the axial direction. However, here we make the assumption that the wave has a component that propagates axially, in the same direction as the ions. The group velocity for a wave with wavenumber k is then[2]

$$\mathbf{v}_g = \mathbf{u}_i \pm \frac{c_s}{\left(1 + k^2\lambda_{de}^2\right)^{3/2}} \hat{u}_i, \tag{7.19}$$

where $\hat{u}_i \equiv \mathbf{u}_i/|\mathbf{u}_i|$ is the unit vector of ion propagation. If the majority of the growth occurs at a single wavenumber and the wavenumber of maximum growth that given by the ion acoustic dispersion relation (i.e. $k^2\lambda_{de}^2 = \frac{1}{2}$), then we arrive at

$$\mathbf{v}_g = \mathbf{u}_i \pm \sqrt{\frac{8}{27}} c_s \hat{u}_i \approx \mathbf{u}_i \pm 0.544 c_s \hat{u}_i. \tag{7.20}$$

We can thus see that there may be some theoretical justification for shifting the profile forward by a factor of c_s. We also note that when implementing the model of Lafleur, Baalrud and Chabert (**Chapters 5** and **6**) we found that the minimum of the collision frequency occurred significantly further downstream than in the empirical reference simulation. That model was similar to ours in all respects except for the assumption that the wave saturated at some fraction of the plasma thermal energy density rather than the electron drift kinetic energy density. Combined with a more complete expression for the group velocity and improved modeling near the anode, it may be that allowing the wave to saturate at a fraction of the electron *total* energy density ($W \propto \frac{3}{2} n_e k_B T_e + m_e n_e v_{de}^2$) might cause the minimum in the anomalous collision frequency to move to a location intermediate between that predicted by the drift and thermal energy assumptions alone. However, the thermal energy typically dominates the drift kinetic energy, so it is possible that the result would be similar to that predicted by Lafleur et al.

Overall, while were unable to quantitatively match ion velocity measurements, we were able to obtain good agreement with the experimental anode efficiency values and qualitative agreement with some key features of the empirically-inferred anomalous collision frequency profile. In the next part of this chapter, we will examine a model which does a better job of matching the ion velocity profiles, but does not predict performance to the same degree.

7.2   Radial effects on anomalous transport models

In **Chapter 3**, we discussed that, by taking different forms of convection term in the wave energy equation, one could arrive at subtly-different models. Two of these forms were:

$$\frac{\partial W}{\partial t} + \nabla \cdot \left(\mathbf{v}_g W \right) = 2\gamma W \tag{7.21}$$

$$\frac{\partial W}{\partial t} + \mathbf{v}_g \cdot \nabla W = 2\gamma W. \tag{7.22}$$

Recall that W is the wave energy density, $\mathbf{v}_g$ is the wave group velocity, and γ is the growth rate. We will refer to the first form (**Equation 7.21**) as the *conservative form* of the wave energy conservation equation, while **Equation 7.22** is the *nonconservative* form. The model of Lafleur et al. employs the conservative form of this equation, while the drift-saturated models we tested in the preceding section used the non-conservative form. We sought to test the non-conservative form of the gradient, in the hopes that it would perhaps not be subject to the same downstream axial displacement that the baseline model of Lafleur et al. exhibited.

We used the same assumptions of the model of Lafleur et al.—namely that the wave saturates due to ion trapping at the electron thermal energy density and that the group velocity of the wave is equal to the ion convection velocity—and implemented this non-conservative model into Hall2De. This was done by modifying the implementation of the conservative form, since they were identical except for the gradient term. Using the definition of the divergence and gradient operators in cylindrical coordinates, these convection terms can be expanded as

$$\nabla \cdot (\mathbf{u}_i n_e k_B T_e) = \frac{\partial}{\partial z} (u_{iz} n_e k_B T_e) + \frac{\partial}{\partial r} (u_{ir} n_e k_B T_e) + \frac{u_{ir} n_e k_B T_e}{r} \tag{7.23}$$

$$\mathbf{u}_i \cdot \nabla (n_e k_B T_e) = u_{iz} \frac{\partial}{\partial z} (n_e k_B T_e) + u_{ir} \frac{\partial}{\partial r} (n_e k_B T_e) , \tag{7.24}$$

where u_{iz} and u_{ir} are the axial and radial components of the ion velocity, respectively. In the first expansion (**Equation 7.23**), we get a third term ($u_{ir} n_e k_B T_e / r$) which is not

present in the second form. We call this a *centripetal* term, as it represents the flux of a quantity out of a control volume due to centripetal acceleration. In implementing the non-conservative form of the gradient, we accidentally left this centripetal term in, which surprisingly led to significantly improved results over the baseline model of Lafleur et al. Removing this erroneous term resulted in poorly-convergent simulations with results worse than the baseline model. We will refer to the model with the non-conservative gradient and included centripetal term as the *centripetal model* throughout the remainder of this chapter. We now present the results of a study we performed to characterize this model.

7.2.1 Simulations

For this study, all simulation-related details are exactly as described in **Chapter 6**. As in that chapter, we simulated the H9 at four different operating conditions—300 V and 15 A on xenon, 300 V and 30 A on xenon, 600 V and 15 A on xenon, and 300 V and 15 A on krypton. We employed a field-aligned mesh for the electrons with 3925 cells and a rectilinear mesh for the ions with 3955 cells. We used a timestep of 15 nanoseconds in all cases. In contrast with our work in **Chapter 6**, we do not perform a calibration step. Instead, we use a value of c_1 = 0.057 for each operating condition. This is the same value that was found in that chapter. This allows us to more clearly assess the ways in which this model differs from the baseline model of Lafleur et al.

7.2.2 Results

In **Tables 7.3**, **7.4**, and **7.5,** we present the discharge current, thrust, and integrated velocity residual obtained for this centripetal model. As in **Chapter 6**, we compare these results to experimental data at the same operating conditions. As an additional point of comparison, we also report the data from the baseline model of Lafleur et al. These are the same results that were previously described in **Chapter 6**. In all cases excepting the 600 V

and 15 A case, which we will discuss momentarily, the centripetal model appears to be a marked improvement over the baseline model. The discharge current is changed by less than 0.5 A, indicating that the coefficient calibrated for the baseline model seems to apply to this model as well. Additionally, the thrust in these three cases is higher, yielding a closer agreement with experiment. Lastly, the integrated velocity residuals have declined substantially. While not as low as those of the the calibrated empirical profiles, the IVRs of the centripetal model appear to be in better agreement with the data than any of the models evaluated in **Chapter 6**.

Unfortunately, the 600 V and 15 A case does not conform to these trends. In this case, the discharge current and thrust are in worse agreement with experiment than the baseline case, and although the integrated velocity residual is somewhat improved, it is still higher than many of the models tested in the previous chapter.

Case	Condition	I_D (Experiment, A)	I_D (Lafleur et al., A)	I_D (Centripetal, A)
1	Xe, 300 V, 15 A	15 A	14.9	15.0
2	Xe, 300 V, 30 A	30 A	26.5	26.2
3	Xe, 600 V, 15 A	15 A	14.5	11.6
4	Kr, 300 V, 15 A	15 A	10.8	10.6

Table 7.3: Discharge current (I_D) predicted by the models compared to experimental discharge current.

Case	Condition	T (Experiment, mN)	T (Lafleur et al., mN)	T (Centripetal, mN)
1	Xe, 300 V, 15 A	292.9 ± 3.5	215.5	241.5
2	Xe, 300 V, 30 A	539.0 ± 4.0	398.2	465.1
3	Xe, 600 V, 15 A	447.2 ± 3.0	273.5	213.7
4	Kr, 300 V, 15 A	235.8 ± 2.5	149.0	156.7

Table 7.4: Thrust (T) predicted by each of the models compared to experimentally-measured thrust.

