

DISSERTATION

AUGMENTED DIAGNOSTICS FOR ELECTRIC PROPULSION

Submitted by

Seth J. Thompson

Department of Mechanical Engineering

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2026

Doctoral Committee:

Advisor: John D. Williams

Azer P. Yalin
Ciprian Dumitrache
Jacob L. Roberts

Copyright by Seth J. Thompson 2026

All Rights Reserved

ABSTRACT

AUGMENTED DIAGNOSTICS FOR ELECTRIC PROPULSION

Hall effect thrusters (HETs) have become the leading electric propulsion technology for satellite constellations and deep-space missions, yet predicting their in-space performance from ground test data remains a central challenge. Extended Kalman filters and other data-driven modeling frameworks offer a possible path toward bridging the ground-to-orbit gap by fusing experimental measurements with physics-based models to estimate poorly understood model parameters, such as anomalous electron transport. However, these frameworks are fundamentally limited by the information content of the measurements supplied to them. This dissertation addresses the challenge of improving the specificity of traditional diagnostic measurements used to inform Hall thruster models.

Three complementary techniques are developed, each targeting a distinct dimension of measurement specificity. The first is ion-species specificity, achieved with a combined energy-velocity analyzer (EVADER) probe that sequentially filters ions by energy per charge and velocity. By collecting ExB spectra at multiple pass energies set by an upstream electrostatic energy analyzer, the EVADER deconvolves overlapping ion energy distributions that can be inseparable with conventional diagnostics. This can be used to produce species-specific ion energy distribution functions for singly, doubly, and triply charged ions across multiple thruster operating conditions.

The second contribution addresses temporal specificity through the application of shadow manifold interpolation (SMI), a nonlinear state-space reconstruction technique based on Takens' embedding theorem. The SMI technique is used to reconstruct time-resolved ion current density profiles from slow-sweep Faraday probe measurements, using synchronous discharge current waveforms as the reference signal, producing spatiotemporal current density distributions that capture breathing-mode oscillation dynamics inaccessible to conventional time-averaged measurements.

The third contribution targets causal specificity through extended convergent cross mapping (eCCM), a causality determination technique applied to HET telemetry signals. The eCCM analysis evaluates causal relationships between cathode current, discharge voltage, and cathode-to-ground voltage across multiple operating conditions and background pressures. The results indicated that discharge voltage retains the most dynamically relevant information and that facility pressure degrades causal relationships among telemetry signals.

Together, these improvements in species, temporal, and causal specificity form a measurement framework that can enhance the data supplied to next-generation data-driven thruster models, supporting both faster qualification of current flight hardware and the development of more capable Hall effect thrusters.

ACKNOWLEDGMENTS

"It takes a village," the adage goes, and I know of nowhere in my life that it is more applicable than in the help and support I have received throughout my graduate career. I am eternally indebted to the people described below and intend to do my best to pay these gifts forward.

I would first like to thank my advisor, Professor John D. Williams. There is no doubt in my mind that, with any other advisor in the same circumstances I have faced, I likely would never have finished my PhD. Several life-changing events affected my family and me throughout this process, and Professor Williams was steadfast and ready to find a way to make it work at every step. Thank you so much.

I would also like to thank my committee members and collaborators. Professor Azer Yalin has been incredibly supportive in helping me navigate my unconventional path to completing my PhD. Professors Ciprian Dumitrache and Jacob Roberts helped clarify my thinking on these topics, and their valuable feedback undoubtedly improved the quality of this dissertation. Professor Kentaro Hara of Stanford University is also greatly responsible for the work presented herein — his guidance and insightful discussions helped fuel the final effort to publish and tie together this research. Thank you all.

Next, there is an extensive list of colleagues over the past decade that I would like to thank for their help in this process. Their contributions and support are bright spots throughout the 10 years that this work spans. My previous and current lab mates: Tucker, Bao, Tom, Ryan, Jack V, Susan, Jack G, Emily, Zach, Tait, Roberto, Hendrix, Alex, Jon, Mitchell, and Matteo. Your company, camaraderie, and enjoyable presence have significantly improved the day-to-day of doing fun, interesting, and often hard research in the lab. Thank you all.

There is a group of individuals who have been omnipresent every step of the way through my graduate career: Casey, Cody, and Shawn Farnell. Casey, your willingness to talk through problems and potential solutions has been invaluable to me throughout everything I have done in the lab. I value your perspective immensely, and I know the lab would not have run without your tireless efforts. Cody, your

patience and humility when I have come to you with blatantly ignorant questions are something I aspire to emulate, and your "brass tacks" engineering approach has been an amazing example. Shawn, your direct support for almost every experiment run in the lab since 2017 has had a direct causal influence on the quality and quantity of the work produced. Both my own work and that of the whole lab have benefited dramatically from your software efforts and constant support, and your willingness to endure my ridiculous coding and math questions directly led to the work presented in this dissertation. Beyond the technical, each of y'all's personalities, senses of humor, and incredible character is something I am eternally grateful for. I have no doubt that without y'all's help, this work would never have been completed. Thank you.

Finally, there are truly an astonishing number of amazing friends and family who have supported me throughout my PhD. Tucker, your sense of humor has always kept me grounded and smiling, even when I managed to break everything at the lab. Brady, your persistence in re-examining my assumptions and clarifying my decision-making has been immensely helpful in both my professional and personal life. To my siblings — Brice, Sean, and Kirsten — without your support on so many fronts, I don't think I would have ever made it through. Susan, thank you for all your help with childcare; without it, this dissertation would certainly have been delayed even further. Padre, if there is one thing I have learned about nonlinear dynamical systems throughout this process, it is that initial conditions matter — thank you for all the hard work you have put into helping us think and operate as we do today. John B, your consistent support for the family and me, and the space we lived for the majority of this journey, is a significant reason we were able to get to where we are, and we are eternally grateful. Mom, it is interesting to grow up and eventually see that which you were blind to your entire childhood. Your conscious decisions to attack and overcome any challenge that comes your way are not only awe-inspiring but fundamental to how you raised your children. You are selfless while also overcoming numerous challenges in support, and I am incredibly grateful for your lifelong support and the amazing example you continue to set every day. Finally, to my amazing wife, Miranda — it is one thing to go

through a PhD program with a child, but to do it with brand-new baby girls, a four-year-old, and a wife who battles through her own challenges with MS and a full-time workload is another thing entirely. Your exhaustive effort in taking on far more than your fair share of the family responsibilities while I hid away to make incremental progress is the only reason this dissertation exists. I most certainly could not have done it without you.

TABLE OF CONTENTS

ABSTRACT		ii
ACKNOWLEDGMENTS		iv
1	Introduction.....	1
1.1	Motivation for Electric Propulsion and Hall-Effect Thrusters	1
1.2	The Modeling Challenge: Bridging Ground Test to Space Performance	3
1.3	The Diagnostic Gap: Limitations of Traditional Measurement Approaches	6
1.4	A Measurement-Specificity Perspective on Plasma Diagnostics	7
1.5	Dissertation Overview and Contributions	8
2	Background.....	10
2.1	Hall effect thruster physics	10
2.2	Plasma Diagnostic Methods for Electric Propulsion	15
2.3	Data-Driven Approaches in Plasma Modeling	18
2.4	Nonlinear Time Series Analysis and Dynamical Systems	20
2.5	Summary: Opportunities for Diagnostic Innovation (Problem Statement)	22
3	Species-Resolved Ion Energy Distributions Using Combined Energy-Velocity Filtering	24
3.1	Introduction.....	25
3.2	Theory of Combined Energy-Velocity Probe.....	26
3.2.1	EVADER Probe.....	26
3.2.2	ESA Theory	27
3.2.3	ExB Theory	29
3.3	Experimental Setup.....	30
3.3.1	Test Facility	30
3.3.2	Hall Effect Thruster Operation.....	31
3.4	Analysis Methodology.....	33
3.4.1	Standalone Energy and Species Spectra.....	33
3.4.2	EVADER Analysis.....	35
3.4.3	Plume CEX correction.....	39
3.4.4	ExB Spectra Analysis	39
3.5	Results of Combined Probe Measurements	40
3.5.1	CSU 1.5-kW Test.....	40
3.6	Discussion.....	47
3.7	Summary.....	48
4	Temporal Reconstruction of Plasma Properties via Shadow Manifold Interpolation	50
4.1	The Temporal Information Gap.....	51
4.2	Experimental Configuration	53
4.3	Shadow Manifold Interpolation.....	54
4.3.1	Signal reconstruction used for cross-convergent mapping.....	55
4.3.2	Reconstruction process for ion-current measurements	57
4.4	Spatiotemporal Ion Current Density Reconstruction.....	59
4.5	Discussion.....	61
4.6	Summary.....	63
5	Causal Analysis of Hall Thruster Telemetry.....	65

5.1	Introduction: Signal Selection for Data-Driven Models.....	65
5.2	Extended Cross-Convergent Mapping.....	68
5.2.1	Background of extended cross-convergent mapping	69
5.2.2	Extended cross-convergent mapping.....	71
5.2.3	Embedding parameter selection	72
5.2.4	Example Problem	72
5.3	Experimental Configuration	75
5.4	Embedding Parameter Analysis.....	76
5.4.1	4.5-kW Operation.....	77
5.4.2	6-kW Operation.....	82
5.5	Extended Cross-Convergent Mapping Results.....	85
5.5.1	Extended cross-convergent mapping results of the 4.5-kW condition.....	85
5.5.2	Extended cross-convergent results of the 6-kW operation.....	89
5.5.3	Pressure effects on cross-mapping	90
5.6	Discussion.....	93
5.7	Summary.....	94
6	Broader Implications	96
6.1	Measurement Specificity as a Unifying Framework	96
6.2	Toward Integrated Diagnostic Frameworks.....	97
6.3	Implications for Hall Thruster Modeling.....	99
6.4	Limitations and Scope	100
7	Conclusion	103
7.1	Summary of Contributions	103
7.2	Future Work	105
	References	108

1 INTRODUCTION

“To measure is to know. If you cannot measure it, you cannot improve it.” – Lord Kelvin

1.1 Motivation for Electric Propulsion and Hall-Effect Thrusters

Electric propulsion (EP) is a highly enabling propulsion technology that allows consideration of missions that would not be feasible with their chemical-based counterparts. This is due to the order-of-magnitude improvement in fuel efficiency, referred to as specific impulse in propulsion parlance, which dramatically reduces the amount of propellant required. The currency for any maneuver in space is Δv , or the change in velocity, required to make the maneuver. We can relate the mass of our delivered payload to this change in velocity through the rocket equation [1]

$$m_d = (m_d + m_{prop}) \exp\left(-\frac{\Delta v}{\eta_m v_{ex}}\right) \quad (1)$$

where m_d is the mass of the delivered payload, m_{prop} is the mass of the propellant, η_m is the propellant mass utilization, and v_{ex} is the exhaust velocity of the propellant. This exhaust velocity and propellant utilization product is related to a performance metric known as specific impulse, or Isp, by

$$\eta_m v_{ex} = I_{sp} * g_{e0} , \quad (2)$$

with g_{e0} representing Earth’s gravitational acceleration at sea-level (9.81 m/s²). From here, equation (1) can be rewritten as

$$m_d = (m_d + m_{prop}) \exp\left(-\frac{\Delta v}{I_{sp} g_{e0}}\right), \quad (3)$$

highlighting the exponential dependence of the propellant mass required to perform a given mission, which one wants to be small compared to the mass of the delivered payload. Solving directly for the propellant mass required for a given Δv mission yields:

$$m_{prop} = m_d \left(\exp \left(\frac{\Delta v}{I_{sp} g_{e0}} \right) - 1 \right). \quad (4)$$

Typical state-of-the-art electric propulsion devices, such as Hall effect thrusters (HETs), offer specific impulses in the range of 1500-3000 seconds for many commonly used propellants. This order-of-magnitude higher I_{sp} , relative to the commonly achieved 200-300 seconds of chemical rockets, in the exponential term defining the required propellant mass for a given Δv , is what makes EP devices so attractive. In addition to HETs, there are other EP devices, like gridded ion thrusters with an I_{sp} around 2500-4000 seconds, and arcjets with I_{sp} around 150-600 seconds, that have a long history with spacecraft propulsion spanning more than 50 years [1]. The improved I_{sp} of EP devices over chemical systems comes at the cost of requiring a power supply system to generate (and accelerate) the plasma that is fundamental to their operation. In most cases, the total available spacecraft power limit the devices to lower thrust (mN-N) than typically achievable using chemical rockets (N-MN), meaning that EP device use is limited to performing in-space maneuvering, such as orbit raising or N-S and E-W transfers, as well as station keeping for geostationary satellites when short mission duration times aren't critical. Historically, power limitations on spacecraft have prevented the use of EP in systems that need to respond rapidly, because the available power limited the use of higher-power EP thrusters capable of delivering the thrust needed to maneuver quickly. Common alternatives, such as chemical-based monopropellant hydrazine thrusters, are, however, now being reexamined as part of hybrid propulsion systems that combine the best of both worlds (high thrust chemical, low thrust electric) along with a growing focus on increasing spacecraft power to enhance EP for use in avoidance maneuvering [2], [3].

In recent years, HETs have emerged as leaders of existing EP technologies. They are the primary satellite propulsion technology for many constellations, such as SpaceX's Starlink, Amazon LEO, OneWeb, and AST Mobile, and are the primary propulsion system for the SDA's multi-tranche satellite constellations. Apart from the commercial interest that has led to orders-of-magnitude increases in the number of near-Earth satellites equipped with HETs, there is also growing interest in this technology for

Exploration-class spacecraft, with a recent example being the Psyche mission [4], [5]. The HETs in commercial space are in the few-hundred-watt to the 6kW range for most systems, although 13-20 kW thrusters are actively being developed [6], [7]. The growing demand for low-power HETs, coupled with a renewed interest in higher-power systems, has placed renewed emphasis on how best to design, develop, understand, test, and qualify these devices.

Although Hall effect thrusters must be operated in space for thousands to tens of thousands of hours, they must first be qualified in ground-based vacuum test facilities. This introduces a challenge, as ground-based test facilities create an environment that interacts with a thruster, resulting in behavior that differs from what can be expected in space. These so-called “facility effects” can be categorized into three broader categories. First, pressure effects arising from limited pumping speed result in higher neutral-density environments than expected in orbit, which can affect the thruster plasma discharge and neutralizer performance, as well as the thrust, Isp, and power of the thruster system [8], [9], [10], [11]. Second, most large vacuum facilities have conductive walls that intercept the beam plasma, resulting in electrical boundary conditions that differ from those in space [12], [13], [14]. Third, the highly energetic ion beam that produces thrust must collide with surfaces at the beam target or on the chamber walls, resulting in sputtering of these materials that can redeposit on the thruster and obfuscate true erosion rates and potentially cause arcing and flare events [15], [16]. As discussed in 2.3 the Hall effect thruster discharge relies on several coupled nonlinear processes [17], and as a result, extrapolation to space is difficult to predict without properly capturing the key physics of the device under the influence of its test environment.

1.2 The Modeling Challenge: Bridging Ground Test to Space Performance

Predictive modeling of EP devices serves two main goals. First is the accurate prediction of in-space performance, and the second is estimating the lifetime. As discussed above, there are several reasons why we may not replicate the conditions the thruster will experience in space here on the ground, all of which are expected to worsen as power levels increase. Properly translating ground-test

performance to in-space performance is crucial to enabling high-power thruster qualification and would dramatically reduce qualification costs for lower-power thrusters used commercially. Furthermore, predictive models can serve as a design tool for developing next-generation thrusters with improved performance. Present modeling capabilities range from 0D global plasma models to multidimensional fluid and even hybrid fluid-PIC models [17], [18], with varying degrees of utility and empirical closure. The HET discharge plasma presents a unique modeling challenge in that it is a low-temperature, partially magnetized plasma whose characteristics are primarily driven by electron transport phenomena that are not fully understood. Classical descriptions fail to match experimentally observed phenomena without use of adjustable parameters, and temporal behavior both from models and experiments can vary significantly from that predicted by non-temporal models [19].

In HET devices, classical electron transport underpredicts observed transport behavior by 1-2 orders of magnitude. This discrepancy is described as “anomalous electron transport” and is attributed to both plasma instabilities, such as azimuthal drift waves, the electron cyclotron drift instability, breathing-mode oscillations, as well as electron-wall interactions [17], [18], [19]. This results in the effective electron mobility parameters being treated as free parameters, fitted from the comparison of model output with experimental data, making these models descriptive rather than predictive [20]. Temporal fluctuations further complicate the experimental calibration process, as many of the measures used to determine the spatial electron mobility parameter are time-averaged, even though experimental measurements indicate that it varies in time [21].

The HET plasma discharge is highly nonlinear, and its key model parameters, such as anomalous mobility, ionization rates, and wall-interaction coefficients, are not constant but time-varying [21], [22]. The ubiquitous breathing mode, present in the 10-30 kHz range, can exhibit plasma property oscillations that exceed 100% of average values [23], [24], [25]. Time-average measurements and non-temporal models miss this dynamic information, which may be critical to accurate plasma and model parameter estimation. Simply adding temporal capabilities to existing steady-state models may be too

computationally intensive, while providing the key calibration data may also be difficult. To meet this challenge, data-driven frameworks have begun to emerge that fuse experimental data into the model to constrain or inform poorly understood model parameters [22], [26], [27]. These models link key parameters to more easily measurable dynamic data, such as discharge current waveforms, enabling live-model updates. Extended Kalman filters and their physics-constrained variants have been demonstrated as effective tools for estimating unknown model parameters of the HET discharge, such as the time-resolved bulk electron temperature [28] or anomalous electron mobility [22]. These data-driven approaches can act to bridge the gap between fully predictive, high-fidelity models and calibration-reliant non-temporal models commonly employed today.

The central challenge motivating this dissertation is to improve the quality and the specificity of the experimental data supplied to predictive Hall thruster models. The data-driven modeling frameworks described above are only as good as the measurements fed into them. The measurement's ability to constrain model parameters, reduce state-space uncertainty, and produce reliable predictions depends fundamentally on the information content of the input measurements. This doesn't necessarily mean the improvement of signal-to-noise, but rather the specificity of the measurements – ensuring they contain species-specific, temporally resolved, and causally relevant information that maps directly onto the parameters and states these models are trying to estimate. Higher-specificity measurements enable tighter model constraints, which in turn improve predictions of ground-to-orbit performance, addressing a major bottleneck in HET qualification. Beyond this near-term application, improving our fundamental understanding of the discharge physics (anomalous transport, oscillation dynamics, erosion) through better-constrained models opens the path to model-based designs for the next generation of Hall thrusters, rather than relying on empirical scaling and heritage. Despite the promise of data-driven modeling, the diagnostic measurements commonly available today fall short of the specificity these frameworks require.

1.3 The Diagnostic Gap: Limitations of Traditional Measurement Approaches

A host of diagnostics are commonly employed to characterize EP plasma discharges. They range from electrostatic probes, some of which can be perturbative, to optical techniques used to non-invasively explore the relatively higher energy and density regions of the discharge plasma. An electrostatic energy analyzer (ESA) [29] or a retarding potential analyzer (RPA) [30] is often used to characterize ion-energy distributions in the downstream far-field of the plume. An ExB probe is frequently co-located in the downstream plume and is used to provide species distributions for multiply charged ions [31]. Langmuir probes can provide information on plasma potential, plasma density, and electron temperature in the plume [32] but can be perturbative if placed too close to the discharge [33]. A Faraday probe is often used to quantify the ion current density distribution, and the resulting data, along with data from other probes, are frequently employed in phenological efficiency analysis [34]. Measurements of the discharge channel properties require non-invasive optical techniques, such as laser-induced fluorescence, to obtain IVDF profiles in the channel [35]. Optical emission spectroscopy (OES) can also be used in conjunction with collisional-radiative models (CRMs) to estimate electron temperature and plasma density from relative emission spectra of the plasma [36]. These diagnostics have long served as workhorses for characterizing EP devices [37]. Most test facilities will have access to electrostatic diagnostics due to their relatively low cost, ease of installation, and straightforward operation, while optical measurements can be more specialized and used only in specific test campaigns.

A challenge with many of these diagnostics is how they measure aggregate or convoluted quantities. Energy measurements from ESA or RPA probes show the energy-per-charge (E/q) distribution and do not distinguish by species [29]. The species distributions inferred from ExB spectra can overlap if each species exhibits a wide energy distribution [38]. Faraday probe measurements detect ions of all energies at a given location and therefore include low-energy ions generated through charge-exchange collisions (CEX) [39]. These convolutions can limit the specificity of the information needed to constrain model parameters. In stronger terms, one could say that if a model could be cast to provide species-

specific information to constrain the discharge model parameters, then the convoluted measurements would lead to ambiguity in those parameters.

Furthermore, many experimental measurements are sweep-based, meaning that the diagnostic builds a distribution by sweeping over a range of measurement states adjusted by applying new voltages, wavelengths, or adjusting measurement locations. This scanning can be quite slow relative to the time scales of interest. Even the lowest-frequency (most accessible) oscillations of the breathing mode, at 10-30 kHz, are well beyond the sweep rate of an ESA or ExB probe. There have been attempts to increase the temporal capabilities of existing Langmuir probe measurements [40] and LIF measurements [41], which have shown frequency responses fast enough to access breathing-mode oscillations. The additional techniques discussed below have been developed to address these limitations [42].

The challenge is that models benefit from species-specific, time-resolved, spatially distributed data about the plasma, while many of today's most common diagnostics yield time-averaged, spatially local, and species-convolved measurements. Data-driven models are, to some degree, fundamentally limited by the measurements that can be fed into them [22], [27]. Closing this gap will require improvements in diagnostic techniques and in how we extract useful information from available measurements.

1.4 A Measurement-Specificity Perspective on Plasma Diagnostics

Improving measurement specificity is an important step toward meaningfully constraining model parameters. In the context discussed herein, measurement specificity refers to the degree to which a diagnostic measurement provides unambiguous, targeted information about plasma properties or model parameters. In this context, a low specificity measurement could be a convoluted measurement, *e.g.*, overlapping ion energy distributions of multiple species, whereas a high-specificity measurement could include species-specific information, time-resolved measurements, or spatially defined temporal changes.

The work presented herein focuses on improving measurement specificity through probe-level deconvolution, spatiotemporal reconstruction, and causal signal selection.

The purpose of this improved measurement specificity is to enable more capable modeling. Each experimental measurement included in a model should help constrain the set of possible model states (the state space). The more specific the experimental measurements are, the more the available state space can be reduced. In extended Kalman filter approaches [22], [26], [27], this corresponds to a reduction in the uncertainty (as measured by the covariance) of the estimated states and model parameters. For example, if ion-energy distributions are included in a model, species-specific energy distributions constrain the ionization and acceleration model parameters more tightly than species-averaged distributions. Similarly, temporal models that implement time-resolved ion measurement data constrain dynamic model behavior, *e.g.*, time-dependent mobility parameters. In addition, causally relevant signals provide more informative measurement updates than arbitrary or redundant signals, thereby improving computational cost.

The diagnostic augmentation efforts described in this dissertation center on three implementations that improve measurement specificity. The first is serial filtering for probe specificity, discussed in Chapter 2, by combining energy and velocity filters to deconvolve overlapping species distributions. The second augmentation of measurements comes through the addition of spatiotemporal reconstruction of state-based measurements in Chapter 4, where, by way of example, shadow manifold interpolation techniques are applied to generate time-resolved ion-current density measurements. Finally, in Chapter 5, improvement in signal selection is evaluated through a causality analysis framework to determine which signals retain the most dynamically relevant information.

1.5 Dissertation Overview and Contributions

This dissertation demonstrates the ability to improve the specificity of diagnostic measurements—through species-resolved energy analysis, temporal reconstruction, and causal selection of dynamic signals—to better inform data-driven predictive models of Hall effect thruster discharges,

with the ultimate aim of improving ground-to-orbit performance prediction and enabling model-based thruster design. Chapter 2 introduces some of the concepts and physics of the devices used in the testing, as well as some of the concepts that underpin the analysis techniques. Chapter 2 describes the demonstration of an inline energy-velocity filter for deconvolving species-specific ion-energy distributions. Chapter 4 demonstrates the ability to use non-linear state-space reconstruction to produce spatiotemporally resolved ion-current density profiles via shadow manifold interpolation. Chapter 5 discusses the use of causality determination techniques as a framework for improving signal selection in time-resolved data-driven models. Chapter 6 will present how future extensions of these concepts can further enable difficult measurements.

2 BACKGROUND

Hall effect thrusters are considered the electric propulsion device of choice for many satellites, selected for use in low Earth orbit commercial constellations and for discovery-class, deep-space missions [5]. Despite their growing use, fundamental gaps in understanding physical phenomena still limit our ability to deploy predictive codes as design tools or to predict on-orbit performance [19]. This chapter introduces key concepts needed to clearly articulate the diagnostic innovations for improving measurement specificity, detailed in the following chapters. To do this, background information on Hall effect thrusters, traditional EP diagnostics, and the temporal domain of these devices is explored to support the more detailed information introduced in the following chapters.

2.1 Hall effect thruster physics

The Hall effect thruster is described most prominently by its crossed axial electric and radial magnetic fields that produce an electron drift azimuthally, in the $\mathbf{E} \times \mathbf{B}$ direction, throughout the annular channel. A diagrammatic representation of a 1.5kW research Hall thruster is shown in Figure 1 next to a photograph of it operating on krypton. The majority of the propellant (neutral gas particles) is distributed in the back of the annular channel through the anode, while a small fraction of the total gas is commonly fed to the hollow cathode used for discharge seeding and ion beam neutralization [37].




Figure 1: Cross-section view of a Hall effect thruster with schematic representation of electric and magnetic fields (left) and photograph of the CEPPE's 1.5kW research thruster operating on krypton (right).

In these devices, the electrons are highly magnetized by the radial magnetic field, with the Larmor radius typically much smaller than the channel width, $r_L \ll w_{ch}$. Ions are relatively unmagnetized by these magnetic fields due to their significantly greater mass. The good electron confinement enables most of the voltage drop between the anode and cathode to be concentrated at the downstream end of the discharge channel, as shown in Figure 2. This region of a strong electric field heats the electrons, which then efficiently ionize the propellant flowing into this region of the channel. Generally, the resulting plasma discharge is modeled with three distinct plasma regimes [17], [43]. The first is the ionization region, just upstream of the highest radial magnetic field strength, where the neutral density is highest. Immediately downstream, and sometimes overlapping with, the ionization region is the acceleration region near the channel exit. The acceleration region occurs in the region of highest magnetic field strength and is characterized by a strong axial electric field that electrostatically accelerates ions downstream. As a result, although Hall thrusters utilize magnetic fields, they are often characterized as electrostatic electric propulsion devices [37]. The next downstream region is the near-field plume, which has significantly reduced neutral and plasma densities. This is where the exiting ion beam current density and the neutralizing electron current density balance each other, in the absence of any influence from the vacuum facility. Due to the dependence of electron confinement on the magnetic field, the thruster

magnetic field topology can significantly impact the spatial distribution of electron temperature, ionization rate, and plasma potential [37]. It is the potential structure experienced by ions born and accelerated through these regions that dictates their ion energy distribution, thereby imparting information about the discharge in the unique distribution of those ions.

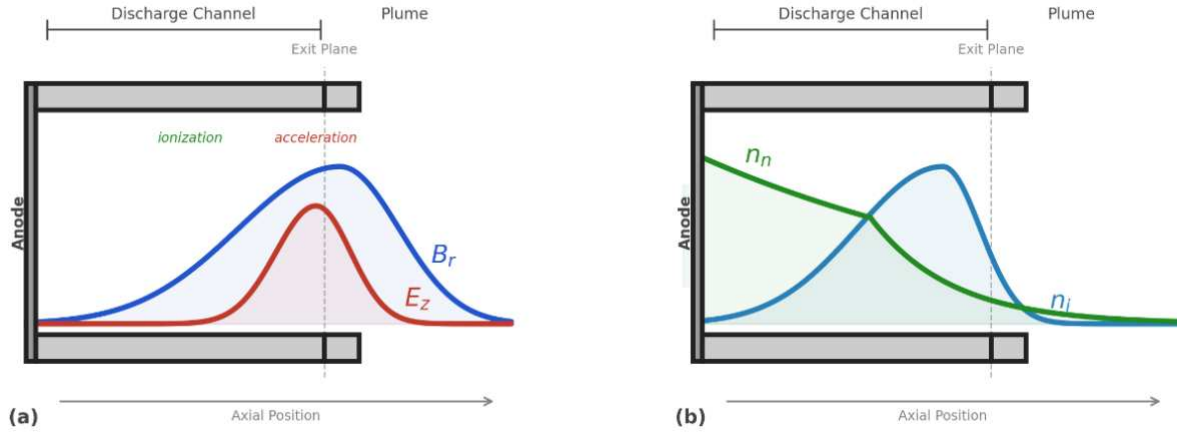


Figure 2: Notional diagram of field strengths (a) and particle densities (b) generated for illustrative purposes.

The HETs are characterized by key performance parameters that translate their operation into metrics for mission planning, namely I_{sp} , thrust, and efficiency. The first, and key proposition of all-electric propulsion thrusters, is the specific impulse. Specific impulse is related to thrust, T , by

$$I_{sp} = \frac{T}{\dot{m}_p g_{e0}}, \quad (5)$$

where \dot{m}_p is the mass flow rate of the propellant and g_{e0} is Earth's gravitational constant. As shown in equation (3) this can also be related to exhaust velocity, but in electric propulsion, an additional term is also considered, known as mass utilization η_m , which accounts for the fraction of the propellant that is actually ionized and accelerated away—typically in the 0.8 range. This relatively high utilization efficiency is due to the lower velocity of the unionized neutral particles, which is about a magnitude lower than the ions entering the acceleration region. As a result, we consider

$$I_{sp} = \frac{v_{ion} \dot{m}_{ion}}{g_{e0} \dot{m}_p} = \frac{v_{ion}}{g_{e0}} \eta_m \quad (6)$$

where \dot{m}_{ion} is the mass flow rate of the ion accelerated away, and v_{ion} is the velocity of the accelerated ions. This mass utilization term is a crucial efficiency term that has a strong impact on the thruster's overall performance. The other efficiency often evaluated is the electrical efficiency η_e of the device which relates the fraction of the total input power P_T to the amount of power transferred to the ion beam P_b , through,

$$\eta_e = \frac{P_b}{P_T} = \frac{I_b V_b}{P_d + P_0} \quad (7)$$

where I_b is the current and V_b is the average beam potential of the ions accelerated from the thruster. The power required for the discharge is P_d , and the other input power requirement, P_0 , to produce the ion beam, such as magnetic coils, keeper power, cathode heater, etc. We can relate the input power required for a given amount of thrust to these efficiencies and specific impulse through

$$\frac{T}{P_T} = \frac{2\eta_m \eta_e}{g_{e0} I_{sp}} = \frac{2}{g_{e0}} \frac{\eta_T}{I_{sp}} \quad (8)$$

where η_T is the total thruster efficiency. The relationship between thrust, efficiency, and specific impulse, as given in equation (8), is used to characterize the device's performance for any given mission set.

Typical values for the thruster efficiency of HETs are 50-75% with specific impulses between 1500-3500 seconds, depending on propellant and operating voltage [6], [37].

The magnetic field strength and topology play a pivotal role in determining properties of the thruster discharge, including electron confinement, ionization efficiency, and wall erosion. In 2012, the concept of magnetic shielding was introduced into the broader EP literature [44], with the primary goal of reducing channel-wall bombardment by tailoring field lines to avoid wall intersections. This reduction in wall bombardment increased thruster lifetime from roughly 8,000 hours to greater than 20,000 hours [37]. As a result of the adjusted magnetic field shape, magnetic shielding was found to affect plasma discharge

physics, leading to changes in potential structures and ion energy distributions, as well as in the types of temporal phenomena observed.

Hall effect thrusters exhibit rich oscillatory and turbulent behavior characteristic of their operation [44]. These span orders of magnitude in time scales and are attributed to many different types of instabilities [19] that dictate much of the thruster's steady-state performance [18]. Most of the temporal diagnostic work has focused on evaluating breathing mode oscillations [23], [40], [41] due to their accessibility in the 10-30kHz range. Additional imaging work on rotating spokes in the 1-25 kHz range seen in unshielded thrusters has also been evaluated [45]. Higher-frequency oscillations are also observed in the MHz range, consistent with ion acoustic turbulence and electron-cyclotron drift instabilities, which have been confirmed via coherent Thomson scattering studies [46]. Therefore, when a large degree of temporal variation is observed in a given diagnostic measurement, it is not merely noise; rather, the oscillations are intrinsic to the discharge physics and affect transport, performance, and model fidelity.

A key process in these devices that should be emphasized is the anomalous electron transport. This is the central physics problem underlying predictive modeling challenges. The cross-field transport described by classical collisional or even Bohm-type cross-field diffusion underpredicts the electron current to the anode by more than a factor of 10 relative to the measured current observed in thruster operation. Several candidate mechanisms have been proposed: azimuthal turbulence-driven transport, near-wall conductivity with secondary electron emission, and cyclotron drift instabilities [19], [37]. The effective anomalous mobility is the primary "knob" to turn in models to match them to experimental data, which, unfortunately, yields descriptive rather than predictive models.

Further complicating the matter of prediction is the altered conditions under which the thruster must be tested on the ground, thereby introducing additional effects of the facility. These come in three primary forms, as mentioned briefly above. The first is a finite background pressure that is orders of magnitude higher than space conditions. This can cause the production of CEX ions, modify ionization rates, and change the breathing mode frequency, and some indications suggest changes in the thruster's

acceleration region as well [8], [9], [10], [11]. The second effect results from the need to test the thrusters in a metal, electrically conducting, facility that intercepts the plasma beam and may alter the thruster behavior due to electrical interactions with the facility [12], [13], [14]. Finally, there is also the problem of back-sputtered material from the chamber walls redepositing on the thruster, leading to uncertainties in erosion rates and causing operational problems such as arcing. All three of these problems are expected to increase as higher-power thrusters are explored [47] further motivating the need for accurate predictive models that enable one to extrapolate from ground test measurements to in-space performance.

2.2 Plasma Diagnostic Methods for Electric Propulsion

Several electrostatic diagnostics are commonly used to characterize HETs. The focus of this dissertation is on electrostatic energy analyzers (ESAs), Wien velocity filters or ExB probes, and Faraday probes; however, Langmuir probes [32], retarding potential analyzers [30], and emissive probes [48] are also commonly used in these tests as well, although they are not covered in detail herein. Electrostatic analyzers are used to evaluate the ion energy distribution function of ions in the region $\sim 1\text{-}2$ m downstream of the thruster. There are several configurations of these devices [29]. The method used in this work is a 120° segment of a spherical deflector analyzer, the details of which are discussed further in Chapter 2. A key aspect of most energy-discriminating diagnostics is that they slowly build up an ion energy distribution as the probe structures are swept with applied voltages. The current collected at a given applied voltage is representative of the number of ions that pass through the analyzer at a given energy-per-charge (E/q). However, it is important to note that this is a cumulative current and therefore distributions of all species (e.g., singly, doubly, triply, ... charged ions) at a given E/q ratio will arrive at the collector, resulting in a convolution of energy distributions for all ion species leaving the thruster.