To further investigate these (largely) fortuitous results, we compare in **Figure 7.3** the axially-resolved plasma properties of the baseline model of Lafleur et al. to those of the new centripetal model for all four conditions. Th anomalous collision frequency profiles now more closely resemble the empirically-inferred profiles. As a result, the ion velocity
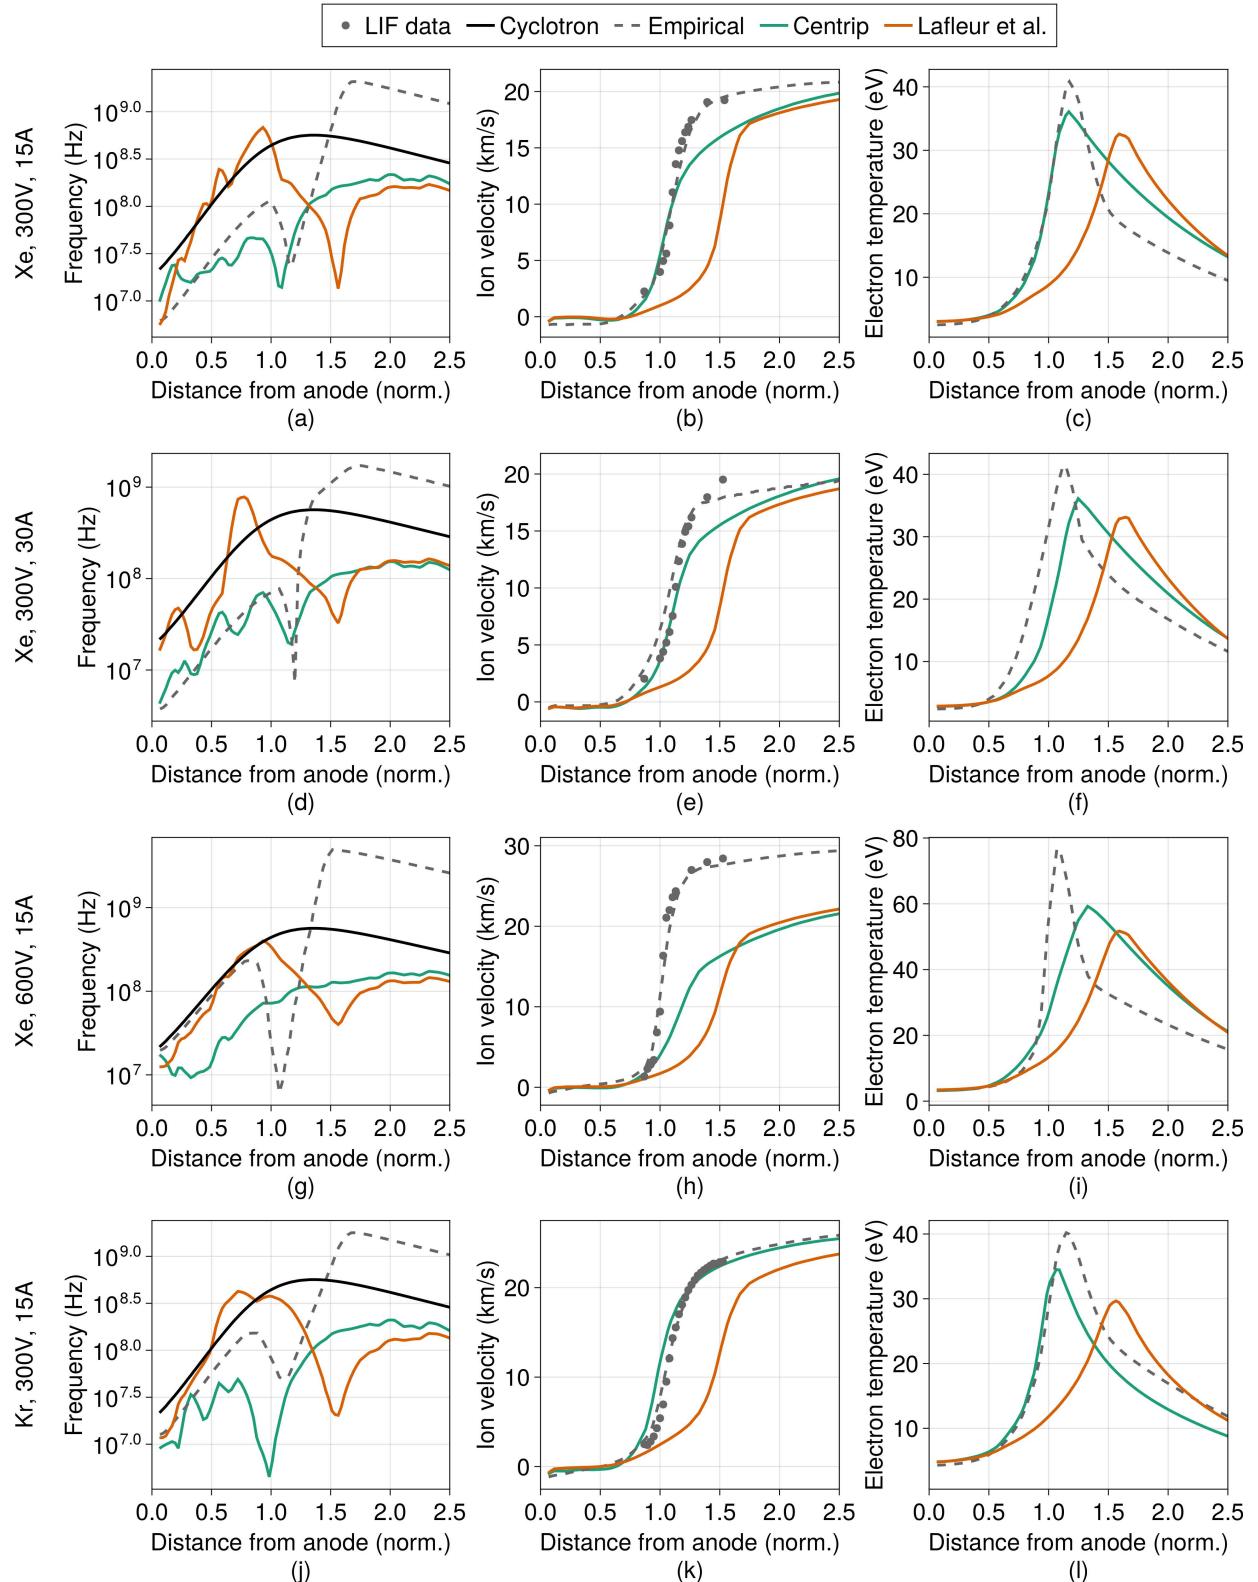
Figure 7.3: Comparison of the baseline model of Lafleur et al. to the model with the erroneous centripetal term (labelled Centrip) at the four test operating conditions.

Case	Condition	IVR (empirical)	IVR (Lafleur et al.)	IVR (Centripetal)
1	Xe, 300 V, 15 A	0.069	0.761	0.204
2	Xe, 300 V, 30 A	0.095	0.730	0.184
3	Xe, 600 V, 15 A	0.065	0.793	0.543
4	Kr, 300 V, 15 A	0.059	0.669	0.168

Table 7.5: Integrated velocity residual (IVR) the models compared to the IVR of the empirical reference simulation.

profiles now match the experimental data more closely than the baseline model. In the case of 300 V and 15 A on krypton, this match is almost as good as the empirical reference simulation. Unfortunately, as suggested by the global metrics for the 600 V and 15 A case, the anomalous collision frequency and ion velocity profiles are in very poor agreement with the empirical profiles and the data. One interesting feature shared by both the baseline model and the centripetal model at this condition is that the ion velocity seems to "plateau" very early compared to the other cases. The minima of the anomalous collision frequency profiles for both models are much shallower than the empirical profile, which could explain this behavior. Despite this poor performance at higher discharge voltages, it appears that the centripetal model represents an improvement over the baseline model of Lafleur et al.

7.2.3   Discussion

In our discussion of the model of Lafleur et al. in Chapter 6, we suggested that the model's convection term has the effect of modifying the electron energy equation in a manner which leads to the displacement of the location of peak electron temperature, and thus the minimum in the anomalous collision frequency profile. We posited that the balance between the convection term and the source terms in the energy equation was responsible for the ultimate location of the ion acceleration region, which corresponds to this minimum. These results support that hypothesis. By adjusting the form of the convection term, we were able to improve the model results in most cases. However, while a modification of the

convection term to a non-conservative form can be justified, the presence of the centripetal term cannot be easily explained.

One possible interpretation is that, since this term depends on the ion radial velocity, it may represent a loss of wave energy in the radial direction. Villafana et al.[39] recently performed three-dimensional particle-in-cell simulations which found that energy loss in the radial direction alters the character of transport when compared with equivalent two-dimensional simulations. Exploring alternate forms of this term may help incorporate an approximation of this radial wave transport into future self-consistent models. As we discussed in **Chapter 3**, the character of the ECDI changes based on the radial component of the wave-vector. The factor of $1/r$ in the centripetal model has the same units as a wavenumber, and could be on the right wavelength scale to explain the experimentally and numerically-observed results. The true form of the centripetal model could then be written as

$$\nu_{an} = c_1 q_e \frac{\left| \mathbf{u}_i \cdot \left(\nabla \left(n_e k_B T_e \right) + n_e k_B T_e \mathbf{k} \right) \right|}{m_e n_e c_s v_{de}}. \tag{7.25}$$

Here, $\mathbf{k}$ is the dominant wavenumber of the instability. If this has a finite radial component, then $\mathbf{u}_i \cdot \mathbf{k}$ would take on a form similar to our erroneously-included centripetal term.

Despite the good agreement with the ion velocity profile obtained in the 300 V and 15 A krypton case, the discharge current is still 33% lower than the experimental value. The anomalous collision frequency profiles for this case (**Figure 7.3j**) suggest that the overall value of the anomalous transport is too low in this case. Indeed, the simulated anomalous collision frequency curve in this case lies below the empirical curve at all points. In the other cases, the curves intersect. We previously hypothesized that the low anomalous transport exhibited by krypton in the model of Lafleur et al. might mean that the assumption that the group velocity scales with the ion convection velocity might need to be revisited. As the discharge current predicted for krypton by the centripetal model was the same as in the baseline case, the present results do not provide much new insight into this question.

The degraded performance at high voltage observed in the centripetal model suggest that something additional needs to be added to the model to properly scale to high voltages. As discussed in **Chapter 3**, it is possible that adding a shear-suppressing factor could help steepen the ion velocity gradients in all cases, with an increasing effect at higher discharge voltages.

7.3   Summary

In this chapter, we evaluated five algebraic closure models of the Hall thruster anomalous transport. We assumed that the instability responsible for this transport saturated at the electron drift kinetic energy, and that the primary loss mechanism was convection of the wave energy out of the thruster. We investigated the effect of several additional damping mechanisms, including classical electron-ion and ion-neutral charge exchange collisions, in addition to ion Landau damping at both the plasma frequency and at cyclotron resonances. We found that these secondary damping mechanisms were insignificant compared to convection, and that as a result all of these models produced similar estimates of the thruster anode efficiency and the error in the ion velocity compare to experiment.

Despite significant differences in the ion velocity profiles between our simulations and experimental measurements obtained via laser-induced fluorescence measurements of the H9 Hall thruster, the anode efficiencies predicted by our simulations were very close to the experimental values, something which previous closure models have not yet been able to attain. In addition, the predicted anomalous collision frequency profiles shared many features in common with empirically-inferred collision frequency profiles for this thruster.

We then examined a new algebraic closure model derived by replacing the conservative gradient in the wave energy conservation equation with a non-conservative form. Through an oversight, we erroneously left in the centripetal term of the divergence operator from the conservative form of the equation. We evaluated this model against four operating

conditions and found that it produced improved results over a version of the model with a conservative gradient, in all cases except for the high voltage test case. Finally, we suggested that this centripetal term may represent a loss due to a long-wavelength radial mode and proposed an alternate form which better captures this hypothesized physical explanation.

CHAPTER 8

Conclusions and Future Work

> I want to [model anomalous transport], and there's no more time.
> Kurosawa, when he got his Oscar, when George [Lucas] and Steven
> [Spielberg] gave it to him, he said, "I'm only now beginning to see
> the possibility of what [anomalous transport modeling] could be, and
> it's too late." He was 83. At the time, I said, "What does he mean?"
> Now I know what he means.

> Martin Scorcese (kind of)

8.1   Summary of work

In this dissertation, I described our efforts to develop, calibrate and evaluate new models of anomalous electron transport in Hall thrusters.

In **Chapter 1**, I introduced the fundamentals of spacecraft propulsion. I then described electric propulsion in general, and Hall thrusters in specific. I discussed the history of these devices and the problem of anomalous cross-field electron transport which has been known to exist in Hall thrusters for over 50 years. I then overviewed some of the proposed causes of this important phenomenon.

Next, in **Chapter 2**, I gave an overview of the plasma physics relevant to Hall thrusters. I then showed how we derive the basic equations of motion for a fluid plasma and discussed the types of models we typically apply to Hall thrusters. Lastly, I described how anomalous electron transport can be incorporated in models which treat electrons as a continuum rather than kinetically.

Then, in **Chapter 3**, I introduced previous efforts to model anomalous transport. I introduced the classification of models based on the number of differential equations present, and described several different categories of models. These included semi-empirical models, models derived from first principles, and shear-suppressed models. Finally, I showed how altering the assumptions underlying these models can produce new models to test.

In **Chapter 4**, I then described our simulation approach. I introduced Hall2De, the code we used to simulate Hall thrusters throughout this dissertation. I surveyed its governing equations and boundary conditions, and presented the results of grid convergence and parameter sensitivity studies. Next, I introduced the H9 Hall thruster, the device we simulated in the remainder of the dissertation, and described the data to which we compare our simulation results. Lastly, I introduced the metrics we employed to compare our simulations to data.