The ExB probe, also discussed in greater detail in Chapter 2, is often used as a velocity analyzer and is primarily leveraged to distinguish ion species in the plumes of Hall thrusters [31]. However, in the same fashion as the ESA all current measured at a given sweep voltage represent a combination of ions that were able to arrive at that plate voltage, and when the ion-energy distributions for the species are

wide, the species distributions overlap enough to add uncertainty in the fraction of ions that belong to each species [31], [49], [50]. Analyses have been proposed to better handle the overlapping distributions by fitting assumed velocity distribution functions VDFs [38]; however, this requires assumptions about the shapes of these VDFs.

Faraday probes are the standard for providing ion current density distributions of EP devices. The data are often in the form of current density plotted as a function of angle as the Faraday probe is moved on a circular arc that has its origin centered on the HET. These ion current distributions are used to derive the beam current, which is central to the efficiency analyses of these devices and is directly tied to mass utilization and other key efficiencies. However, due to the need to test in vacuum facilities, finite background pressure can alter the measured ion current density profiles from those expected in space, leading to significant efforts to remove these effects [39]. Furthermore, these measurements are almost always reported in steady-state terms, although it has been shown that the plasma ejected from the thruster exhibits rich temporal oscillations with varying densities and profiles [40].

These electrostatic diagnostics are useful in their application to downstream measurements; however, their ability to measure properties near the plasma discharge is limited due to the life of the probe in those conditions, as well as the perturbative effect of the probe, which can alter quantities of interest that the probe seeks to measure [33]. As a result, non-perturbative optical measurements of discharge properties have been the primary means of probing discharge plasma regions. The two most used techniques are laser-induced-fluorescence (LIF) and optical emission spectroscopy (OES). Laser-induced fluorescence has the ability, through measurements of Doppler-shifted light, to couple laser light into metastable ions and neutrals that will then decay at a wavelength, usually non-resonant with the excitation light, providing information not only about the number of particles in the metastable state, but also, based on the Doppler shift, the relative velocity of the ion or neutral [35]. This enables measurement of ion-velocity distribution profiles in the acceleration regions of a Hall thruster. This profile has become the primary means of calibrating anomalous electron transport profiles in simulations to match

experimentally observed results [20], [51]. Optical emission spectroscopy is often used for its ease of integration. No probing laser light is required; instead, an entire spectrum of light is acquired, or occasionally specific bands, and this data is combined with a collisional radiative model (CRM) to constrain the information about the electron energy, plasma density, and neutral density that would have been required to produce such a spectrum [36]. Two other critical optical measurements have been used to advance understanding of the physical properties within the thruster discharge: two-photon absorption laser-induced fluorescence (TALIF) and Thomson scattering. The TALIF approach is similar to LIF but can be used to measure ground states rather than relying on the plasma to populate metastable states, and it provides density estimates typically obtained by mapping flow profiles [52], [53]. Thomson scattering comes in two flavors, incoherent and coherent. The coherent Thomson scattering measurements have been used to evaluate high-speed density fluctuations, which have been key to experimentally verifying the presence of ECDI in the Hall thruster discharge [54], and have now explored the turbulence experienced by the cathode [55]. Incoherent Thomson scattering has been used to evaluate electron energy distributions in the devices' discharges and is currently prompting a re-evaluation of how electrons are treated along magnetic field lines [56][57]. Both TALIF and Thomson scattering can be more intensive optical setups and therefore their deployment is limited and bespoke to a given test campaign and article, although efforts are underway to produce these measurements in several locations [58], [59], [60], [61].

An additional wide range of diagnostics has been used to characterize various aspects of HET discharges, although their deployment is less common. High-speed cameras have been used in many cases to evaluate various instabilities in the Hall thruster [40], notably for the characterization of the spokes instability [46], and also for drift instabilities observed in shielded devices [62]. Microwave interferometry schemes have been developed for density measurements [63]. Hall current tomography tools have been developed to non-invasively reconstruct the Hall current inside the channel of an

unshielded thruster [64]. Although all of these techniques have yielded insight and interesting measurements, their wide-scale deployment remains limited.

The combination of traditional electrostatic and optical-based diagnostics provides a rich, yet incomplete picture of key properties of the discharge. These key gaps center around deconvolving species-specific information, adding temporal resolution to the measurements, and a systematic understanding of which measurements carry the most relevant information for models. These key gaps are what motivate the diagnostic augmentation discussed in Chapters 2-5.

2.3 Data-Driven Approaches in Plasma Modeling

There has been a wide variation in the methods and models used to capture Hall thruster behavior over the years. Zero-dimensional, so-called 0-D global models, are volume-averaged approaches that capture gross thruster behavior, such as mean performance or breathing-mode behavior, if temporal variation is considered [65]. These models are computationally cheap and offer no spatial resolution. One-dimensional (1-D) fluid models are used to resolve axial profiles of the plasma properties and capture the basic discharge structure. These models require transport coefficient inputs to close [66]. Additional specificity is added for 2-D (r-z) hybrid models, in which kinetic ions and fluid electrons are used to resolve radial and axial structures [20], [67]. These hybrid models are among the most commonly used in the community, but they still fail to accurately capture electron mobility profiles in Hall thrusters. The next, higher fidelity, models are typically 2-D and 3-D PIC simulations [68], [69], [70], [71]. These simulations are fully kinetic and are used to resolve instabilities and turbulence, but are too computationally prohibitive for the thruster-scale simulation using realistic parameters. This leads to shortcuts used to reduce computational cost, such as adjusting particle mass values, but at the cost of model fidelity. Note that each level of fidelity introduces additional unknowns or poorly constrained parameters.

The problem of parameter estimation stems from the fact that all practical HET models require specification of transport coefficients (especially anomalous electron mobility) that cannot be determined from first principles [17], [18], [37]. Traditionally, these are tuned to match specific experimental data, particularly LIF-derived mobility profiles for many of the Hybrid models [20], resulting in descriptive models rather than predictive modes. The primary goal of data-driven modeling is to systematically estimate these manually tuned parameters from an experimental dataset in a generalizable way [28].

A method for systematically estimating internal states using an extended Kalman filter (EKF) was demonstrated for the EP community by Greve *et al.* [26] and further developed in [28]. Extended Kalman filters are not new to aerospace and were even leveraged in the original navigation processes of the Apollo mission to gauge spacecraft trajectory in sparse, noisy tracking data [72]. The basic concept is that a physics-based model propagates the state forward in time, called a prediction state. Periodic experimental data are then used to update the state estimate, referred to as a correction step, and the process repeats. The Kalman gain determines how much to trust the model prediction vs. the measurement, based on the respective uncertainties, referred to as process noise covariance Q and measurement noise covariance R . The continuous-discrete form was used, whereby the continuous model propagation is updated with discrete measurements when available [73]. The innovation of Greve's approach was to use a set of physics-informed constraints on the process and measurement noise covariance to prevent unphysical solutions, such as negative density, temperatures, or uncertainty. Furthermore, the framework can describe a process noise, Q , that is related to model imperfections and how it varies in time. These physics-constrained extended Kalman filter (PC-EKF), augmented the EKF approach to be robust to sparse, noisy measurements that are critical in leveraging experimental data [28].

The EKF framework has now been applied to multiple thruster data sets with various modifications. In 2021, Greve first applied the EKF framework for real-time state estimation of low-frequency oscillations in HETs through the use of a 0D global model [26]. This model was further extended in 2023 by Troyetsky *et al.* to evaluate a wide range of discharge voltages and the resulting

temporal dynamics of the discharge current for a 1.5-kW Hall thruster [22]. A similar 0-D global model was used, but with the addition of the electron energy equation to better capture the effects of the electron dynamics on the discharge oscillations. Another important extension of the Greve EKF framework for Hall thrusters was brought to the floor by Zivre and Hara in 2025 [27], where changes to the framework were made to allow multi-modal data sets, where the EKF was more frequently updated by data like the discharge current, but less frequently included additional data such as thrust. The result of this estimation effort revealed a nonlinearity in the change in thruster oscillations under space-like conditions that is not captured by historical linear regressions of pressure on performance data [8], [9], [74]. This underpins a key finding for prediction of in-space performance, that true and actionable knowledge cannot be derived from observing the macroscopic performance properties of a complex system operating in a compromised environment.

A key limitation of all these data-driven models thus far is a limitation in the measurements available for model integration. Presently, the only dynamic data that has been utilized is the discharge current, and the primary macroscopic data to compare against has been thrust. This is a very limited window of data to view about the thruster. This raises the question: what measurements should be supplied to these data-driven frameworks, and how is maximal information extracted from the available diagnostics? Note that current efforts have primarily focused on discharge current because it can be measured easily at high speed and is directly integrated into model parameters, such as bulk electron current. But as stated previously, both discharge current and thrust are convolutions of other states in the plasma dynamics system and may not provide the best observability of all important model states. Richer, more specific measurements could provide better state-space constraints, as discussed in Chapter 5.

2.4 Nonlinear Time Series Analysis and Dynamical Systems

A key aspect to discuss when analyzing temporal phenomena in a Hall thruster is its characterization as a dynamical system. As discussed previously, the underlying physical processes are nonlinear and span orders of magnitude in dimensional complexity [17], [18], [37]. Temporal fluctuations

ranging from pico- to millisecond timescales that influence one another, with length scales on the order of centimeters to sub-millimeters, result in an environment ripe for deterministic chaos (strong sensitivity to initial conditions) [19], [74], yielding complex time series in the thrusters' signals. The breathing mode is a limit-cycle-like oscillation that can become chaotic under certain conditions [75]. It is essential to understand this characterization of the system as likely being chaotic, as it can preclude the use of many linear-based statistical analysis techniques, such as FFT correlation methods, from fully capturing thruster behavior [42], [75].

In evaluating Hall thrusters as deterministic dynamical systems, a different set of tools is needed to characterize and analyze the signals. One of the key concepts to understand is Takens' theorem [76], a foundational proof in dynamical systems theory, which demonstrates that a time series of a single observable variable from a dynamical system can be used to reconstruct a topologically equivalent representation of the system's full attractor via time-delay embedding. This essential states that all participating parts of the dynamic system will leave their "fingerprints" on all the other interacting variables of the system. The details of how these time-delay embeddings are generated will also be covered in detail in Chapter 4. However, a key caveat is that they are causally linked variables. If so, a single measurement can recover all of the underlying dynamics, a key aspect of the reconstructions discussed in Chapters 4 and 5.

It is this concept that is extended to the embedding of multiple observables: if X and Y are both variables in the dynamical system, their respective time-delay embeddings reconstruct "shadow manifolds" that are topologically related (diffeomorphic). If X causally influences Y , then the shadow manifold of Y contains information about X , and therefore X can be predicted from the neighborhood structure of M_Y , the manifold of the Y timeseries data [42]. This is the foundation of cross-convergent mapping, originally proposed by Sugihara *et al.* [77] and later extended by Ye *et al.* [78].

Cross-convergent mapping (CCM) algorithmically uses the near neighbors in one manifold, like M_Y , to pick the corresponding location in the other variable's manifold, say M_X . A key property of the

original CCM framework is *convergence* of the reconstruction [77]. If a causal relationship exists, the cross-mapping skill, most often defined as the Pearson correlation (ρ) between the reconstructed and true time-series signals, improves as the library length L increases, because the manifold becomes denser and nearest neighbors become closer. This is an improvement to the Granger causality paradigm, which was previously used as the gold standard for determining causality [79]. However, the Granger method assumes a separability of variables that, by nature of dynamical systems, cannot be achieved. Cross-convergent mapping fundamentally leverages this inseparability that limits the use of Granger causality to evaluate dynamic causality. An extension of CCM has been developed to better determine causal directionality in systems under strong unidirectional forcing that leads to *generalized synchrony*, while also identifying indirect causal chains [78].

As a result, the dynamic time series measured from HETs (discharge current, voltage, floating potential, ion current, etc.) contain embedded information about the underlying discharge dynamics. Takens' theorem tells us that sufficiently rich single-point measurements can reconstruct the attractor, but not all signals are equally informative. Causality analysis can identify which signals carry the most dynamically relevant information about the discharge physics, guiding the selection of measurements for data-driven models. These concepts are explored in further detail in Chapter 5 which examines the dynamic data collected from circuit measurements of a magnetically shielded Hall thruster tested in a varying-pressure environment.

2.5 Summary: Opportunities for Diagnostic Innovation (Problem Statement)

To further advance our understanding of Hall effect thrusters, predictive codes are needed to capture scaling and ground-to-space translation challenges. Predictive modeling of HETs is currently limited by the lack of known electron transport parameters. Data-driven approaches have been shown to estimate mobility parameters by incorporating dynamically varying data into a reduced-order model; these models are fundamentally limited by the information content of the measurements used to constrain the parameter space.

This work focuses on opportunities to meet this challenge with existing diagnostic tools common in most EP research facilities. This dissertation provides a framework for each opportunity along with a demonstration of the technique. The first opportunity is to improve probe specificity by serially operating filtering probes. In Chapter 2 this is demonstrated by the combination energy-velocity probe, which can deconvolve overlapping species distributions to provide species-specific ion energy distributions. The second opportunity is to add temporal information to many of the common diagnostics used today. Chapter 4 demonstrates the ability to use nonlinear dynamic reconstruction methods to provide these measurements for traditionally time-averaged diagnostics that are swept through measurement states. The third opportunity is to improve the amount of relevant dynamic information captured in signals selected for data-driven frameworks.

Improved probe specificity enhances the uniqueness of the information contained in the measured ions, while temporal reconstruction provides a comprehensive time-resolved view of all measured states at once with respect to a linked input signal, and finally, causality analysis ensures that improved dynamic measurements are selected for model integration. Together, they expand the information available to data-driven predictive models, with the ultimate goal of improving ground-to-orbit performance prediction and enabling model-based thruster design.

3 SPECIES-RESOLVED ION ENERGY DISTRIBUTIONS USING COMBINED ENERGY-VELOCITY FILTERING¹

Ion energy distributions and species fractions are common measurements used to characterize Hall-effect thrusters. They are often used to calculate some of the phenomenological sub-efficiencies that comprise the overall thruster efficiency, but they also contain information about the thruster discharge itself. Energy distributions and species fractions are dependent on the properties of the discharge plasma, and they are measured downstream of the thruster in the far-field plume with (1) minimal perturbation of the thruster and (2) low loss of fidelity due to the small number of collisions the ions suffer before reaching the probe. However, direct measurements of the energy and species fractions can be confounded by the diagnostics used to obtain them. Wide velocity distribution functions can cause overlap in Wien filter (ExB probe) spectra that can increase uncertainty in species fractions [31], [49], [50]. Although this is widely acknowledged, it is worth noting that traditional ExB probe analysis and analysis based on species-specific IEDF measurements reported on herein do not result in significantly different charge utilization sub-efficiency calculation results. Rather, what is lost in traditional ExB probe spectra is discharge plasma ionization and acceleration processes—this is because it is hard to recover discharge plasma effects on multiply charged ions due to the overlap of spectra features from the wide ion energy distributions and the relatively poor energy resolution of ExB probes. At the same time, most common ion energy measurements are made in terms of energy-per-charge (E/q) for all ions entering the probe [29], [30], resulting in a cumulative energy-per-charge distribution that no longer contains information about multiply charged species. However, because both energy and velocity-based diagnostics use filtering to obtain their

¹ Part of this chapter is reprinted with permission from “Combined electrostatic analyzer—Wien filter probe for characterization of species distributions in Hall thrusters” by Thompson *et al.* 2021, *Journal of Applied Physics*, 130, COPYRIGHT 2021 by AIP Publishing.

cumulative distribution function measurements, they can be used in a serial configuration where the specificity of ion measurements is improved by limiting the influence of overlapping distributions.

3.1 Introduction

The species distribution of ions in Hall-effect and ion thrusters has long been measured to provide information on thruster performance, predict erosion and lifetime, and offer insight into the devices' fundamental operation. One traditional method for measuring the species distribution in the thruster plume uses a Wien velocity filter, often called an ExB probe in electric propulsion (EP) literature [38]. Determination of species fractions of multiply charged ions is relatively straightforward in gridded ion thrusters (GITs), which have narrow ion energy distributions. However, the analysis of ExB spectra to determine species fractions becomes more difficult as the ion energy distribution widens, which is commonly observed in Hall-effect thrusters. Even wider ion energy distributions are created in magnetically shielded Hall thruster configurations, making ExB spectra analysis even harder [49], [50]. Previous methods utilizing peak heights [80] and areas under the curve [81] have been shown to be inaccurate for analyzing the ion species fractions in a Hall thruster plume [38]. Huang and Shastry presented an improved method for analyzing ExB spectra collected from HETs using velocity distribution functions (VDFs) to approximate the current associated with each ion species [50]. However, that method requires assumptions about the shape of the VDF for a specific ion species, which are difficult to verify with information typically collected with individual Faraday probes, ExB probes, and energy analyzers, and may therefore increase uncertainty in the measurement of ion species fractions.