After that, in **Chapter 5**, I presented the results of two studies which addressed the question of model calibration. These studies investigated the validity of the common practice of validating models of the anomalous collision frequency to empirically-inferred collision frequency profiles. In both cases, we found that a positive result using this method does not imply that a given model will perform well when implemented self-consistently into a Hall thruster code. I then suggested a few reasons for this discrepancy, including artifacts introduced by time averaging as well as lack of feedback between the physics of the closure model and the physics of the simulation. These results underscore the importance of evaluating models by implementing them directly into a Hall thruster code.

In **Chapter 6**, I followed up these results by implementing four of the models introduced in **Chapter 3** into Hall2De and evaluated their performance against data from four operating conditions of the H9 Hall thruster. I found that while some of the models could capture bulk performance trends in a few of the operating conditions, none was able to accurately predict performance in all cases. Additionally, the spatially-resolved plasma properties obtained from the simulations showed that none of the models was able to adequately matched ion velocity data obtained via laser-induced fluorescence. I then discussed these results in the context of the functional dependencies of the model, providing possible physical explanations for the models' behavior in each case. We also highlighted the role of the electron thermal energy equation in establishing the location of the ion acceleration region.

Lastly, **in Chapter 7**, I investigated a few new models which were derived by altering the underlying assumptions of a model presented in **Chapter 3**. I first examined a model in which the wave energy saturates at the electron kinetic—rather than thermal—energy density. Additionally, we tested the effects of incorporating collisional and wave-driven damping mechanisms in the closure model. We found that while this drift saturation model could adequately predict the thrust and efficiency at the single test condition, it produced a poor fit with ion velocity profiles measured using LIF. Additionally, the incorporation of these additional loss mechanisms did not improve performance. After that, I examined the performance of a model derived by modifying the form of the wave energy convection equation. In all cases except for the high voltage test case, this altered model produced better results than the model with the baseline form of the convection term. I suggested that these encouraging results may stem from enhanced radial energy losses from the wave-driven turbulence present in Hall thrusters.

Overall, I have reviewed and tested nearly a dozen different closure models of the Hall thruster anomalous transport in this dissertation. These represent a small subset of the models we tested and discarded over the course of my thesis work. While I was unable to

completely resolve the problem of anomalous transport in Hall thrusters, my results give into the proper way to evaluate transport closures in Hall thrusters, and suggest some new avenues for exploration.

8.2   Next steps

First, I have shown that closure models should not be calibrated or evaluated against steady-state or time-averaged data. Instead, models should be validated by implementing them directly into a Hall thruster code. Running the number of simulations needed to properly calibrate a model with more than two coefficients is time-consuming and computationally expensive. Additionally, I have highlighted the importance of directly evaluating model extensibility by simulating multiple operating conditions, and ideally multiple thrusters. To test more complex models at more operating conditions, it would be ideal to make use of acceleration techniques to reduce time to convergence of the individual simulations, or to reduce the number of simulations needed. In one of our conference papers, we applied a multi-fidelity surrogate optimization process[104] to attempt to speed up the calibration of empirical transport profiles. This work was not very successful, but I believe the core of the idea is good. By using lower-fidelity methods, whether 1D codes or simulations with reduced grid resolution, we may be able to speed up searches for optimal model coefficients over large parameters spaces. Such methods could also pave the way for automated, data-driven model discovery techniques. These would be analogous to the methods employed by Jorns,[3] but we would have more confidence in the performance of the models generated.

Next, while this work focused primarily on fluid models, the role of kinetic simulations should not be understated. As the maturity of these simulations increases, we will gain more knowledge about the origins and scaling of turbulence in Hall thrusters, and hopefully be able to derive insights which can be applied to closure modeling. In particular, the role

189

of anomalous transport on the conductive heat flux in Hall thrusters should be reviewed, both theoretically and computationally. Wave-driven cross-field heat flux is an important factor in many magnetized plasmas, from magnetic confinement fusion devices to the atmospheres of stars. An improved theoretical understanding of this phenomenon may help improve the match between simulations and experiment. Particle-in-cell simulations may help with this by giving us insight into how the cross-field heat flux and thermal conductivity scale with other plasma properties.

Lastly, closure models of the anomalous collision frequency like the ones I discussed should continue to be tested and evaluated. I was unable to evaluate many of the shear-suppressed transport models in time for this dissertation, but work toward that is ongoing. Looking beyond algebraic closure models, we should develop and test new multi-equation models. By explicitly solving for the transport of wave energy (or other key wave parameters), we may reduce the number of assumptions needed to capture the anomalous electron transport, and thus improve the extensibility and physical fidelity of the resulting models.

8.3 Extensibility to other plasma devices

Anomalous electron transport is a problem not just in Hall thrusters, but also in magnetized plasmas of many types in many fields of plasma science. In this thesis, we have primarily considered the problem of cross-field electron transport, which arises when the electric field and magnetic field are perpendicular to each other. In other plasma devices, this may not be the case. Indeed, in **Chapter 4**, we briefly discussed a different type of anomalous transport—namely, the problem of enhanced field-aligned resistivity in the cathode plume—where the electric and magnetic fields are approximately parallel to each other. Here, anomalous resistivity is needed to explain the observed gradients in plasma properties along the magnetic field lines. A similar problem has been observed in ECR

magnetic nozzle thrusters, where instability-induced electron transport may play a key role in moderating the field-aligned heat flux.[114]

In each plasma device where anomalous transport plays a role, it may be possible to resolve the transport in a fluid simulation using an appropriate closure model. We have already seen one such model for anomalous effects in hollow cathode plumes (the Sagdeev model, **Equation 4.10**), but other have been proposed.[86] Evaluating these models will require researchers to face challenges similar to those encountered by Hall thruster modellers. It is unlikely that a solution to the transport problem in Hall thrusters will lead to a solution of the same problem in these other devices, as the dominant transport mechanisms may differ from device to device. Nevertheless, I hope some of the methods developed in this work may be applied to evaluating and validating anomalous transport models in these other fields.

APPENDIX A

HallThruster.jl: An Open-Source One-Dimensional Fluid Hall Thruster Code

Everything should be made as simple as possible, but no simpler.

Unknown, commonly attributed to Albert Einstein

Some of this work was previously published as Ref. 115

As part of my work in this thesis, I worked with Paul Schedler, an Austrian visiting scholar from ETH Zurich, to develop a one-dimensional fluid Hall thruster code. The goal of this effort was to accelerate the calibration and testing of anomalous transport models. If we had a code which ran in seconds to minutes, it would be vastly faster to search large regions of model parameter space than if we used Hall2De, which takes hours to days to run. While hybrid-PIC codes may provide more information about the ion distribution functions, it was hoped that by using a fluid model, we could get an improved correspondance between our 1D results and the results from Hall2De. In this appendix, I will introduce the resulting open-source one-dimensional fluid code, HallThruster.jl, before discussing its ultimate utility for Hall thruster anomalous transport modeling.

A.1   Statement of need

`HallThruster.jl` was designed as a tool to more rapidly assess different models for this electron transport through parameter estimation, uncertainty quantification, and surrogate optimization. It allows users to alter many parts of the physics model for electron transport (and related processes) without modifying the source code of the package, and the one-dimensional domain only requires a few seconds to a few minutes to run a simulation, depending on the grid resolution. With that said, a central premise of `Hallthruster.jl` is that the dynamics in the Hall thruster are suitably one-dimensional that the one-dimensional code is usually sufficient for an accurate approximation of a two- or three-dimensional code when run on the scale of the entire thruster. But, with appropriate calibration, errors may be corrected by training a surrogate model correction for higher dimensional codes. Indeed, this is one of the anticipated use cases.

While one-dimensional codes are commonly used in Hall thruster research,[116,117,118,119,120] `HallThruster.jl` is unique in that, in addition to its primary goal of informing closure models, it is written from the ground up with a strong emphasis on code verification, collaboration, and extension. As an open source code, our goal is that it will allow the community to reproduce computational results obtained and to compare and contrast with parallel studies.

A.2   Physics model

Figure A.1 depicts the one-dimensional simulation domain for `HallThruster.jl`. In this model, we assume quasineutrality in the entire domain and treat all species—neutrals, ions, and electrons—as fluids. Neutrals are assumed to have constant temperature and velocity and are solved using the continuity equation, while ions are assumed isothermal and are modeled using the isothermal Euler equations. Electron inertia is neglected, so
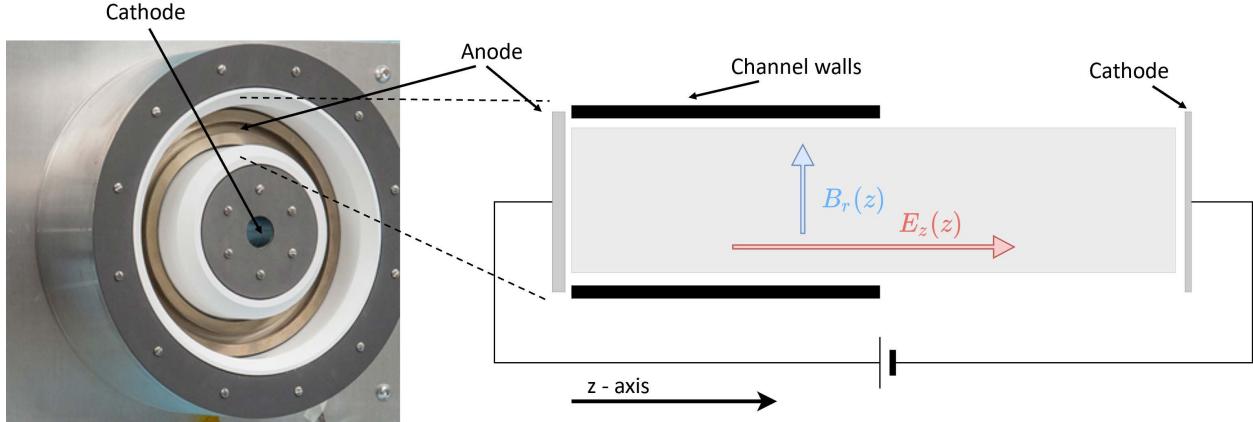
Figure A.1: 1D simulation domain of HallThruster.jl

the electron momentum equation reduces to a generalized Ohm's law. As discussed in the preceding section, electrons drift primarily in the $\hat{\theta} = \hat{z} \times \hat{r}$ direction but have a small axial drift enabled by collisions. This cross-field electron current observed in experiment is much higher than that which would be predicted from classical collisions alone, so we apply an additional "anomalous" collision frequency in the Ohm's law to enable the user to better tune the simulation to match experiment. Different forms for this collision frequency can be specified to explore different closure models. We combine this Ohm's law with current conservation to obtain a second-order partial differential equation for the electrostatic potential. Lastly, we solve a partial differential equation for the transport of electron internal energy, taking into account losses due to ionization, excitation, and wall losses. A detailed listing of the equations employed in the physics model is available in the code documentation.