The probe discussed in this chapter, the **Energy and Velocity Analyzer for Distributions of Electric Rockets (EVADER)** probe manufactured by Plasma Controls, LLC, combines an energy filter/analyzer placed in-line with a velocity filter to provide energy-selective species analysis. Specifically, the EVADER uses an electrostatic analyzer (ESA) section to pass a narrow window of ions having a particular energy-per-charge (E/q) ratio to a Wien filter, where a standard ExB spectrum is taken. Then, by recording spectra at multiple, ESA-stage-selected energies, it is possible to deconvolve the

overlapping peaks of different species observed in a typical ExB probe spectrum collected from a Hall thruster plume and create species-specific ion energy distribution function (IEDF) measurements. Additionally, the probe may be used to differentiate charged particles of different mass, as found in alternative propellants such as water, iodine, and nitrogen, distinguishing H^+ vs O_2^+ , or I_2^+ versus I^+ , for example [82], [83], [84].

In this work, the EVADER was utilized to characterize a 1.5 kW Hall thruster operating on krypton [64], [85], [86], and a BHT-200 operated on xenon. Our group has used standalone ESA and ExB probes in a combined manner for over a decade prior to other similar diagnostics reported in the literature, such as an energy-selective mass spectrometer (ESMS) device that was used to take data downstream of an SPT-100D EM1 HET and RIT-22 GIT, though the reported data and its analysis were very limited [87].

Details of the EVADER probe and the experimental setup used in its initial demonstration, along with an improved data analysis technique, are discussed. Important aspects of the analysis and results are compared to those from VDF-fit and integration-based analyses of traditional ExB spectra. Results from two thruster test campaigns are shown, with improvements in probe operation with the BHT-200 that demonstrate better resolution of species-specific IEDF information, yielding insights into underlying physics hitherto unavailable to EP researchers.

3.2 Theory of Combined Energy-Velocity Probe

The combined probe leverages filtering of ion energy and velocity using two probe segments. Details of the probe segments and their integration are discussed, along with the theory of operation for each segment. These concepts are then used to explain the EVADER analysis approach.

3.2.1 EVADER Probe

The EVADER probe consists of an inline electrostatic analyzer (ESA) and a Wien velocity filter (or ExB). Figure 3 shows a schematic of the EVADER probe with the primary components of both the ESA and ExB sections labeled. Those familiar with plasma diagnostics will note the components of two,

nested, spherical segment electrostatic deflection plates for energy per charge differentiation (E/q), ion collimators, and the use of an ExB electric and magnetic field section for mass (m), velocity (v), and charge state (q) differentiation.




Figure 3. A top-view schematic of the EVADER energy and velocity filter.

Note that by measuring current at the ExB collimator / ESA collector, the EVADER can operate as a standalone ESA probe. Details of the ESA can be found in [29]. Both stages used in the probe are derivatives of commercially available standalone ESA and ExB probes [50], [88], [89].

3.2.2 ESA Theory

The ESA is a deflector-type analyzer designed to separate ions according to their energy-to-charge ratio, E/q . The geometry of the concentric electrodes and applied voltages is chosen such that charged particles of a particular transmission energy, E_0/q , curve along the centerline of the analyzer. Particles with lower or higher energy are vectored off-axis and not transmitted through the exit aperture.

A particle beam is analyzed by varying either the electric field (i.e. transmission energy) or an accelerating/decelerating voltage V_s with a fixed electric field (constant energy mode) and measuring the fraction of transmitted particles at a downstream detector. The resulting (ion) current-versus-energy plot is

called an ion energy distribution function (IEDF), or energy transmission function. The electrostatic sector is used as a filter by removing the (blocking) collector and allowing the selected ions to pass beyond.

The analyzer will selectively pass ions of energy E/q based on the applied accelerating voltage and transmission energy according to equation (9).

$$E/q = V_s + E_0/q = V_s + \Delta V/C \quad (9)$$

The voltage difference between the plates ΔV , transmission energy, and geometry are related through equation (10), where C is the analyzer's geometrical constant (for spherical segments) [29]. The particle charge is $q = ze$, where z is the integer charge state (1,2,3, etc.).

$$\Delta V = (E_0/q) C = (E_0/q) \left(\frac{r_2}{r_1} - \frac{r_1}{r_2} \right) \quad (10)$$

The field of view of the device can be limited by an entrance and exit collimator placed at each end of the spherical segments. Herein, the entrance collimator was equipped with 0.152 cm diameter orifices separated by 3.2 cm, giving a maximum acceptance angle of ± 2.75 degrees. The exit collimator had one 0.152 cm orifice. The collector of the ESA also acts as a collimator for the ExB section of the probe. The inner and outer segment radii were 4.65 and 5.35 cm, respectively, and the (mean) optic axis radius was 5.0 cm, giving $C = 0.2814$ (the analyzer constant, also called the calibration constant = $1/C = 3.554$).

In electric propulsion applications, the kinetic energy of a particle comprises its thermal energy plus energy gained through acceleration in electric and/or magnetic fields. The thermal velocity is typically small compared to the velocity gained due to the body forces. The velocity of the particle upon entering the analyzer (E_p) is given by equation (11), where we only consider the particle velocities to be much less than the speed of light. The particle energy E_p is in Joules, ΔV_p is the acceleration potential from the plasma thruster in eV, m is the ion mass in kg, and v is velocity in m/s.

$$E_p = q\Delta V_p = \frac{1}{2}mv^2 \text{ or } v = \sqrt{\frac{2q\Delta V_p}{m}} \quad (11)$$

Note that the ESA does not distinguish between singly charged and multiply charged ions born at the same potential (equal ΔV_p) measured relative to the probe reference potential (often chamber ground), as the ions will have the same energy-per-charge ratio when passing through the electrostatic section. They will, however, have different velocities due to differences in mass-to-charge ratio, which is the basis for subsequent discrimination using the Wien filter.

3.2.3 ExB Theory

The ExB sector uses magnetic and electric fields oriented orthogonal to each other and to the axial path of the charged particles, such that a particle stays on axis when the electric and magnetic forces are equal, i.e. $\vec{E} = -\vec{v} \times \vec{B}$, as shown in Figure 4. Here, with krypton singly and multiply charged ions having equal mass and a narrow energy band from the ESA sector, the Wien velocity filter physically separates ions according to charge state. The magnetic field is constant, created by permanent magnets. The electric field is varied via the voltage difference $\Delta\phi$ applied between the parallel plates, generating a current-versus-voltage spectrum with peaks associated with charge state.

$$\frac{\Delta\phi}{d} = B \sqrt{\frac{2q\Delta V_p}{m}} \quad (12)$$

Here, d is separation distance between the plates, and ΔV_p is the potential difference between the ion creation point in the plasma and the magnetic sector of the ExB.




Figure 4. ExB probe schematic showing the electric field and magnetic field separation section. The amount of deflection an ion undergoes depends on its entrance velocity and the applied electric and magnetic field strength.

Charge fractions can be determined from the ExB scan by methods such as integration under each curve, triangle area fitting at full width half maximum (FWHM) or simply determining the ratios of peak heights [38], [50], [81]. In EP literature, it is common to use Omega (Ω) to represent the collector current fractions. For determining the total number of ions for calculating charge state number fractions, the multiply charged particle current measurements need to be reduced by correcting for their charge state.

3.3 Experimental Setup

Although both ESA and ExB probes are commonplace for ion source characterization, the description herein focuses on the use of the EVADER to characterize the ions in the plume of HETs. This testing requires the use of high-pump speed vacuum facilities to produce relevant thruster environments. Common propellants used in the devices are xenon and krypton, although there is growing interest in cost-saving propellants like argon [90] and capability-enhancing propellants like molecular gases [91].

3.3.1 Test Facility

Testing of the EVADER probe was conducted at Colorado State University's (CSU) Center for Electric Propulsion and Plasma Engineering (CEPPE) in the Orion vacuum chamber measuring 1.7 m in diameter and 4.6 m in length. The chamber is equipped with a Leybold DryVac roughing pump, two Varian HS20 diffusion pumps, and two internal cryosails, as shown in Figure 5. The chamber can reach a base pressure of 5×10^{-7} Torr with a maximum pumping speed of 24,000 L/s for krypton. A graphite beam dump, located 2.8 m downstream of the thrust stand, separates the cryosails from direct exposure to

the thruster plume. The four-degrees-of-freedom motion stage was used to align either the EVADER or a standalone ExB probe with the centerline of the thruster being tested.




Figure 5. Schematic of the vacuum chamber used in this study.

3.3.2 Hall Effect Thruster Operation

Two Hall effect Thrusters were evaluated with the EVADER probe, the first a laboratory model 1.5kW Stationary Plasma Thruster (SPT)-style design developed at the CSU CEPPE with a center-mounted hollow cathode and a ceramic borosil (BN-SiO_2) channel [64], [85], [86] (Figure 6). The second is a Busek BHT-200, which has a long history of performance evaluation in EP literature [92], [93], [94], [95], [96], [97]. The 1.5 kW CEPPE laboratory thruster channel dimensions are 104 mm O.D., 17 mm wide, and 32 mm long. Four outer coils coupled with a center coil generate the channel magnetic field, which features a plasma lens topography. The CSU HET can operate at power levels ranging from 0.5 kW to 2.5 kW with krypton or xenon. A nominal operating condition of 1.5 kW and 300 V on Kr was used as the reference operating condition, at which the thruster is capable of 1600 seconds Isp and 35% efficiency.




Figure 6. Photo of the 0.5-to-2 kW laboratory Hall thruster, developed at CSU, operating on krypton.

The CSU HET was tested at three operating points. The “Nominal” and “High-Performance” conditions were in the standard thruster configuration, while the third condition was labeled a “Magnetically Shielded” configuration. These three operating conditions/configurations are discussed by Andreano in [85] with in-depth reporting of thruster performance and phenomenological efficiencies that are summarized in Table 1.

Table 1. Summary of thruster conditions.

Condition/Configuration	Discharge Voltage	Anode Mass Flow Rate	Anode Power
Nominal	300 V	4.50 mg/s Kr	1500 W
Magnetically Shielded	300 V	4.25 mg/s Kr	1425 W
High-Performance	300 V	4.25 mg/s Kr	1615 W

The nominal condition was selected as the reference case because it was the most frequently replicated operating condition in past studies. However, both magnetically shielded and high-performance conditions had subtle effects on the IEDFs collected with the ESA and on standard ExB spectra.

The Busek BHT-200 used in evaluation with the EVADER diagnostic was operated with a separate heaterless hollow cathode, as shown in Figure 7. This was done to demonstrate the heaterless cathode's functionality and to reduce the cycle time from vented-chamber operation to thruster operation. It was placed at a location similar to the existing cathode, but mirrored across the thruster centerline. No performance measurements were made, but initial operating conditions were selected to match those used

in previous measurements. The thruster was operated at three anode voltages: 200V, 250V, and 300V at 250W to match ExB measurements made by others. Similar, slightly lower magnet currents were used in our tests to maintain stability as discussed by Ekholm and Hargus [97].



Figure 7: Photograph of the BHT-200 HET operated on xenon with a channel-aligned EVADER probe mounted on an in-vacuum motion stage.

3.4 Analysis Methodology

Here, we discuss the method for collecting and analyzing data from the combined probe. The differences between standard ESA and ExB measurements that can be made with stand-alone probes are discussed, with attention to the difficulty of analyzing them. The effects of the CEX in the plume are discussed, along with the CEX correction method used to analyze the standard ExB probe and EVADER data.

3.4.1 Standalone Energy and Species Spectra

In using the ESA sector alone, care was taken to first collect a unique IEDF at each condition to characterize the energy ranges that occur and to use as references when discussing ExB spectra. The IEDF data shown in Figure 8 were centered at nearly the same energy measured with respect to chamber ground that were about equal to the voltage applied between the anode and cathode. The collector currents were normalized to the most probable value of each spectrum to allow for easier visual comparison of the IEDF shapes. The nominal and magnetically shielded IEDFs are similar, with the magnetically shielded

configuration having slightly fewer high-energy ions and a larger fraction of lower-energy ions than the nominal case. The so-called high-performance mode exhibited a significantly wider spread of ion energies, with ion populations at energies both higher and lower than the discharge voltage.

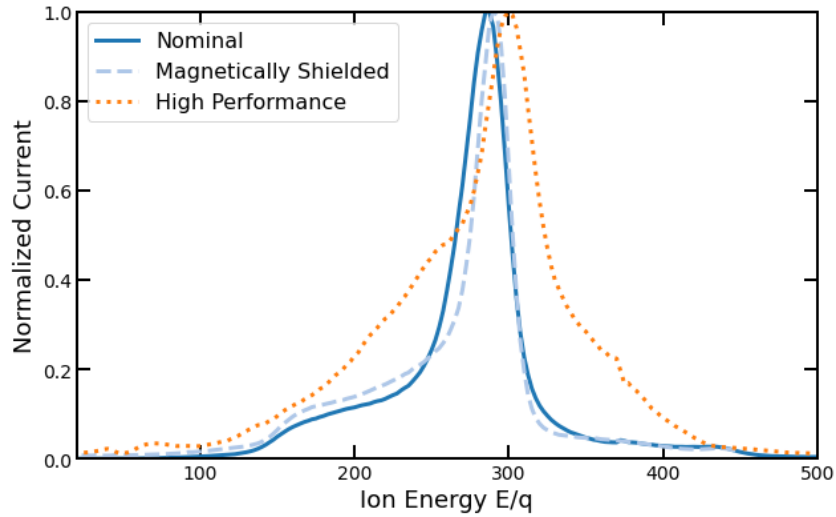


Figure 8. Ion energy distributions at the three operating conditions measured with the ESA sector of the EVADER probe.

The ExB spectra collected for the three thruster operating conditions are shown in Figure 9. These were measured using a standalone ExB probe (without energy filtering), and they reveal substantial qualitative differences. Note that there is visible overlap between the current corresponding to the singly charged ions, signaled by the large peak at ~ 28 V, and the current corresponding to doubly charged ions, signaled by the peak at ~ 40 V, in both the nominal and magnetically shielded conditions.

The overlap is much greater for the high-performance condition, which is a clear indication of how the wide IEDF of a Hall thruster can make analysis of traditional ExB probe spectra difficult. There is an apparent high-energy shoulder to the singly charged ion signal that could also be interpreted as a low-energy doubly charged ions and is likely to be a combination of both. This complicates fitting of VDFs to the ExB spectra for this condition.

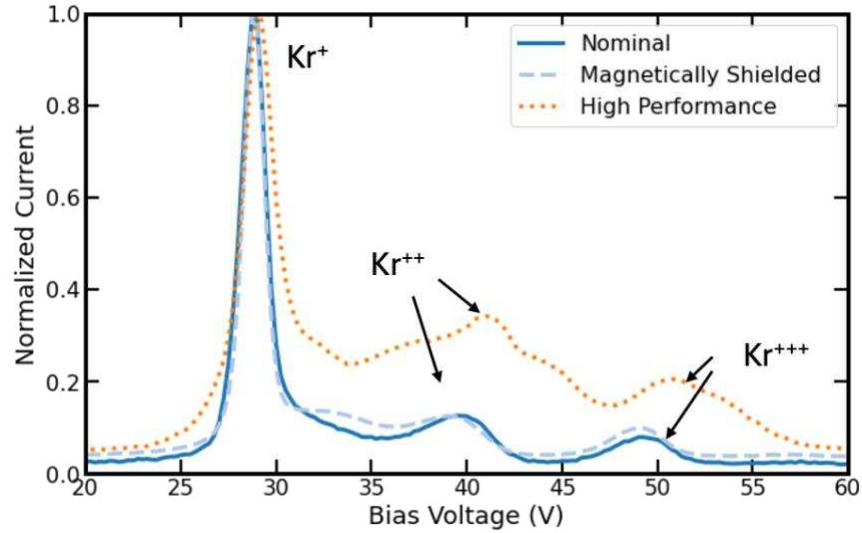


Figure 9. Standard ExB spectra collected for all three thruster configurations without the energy filter (ESA).

3.4.2 EVADER Analysis

First, an IEDF was measured using the ESA sector with the ExB entrance collimator acting as the ion current collector, e.g. those of Figure 8 and Figure 10 for the nominal condition. Discrete pass energies E_1, E_2, \dots, E_k (E/q) were then selected across the domain of the IEDF at which ExB spectra were to be collected, represented by the blue dots in Figure 10.

The transmission energy of the ESA was fixed at 375 eV by setting a bias of 105.5 V across its curved segments; the pass energy (E/q) was controlled by adjusting the collimator voltage. An ExB spectrum was then collected at the selected pass energy by sweeping the plates in the ExB sector and measuring the ion current to the ExB collector.

Importantly, the relative current fractions of species change with each selected pass energy as one would expect, and Figure 10 illustrates this by showing the ExB spectra at the 250 and 290 eV per-charge pass energies. Note that the ExB spectra collected at a given pass energy contain no overlap of species data, in contrast with a traditional ExB spectrum taken without an energy filter for the same operating condition as was shown in Figure 9. The energy filtering step of the ESA allows one to completely distinguish the charge species signals. Most Hall-effect thrusters do not have a significant fraction of ions

with energies below $\frac{1}{2}$ of the discharge voltage, and measurements at those ion energies were only taken at relatively wide ion energy steps for the CSU HET tests.