While many aspects of Hall thruster physics are well-represented by a one-dimensional model, certain important features of real Hall thruster discharges will not be captured in such a simplified code. Among other effects, the most important of these missing features are plume divergence, impacts of magnetic field curvature, and effects of radial plasma non-uniformity on wall losses.

A.3 Numerics

The finite volume method is applied to discretize the ion and neutral equations, transforming them into a system of ordinary differential equations. These are integrated in time using a second order strong stability preserving Runge-Kutta (SSPRK22) scheme. Several choice of numerical flux are available, including the Rusanov, global Lax-Friedrichs, and Hartford-Lax-van-Leer-Einfeld (HLLE) fluxes. The elliptic equation for the potential is transformed by finite differences into a tridiagonal linear system and solved using Thomas' algorithm.[121] The electron energy equation is solved semi-implicitly to ease timestep restrictions and is discretized in space using first-order upwind differences. We compute reaction and collision rate coefficients using look-up tables. The method of manufactured solutions is used to verify that we have correctly discretized the PDEs and obtain the design order of accuracy. We are aided in this by the Symbolics.jl package,[122] which automatically computes the needed source terms.

A.4 Functionality

`HallThruster.jl` provides extensive options to allow the user to customize their simulation. These include

1. Custom (one-dimensional) thruster geometry

2. Anode and wall sheath model options

3. Multiple propellants: the code has built-in capabilities for krypton and xenon. Additional propellants can be added by introducing the appropriate reaction rates and mass

4. Restarts

5. Custom anomalous transport models

6. User-specified source terms

7. Custom additional collisions and reactions

Detailed documentation about all of these features and mode is available at the Github repository.

A.5 Validation

We show in **Figures A.2-A.4** comparisons of our results to cases 1-3 of the LANDMARK benchmark,[123] respectively. This benchmark provides reference one-dimensional Hall thruster simulations to which code authors can compare their own simulations. This benchmark has three cases, distinguished by the amount of electron energy lost to the walls. We show all results with 1024 cells, time averaged after 2 ms of simulation time. We compare three included fluxes: HLLE, local Lax-Friedrichs / Rusanov (LLF), and global Lax-Friedrichs (GLF) to the three sub-cases of the LANDMARK benchmark - two fluid models with adjustable viscosity parameter δ and a hybid-particle-in-cell case. Note that the benchmark does not include plots of the ion and electron velocities, nor does it include details about the numerics of the example case. This makes a quantitative comparison difficult. Nevertheless, we see good qualitative agreement. Differences between our model and the LANDMARK fluid model may be related to differences in boundary condition implementation, order of accuracy of method, or choice of numerical flux.

A.6 Use in closure modeling

HallThruster.jl is able to incorporate many of the closure models introduced throughout this dissertation and obtain results qualitatively similar to those produced using Hall2De. Additionally, we were able to use the code to tune the coefficients of the PI controller we

employed to accelerated model calibration in **Chapters 6** and **7**. Unfortunately, we found through extensive testing that the correspondance between the 1D results from our code and the 2D results from Hall2De is not as straightforward as we hoped.

This was exemplified by our attempt to use this code as part of a multi-fidelity framework in 2022.[104] The goal of this effort was to use the results of the 1D code to suggest new stationary spatially-varying anomalous collision frequency profiles for use in Hall2De. The hypothesis underlying this is that we could treat the 1D code as a "surrogate model" of the Hall2De, and use only a few evaluations of Hall2De to "correct" the predictions of HallThruster.jl. In this case, Using HallThruster.jl did not measurably speed up the calibration procedure in this case. While it is possible that improved surrogate modeling efforts could fix this, it is also possible that radial effects are more significant than we appreciated.

Ultimately, while we have not been able to use the code in the manner we originally intended, it has still been a valuable tool. Its short runtime (10-60 seconds for a typical simulation) makes it possible to rapidly assess the effects of changing physics assumptions on Hall thruster simulations. Additionally, while the results to not always quantitatively agree with Hall2De, they often match Hall2De well enough that we can use the 1D code to "pre-screen" closure models before we implement the model in the high-fidelity code. Lastly, the code is to our knowledge the only open-source Hall thruster code available. As the code becomes more mature, it is hoped that other researchers will use this code instead of writing their own, reducing reduplication of effort. As these researchers contribute back to the code, HallThruster.jl's abilities will grow and the model's fidelity will continue to increase.
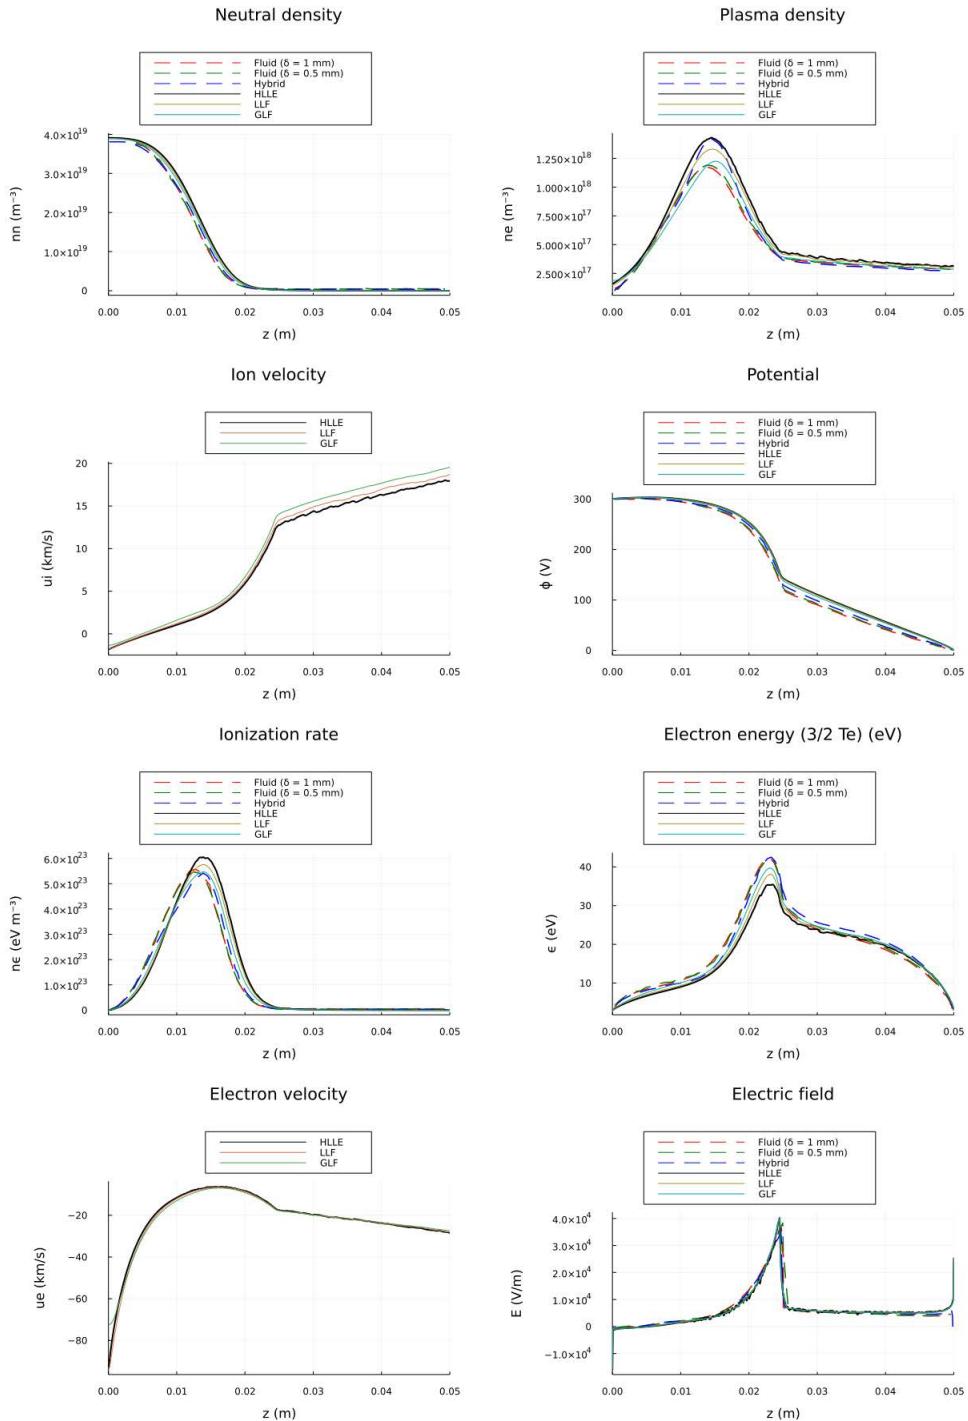
Figure A.2: Comparison of HallThruster.jl results with different numerical fluxes to LANDMARK benchmark case 1.
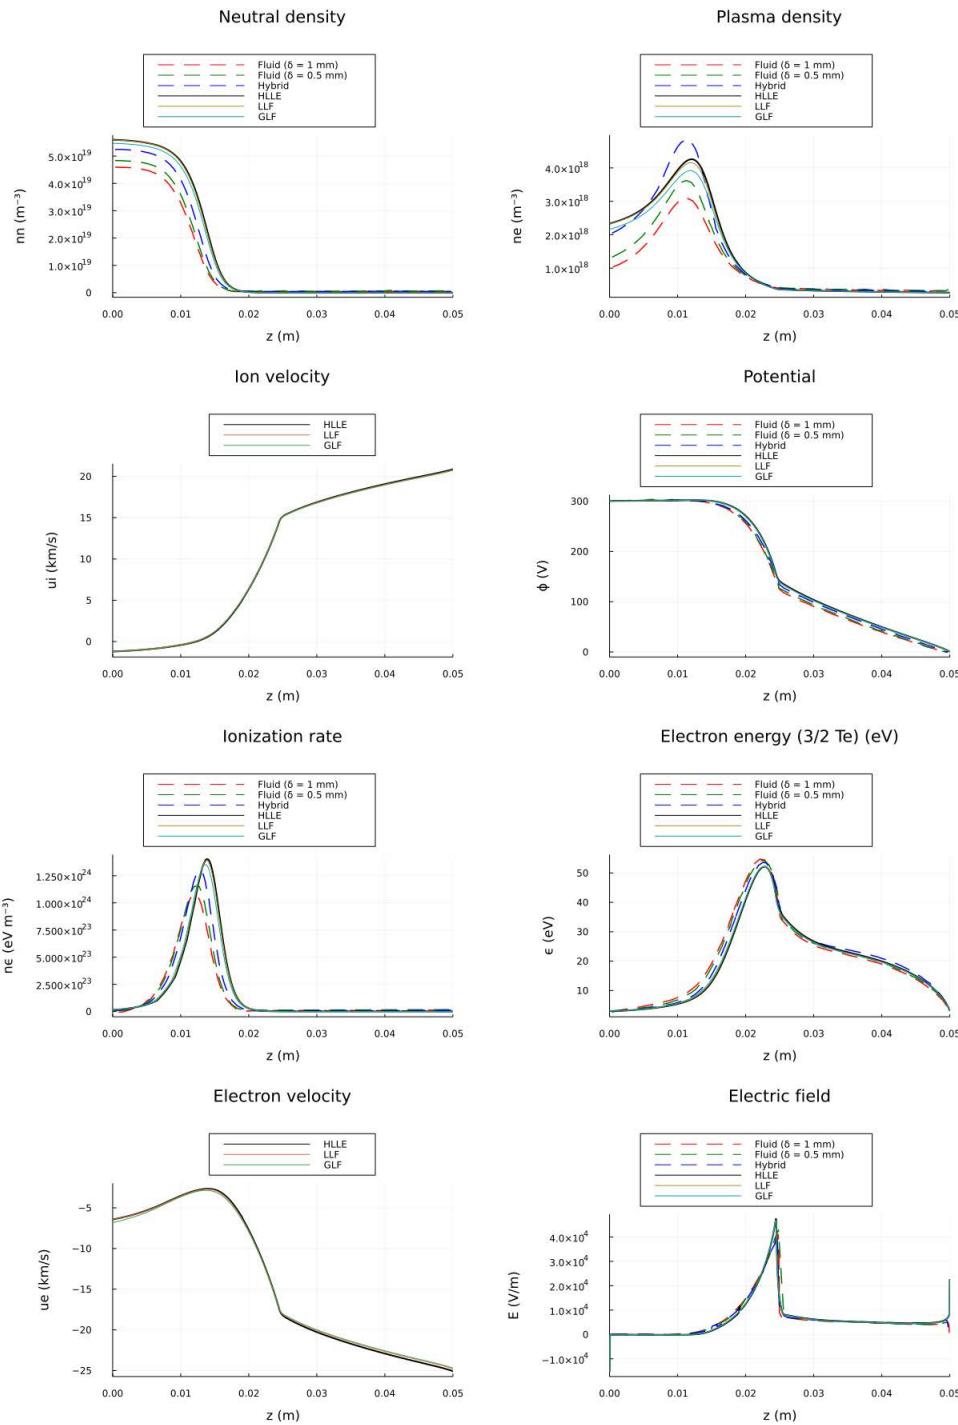
Figure A.3: Comparison of HallThruster.jl results with different numerical fluxes to LANDMARK benchmark case 2.
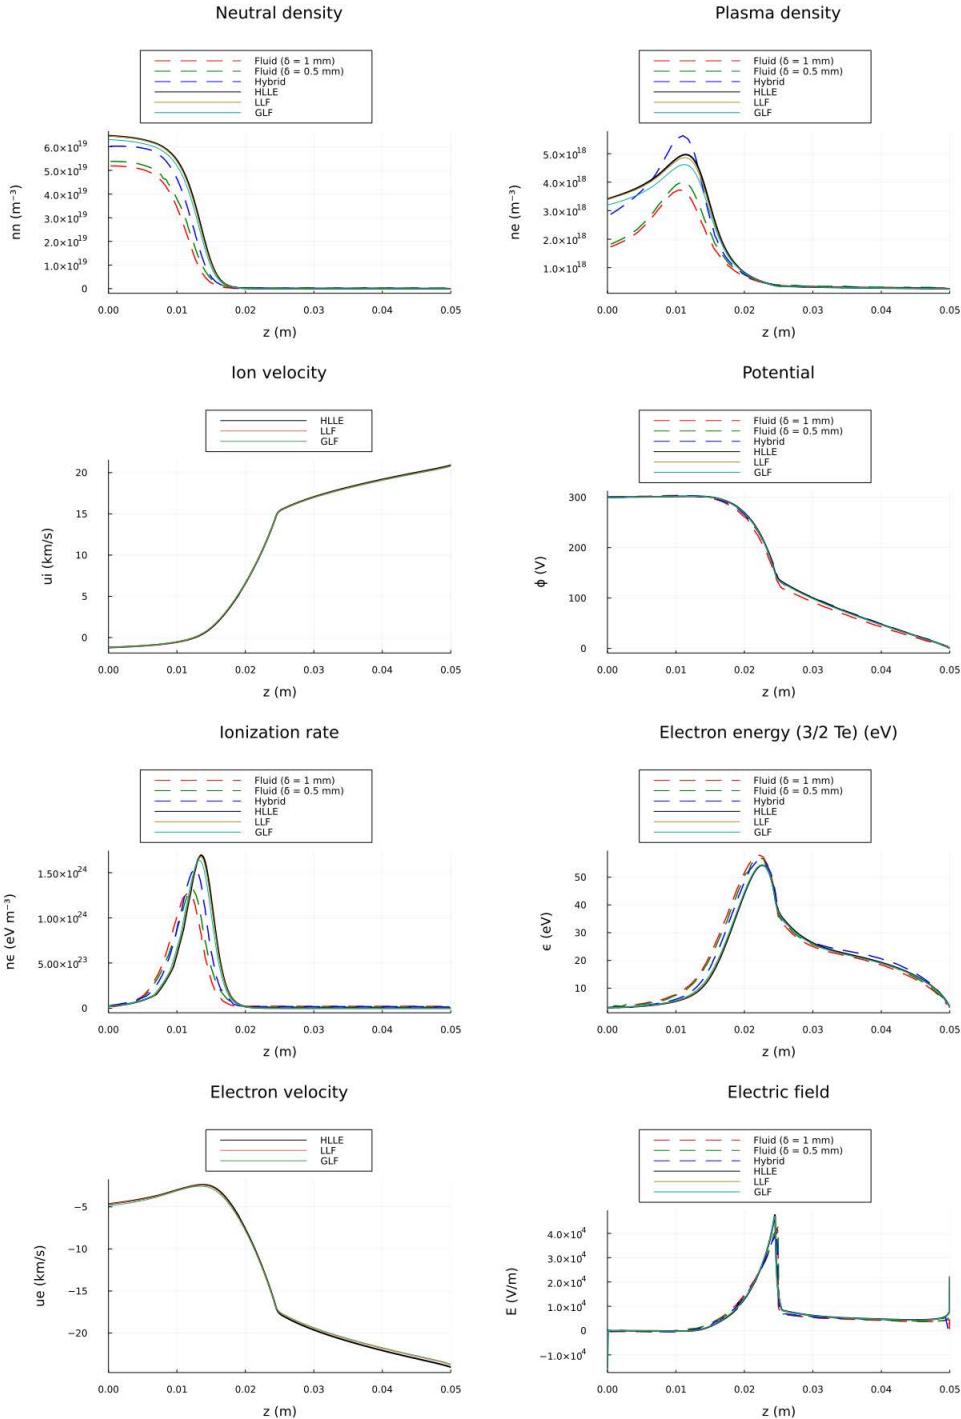
Figure A.4: Comparison of HallThruster.jl results with different numerical fluxes to LANDMARK benchmark case 3.

BIBLIOGRAPHY

[1] L. L. Su and B. A. Jorns. Performance comparison of a 9-kw magnetically shielded hall thruster operating on xenon and krypton. *Journal of Applied Physics*, 130: 163306, 10 2021. ISSN 0021-8979. doi:10.1063/5.0066849. URL https://aip.scitation.org/doi/abs/10.1063/5.0066849.

[2] T. Lafleur, S. D. Baalrud, and P. Chabert. Theory for the anomalous electron transport in hall effect thrusters. i. insights from particle-in-cell simulations. *Physics of Plasmas*, 23, 5 2016. ISSN 10897674. doi:10.1063/1.4948495.

[3] Benjamin Jorns. Predictive, data-driven model for the anomalous electron collision frequency in a hall effect thruster. *Plasma Sources Science and Technology*, 27, 10 2018. ISSN 13616595. doi:10.1088/1361-6595/aae472.

[4] John D. Clark. *Ignition!: an informal history of liquid rocket propellants*. Rutgers University Press classics. Rutgers University Press New Brunswick, New Jersey, New Brunswick, New Jersey, 1972. ISBN 9780813507255; 0813507251; 9780813595832; 0813595835; 0813599172; 9780813599175.

[5] Dan M. Goebel and Ira Katz. *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*. Jet Propulsion Laboratory, 2008. ISBN 9780470429273. doi:10.1002/9780470436448.

[6] John Brophy. *The Dawn Ion Propulsion System*, pages 251–261. Springer New York, New York, NY, 2012. ISBN 978-1-4614-4903-4. doi:10.1007/978-1-4614-4903-4_11. URL https://doi.org/10.1007/978-1-4614-4903-4_11.

[7] Joakim Kugelberg, Per Bodin, Staffan Persson, and Peter Rathsman. Accommodating electric propulsion on smart-1. *Acta Astronautica*, 55(2):121–130, 2004. ISSN 0094-5765. doi:https://doi.org/10.1016/j.actaastro.2004.04.003. URL https://www.sciencedirect.com/science/article/pii/S0094576504001328.

[8] Hitoshi Kuninaka, Kazutaka Nishiyama, Ikko Funaki, Tetsuya Yamada, Yukio Shimizu, and Jun'ichiro Kawaguchi. *Powered Flight of HAYABUSA in Deep Space*. AIAA, 2006. doi:10.2514/6.2006-4318. URL https://arc.aiaa.org/doi/abs/10.2514/6.2006-4318.

[9] Jerry Jackson, May Allen, Roger Myers, Erich Soendker, Benjamin Welander, Artie Tolentino, Sam Hablitzel, Chyrl Yeatts, Chris Sheehan, Joseph Cardin, John Steven Snyder, Richard R Hofer, Todd Tofil, and Dan Herman. 13kw advanced electric propulsion flight system development and qualification.

The 36th International Electric Propulsion Conference, page 18, 2019. URL https://iepc2017.org/sites/default/files/speaker-papers/iepc-2017-223_aeps_development_and_qualification.pdf.

[10] John Steven Snyder, Vernon Chaplin, Dan M. Goebel, Richard R. Hofer, Alejandro Lopez Ortega, Ioannis G. Mikellides, Taylor Kerl, Giovanni Lenguito, Faraz Aghazadeh, and Ian Johnson. Electric propulsion for the psyche mission: Development activities and status. *AIAA Propulsion and Energy 2020 Forum*, pages 1–15, 2020. doi:10.2514/6.2020-3607. URL https://arc.aiaa.org/doi/10.2514/6.2020-3607.