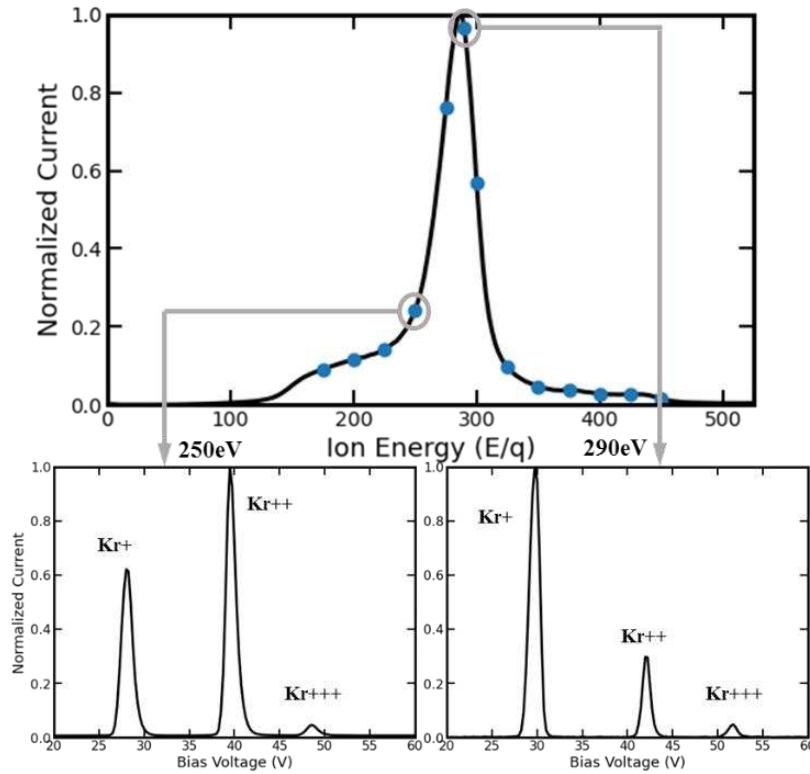


Figure 10. IEDF at the nominal 1.5 kW, 300 V, 4.50 mg/s on Kr condition measured using the ESA section of the EVADER probe. ExB scans at selected pass energies of 250 eV and 290 eV.

In computing current fractions for a given EVADER ExB spectra, the area fraction method in [81] was compared to the sequential VDF fit method of [38]. For a given fitting function, no difference between current fractions was observed when using the two methods, a result predicted by Huang and Shastry [38] when explaining the validity conditions related to area integration methods, where they deemed the area fraction method suitable if ions arriving at the probe had all originated from the same potential. The filter in front of the ExB sector in the EVADER ensures that this condition is satisfied; therefore, the simpler area-fraction method was used to calculate the current fractions at each transmission energy.

To get the total current fractions in the plume, a weighted sum of the current fractions at the selected pass energies is used, weighted by the fraction of the observed current in the plume in a band

around the selected pass energy. We let the weight factor α_i for the section of the IEDF corresponding to the selected energy E_i be

$$\alpha_i = \frac{\int_{E_{i-1/2}}^{E_{i+1/2}} I_{IEDF} dE}{\int_0^{\infty} I_{IEDF} dE} \quad (13)$$

where

$$E_{i-1/2} = \begin{cases} 0 & i = 1 \\ \frac{E_{i-1} + E_i}{2} & 1 < i \leq k \end{cases}$$

and

$$E_{i+1/2} = \begin{cases} \frac{E_i + E_{i+1}}{2} & 1 \leq i < k \\ \infty & i = k \end{cases}$$

These represent the halfway distances in the IEDF domain between the neighboring selected transmission energies (except at the upper and lower ends).

The current fraction $\omega_{z,i}$ of ions with charge z at the i^{th} transmission energy was calculated using

$$\omega_{z,i} = \frac{I_{z,i}}{I_{total}} = \frac{\int_{D_z} I_i dV}{\int_0^{\infty} I_i dV}, \quad (14)$$

where I_i was the ExB spectrum current at energy E_i and where for a given z , $\int_{D_z} I_i dV$ was integrated over the portion of the domain corresponding to ions with charge state z .

Finally, the total current fraction of the species with charge z across the entire IEDF was represented by equation (15).

$$\Omega_z = \sum_{i=1}^k \alpha_i \omega_{z,i} \quad (15)$$

A graphical representation of the analysis process is depicted in Figure 11 for the 300 V nominal condition, and the numerical analysis of relative weights and current fractions is listed in Table 2. The fidelity of the species fraction calculation improves as more pass energies are selected.




Figure 11. Graphical representation of area weight to the current fraction measured at each pass energy in the nominal configuration.

Table 2. Numerical analysis of the nominal thruster condition.

Band	Selected Energy (eV)	Area Weight Factor	Species Current Fractions			
i	E_i	α_i	$\omega_{1,i}$	$\omega_{2,i}$	$\omega_{3,i}$	
0	175	0.07	0.33	0.66	0.01	
1	200	0.05	0.39	0.60	0.01	
2	225	0.06	0.44	0.55	0.01	
3	250	0.11	0.57	0.42	0.01	
4	275	0.24	0.84	0.15	0.01	
5	290	0.21	0.88	0.11	0.01	
6	300	0.13	0.88	0.11	0.01	
7	325	0.05	0.96	0.04	0.00	
8	350	0.02	0.97	0.2	0.01	
9	375	0.02	0.92	0.07	0.01	
10	400	0.01	0.73	0.26	0.01	
11	425	0.01	0.46	0.54	0.00	
12	450	0.01	0.52	0.48	0.00	
Total Current Fraction			Ω_z	$\Omega_1 = 0.74$	$\Omega_2 = 0.25$	$\Omega_3 = 0.01$

3.4.3 Plume CEX correction

The EVADER and standalone ExB probes were located 1.5 m downstream of the CSU HET. The chamber pressure, measured with a gauge behind the thruster at the wall, was near 5×10^{-5} Torr during operation. Background neutrals interact with both the thruster and plasma diagnostics [98], [99], [100].

The basic methods of [81] and [38] were used to correct current fractions for CEX in the ExB VDF-fit and integration analysis, where the attenuation of each species was accounted for by the background neutral density, the charge-exchange cross section, and the distance from the thruster. For the EVADER measurements, the variation in the CEX cross-section with ion energy was used. In the VDF-fit and integration method, the energy for the CEX cross-section was assumed to be the average energy of the IEDF. This results in an underestimation of CEX attenuation for lower-energy ions, as the CEX cross-section is larger at lower ion energies. Similarly, CEX attenuation of higher-energy ions is overestimated.

Hause et al [101] provide singly and doubly charged krypton ion-neutral CEX cross-sections that are within $\sim 15\%$ of the CEX cross-sections of xenon. As a result, this scaling factor was applied to the triply charged xenon cross-section to supplement the gap in krypton cross-section information. Asymmetric CEX reactions could contribute up to 3% error at the pressure ranges and distances of interest, past the $\sim 2\%$ threshold discussed by Shastry and others [81], but were neglected in our study due to a lack of information about these CEX cross-sections for krypton.

3.4.4 ExB Spectra Analysis

The standard ExB spectra of Figure 9 were analyzed using the VDF-fit and integration method proposed by Huang and Shastry [38]. Several fit functions were explored, including triangle, Gaussian, and skew normal. In line with Huang and Shastry, the results varied only slightly across the VDF functional forms, but the skew normal generally provided the best fit, with R^2 consistently above 0.95. As recommended by Huang and Shastry, charge state peaks were fit and subtracted sequentially. Charge-

exchange corrections, discussed above, were performed using the simplified method introduced in [21] using krypton charge-exchange cross-sections from [101].

3.5 Results of Combined Probe Measurements

The combined diagnostic was used to evaluate energy and species distributions of the 1.5-kW HET operated on Kr at two conditions and in a magnetically shielded configuration. This test used a limited number of pass energies for the ESA segment of the probe; as a result, a follow-on study was conducted with a BHT-200 operated on xenon to obtain species specific IEDFs at higher resolution. Comparisons of VDF-fit methods and between the two tests are discussed.

3.5.1 CSU 1.5-kW Test

The following sections discuss the results of testing a 1.5-kW CSU HET on krypton. As mentioned above, these experiments had fewer pass energies selected for the area-weighted IEDF section of the analysis, resulting in limited species-specific energy resolution relative to the follow-on BHT-200 evaluation. The results are compared to the existing VDF-fit methodology applied to traditional ExB probe spectra collected at the same conditions.

3.5.1.1 *VDF-fit Comparison to EVADER Analysis*

The skew normal R^2 values for the VDF fit of the nominal, magnetically shielded, and high-performance operating conditions were 0.97, 0.972, and 0.9717, respectively. The results of both the VDF fits and the EVADER measurements are summarized in Table 3. Mass utilization is affected to a lesser degree and involves more information about the beam current; this was reported in detail in [85]. Finally, the equations for both phenomenological efficiencies are covered in [81].

In all conditions, a decrease in the fraction of triply charged ions was observed via EVADER analysis. In the nominal and magnetically shielded cases, which had similar IEDF shapes and typical ExB spectra, the current fraction of the doubly charged ion current was significantly higher. The high-performance condition, which showed a large discrepancy in the singly and doubly charged ion current

fractions, exhibited markedly different IEDF and ExB spectra. In all cases, the impact on charge utilization is on the order of 0.1-0.2%, which is well below the resolution of overall thruster efficiency. We note that the impacts of the species fraction propagate out to both current utilization and to some degree, mass utilization [34]. As mentioned in the introduction to this chapter, the more consequential result of more accurately constructed species-specific IEDFs is the plasma condition they imply must exist in the thruster discharge [17], [102].

3.5.1.2 Ion Energy Distribution of Species

An advantage of the EVADER probe is that ExB spectra collected at a sufficient number of pass energies along the overall IEDF can be used to approximate species-specific IEDFs. The sum of these should resemble the overall IEDF collected for all species in the ESA sector. Figure 12 shows a plot of the current fractions at each selected pass energy for the high-performance condition shown with the summed current fractions, as well. These help clarify the discrepancy between the tabulated current fractions obtained from the VDF-fit and EVADER methods.

Table 3. Comparison of EVADER and VDF-fit analysis current fractions.

Operating Condition	Measurement Method	Kr⁺ Current Fraction - Ω_1	Kr⁺⁺ Current Fraction - Ω_2	Kr⁺⁺⁺ Current Fraction - Ω_3	Charge Utilization
Nominal 300V – 4.5 mg/s	ExB – VDF Fit	0.766	0.184	0.050	0.978
	EVADER	0.738	0.252	0.010	0.980
Magnetically Shielded 300V – 4.25mg/s	ExB – VDF Fit	0.767	0.197	0.037	0.980
	EVADER	0.746	0.245	0.009	0.981
High Performance 300V – 4.25mg/s	ExB – VDF Fit	0.378	0.574	0.048	0.967
	EVADER	0.625	0.336	0.035	0.969




Figure 12. Current fraction contributions, by each ion species as a function of selected pass energy, for the high-performance operating condition.

In Figure 12, there is a notable contribution of the singly charged ions with energies above 300eV (near the discharge voltage), which, when considered with the wide and lower energy distribution of the doubly charged ion current, results in a significant overlap of the singly and doubly charged ion current peaks in the standard ExB spectrum. Only with this species-specific IEDF information can the constituents of the standard ExB spectrum's singly charged shoulder be properly deconvolved. In this case, a 65% increase in the singly charged ion current fraction and a 41.3% decrease in the doubly charged ion current fraction resulted when comparing the VDF and EVADER analyses.

The ExB spectrum for the high-performance condition was particularly unique. Even if a purposefully over-wide VDF was applied to the singly charged peak (with a reduced R^2 value), the resulting singly charged ion current fraction was still lower than that given by EVADER. Such discrepancies were less pronounced for the nominal and magnetically shielded cases, where the ExB spectra were more typical.

The nominal and magnetically shielded IEDFs were less refined than the high-performance condition due to a reduction in the number of selected pass energies in our demonstration testing. For instance, the magnetically shielded condition had the first selected pass energy at 225eV, to represent the current fraction of the entire low-energy shoulder in the overall IEDF, see Figure 8. That area, with a

weight factor of $\alpha_i > 0.25$, is likely to cause an artificially high current fraction contribution at that singular energy, as seen in Figure 13, with a significant contribution at 225 eV per charge.




Figure 13. Current fraction contributions for the magnetically shielded condition.

Although there were fewer selected pass energies at the nominal operating condition, the current fraction contribution with respect to ion energy, as shown in Figure 14, has a shape similar to that of the ESA-sector-collected IEDF in Figure 8. Both magnetically shielded and nominal operating conditions showed a higher fraction of doubly charged ions in the EVADER measurement than in the VDF-fit and integration analysis. In the nominal condition, this could be explained in part by a substantial contribution from low-energy doubly charged ions that have the potential to widen the singly charged ion current peak in a standard ExB spectrum.




Figure 14. Current fraction contributions for the nominal condition.

Using the ion separation equation (12) for the ExB probe and applying dimensional and magnetic field specific information to the ExB probe, we could calculate that the low-energy doubly charged current (~ 175 eV per charge) and singly charged ion current with energies around 300 eV per charge will have a < 2 V separation on an ExB spectrum. Furthermore, any singly charged ion current with an energy of 350 eV will appear at the same probe voltage as this lower energy doubly charged ion current.

3.5.1.3 BHT-200 Testing

The BHT-200 was operated at three discharge voltages: 200V, 250V, and 300V at a power level of ~ 250 W. In previous studies, these operating conditions have produced plumes with varying species fractions [97]. A major goal of this testing was to evaluate the impact of increasing the number of selected pass energies for which to collect a subsequent ExB spectrum. Figure 15 shows an example of the number of pass energies for which ExB spectra were collected projected onto the IEDF of a BHT-200 operating at 250V. This stands in contrast to the number of pass energies selected for the CSU HET testing shown in Figure 11.




Figure 15: Representation of each pass energy where the ExB probe trace was taken. The area of each smaller section is used to weight the species fraction calculated at that area.

Figure 16 shows the IEDFs collected with the ESA sector of the EVADER at each anode voltage. The complete overlay of all ExB spectra collected along the IEDF of the 250V anode voltage condition is

shown in Figure 17, along with the same complete overlay plot corresponding to the 200 V anode voltage condition.




Figure 16: IEDF collected with the first stage of the EVADER at the three operating conditions explored for the paper.

The shifting location of the peaks is easily discernible as the selected energy changes. This figure, although very busy, explicitly shows the need to replace fitting traditional ExB spectra with individual fits for each species and each energy band, and also highlights the complexity involved in doing so. For example, the same plot of the 200V condition ExB traces, shown in Figure 17 indicates a high energy shoulder of singly charged ions is present that overlaps with the lower energy distribution of doubly charged ions that would be nearly impossible to discern with a standalone ExB probe.




Figure 17: ExB spectra collected at an incremented pass energy of the ESA sector every 5eV from 100eV-300eV for the 250V anode voltage (left) and the 200V anode voltage condition (right).

The overlapping of the species IEDFs can be better visualized with a contour plot of the collector current with ESA pass energy and the bias voltage applied to the ExB sector, as shown in Figure 18. The location of the singly and doubly charged xenon ions moves with the energy of the incoming ions. The tails of these distributions can often lead to overlapping distributions previously discussed (Figure 17).

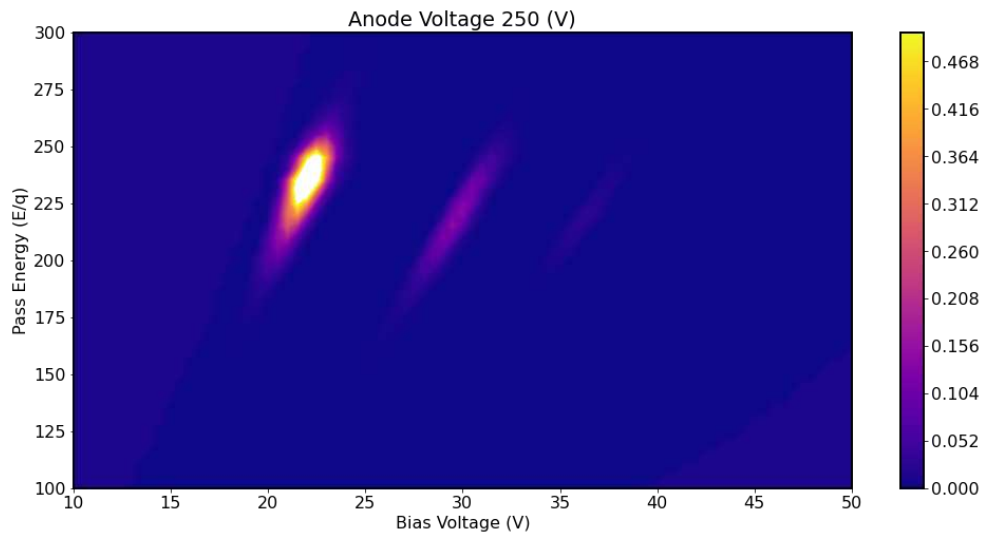


Figure 18: Contour plot of collector current of the EVADER probe with respect to the ESA pass energy and the voltage applied to the ExB sector, peak normalized to one with scale change for clarity.

An ExB spectrum is collected at each pass energy set by the ESA, enabling species-specific ion energy distributions to be reconstructed over the energy range swept by the ESA. These species-specific ion energy distributions for each thruster operating condition are shown in Figure 19 with the primary ordinate representing the current fraction contributed by each species at a given energy (the abscissa). These contributions are summed together at each energy and compared to the normalized IEDF shown on the secondary ordinate. These summed current fraction contributions accurately recreate an IEDF specific for each ion species.