[11] Tom Percy, Melissa McGuire, and Tara Polsgrove. Combining solar electric propulsion and chemical propulsion for crewed missions to mars. In *2015 IEEE Aerospace Conference*, pages 1–10, 2015. doi:10.1109/AERO.2015.7119289.

[12] Sean K. Patrick and Alfred Lynam. *Optimal SEP Trajectories from Earth to Jupiter with Triple Flyby Capture*. AIAA, 2014. doi:10.2514/6.2014-4218. URL https://arc.aiaa.org/doi/abs/10.2514/6.2014-4218.

[13] A. I. Morozov. The conceptual development of stationary plasma thrusters. *Plasma Physics Reports*, 29(3):235–250, Mar 2003. ISSN 1562-6938. doi:10.1134/1.1561119. URL https://doi.org/10.1134/1.1561119.

[14] Jean Pierre Boeuf. Tutorial: Physics and modeling of hall thrusters. *Journal of Applied Physics*, 121, 1 2017. ISSN 10897550. doi:10.1063/1.4972269.

[15] Colleen M. Marrese-Reading, Robert Frisbee, Anita Sengupta, Mark A. Cappelli, Sergey Tverdoklebov, Sasha Semenkin, and Iain Boyd. Very high Isp thruster with anode layer (VHITAL) : an overview. Technical report, 2004. URL https://hdl.handle.net/2014/40499.

[16] Ioannis G. Mikellides and Ira Katz. Numerical simulations of hall-effect plasma accelerators on a magnetic-field-aligned mesh. *Physical Review E*, 86, 10 2012. ISSN 15393755. doi:10.1103/PhysRevE.86.046703.

[17] Leanne L. Su and Benjamin Jorns. *Performance at High Current Densities of a Magnetically-Shielded Hall Thruster*. AIAA, 2021. doi:10.2514/6.2021-3405. URL https://arc.aiaa.org/doi/abs/10.2514/6.2021-3405.

[18] Leanne L. Su, Parker J. Roberts, Tate Gill, William Hurley, Thomas A. Marks, Christopher L. Sercel, Madison Allen, Collin B. Whittaker, Matthew Byrne, Zachariah Brown, Eric Viges, and Benjamin Jorns. *Operation and Performance of a Magnetically Shielded Hall Thruster at Ultrahigh Current Densities on Xenon and*

Krypton. AIAA, 2023. doi:10.2514/6.2023-0842. URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-0842.

[19] Jacob Simmonds, Yevgeny Raitses, and Andrei Smolyakov. A theoretical thrust density limit for hall thrusters. *Journal of Electric Propulsion*, 2(1):12, Mar 2023. ISSN 2731-4596. doi:10.1007/s44205-023-00048-9. URL https://doi.org/10.1007/s44205-023-00048-9.

[20] Jacob Simmonds and Yevgeny Raitses. Ion acceleration in a wall-less Hall thruster. *Journal of Applied Physics*, 130(9):093302, 09 2021. ISSN 0021-8979. doi:10.1063/5.0062607. URL https://doi.org/10.1063/5.0062607.

[21] Yevgeny Raitses, Enrique Merino, and Nathaniel J. Fisch. Cylindrical Hall thrusters with permanent magnets. *Journal of Applied Physics*, 108(9):093307, 11 2010. ISSN 0021-8979. doi:10.1063/1.3499694. URL https://doi.org/10.1063/1.3499694.

[22] Tommaso Andreussi, Eugenio Ferrato, and Vittorio Giannetti. A review of air-breathing electric propulsion: from mission studies to technology verification. *Journal of Electric Propulsion*, 1(1):31, Dec 2022. ISSN 2731-4596. doi:10.1007/s44205-022-00024-9. URL https://doi.org/10.1007/s44205-022-00024-9.

[23] John Brophy, James Polk, and John Dankanich. *Lifetime Qualification Standard for Electric Thrusters*. AIAA, 2009. doi:10.2514/6.2009-5095. URL https://arc.aiaa.org/doi/abs/10.2514/6.2009-5095.

[24] Harold R. Kaufman. Explanation of bohm diffusion. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 8:107, 1 1990. ISSN 0734211X. doi:10.1116/1.584855.

[25] A. I. Morozov and V. V. Savelyev. *Fundamentals of Stationary Plasma Thruster Theory*, pages 203–391. Springer, Boston, MA, 2000. doi:10.1007/978-1-4615-4309-1_2. URL https://link.springer.com/chapter/10.1007/978-1-4615-4309-1_2.

[26] S. Barral, K. Makowski, Z. Peradzyński, N. Gascon, and M. Dudeck. Wall material effects in stationary plasma thrusters. ii. near-wall and in-wall conductivity. *Physics of Plasmas*, 10:4137–4152, 10 2003. ISSN 1070664X. doi:10.1063/1.1611881. URL http://aip.scitation.org/doi/10.1063/1.1611881.

[27] Richard Robert Hofer, Hani Kamhawi, Daniel A. Herman, James E. Polk, John Steven Snyder, Ioannis G. Mikellides, Wensheng Huang, James L. Myers, John T. Yim, George J. Williams, Alejandro Lopez Ortega, Benjamin A. Jorns, Michael J. Sekerak, Christopher Griffith, Rohit Shastry, Thomas W. Haag, Timothy

Verhey, Bradley Gilliam, Ira Katz, Dan M. Goebel, John R. Anderson, James H. Gilland, and Lauren Clayman. Development approach and status of the 12.5 kw hermes hall thruster for the solar electric propulsion technology demonstration mission. *34th International Electric Propulsion Conference*, 2015.

[28] Sarah E Cusson, Richard R Hofer, Robert B Lobbia, Benjamin A Jorns, and Alec D Gallimore. Performance of the h9 magnetically shielded hall thrusters. In *2017 International Electric Propulsion Conference*, 2017.

[29] L. Garrigues, G. J.M. Hagelaar, C. Boniface, and J. P. Boeuf. Anomalous conductivity and secondary electron emission in hall effect thrusters. *Journal of Applied Physics*, 100, 2006. ISSN 00218979. doi:10.1063/1.2401773.

[30] Richard Hofer, Ira Katz, Dan Goebel, Kristina Jameson, Regina Sullivan, Lee Johnson, and Ioannis Mikellides. *Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations*. AIAA, 2012. doi:10.2514/6.2008-4924. URL https://arc.aiaa.org/doi/abs/10.2514/6.2008-4924.

[31] Ethan T. Dale and Benjamin A. Jorns. Non-invasive time-resolved measurements of anomalous collision frequency in a hall thruster. *Physics of Plasmas*, 26, 2019. ISSN 10897674. doi:10.1063/1.5077008. URL http://dx.doi.org/10.1063/1.5077008.

[32] Yu V. Esipchuk and G. N. Tilinin. Drift instability in a hall-current plasma accelerator. *Sov Phys Tech Phys*, 21:417–423, 1976.

[33] S. Tsikata, N. Lemoine, V. Pisarev, and D. M. Gŕsillon. Dispersion relations of electron density fluctuations in a hall thruster plasma, observed by collective light scattering. *Physics of Plasmas*, 16:33506, 3 2009. ISSN 1070664X. doi:10.1063/1.3093261. URL https://doi.org/10.1063/1.3093261.

[34] S. Tsikata, A. Héron, and C. Honoré. Hall thruster microturbulence under conditions of modified electron wall emission. *Physics of Plasmas*, 24:053519, 5 2017. ISSN 10897674. doi:10.1063/1.4984255. URL https://aip.scitation.org/doi/abs/10.1063/1.4984255.

[35] Benjamin Vincent, Sedina Tsikata, Stéphane Mazouffre, Tiberiu Minea, and Jérôme Fils. A compact new incoherent thomson scattering diagnostic for low-temperature plasma studies. *Plasma Sources Science and Technology*, 27:055002, 5 2018. ISSN 0963-0252. doi:10.1088/1361-6595/AABD13. URL https://iopscience.iop.org/article/10.1088/1361-6595/aabd13https://iopscience.iop.org/article/10.1088/1361-6595/aabd13/meta.

[36] Zachariah A. Brown and Benjamin A. Jorns. Spatial evolution of small wavelength fluctuations in a hall thruster. *Physics of Plasmas*, 26:113504, 11 2019. ISSN 10897674. doi:10.1063/1.5116708. URL https://doi.org/10.1063/1.5116708.

[37] Zachariah A. Brown and Benjamin A. Jorns. Growth and saturation of the electron drift instability in a crossed field plasma. *Phys. Rev. Lett.*, 130:115101, Mar 2023. doi:10.1103/PhysRevLett.130.115101. URL https://link.aps.org/doi/10.1103/PhysRevLett.130.115101.

[38] J. Cavalier, N. Lemoine, G. Bonhomme, S. Tsikata, C. Honoré, and D. Grésillon. Hall thruster plasma fluctuations identified as the e×b electron drift instability: Modeling and fitting on experimental data. *Physics of Plasmas*, 20:082107, 8 2013. ISSN 1070-664X. doi:10.1063/1.4817743. URL https://aip.scitation.org/doi/abs/10.1063/1.4817743.

[39] W. Villafana, B. Cuenot, and O. Vermorel. 3d particle-in-cell study of the electron drift instability in a hall thruster using unstructured grids. *Physics of Plasmas*, 30, 3 2023. ISSN 10897674. doi:10.1063/5.0133963.

[40] Giancarlo Alfonsi. Reynolds-Averaged Navier–Stokes Equations for Turbulence Modeling. *Applied Mechanics Reviews*, 62(4):040802, 06 2009. ISSN 0003-6900. doi:10.1115/1.3124648. URL https://doi.org/10.1115/1.3124648.

[41] Michelle K. Scharfe, Cliff A. Thomas, David B. Scharfe, Nicolas Gascon, Mark A. Cappelli, and Eduardo Fernandez. Shear-based model for electron transport in hybrid hall thruster simulations. *IEEE Transactions on Plasma Science*, 36:2058–2068, 2008. ISSN 00933813. doi:10.1109/TPS.2008.2004364.

[42] Mark A. Cappelli, Christopher V. Young, Eunsun Cha, and Eduardo Fernandez. A zero-equation turbulence model for two-dimensional hybrid hall thruster simulations. *Physics of Plasmas*, 22, 11 2015. ISSN 10897674. doi:10.1063/1.4935891.

[43] P. L. Bhatnagar, E. P. Gross, and M. Krook. A model for collision processes in gases. i. small amplitude processes in charged and neutral one-component systems. *Phys. Rev.*, 94:511–525, May 1954. doi:10.1103/PhysRev.94.511. URL https://link.aps.org/doi/10.1103/PhysRev.94.511.

[44] L. D. Landau. Kinetic equation for the case of coulomb interaction. *Phys. Z. Sowjetunion*, 10:154–164, 1936.

[45] S. Chapman and T. G. Cowling. *The Mathematical Theory of Non-uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction and Diffusion in Gases*. Cambridge University Press, 1970. ISBN 052140844X.

[46] Natalia MacDonald-Tenenbaum, Quinn Pratt, Michael Nakles, Nickolas Pilgram, Michael Holmes, and William Hargus. Background pressure effects on ion velocity distributions in an spt-100 hall thruster. *Journal of Propulsion and Power*, 35:403–412, 2019. doi:10.2514/1.B37133. URL https://doi.org/10.2514/1.B37133.

[47] John Michael Fife. *Hybrid-PIC modeling and electrostatic probe survey of Hall thrusters*. PhD thesis, Massachusetts Institude of Technology, 1998.

[48] Adnan R. Mansour, Christopher L. Osgood, Rupali Sahu, and Kentaro Hara. Development of a one-dimensional full-fluid moment model for hall effect thrusters. *AIAA Propulsion and Energy 2020 Forum*, pages 1–8, 2020. doi:10.2514/6.2020-3623. URL https://arc.aiaa.org/doi/abs/10.2514/6.2020-3623.

[49] Alejandro Lopez Ortega, Ioannis G. Mikellides, Vernon H. Chaplin, Wensheng Huang, and Jason D. Frieman. Anomalous ion heating and pole erosion in the 12.5-kw hall effect rocket with magnetic shielding (hermes). pages 1–23. American Institute of Aeronautics and Astronautics Inc, AIAA, 2020. ISBN 9781624106026. doi:10.2514/6.2020-3620. URL http://arc.aiaa.org.

[50] Kristi De Grys, Alex Mathers, Ben Welander, and Vadim Khayms. Demonstration of 10,400 hours of operation on a 4.5 kw qualification model hall thruster. 2010. doi:10.2514/6.2010-6698. URL http://arc.aiaa.org.

[51] John M Sankovic, John A Hamley, Thomas W Haaa, and Thomas W Haagt. Performance evaluation of the russian spt-100 thruster at nasa lerc. page 094, 1993.

[52] Jonathan A. Walker, Jason D. Frieman, Mitchell L.R. Walker, Vadim Khayms, David King, and Peter Y. Peterson. Electrical facility effects on hall-effect-thruster cathode coupling: Discharge oscillations and facility coupling. *Journal of Propulsion and Power*, 32:844–855, 1 2016. ISSN 15333876. doi:10.2514/1.B35835/ASSET/IMAGES/LARGE/FIGURE14.JPEG. URL https://arc.aiaa.org/doi/10.2514/1.B35835.

[53] Leanne L Su, Thomas A Marks, and Jorns. Investigation into the efficiency gap between krypton and xenon operation on a magnetically shielded hall thruster. *International Electric Propulsion Conference*, 2022.

[54] Ioannis G. Mikellides, Ira Katz, Richard R. Hofer, and Dan M. Goebel. Magnetic shielding of a laboratory hall thruster. i. theory and validation. *Journal of Applied Physics*, 115:043303, 1 2014. ISSN 10897550. doi:10.1063/1.4862313. URL http://aip.scitation.org/doi/10.1063/1.4862313.

[55] Alejandro Lopez Ortega, Ioannis G Mikellides, Vernon H Chaplin, Wensheng Huang, and Jason D Frieman. Anomalous ion heating and pole erosion in the 12.5-kw hall effect rocket with magnetic shielding (hermes). *AIAA Propulsion and Energy Forum*, pages 125–109, 2020. doi:10.2514/6.2020-3620. URL http://arc.aiaa.org.

[56] Michelle K. Scharfe, Nicolas Gascon, Mark A. Cappelli, and Eduardo Fernandez. Comparison of hybrid hall thruster model to experimental measurements. *Physics of Plasmas*, 13, 2006. ISSN 1070664X. doi:10.1063/1.2336186.

[57] Justin William Koo. Hybrid pic-mcc computational modeling of hall thrusters. *PhDT*, 2005. URL https://ui.adsabs.harvard.edu/abs/2005PhDT........35K/abstract.

[58] Kentaro Hara. Non-oscillatory quasineutral fluid model of cross-field discharge plasmas. *Physics of Plasmas*, 25(12):123508, 12 2018. ISSN 1070-664X. doi:10.1063/1.5055750. URL https://doi.org/10.1063/1.5055750.

[59] Lewis F Richardson. *Weather prediction by numerical process*. University Press, 1922.

[60] Benjamin A. Jorns. Two equation closure model for plasma turbulence in a hall effect thruster. *36th International Electric Propulsion Conference*, pages 1–12, 2019.

[61] Joseph Plemelj. *Problems in the sense of Riemann and Klein*. Interscience Publishers, 1964.

[62] Daniel L Brown, C William Larson, Brian E Beal, and Alec D Gallimore. Methodology and historical perspectiveof a hall thruster efficiency analysis. *Journal of Propulsion and Power*, 25, 2020. doi:10.2514/1.38092.

[63] T. Lafleur, S. D. Baalrud, and P. Chabert. Theory for the anomalous electron transport in hall effect thrusters. ii. kinetic model. *Physics of Plasmas*, 23:11101, 5 2016. ISSN 10897674. doi:10.1063/1.4948496. URL http://dx.doi.org/10.1063/1.4948496].

[64] Ira Katz, Alejandro Lopez Ortega, Benjamin A. Jorns, and Ioannis G. Mikellides. Growth and saturation of ion acoustic waves in hall thrusters. 7 2016. URL https://dataverse.jpl.nasa.gov/dataset.xhtml?persistentId=hdl:2014/46151.

[65] W. Manheimer and J. P. Boris. Marginal stability analysis. *Comments Plasma Phys Controlled Fusion*, 3(1):15–24, 1977. URL http://inis.iaea.org/search/search.aspx?orig_q=RN:08336117.

[66] P. W. Terry. Suppression of turbulence and transport by sheared flow. *Rev. Mod. Phys.*, 72:109–165, Jan 2000. doi:10.1103/RevModPhys.72.109. URL https://link.aps.org/doi/10.1103/RevModPhys.72.109.

[67] Eunsun Cha, David B. Scharfe, Michelle K. Scharfe, and Eduardo Ramos Fernandez. Hybrid simulations of hall thrusters operating on various propellants. 2009. URL https://api.semanticscholar.org/CorpusID:99272674.

[68] R. Chodura. A hybrid fluid-particle model of ion heating in high-mach-number shock waves. *Nuclear Fusion*, 15(1):55, feb 1975. doi:10.1088/0029-5515/15/1/008. URL https://dx.doi.org/10.1088/0029-5515/15/1/008.

[69] A. G. Sgro and C. W. Nielson. Hybrid model studies of ion dynamics and magnetic field diffusion during pinch implosions. *The Physics of Fluids*, 19(1):126–133, 01 1976. ISSN 0031-9171. doi:10.1063/1.861309. URL https://doi.org/10.1063/1.861309.

[70] R. D. Milroy and J. U. Brackbill. Numerical studies of a field-reversed theta-pinch plasma. *The Physics of Fluids*, 25(5):775–783, 05 1982. ISSN 0031-9171. doi:10.1063/1.863832. URL https://doi.org/10.1063/1.863832.

[71] M. E. Kayama. Resistivity in the dynamic current sheath of a field reversed configuration. *Physics of Plasmas*, 19(3):032511, 03 2012. ISSN 1070-664X. doi:10.1063/1.3698405. URL https://doi.org/10.1063/1.3698405.

[72] F. I. Parra, E. Ahedo, J. M. Fife, and M. Martínez-Sánchez. A two-dimensional hybrid model of the Hall thruster discharge. *Journal of Applied Physics*, 100(2):023304, 07 2006. ISSN 0021-8979. doi:10.1063/1.2219165. URL https://doi.org/10.1063/1.2219165.

[73] G. J.M. Hagelaar, J. Bareilles, L. Garrigues, and J. P. Boeuf. Two-dimensional model of a stationary plasma thruster. *Journal of Applied Physics*, 91:5592–5598, 5 2002. ISSN 00218979. doi:10.1063/1.1465125. URL https://doi.org/10.1063/1.1465125.

[74] Adrian Domínguez-Vázquez. *Axisymmetric simulation codes for Hall effect thrusters and plasma plumes*. PhD thesis, Universidad Carlos III de Madrid, 2019.

[75] Ira Katz and Ioannis G. Mikellides. Neutral gas free molecular flow algorithm including ionization and walls for use in plasma simulations. *Journal of Computational Physics*, 230:1454–1464, 2 2011. ISSN 0021-9991. doi:10.1016/J.JCP.2010.11.013.

[76] J. Perales-Díaz, A. Domínguez-Vázquez, P. Fajardo, E. Ahedo, F. Faraji, M. Reza, and T. Andreussi. Hybrid plasma simulations of a magnetically shielded hall thruster. *Journal of Applied Physics*, 131(10):103302, 03 2022. ISSN 0021-8979. doi:10.1063/5.0065220. URL https://doi.org/10.1063/5.0065220.

[77] G D Hobbs and J A Wesson. Heat flow through a langmuir sheath in the presence of electron emission. *Plasma Physics*, 9(1):85, jan 1967. doi:10.1088/0032-1028/9/1/410. URL https://dx.doi.org/10.1088/0032-1028/9/1/410.

[78] Marcel P Georgin, Benjamin A Jorns, and Alec D Gallimore. Transient non-classical transport in the hollow cathode plume i: measurements of time-varying electron collision frequency. *Plasma Sources Sci. Technol.*, 29:13, 2020. doi:10.1088/1361-6595/abb0ce. URL https://doi.org/10.1088/1361-6595/abb0ce.

[79] Ioannis G Mikellides, Ira Katz, Dan M Goebel, and Kristina K Jameson. Evidence of nonclassical plasma transport in hollow cathodes for electric propulsion. *J. Appl. Phys*, 101:63301, 2007. doi:10.1063/1.2710763. URL https://doi.org/10.1063/1.2710763.

[80] Benjamin A. Jorns, Ioannis G. Mikellides, and Dan M. Goebel. Ion acoustic turbulence in a 100-a lab6 hollow cathode. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 90, 12 2014. ISSN 15502376. doi:10.1103/PhysRevE.90.063106.

[81] Benjamin A. Jorns, Sarah E. Cusson, Zachariah Brown, and Ethan Dale. Non-classical electron transport in the cathode plume of a hall effect thruster. *Physics of Plasmas*, 27:22311, 2020. ISSN 10897674. doi:10.1063/1.5130680. URL https://doi.org/10.1063/1.5130680.

[82] Alejandro Lopez Ortega and Ioannis G. Mikellides. The importance of the cathode plume and its interactions with the ion beam in numerical simulations of hall thrusters. *Physics of Plasmas*, 23:1411, 4 2016. ISSN 10897674. doi:10.1063/1.4947554. URL https://doi.org/10.1063/1.4947554.