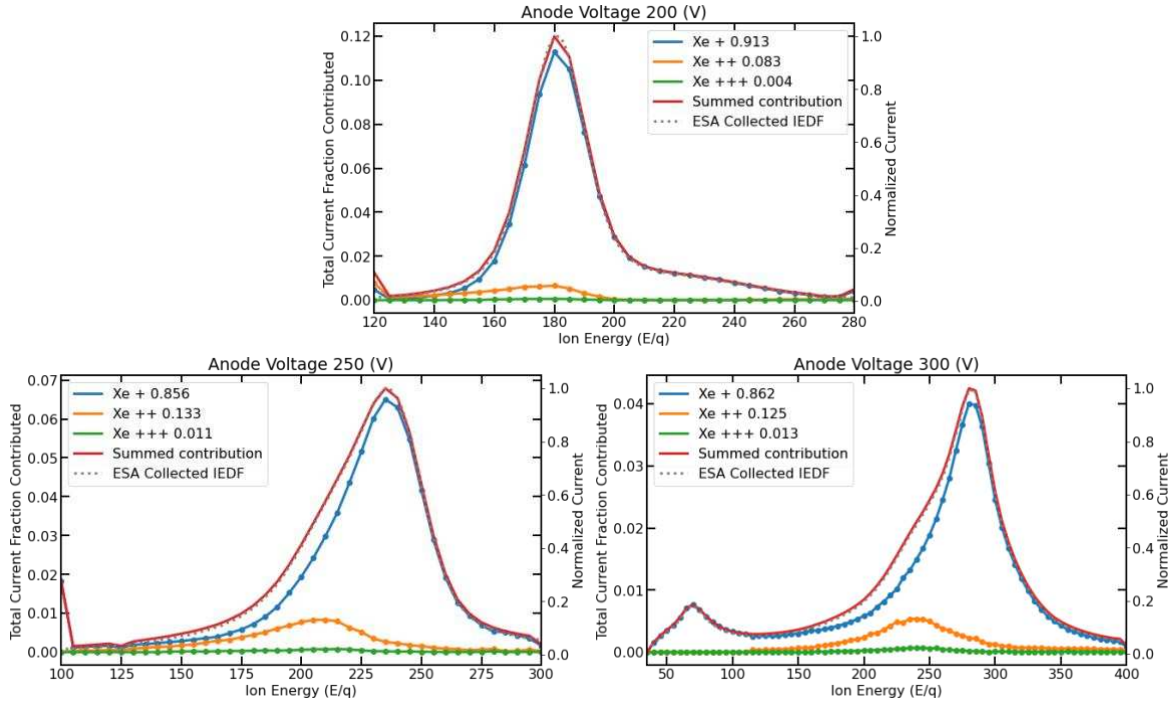


Figure 19: Plot of species-specific ion energy distribution with the summed contribution (total IEDF) outlined over the top for 200V (top), 250V (left) and 300V (right).

3.6 Discussion

The combined energy and velocity filter diagnostic has been shown to be particularly effective at deconvolving overlapping velocity distribution functions arising from wide IEDFs between multiply charged ion species in a HET plume. The combined diagnostic can therefore better estimate the correct species fraction, which, apart from a minor impact on efficiency analysis, provides more detailed measurements of the species fractions generated by the thruster discharge plasma. Furthermore, energy filtering performed prior to ExB probe spectra collection enables measurement of species-specific ion energy distributions. With a sufficient number of pass energies selected, these distributions can also be obtained without the need to assume a specific shape to the VDF or the IEDF. These two improvements in species-fraction identification and species-specific ion energy distributions provide more information about plasma discharge properties than can traditionally be obtained from a stand-alone ESA or ExB, especially when the ion energy distributions, in terms of energy-per-charge, differ significantly. As shown in the BHT-200 measurements for example, there is a disparity in the energy distribution of doubly charged ions relative to the IEDF of singly charged ions.

Different E/q ion energy distribution profiles imply information about the plasma potential in relation to the generation of the ions being measured. Many models assume equipotential magnetic field lines in certain regions of the discharge [103], and, as a result, a wider ion energy distribution of doubly charged ions with a lower average indicates a larger spatial domain in which the ions are generated. Outside of these assumptions, small-scale physics concerning the coupling of instabilities has also predicted disparity of the IEDF between doubles and singles [71]. Depending on the format and fidelity of the modeling used, the species-specific information can be used, as discussed previously, in more intensive data-driven or data-fusion frameworks to better limit the available state space for given model properties, and, in frameworks with even larger datasets, may fill a crucial role in informing such models.

It is also important to note that, as these data are implemented in data-driven frameworks, a key metric to report with such measurements is the uncertainty. Historically, only the fit-parameter uncertainty, or the random error of the measurement instrument has been reported, but these do not account for systemic errors that may result from probe effects, *e.g.*, fringe fields, finite energy pass-band widths, off-axis ions, etc. Although it is common to compare ion energy and species measurements made in the probe to a well-characterized ion source, such as an ion beam. This does not provide a full calibration of what may be shifted in a wide-ranging energy source with multiple species (typically at higher species fractions than observed in such ion sources). As a result, the measurement uncertainty of the EVADER probe might be possible to estimate from models that capture such probe effects that can be operated in reverse to correct collected spectra. We note that efforts have begun in attempting to account for the measurement uncertainty of the EVADER probe [104].

3.7 Summary

Initial work using existing electrostatic probe techniques, such as ESA and ExB probes, has demonstrated increased accuracy in species estimation, particularly when multiply charged ion species have different energy distributions, as observed in Hall effect thruster plumes. In addition to improved species fraction identification, the use of sufficient pass energies to collect energy-per-charge filtered ExB

spectra can generate species-specific ion energy distribution functions that don't require assumptions about the VDF-fit. This improved measurement specificity provides previously inaccessible information from a non-perturbative measurement made downstream of the plasma device. The improved measurement specificity lends itself well to thruster modeling efforts, as it limits the possible state space for the discharge properties required to produce such species and energy distributions. The details of the integration are model-dependent, but in all cases, will benefit from accurate estimation of the measurement error. Efforts are currently underway at the CEPPE Lab to address this need.

4 TEMPORAL RECONSTRUCTION OF PLASMA PROPERTIES VIA SHADOW MANIFOLD INTERPOLATION

Following the theme of improving the information gained from commonly employed diagnostic techniques for plasma thruster characterization, we can investigate the ability to add temporal dimensionality to traditional measurements. Hall thrusters exhibit rich temporal behavior; they are known to exhibit a large number of oscillations and instabilities, which are fundamental to their operation [49], and the relative importance of these oscillations and instabilities in producing accurate predictive models is an active area of research [19]. Many of these oscillations can result in variations in plasma properties by nearly 100% of the steady-state values [22], [105]. In addition to HET discharge oscillations, a significant number of cathode instabilities influence their time-dependent and steady-state behavior [106], [107], [108], [109]. As a result, the coupled system can exhibit a slew of non-linear dynamical phenomena that are difficult to capture properly in steady-state plasma models. By augmenting existing diagnostics with temporal information, greater model parameter specificity can be achieved in data-driven frameworks [22], [26], [27]. However, due to the approach used in many diagnostic measurements, temporal measurements cannot always be simply achieved by increasing the rate at which current, voltage, or light intensity is measured. This is often because many diagnostics are measurement-state-dependent, meaning the probe measurement state is swept through voltage, current, position, or wavelength to build a distribution across the states represented by each. Practical limitations of how fast the diagnostics states can limit measurement speeds well below the frequencies of interest. Methods for overcoming these challenges and their limitations are discussed, and an example application of a nonlinear reconstruction technique, Shadow Manifold Interpolation (SMI), is used to obtain time-resolved ion current density distributions from Faraday probe measurements.

4.1 The Temporal Information Gap

As mentioned above, often diagnostic tools used in EP device characterization are scanned through a series of measurement states to build a measurement-state distribution. As discussed in Chapter 4, for ESA and ExB probes, the measurement state distribution is a function of E/q or the applied plate voltage, which can be translated into a velocity for a given species. The scan through measurement states is accomplished by applying consecutive voltages to the internal probe surfaces, and the current measured at each voltage is used to determine the quantity of ions in that discrete distribution bin. This process is not unique to electrostatic energy and velocity probes, as laser-induced fluorescence (LIF) and optical emission spectroscopy (OES) scan wavelengths, while Faraday probe measurements are made at many angles relative to the thruster centerline to provide spatial information. As a result, because successive measurement states must be scanned to obtain the desired distribution, most measurements are necessarily steady-state or DC. As mentioned above, EP devices, on the other hand, often exhibit highly dynamic, rich, and often non-periodic oscillatory behavior.

The “breathing mode” oscillation in HETs is one of the most studied and characterized due to its ubiquity and the accessibility of measurement frequencies (10-30kHz). The breathing-mode oscillation is typically detected and characterized by temporal measurements of the discharge current, yet even its lower bound of 10 kHz is often well beyond the scan speed of many diagnostics. Furthermore, if spatial and measurement-state information is required, a further complication of the fast measurement technique is how to provide temporal measurements at all spatial locations in a synchronous way that can be used for device characterization and model validation.

To address the challenge of spatial and temporal synchronization, many approaches have been used. Sample-and-hold techniques, based on zero-crossing of the discharge current oscillation [110], [111], and heterodyning of laser light [112] were used in time-resolved measurements around the breathing mode oscillation. Although these can provide relatively good temporal resolution and signal-to-noise, they require a highly periodic breathing mode; however, the breathing mode does not often exhibit

strong coherence around a single frequency. Rather, it often exhibits jitter and even non-periodic behavior [22], [105]. In 2010, to address this problem, Lobbia pioneered the use of an FFT-based transfer-function method to fuse temporal discharge-current measurements with temporally synchronized Langmuir-probe measurements at several locations in the downstream plume. In this way, he was able to reconstruct the time-varying oscillations of plasma properties measured downstream of a HET [40]. The spatiotemporal map of the plasma ejecta showed rich variations in the downstream plasma throughout the HET discharge current oscillation. Work by Durot demonstrated the potential for this temporal reconstruction technique with LIF [41], and it has been used more recently to characterize the temporal IVDF measurements of the NASA HERMeS 12kW thruster [113].

The fundamental difference of Lobbia’s technique is that each measurement state and location can be produced as a function of an input signal, which is measured synchronously with each independent state and location measurement. After this “construction” phase, a typical single signal input is used to “reconstruct” all state and location measurements in a pseudo-simultaneous temporal manner. Although the transfer function approach can handle non-periodic data, there are a couple of challenges in implementing it for all temporal and state-based measurements. First, a significant amount of data is needed to produce clean reconstructions. Although data recycling can be implemented, the data files themselves can grow quite large. Second, there remain questions about how well artifact-prone, FFT-based methods can capture non-linear behavior observed in the dynamics of the HETs [75], [114].

In 2019, Eckhardt *et al.* proposed another method for nonlinear spatiotemporal reconstruction [63], inspired by the cross-convergent mapping method demonstrated by Sugihara *et al.* [64], leveraging Taken’s theorem [65], to map disparate signals of a unified dynamical system to one another. This was done by Sugihara *et al.* to determine causality between dynamic signals; however, this state-space reconstruction method, used as part of the cross-mapping process, can also handle quasiperiodic nonlinear temporal data. The method put forth by Eckhardt *et al.*, described as shadow manifold interpolation (SMI), was shown to improve temporal reconstruction performance compared with FFT-based transfer

function methods for reconstructing temporal beam-dump and floating-potential signals collected in an experimental HET test. Specifically, the SMI-based technique required a smaller data file size (record length) and did not produce non-physical artifacts in the reconstructed signals. Both FFT and SMI techniques have been recently applied to retarding potential analyzer (RPA) measurements of an ion source operated with a time-varying beam voltage [115], [116] and to sub-kilowatt and multi-kilowatt Hall thrusters [117], [118]. This method has been extended further to LIF [119] using the same data collected in [113].

In the sections below, the process for performing SMI will be outlined, along with a discussion of key considerations during the process and the approach's limitations. The technique is then demonstrated using Faraday probe measurements to reconstruct the spatiotemporal profile of the ion current density of the 1.5-kW CSU HET described in Chapter 1.

4.2 Experimental Configuration

The thruster used for this study was a modified version of the 1.5kW CSU HET described in Chapter 1, operated with an external hollow cathode. In this study, the HET had been modified to enable single-power-supply operation. Because of these changes and the use of a lower than optimal magnetic field strength, the thruster was operated in a condition that led to large discharge current oscillations, with peak-to-peak currents that were nearly 400% of the steady-state current. Details of the thruster operating in this regime are discussed in [120]. The Orion vacuum facility, used for testing, was also described in Chapter 1. A graphite Faraday probe was used to measure current density. To generate the data needed for the temporal reconstruction, a high-speed measurement of the diagnostic at a given measurement state must be collected synchronously with an input signal containing the mutual dynamics. In this example and in previous studies [116], [118], [119], discharge current was used as the input signal, collected with each high-speed measurement state signal. A low-voltage differential probe with a 50 MHz bandwidth was used to measure the voltage across a 1 k Ω resistor in series with the Faraday probe collector to capture the high-speed ion-current signal to the probe when it was placed at each angle of its sweep range.

For the common input signal collected synchronously with the ion current, a 1 MHz Bandwidth (AC + DC) current probe was used to capture the anode current oscillations. All measurements were collected at 50 MHz and downsampled to 10 MHz for data processing.

4.3 Shadow Manifold Interpolation

Shadow manifold interpolation relies on a state-space reconstruction method that was leveraged in cross-convergent mapping work by Sugihara *et al.* [78]. In cross-convergent mapping (CCM), one signal is used to reconstruct another signal suspected to be causally related in a the dynamic system. A time-delay embedding is generated from one signal and mapped to the other. If the signals are causal members of the same dynamic system, then by virtue of Taken's theorem [76], the dynamics of the system are recoverable from the time delay embeddings of each of the variables. In CCM, the ability to reconstruct this signal with increasing accuracy as more dynamic data are included is taken as a sign of the causal relation between the variable and the other variable in the system. This sets aside special cases of strong unidirectional forcing and indirectly driven variables that give rise to a phenomenon known as *generalized synchrony*. An extension of the technique for improving the detection and identification of the generalized synchrony cases was later described by Ye *et al.* [78], which is discussed in further detail in Chapter 5.

The process of reconstructing one signal from another using CCM can be used to reconstruct multiple signals from a common input signal [42]. If a measured dynamic signal is accurately reconstructed by the reference signal and the measured signal is collected across all measurement states of interest, then a temporal signal reconstruction can be performed simultaneously across all measurement states. This provides a means to generate a spatiotemporal reconstruction of the measurement distribution relative to a singular input signal. The process of cross-convergent mapping is discussed next, which is relevant to the SMI process discussed in this chapter but also in the discussion of causality in Chapter 5.

4.3.1 Signal reconstruction used for cross-convergent mapping

The premise of CCM relies on nonlinear state-space reconstruction (SSR) that leverages Taken's theorem [76]. To study if one signal $x(t)$ is causal on another signal $y(t)$, the process attempts to reconstruct the dynamic signals of the time series $x(t)$ from a delayed embedding (which is sometimes called a shadow manifold) of the timeseries $y(t)$. If this reconstruction can be performed with increasing accuracy, as measured by the Pearson correlation coefficient, ρ , between the reconstructed and real time series, when more data are included in the analysis, then it can be understood that x is causal on y .

The basic flow of the CCM process is outlined below. Time-resolved signals are collected synchronously for all variables of interest. The signals are split into test-train segments, i.e., $x(t = 1, \dots, N) = [x_{train}(1, \dots, N_{train}), x_{test}(N_{train} + 1, \dots, N)]$ and $y(t = 1, \dots, N) = [y_{train}(1, \dots, N_{train}), y_{test}(N_{train} + 1, \dots, N)]$, where N is the number of total timesteps and N_{train} is the number of timesteps for the training data. The test signals are reconstructed using the shadow manifold, which will be defined shortly, and the reconstructed signals are compared with the original test data signals to assess the reconstruction quality. The reconstruction is performed by considering a shadow manifold with a reduced number of signals: $M_y(t) = [y(t), y(t - \tau), \dots, y(t - (m - 1)\tau)]$, where τ is the time delay and m is the number of embedding dimensions. The time-delayed phase portrait was first introduced to the EP community by Greve et al. [121] as a tool for model calibration of nonlinear dynamical systems.

Shadow manifold reconstruction is a two phase process, broken up into a “training phase” and a “test phase”. During the training phase the signals that were simultaneously sampled $x(t)$ and $y(t)$ are split in the middle following Eckhardt *et al.* [42] so that the first half may be used for the training, while the second half is used for test. The training data signals are used to create the time delay manifolds following the process above for $M_{y,train}$ and $M_{x,train}$. Together these constitute a nearest-neighbor library for any point in the input signal's state-space we can look up a corresponding point in the output

signal state-space. This is a non-parametric mapping, unlike transfer function methods, no model parameters are fit, the training data itself acts as the model.

In the test phase for SMI only the test input is used for reconstruction, where near neighbors of points input manifold are used to find near-neighbor points in the output manifold's state-space. This is done at each time step of the test data signals, the time-delay embedding test data manifold, *i.e.*, $M_{y,test}(t = t_0) = [y_{test}(t = t_0), y_{test}(t = t_0 - \tau), \dots, y_{test}(t = t_0 - (m - 1)\tau)]$. In addition, the train data manifold is constructed as $M_{y,train}(t = t_1) = [y_{train}(t = t_1), y_{train}(t = t_1 - \tau), \dots, y_{train}(t = t_1 - (m - 1)\tau)]$. Then, we find a finite number of t_1 values, *e.g.*, $m + 1$ values [122], of the train data manifold that result in the smallest Euclidian distance between $M_{y,test}(t_0)$ and $M_{y,train}(t_1)$. These datasets are called the near neighbors. Now, a shadow manifold of x is constructed as,

$$\hat{M}_{x,test}(t_0) = \sum_{i=1}^{m+1} w_i M_{x,train}(t_i), \quad (16)$$

where $w_i = u_i / \sum_i^{m+1} u_i$ is the normalized weight, $u_i = \exp\left(-\frac{d_{y,i}}{d_{y,1}}\right)$ is the weight, and $d_{y,i}$ is the Euclidean distance between $M_{x,train}(t_0)$ and $M_{x,train}(t_i)$. Note that $i = 1$ is the nearest neighbor. Here, the constructed shadow manifold can be considered to have m discrete numbers such that $\hat{M}_{x,test}(t_0) = [x_{sm}(t_0), x_{sm}(t_0 - \tau), \dots, x_{sm}(t_0 - (m - 1)\tau)]$, where the x_{sm} values are estimated from Eq. (16). Finally, it should be noted that the x signal can be reconstructed, which we call $\hat{x}(t)$, by taking an average over m values of $x_{sm}(t)$. The cross-mapping skill is evaluated using the Pearson correlation coefficient ρ between the reconstructed signal $\hat{x}(t)$ and the test signal $x_{test}(t)$.

In SMI the test-train split serves two purposes. This first is that it serves as an honest evaluation of the reconstruction quality. The Pearson correlation coefficient (ρ) between the reconstructed signal $\hat{x}(t)$ and the unseen test data of $x_{test}(t)$, reflecting the processes ability to generalize rather than memorize. The second is that the success of the reconstruction acts as an additional test of stationarity. If the

underlying dynamics between the test and training data had changed, then the nearest-neighbor mapping would degrade. Eckhardt *et al.* found that reconstructions could be sensitive to data length [42], so sufficient data must be used to appropriately sample the attractor.

4.3.2 Reconstruction process for ion-current measurements

As discussed above, the temporal reconstruction process attempts to reconstruct the high-speed signal of a given diagnostic, collected at each measurement state, using a common input signal across all measurement states. In the case of Faraday probe measurements for HET characterization made at a fixed radius from the thruster origin, the measurement state is the angle at which the ion current is measured across a spatial distribution (e.g., from -90° to $+50^\circ$). This ion current density distribution is related to key performance metrics of the Hall thrusters [1], but is almost always reported as a time-averaged parameter, due to the need to slowly scan the probe through the thruster plasma plume.

To collect the desired state-dependent high-speed signals, an oscilloscope was used to collect synchronous anode current and ion-current density measurements. The ion current density reconstruction follows the approach of Eckhardt *et al.* [42] and generates a test-train split at the midpoint of the waveform. The training manifold is constructed from the first half of the synchronous discharge-current and ion-current waveform pair, while the second half is used exclusively for reconstruction and evaluation. The blue plots in Figure 20(a) and Figure 20(b) show measured ion current density temporal behavior and discharge current oscillation over a period of $250\mu\text{s}$ at two different measurement angles, while Figure 20(a-1) and Figure 20(b-1) show comparisons of the reconstructed ion current density to the test data for 0° and 30° , respectively. The Pearson correlation coefficient is used to gauge the accuracy of the reconstructed time-varying signal compared to the measured signal. This works well when the signal far exceeds noise, as shown in ion current measured at 0° shown in the left plot of Figure 20, where the correlation coefficient $\rho = .96$, but in cases where noise makes up a significant portion of the signal shown in the ion current measurement at 30° , the Pearson coefficient can drop significantly, $\rho = 0.56$, even though the qualitative comparison between the reconstructed and measured signals may appear good. This

has implications in denoising and signal decomposition, which have been discussed for SMI elsewhere [123]. However, after applying a small moving average of 10 data points to the lower-SNR traces at 30° and higher, we observed improvements in the Pearson correlation exceeding 0.7, even for the off-angle measurements. Another way to improve data collection would have been to utilize a low-noise high-speed current measurement, especially in the regions of lowest current density.

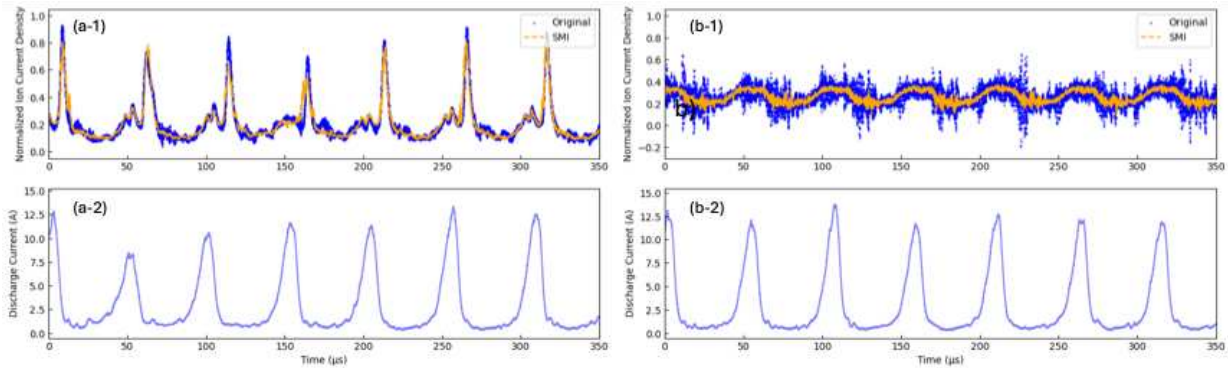


Figure 20: Example of SMI reconstructed signal to the actually measured signal of ion current density at 0° (a-1) and 30° (b-1), along with the input signal, anode current, used for the reconstruction (a-2 and b-2, respectively).

The embedding parameters, both time delay τ and dimension m , play a critical role in the accuracy of the signal reconstruction. The first parameter to evaluate is the time delay parameter τ , which has received much attention in the time series analysis literature [75]. Mutual information of delayed signals is a common basis for finding a preferred delay value [124]. Recent work has extended this approach by using transformed coordinates to more evenly spread the attractor defined by the delayed signal series [125]. For the work presented on ion-current density reconstruction, this method was not employed; instead, we used a mutual information measure based on a fixed bin size, which yielded a delay of 8.6 μs , which is in line with delays used by others to create time delay phase portraits of Hall thrusters [114]. In Chapter 5, the limitations of this selection are demonstrated, as is its influence on signal reconstruction in the work presented. The dimensionality was selected using the modified false nearest neighbors approach proposed by Cao [126]. This method has been used in many other reconstruction and cross-convergent mapping works using HET test data [42], [118], [127]. The Cao method resulted in the selection of 10-12 dimensions, although the diminishing distance and false nearest-neighbor points made

it difficult to determine whether saturation was fully reached, suggesting that increasing the dimension may have fully saturated the false nearest-neighbor detection metric. Reconstruction accuracy remained quite high for the train-test splits that were evaluated with the measurements with dimension selections of 10 to 12; and so the same embedding dimension of 10 was used for all reconstructions of ion current to reduce computation time.

4.4 Spatiotemporal Ion Current Density Reconstruction

The ion current to the Faraday probe was measured with an oscilloscope while simultaneously measuring discharge current oscillations of the thruster with the Faraday probe parked at a single angular position. This measurement was then repeated at each angle from -90° to 45° to build a standard current-density profile. The normalized, time-averaged current density profile is shown on the left in Figure 21 with the transient current density measurements of 0° , 15° , and 30° shown on the right. These transient ion current density measurements at fixed probe locations are shown over the same $250 \mu\text{s}$ period the discharge current waveform (bottom right plot of Figure 21). As discussed previously, the sub-nominal magnetic field condition used to operate the thruster produced large oscillations in discharge current, which in turn caused significant fluctuations in ion current, especially along the centerline (0°). Ion current density signals measured at large angles ($>30^\circ$) began to become more impacted by noise as the angular position was increased.

Train-test splits were applied to each pair of signals after measuring the transient ion current density and discharge current at each angle, and SMI was performed to reproduce the known ion current from the test split in the data. As discussed previously, the higher-angle measurements (corresponding to lower current density) were subject to greater relative noise, although qualitative assessments of these reconstructions appeared to perform reasonably well. The Pearson correlation coefficient for reconstructed test ion current density signals at near-centerline angles $\pm 15^\circ$ was greater than 0.94. This measure dropped to below $\rho = \sim 0.5$ for angles greater than 30° , even though the angular profiles remained similar to those expected in a “denoised” ion current density signal.




Figure 21: Time average current density profile (left) created by averaging current density signals measured at each angle. Three temporal plots of the transient ion current density measurement at 0°, 15°, and 30° as well as a representative discharge current signal are shown on the right across a 250 μs period.

After confirming that relatively accurate reconstructions of the ion current density profile could be achieved with each test-train pair, a representative discharge current signal was taken from the 0° measurement pair and used to reconstruct the time-dependent ion current density at each angle across the entire length of the input discharge current using the full measurement pairs as the training data for each angle. The dimensions and time delays were kept constant across all reconstructed signals and test-train splits. The reconstructed ion current density profile at each angle over a 250 μs period of the input discharge current is shown in Figure 22. The largest peak in ion current appears at a delayed time corresponding to the transit time of a singly charged ion with an energy of ~270 eV traveling roughly 75cm from the thruster to the Faraday probe face, accounting for the disparity in ion velocity peak and current peak that are often assumed to be out of phase [128]. The reconstruction also produces two other distinct peaks in the current density measurements that appear before the larger ion current density signal, as seen in the individual traces for 0° and 15° and shown in Figure 21. These peaks correspond to the transit-delay difference expected for doubly and triply charged krypton ions, with slightly lower E/q distributions (~250 E/q).

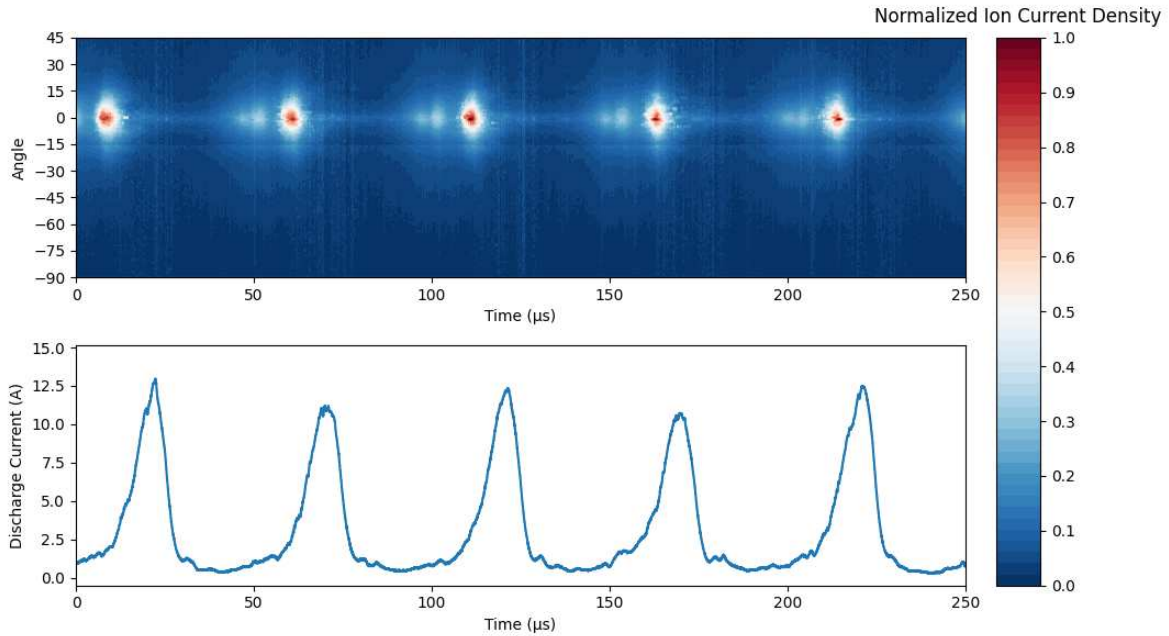


Figure 22: Spatiotemporal reconstructed ion current density profile relative to the input discharge current measurement across a $250\mu\text{s}$ period.

4.5 Discussion

The use of SMI to reconstruct the spatiotemporal maps of ion current with a Faraday probe allows for a more complete picture to be constructed of the temporal behavior of a Hall effect thruster. The spatiotemporal map of the Faraday probe shows that the strongly oscillatory off-nominal operation of the HET at 1 kW leads to significant oscillations in the ion current density measured 75 cm downstream of the thruster. The delay between the discharge current peak and the peak of the ion current density is close to the transit time of a $\sim 275\text{eV}$ krypton ion traveling 75cm between the thruster exit plane and the Faraday probe collector. Proceeding these peaks in the spatiotemporal current density plot shown in Figure 22 are small groups of ions at lower current density with transit times commensurate with that of doubly and triply ionized krypton ions traveling at slightly lower energy per charge ratio. This is consistent with the commonly accepted current fractions of singles and doubles previously reported for this thruster, albeit at higher power and under different magnetic field conditions [129]. This feature of the temporal ion current density measurement may be argued to contain information about the ion species leaving the thruster if measurements are made at a known distance, similar to time-of-flight (TOF)

diagnostics but without the onerous gate timing electronics needed to operate TOF-based instruments. If the temporal measurements are fused with other energy, or especially velocity measurements (e.g. ExB or LIF), we believe one would be able to further reduce uncertainty in the measured quantities.

There are practical experimental challenges in using these approaches to produce high-quality temporal reconstructions of existing diagnostic measurement systems. One of the first is the measurement-state-specific high-speed signal noise, specifically the signal-to-noise ratio (SNR), present in the lowly populated distribution states. For example, at ion current measurements made at angles $>45^\circ$ that were strongly influenced by noise. Where steady-state measurements benefit from increased averaging over time, higher-speed measurements need to contend with common-mode and other circuit noise introduced by making the measurement at high-speed. Furthermore, although these techniques may improve the information provided by existing diagnostic methods, they may still require modifications to the diagnostic methods to achieve the desired level of accuracy for temporal measurements. For example, line capacitance, resistor size, SNR, feedthrough common-mode coupling, and scope/probe accuracy and bit depth are additional considerations that need to be evaluated to ensure that the high-speed signal used for reconstruction is sufficiently well resolved. The work presented herein indicates that using 12-bit-depth (or better) oscilloscopes with high-CMRR differential probes when possible will significantly improve reconstruction quality through better near-neighbor identification.

There are additional challenges to address when implementing SMI to produce spatiotemporal measurement reconstructions. One challenge is selecting embedding parameters for the shadow manifolds generated during signal reconstruction. These parameters can have a large impact on the quality of the reconstruction, and therefore, ideally, sweeps of the best τ and m values should be conducted to confirm that there is robustness in the selected parameters if similar values are used for all reconstructions, as done in Chapter 5. Another challenge is the care that should be taken to ensure the thruster remains in the same dynamic condition throughout a test campaign, so that the causal relationships underlying the dynamic system do not change significantly over the probing time. This generally means the thruster should have

reached steady state and is not significantly drifting, leading to a lack of stationarity in the underlying system. Tools like extended cross-convergent mapping could be used on signals of interest at the beginning and end to ensure that changes in these dynamic relationships have not occurred.

The ability to fuse multiple temporal diagnostic measurements with SMI opens the possibility of increasing the degree to which transient phenomena downstream and even within the thruster can be measured and synchronized on time scales of interest. Adding temporal information can further improve the information content of standard diagnostics, as temporal current density fluctuations will also carry species information. If fused with improved species information collected from a combined probe as discussed in Chapter 1, then we believe further reductions in uncertainty could be obtained in the measurements, as well as improved discharge plasma state uniqueness. This is a consequence not only of the need for a specified environment to produce such species, but also of information about how these states evolve over time to produce specific ions.

4.6 Summary

The ability to reconstruct spatiotemporal measurements from multiple diagnostics is a significant step towards providing data for validating time-resolved models and identifying gaps between models and experimental data. Furthermore, many electrostatic and optical diagnostics require adjustment to a measurement state to build up profiles of the measured property of interest. Application of SMI to current density measurements made downstream of a 1kW CSU HET showed that not only is there a rich temporal variation in the ion current propagating downstream, but that additional information may also be extracted via time-of-flight-like effects based on the disparity of species-specific ion velocities that arrive at the probe. The reconstructed ion current density profile also shows a temporal change in the plume divergence, although in the data presented, this is very likely dominated by the very ionization events that are commensurate with the very large breathing-mode current oscillations. Shadow manifold interpolation has proven robust, enabling high-speed state measurement and, eventually, diagnostic fusion. However, there are still pragmatic challenges when collecting high-speed measurements at a specific measurement

state, which can affect the utility of the measurement outcome. In this sense, it is critical to ensure the fidelity of high-speed measurements by working to improve measurement accuracy and reduce random noise to a minimum. Additionally, attention needs to be paid to thruster dynamics throughout the test to ensure there are no major changes in dynamic behavior during a measurement collection campaign. Tools that can aid in this process are discussed in Chapter 5.

5 CAUSAL ANALYSIS OF HALL THRUSTER TELEMETRY²

Thus far, Chapters 3 and 4 have discussed improvements in the specificity of common diagnostics through combined probe-and-filter diagnostics and through improved temporal information, thereby enhancing the information provided for HET thruster models. However, one of the most commonly used metrics to tie dynamic data to HET discharge models is telemetry measurements, the circuit voltage and current measurements of the thruster under test [22], [26], [27]. This is due to the low barrier to making high-resolution measurements of these metrics, which, as is the case for anode current, can be easily integrated into models of varying fidelity. The increasing dynamic-systems approach to the temporal behavior of thrusters [12], [54], [81] also indicates that, if the appropriate temporal signals are selected, the dynamics of the overall system should be embedded in the time series. Selecting the appropriate time series for such data-driven frameworks is a key challenge for delivering maximally informative dynamic data.