[83] Roald Z. Sagdeev. The 1976 oppenheimer lectures: Critical problems in plasma astrophysics. i. turbulence and nonlinear waves. *Reviews of Modern Physics*, 1979. ISSN 00346861. doi:10.1103/RevModPhys.51.1.

[84] V. N. Tsytovich and D. T. Haar. *Lectures on non-linear plasma kinetics*. Springer, 1995.

[85] L. C. Woods. *Principles of Magnetoplasma dynamics*. 1987.

[86] Marcel P Georgin and Benjamin A Jorns. Transient non-classical transport in the hollow cathode plume ii: evaluation of models for the anomalous collision frequency. *Plasma Sources Sci. Technol*, 29, 2020. doi:10.1088/1361-6595/abb0cf. URL https://doi.org/10.1088/1361-6595/abb0cf.

[87] Richard R Hofer, Sarah E Cusson, Robert B Lobbia, and Alec D Gallimore. The h9 magnetically shielded hall thruster. 2017.

[88] W. A. Hargus and M. A. Cappelli. Laser-induced fluorescence measurements of velocity within a hall discharge. *Applied Physics B 2001 72:8*, 72:961–969, 3 2014. ISSN 1432-0649. doi:10.1007/S003400100589. URL https://link.springer.com/article/10.1007/s003400100589.

[89] Benjamin A. Jorns, Dan M. Goebel, and Richard R. Hofer. Plasma perturbations in high-speed probing of hall thruster discharge chambers: Quantification and mitigation. *51st AIAA/SAE/ASEE Joint Propulsion Conference*, 2015. doi:10.2514/6.2015-4006. URL https://arc.aiaa.org/doi/10.2514/6.2015-4006.

[90] Thomas A Marks, Alejandro Lopez Ortega, Ioannis G Mikellides, and Benjamin Jorns. Self-consistent implementation of a zero-equation transport model into a predictive model for a hall effect thruster. 2021. doi:10.2514/6.2021-3424. URL https://arc.aiaa.org/doi/abs/10.2514/6.2021-3424.

[91] Frank Hutter, Hoiger Hoos, and Kevin Leyton-Brown. An efficient approach for assessing hyperparameter importance. volume 2, pages 1130–1144, 2014. ISBN 9781634393973.

[92] Matthias Feurer and Frank Hutter. *Hyperparameter Optimization*, pages 3–33. Springer, 2019. ISBN 978-3-030-05317-8. doi:10.1007/978-3-030-05318-5_1. URL http://www.springer.com/series/15602.

[93] Li Yang and Abdallah Shami. On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing*, 415:295–316, 11 2020. ISSN 18728286. doi:10.1016/j.neucom.2020.07.061.

[94] Ioannis G. Mikellides and Alejandro Lopez Ortega. Challenges in the development and verification of first-principles models in hall-effect thruster simulations that are based on anomalous resistivity and generalized ohm's law. *Plasma Sources Science and Technology*, 28:48, 1 2019. ISSN 13616595. doi:10.1088/1361-6595/aae63b. URL https://doi.org/10.1088/1361-6595/aae63b.

[95] Heikki Haario, Marko Laine, Antonietta Mira, and Eero Saksman. Dram: Efficient adaptive mcmc. *Statistics and Computing*, 16:339–354, 12 2006. ISSN 09603174.

doi:10.1007/s11222-006-9438-0. URL https://link.springer.com/article/10.1007/s11222-006-9438-0.

[96] Thomas A. Marks and Benjamin A. Jorns. Challenges with the self-consistent implementation of closure models for anomalous electron transport in fluid simulations of Hall thrusters. *Plasma Sources Sci. Technol.*, 32(4):045016, April 2023. ISSN 0963-0252. doi:10.1088/1361-6595/accd18.

[97] Justin W. Koo and Iain D. Boyd. Modeling of anomalous electron mobility in hall thrusters. *Physics of Plasmas*, 13:3541, 3 2006. ISSN 1070664X. doi:10.1063/1.2172191. URL https://aip.scitation.org/doi/abs/10.1063/1.2172191.

[98] Ethan Dale. Investigation of the hall thruster breathing mode, 2019.

[99] Jaehong Park, Dongho Lee, Guentae Doh, Youngho Kim, and Wonho Choe. Deep learning-based electron anomalous transport prediction in hall thrusters with lif-measurement. *International Electric Propulsion Conference*, 2022.

[100] J. M. Beall, Y. C. Kim, and E. J. Powers. Estimation of wavenumber and frequency spectra using fixed probe pairs. *Journal of Applied Physics*, 53:3933–3940, 1982. ISSN 00218979. doi:10.1063/1.331279.

[101] Nathan N. Meezan, William W. Hargus, and Mark M. Cappelli. Anomalous electron mobility in a coaxial hall discharge plasma. *Physical Review E*, 63:026410, 1 2001. ISSN 1063651X. doi:10.1103/PhysRevE.63.026410. URL https://journals.aps.org/pre/abstract/10.1103/PhysRevE.63.026410.

[102] Christine M Greve, Manoranjan Majji, and Kentaro Hara. Real-time state estimation of low-frequency plasma oscillations in hall effect thrusters. *Physics of Plasmas*, 28: 93509, 2021. doi:10.1063/5.0057751. URL https://doi.org/10.1063/5.0057751.

[103] Kybeom Kwon, Mitchell L R Walker, and Dimitri N Mavris. Self-consistent, one-dimensional analysis of the hall effect thruster. *Plasma Sources Sci. Technol*, 20: 45021–45036, 2011. doi:10.1088/0963-0252/20/4/045021.

[104] Thomas A Marks and Benjamin A Jorns. Modeling anomalous electron transport in hall thrusters using surrogate methods. *37th International Electric Propulsion Conference*, 2022.

[105] Alejandro Lopez Ortega, Benjamin Jorns, Ioannis G. Mikellides, and Richard R. Hofer. *Numerical simulations of the XR-5 Hall thruster for life assessment at different operating conditions*. 2015. doi:10.2514/6.2015-4008. URL https://arc.aiaa.org/doi/abs/10.2514/6.2015-4008.

[106] S. I. Braginskii. Transport processes in a plasma. *RvPP*, 1:205, 1965. URL https://ui.adsabs.harvard.edu/abs/1965RvPP....1..205B/abstract.

[107] Yusuke Yamashita, Raymond Lau, and Kentaro Hara. Inertial and anisotropic pressure effects on cross-field electron transport in low-temperature magnetized plasmas. *J. Phys. D: Appl. Phys.*, 56(38):384003, June 2023. ISSN 0022-3727. doi:10.1088/1361-6463/acdb83.

[108] Thomas A. Marks and Benjamin Jorns. *Evaluation of several first-principles closure models for Hall thruster anomalous transport*. 2023. doi:10.2514/6.2023-0067. URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-0067.

[109] R C Davidson and N A Krall. Anomalous transport in high-temperature plasmas with applications to solenoidal fusion systems. *Nuclear Fusion*, 17:6, 1977.

[110] C. J. McKinstrie, R. E. Giacone, and E. A. Startsev. Accurate formulas for the landau damping rates of electrostatic waves. *Physics of Plasmas*, 6:463, 1 1999. ISSN 1070-664X. doi:10.1063/1.873212. URL https://aip.scitation.org/doi/abs/10.1063/1.873212.

[111] R. C. Davidson and N. A. Krall. Anomalous transport in high-temperature plasmas with applications to solenoidal fusion systems. *Nuclear Fusion*, 17:1313–1372, 1977. ISSN 17414326. doi:10.1088/0029-5515/17/6/017.

[112] J. Scott Miller, Steve H. Pullins, Dale J. Levandier, Yu Hui Chiu, and Rainer A. Dressler. Xenon charge exchange cross sections for electrostatic thruster models. *Journal of Applied Physics*, 91:984–991, 2 2002. ISSN 00218979. doi:10.1063/1.1426246.

[113] Michael L. Hause, Benjamin D. Prince, and Raymond J. Bemish. Krypton charge exchange cross sections for hall effect thruster models. *Journal of Applied Physics*, 113, 4 2013. ISSN 00218979. doi:10.1063/1.4802432.

[114] Shadrach Hepner, Benjamin Wachs, and Benjamin Jorns. Wave-driven non-classical electron transport in a low temperature magnetically expanding plasma. *Appl. Phys. Lett*, 116:263502, 2020. doi:10.1063/5.0012668. URL https://doi.org/10.1063/5.0012668.

[115] Thomas Marks, Paul Schedler, and Benjamin Jorns. Hallthruster.jl: a julia package for 1d hall thruster discharge simulation. *Journal of Open Source Software*, 8(86): 4672, 2023. doi:10.21105/joss.04672. URL https://doi.org/10.21105/joss.04672.

[116] J. P. Boeuf and L. Garrigues. Low frequency oscillations in a stationary plasma thruster. *Journal of Applied Physics*, 84:3541–3554, 10 1998. ISSN 00218979. doi:10.1063/1.368529.

[117] E. Ahedo and M. Martinez-Sanchez. One dimensional model of plasma flow in a Hall thruster. *Physics of Plasmas*, 8(6):3058–3068, jun 2001. ISSN 00218979. doi:10.1063/1.1371519.

[118] Kentaro Hara. Non-oscillatory quasineutral fluid model of cross-field discharge plasmas. *Physics of Plasmas*, 25(12):123508, 2018. doi:10.1063/1.5055750. URL https://doi.org/10.1063/1.5055750.

[119] Rupali Sahu, Adnan R Mansour, and Kentaro Hara. Full fluid moment model for low temperature magnetized plasmas. *Physics of Plasmas*, 27(11):113505, 2020. doi:10.1063/5.0021474. URL https://doi.org/10.1063/5.0021474.

[120] Ioannis Mikellides, Ira Katz, Myron Mandell, and J. Snyder. *A 1-D model of the Hall-effect thruster with an exhaust region*. 2001. doi:10.2514/6.2001-3505. URL https://arc.aiaa.org/doi/abs/10.2514/6.2001-3505.

[121] Llewellyn Hilleth Thomas. Elliptic problems in linear difference equations over a network. *Watson Sci. Comput. Lab. Rept., Columbia University, New York*, 1:71, 1949.

[122] Shashi Gowda, Yingbo Ma, Alessandro Cheli, Maja Gwóźdź, Viral B. Shah, Alan Edelman, and Christopher Rackauckas. High-performance symbolic-numerics via multiple dispatch. *ACM Commun. Comput. Algebra*, 55(3):92–96, jan 2022. ISSN 1932-2240. doi:10.1145/3511528.3511535. URL https://doi.org/10.1145/3511528.3511535.

[123] G. Hagelaar, K. Hara, A. Smolyakov, and JP Boeuf. Landmark 1d fluid test case. https://jpb911.wixsite.com/landmark/test-case-3, 2018. Accessed: 2021-12-23.