5.1 Introduction: Signal Selection for Data-Driven Models

As discussed previously in Chapter 2, the need for predictive modeling capabilities has led to a reliance on data-driven modeling approaches to capture dynamic thruster behavior [22], [26], [27]. This direct integration of data into physics-constrained models can enable the estimation of the plasma state in regions and time scales that may be difficult to determine with non-perturbative experimental techniques. In this framework dynamic behavior of the thruster is assumed to be contained in temporal signal data used in the model, however, the Hall effect thruster is a convolution of multiple plasma regions dominated by different physics [17], [18], [19], and as a result the selection of the dynamic signal to integrate into such a modeling framework must weigh the utility of the information contained in a given signal.

² Part of this chapter is reprinted from a submitted manuscript “Causality analysis of vacuum chamber facility effects for Hall effect thruster operation” by Seth Thompson *et al.* 2026, *Journal of Applied Physics* (Unpublished).

Time-resolved measurements of direct plasma properties, such as species densities, electron energy distributions, ion and neutral velocity distributions, and plasma potential, are ideal inputs for data-driven models. However, measuring these properties directly in the spatial or temporal scales of interest is difficult and not readily scalable to providing measurements at every EP vacuum facility due to the limited demonstration of techniques. Even when time-resolved measurements are available, determining whether dynamics observed in one property or spatial location are dynamically relevant to thruster behavior or specific model physics can remain illusory. Therefore, selecting accessible, time-resolved measurement datasets that maximize information gained in data-driven thruster models requires additional scrutiny to ensure that the signal is dynamically relevant to the system and can be incorporated into the model. Although models can include several streams of measurement data, it is just as critical to avoid irrelevant signals that have no bearing on the device's operation as it is to select signals with the most useful dynamic information.

This work describes the use of a dynamic causality-determination technique, extended cross-convergent mapping (eCCM) [78] as an aid in selecting which signals to include in data-driven and physics-based modeling frameworks. This technique can be used to infer causal relationships between dynamic signals, providing context for selecting data streams. This, in turn, can be used to exclude specific signals or adjust models to account for the most relevant sensor signals. Signals with strong causal links to important plasma properties will provide more unique information for model parameter estimation, thereby reducing uncertainty in key model parameters by limiting the available state space for the estimated parameters. These time-varying signals can be either high-speed diagnostic measurements or periodic measurements when multi-modal data-driven HET model frameworks are used [27]. When evaluating accessibility and scalability, the telemetry signals of HETs, such as discharge voltage and current measurements, are among the top candidates. These signals are readily accessible, relevant to the entire thruster system (*e.g.*, discharge, cathode, nearfield), and can be monitored non-invasively. Previous data-driven frameworks have used applied power and discharge currents for data-driven models that

relate bulk plasma parameters to input power or current [22], [27], [73]. Adapting these models to include additional signals, such as discharge voltage, cathode-to-ground voltage, the potential of floating mounting structures, thruster body potentials, or other accessible circuit elements, may better capture state evolution if shown to have a strong causal relationship to key plasma properties of interest.

In considering the use of dynamic signals of interest, it is critical to understand what dynamic information the measurements capture. A HET plasma discharge exhibits a wide range of oscillations and instabilities across spatial and temporal scales. The relevant timescales span several orders of magnitude, from nanoseconds for electron-cyclotron frequencies to hundreds of microseconds for ionization-related oscillations, and range in spatial dimensions from fractions of a millimeter for Debye lengths to fractions of meters for the device size. Many of these discharge oscillations were reviewed by Choueiri in 2000 [44], and these oscillations remain areas of active research [19]. Breathing mode oscillations have remained of top interest due to their influence on stability and performance [24], [65], [130], while kinetic instabilities, such as electron cyclotron drift instabilities, remain of keen interest due to their suspected influence on anomalous electron transport [131], [132]. Furthermore, the forced coupling of the HET discharge to a hollow cathode discharge through a common DC circuit means that oscillations native to hollow cathode operation [106], [107], [108], [109] may manifest in global circuit measurements of a thruster. Beyond plasma-specific dynamics, slow transient behavior during device operation is observed, ranging from tens of seconds to hours, due to thermal effects and material outgassing [133], [134], [135]. This combination of many non-linear coupled dynamics that manifest in the telemetry makes it challenging to determine the most relevant timescales that capture all important dynamics in the system. For this work, the focus is on dynamics in the 5 kHz-1 MHz range, which encompasses most of the oscillations that capture the state of the plasma, with the possible exception of high-frequency kinetic instabilities.

Further complicating the evaluation of useful dynamic signals for HET modeling is the effect of environmental parameters created in the vacuum test facility on the underlying dynamic system. The

underlying dynamic system is the set of interacting deterministic variables, such as electron density, ion density, electric fields, etc., that are causally participating in a true deterministic system. They are underlying because these global telemetry measurements ultimately serve as indirect signals that can still be causally influenced, but are not the system variables themselves. Identified causal relationships may be influenced by external environmental parameters that alter the underlying system. For example, elevated background pressure resulting from limited pumping speed can alter charge-exchange plumes and modify expected neutral particle profiles near the thruster [31], [32], thereby changing the relationships among dynamic properties in the discharge or nearfield region. Some studies have demonstrated pressure dependence of the electric field and electron temperature in magnetically shielded thrusters [136]. In addition to the facility-grounded surfaces and background pressures, questions remain about the effects of thruster body potential on thruster operation [11] and the constant deposition of wall materials on the thruster from backsputtering [15], which may affect the isolation of thruster body or surface potentials over the length of a test. As a result, casual relationships identified in one environment may change from facility to facility, across a range of pressures within a single facility, or even over time within a single facility, and pressure, not to mention changes that will likely occur over a wear test.

In this work, extended cross-convergent mapping (eCCM) is applied to high-speed telemetry signals from an H9 thruster, including cathode current, I_{cath} , discharge voltage, V_{dis} , and cathode-to-ground voltage, V_{cg} , for different vacuum chamber pressures. Description of the eCCM process and the related reconstruction parameters and calibration processes needed to identify causal relationships and their associated delay responses is presented in 5.2, followed by the description of the experimental setup in 5.3. The result of the selected embedding parameters is discussed in 5.4, and cross-convergent mapping results are presented in 5.5.

5.2 Extended Cross-Convergent Mapping

Although it is often said that correlation does not imply causation, the corollary is often overlooked: a lack of correlation does not necessarily imply a lack of causation. It is often the case that

correlation, or the lack thereof, is used to determine if a given variable or data stream should be included in a model. When there is no noticeable change in the underlying measurement or property that *correlates* with the phenomena of interest, it is all too often overlooked or dropped from consideration. However, by evaluating the causal relationships among two or more variables, we can better determine whether the information being considered is ultimately useful for model inclusion. In other words, causal analysis can quantify the dynamic relevance of a signal used to capture the state of a system as defined by the estimated parameters updated in the model.

5.2.1 Background of extended cross-convergent mapping

Granger causality has been used to evaluate causality among potentially interacting variables, focusing on improving estimates by including or excluding those variables in the prediction of a variable of interest [79]. This approach, however, can struggle to correctly identify causal relationships in dynamic systems where the history of one variable may be redundant due to its inseparability from other system variables as a consequence of their causal relationship in a deterministic dynamic system [76]. As a result, effective inclusion and exclusion of variables cannot be achieved, leading to ambiguity in identifying causal relationships. Cross-convergent mapping (CCM) addresses this challenge by fundamentally leveraging the inseparability of causally related variables in a deterministic dynamic system [77]. Furthermore, this technique is valuable where direct measurements of all underlying variables are difficult or infeasible.

Cross-convergent mapping leverages Taken's Theorem [76], which relates a present state of a dynamic system to its past states to evaluate the degree to which the history of one variable is present in that of another. Sugihara *et al.* demonstrated a causal relation in a system containing dynamic variables x and y as follows. If x acts causally on y , then information about x will be present in the history of y . Additionally, reconstruction of x using the historical data of y will then be convergent to higher accuracy with the inclusion of additional data [77]. The metric used to gauge reconstruction accuracy was the

Pearson correlation coefficient between the reconstructed and test signals. The reconstruction process leverages time-delay embeddings of a signal, coupled with a technique such as simplex projection [122], to perform state-space reconstruction (SSR) and produce a signal estimate of the second signal. This process can be leveraged to identify the reduced set of measurement data required to capture the state of a system if many types of measurement data can be reconstructed from a selected set of signals. The reconstruction of different measurement data streams was demonstrated using the shadow-manifold interpolation method with HET test data by Eckardt et al. [42]. The technique was leveraged to generate time-resolved reconstructions of many measurement-state-based diagnostics [115], [116], [119], [137]. While SSR techniques are useful for time-resolved measurements, the focus of the reconstruction in CCM is on determining the nature of the relationship between the input and reconstructed signals.

By using these reconstructions to evaluate dynamic causality, it is possible to generate accurate reconstructions of a given signal from signals that were not causally related. Strong external driving variables can dominate the dynamics of non-interacting species that appear to respond synchronously, a phenomenon in ecology known as the Moran effect [138]. Additionally, situations arise in which one variable strongly drives another to such an extent that the common dynamics of both variables lead to similar reconstruction behavior, making it difficult to distinguish true unidirectional behavior from bidirectional behavior. These situations can be broadly categorized as forms of generalized synchrony and are known to obfuscate the correct identification of causal variables [77]. To address this challenge, Ye et al. developed an extension of CCM (eCCM) that considered the time shifts between dynamic variables and their peak reconstruction quality to identify the dynamic response and infer directionality, i.e., whether one signal precedes or succeeds another signal, thereby improving the interpretation of cross-convergent mapping [78]. Note that the words “time shifts” are used in eCCM to distinguish from “time delays” used in the construction of time-delayed phase portraits that are used to create potentially quantifiable images of the dynamic characteristics of a system that can be used to determine if a system has changed from one dynamic behavior to a different one. The eCCM technique has been applied to HET

simulations [139] and experimental measurements [140] and has been compared with other causality determination techniques [127]. It is important to note that, since causality is transitive, the reconstruction quality of one variable may be driven by many other variables that are not measured directly. This work employs causal analysis to quantify changes in telemetry characteristics across different test conditions, using analysis limited to the relationships within the selected datasets, that are still subject to transitivity and indirect relationships of the underlying dynamic system.

5.2.2 Extended cross-convergent mapping

The process for cross-convergent was described in detail in Section 4.3.1 when introducing the non-linear state-space reconstruction method for shadow manifold interpolation of dynamic signals. The process requires the selection of embedding delay τ , and embedding dimension m , which can have significant impacts on the quality of the reconstruction. However, in extended cross-convergent mapping, an additional parameter is added to the analysis: a time shift or cross-map lag [78] is utilized to identify directional information flow between signals. The time shift, l , is applied between the two signals used in the reconstruction. Therefore, one can perform the same CCM process but for $y(t)$ cross-mapping $x(t + l)$. If $x(t)$ is causal on $y(t)$ with some delay, then y will have a better cross-mapping skill of previous values of x . For instance, if x acts on y after two-time steps, we would expect the largest cross-mapping skill for y cross-mapping x with a shift of $l = -2$. Hence, a condition for causality is a large cross-mapping skill with a negative time shift, i.e. $l < 0$. For the same example, we can consider the other cross map, i.e. x cross-mapping y with a shift of $l = 2$. Even if the cross-mapping skill is high, indicating a cross-mapped signal is reconstructed accurately, the positive time shift indicates that y is not causal on x . This identification of a delay in response from one variable to another is tested with an example problem used in Ye et al. [78] and in Huerta et al. [127]. The details of this test and the results are shown in 5.2.4. Another way of describing the results is that for some variable x to causally influence another y , then changes in x that cause changes in y must occur before or simultaneously with the changes in y .

5.2.3 Embedding parameter selection

As discussed in the previous section, embedding parameters must be selected to generate time-delay embeddings or shadow manifolds used in the cross-mapping process to reconstruct signals. Specifically, a time delay τ and an embedding dimension m must be selected. There is no mathematical requirement for the selection τ to generate a proper embedding; however, real-world applications demonstrate that both the embedding dimension and the time-delay selection significantly impact reconstruction quality. Several methods for determining the relevant time delay τ and embedding dimension m have been developed [75]. Many of these methods focus on reducing the computational cost of evaluation and on distinguishing deterministic and stochastic signals. Cao's method [126] for the selection of the embedding dimension m , based on a modified form of identifying false-near neighbors, has been used in many applications with thruster-relevant signals [127], [140]. However, this process requires selecting a time delay, for which numerous heuristics have been proposed for given applications [75]. Recently, work by Martin *et al.* [125] has demonstrated a robust framework that leverages the mutual information of transformed coordinates (MITC) to select a value τ that yields high reconstruction quality in a noisy system. In this dissertation, a parametric sweep of τ and m is conducted to compare their respective impacts on reconstruction quality for a given pair of signals. As a result, every pair of signals reconstructing one another used the embedding dimension m and the embedding time delay τ that produced the best reconstruction with no time shift l applied.

5.2.4 Example Problem

To evaluate our eCCM analysis, an example problem is used in which a delayed response between variables can be controlled to assess proper identification via the time shift, l , in the eCCM process. A set of equations that contains two variables X and Y , driven by a third causal variable Z with a delay k between the response of X and Z is used.

$$\begin{aligned}
Z(t + 1) &= Z(t)(A_z - A_z Z(t)) \\
Y(t + 1) &= Y(t)(A_y - A_y Y(t) - B_{yz} Z(t)) \\
X(t + 1) &= X(t)(A_x - A_x Y(t) - B_{xz} Z(t - k))
\end{aligned} \tag{17}$$

To generate the timeseries data used for evaluating the extended cross-convergent mapping framework all coefficients are kept the same, $A_z = 3.8$, $A_y = 3.1$, $A_x = 3.1$, $B_{yz} = 0.9$, $B_{xz} = 0.5$, and the initial values of $X(0), Y(0), Z(0) = 0.3$. The cross-mapping and time shift plots of all variables correctly identified a causal relationship with appropriate directional shifts for Y and X both influenced by Z . However, this process highlighted a key result. In a case where two variables, such as X and Y , are strongly driven by a third common variable Z with differing delay in response (when $k = 2$), an apparent causal relationship between the indirect variables can be observed, with a time shift corresponding to the response disparity between the two variables to the common driver. This is especially relevant to systems where it is possible that not all system variables would be evaluated in the causal network, such as in the use of a limited subset of telemetry signals. In the example problem where all variables are reconstructed by one another, the reconstructions are still lower in quality than those of the correct causal directions $Z \rightarrow X$ and $Z \rightarrow Y$. However, when the direct relationships among all system variables cannot be evaluated, ambiguity may arise in the causal relationships inferred solely from time-shift data and reconstruction quality. As a result, it must be understood that apparent causal (unidirectional or bidirectional) relationships with appropriate time delays may also arise from indirect relationships between variables that share a common driving variable with differing response delays.

For the analysis of the time series generated from equation (17), we also conducted a parametric sweep to identify optimal embedding parameters for each signal set. Interestingly, the optimal value τ was 2 in all reconstructions, while the embedding dimension varied depending on the signals being reconstructed: reconstructions in the correct causal direction were dimension of 4, and those in the opposite direction were 5-6, and sometimes 7 if $k = 2$. This increase in dimensions can broaden out the

width of the higher cross-mapping skill vs. time shift, as seen in Figure 23 and Figure 24; however, the correct delays were still identified.

The first application of the extended cross-convergent mapping on the timeseries generated from equation (1) with no response delay ($k = 0$) is shown in Figure 23. Plots (b) and (c) show the correct identification of the causal direction, with the highest reconstruction quality occurring for X and Y cross-mapping Z , with a negative shift. This indicated a strong unidirectional causal relationship. The time shift in the forward direction appears to slightly over-predict, which may be due to the increased dimensionality of the embeddings used to reconstruct signals in the non-causal direction, as well as the centering of the embedding as discussed in [78]. The reconstructions of X and Y cross-mapping one another are both forward-shifted, indicating an inconclusive result, as may be expected for indirect variables driven by a strong common driver.




Figure 23: Extended cross-mapping results of equation (1) with the delay $k=0$, with accurate identification of Z 's causal effect on X and Y , but ambiguous results of how X and Y relate to each other

The response delay of X to Z in equation (1) was changed to $k = 2$, and all three signals were cross-mapped again, with the optimum embedding dimension found for each signal being cross-mapped, shown in Figure 24. The correct identification of the time delay and causal direction for X and Y cross-mapping Z are shown in plots (b) and (c). However, unlike the results with $k = 0$, when $k = 2$, the result of X cross-mapping Y appears to indicate a causal relationship with no time shift. If not for seeing the tightly coupled unidirectional relationship in plots (b) and (c) with higher reconstruction skill, then one might incorrectly attribute the relationship $Y \rightarrow X$, due to the disparity in the response to Z , between the two variables. This result should be kept top of mind when an incomplete set of system variables is

examined; this is to say that causal relationships identified in a subset of system signals may result from indirectly related variables that share a common driver with a response disparity.




Figure 24: Extended cross-mapping results of equation (1) with the delay $k=2$ used to demonstrate how information delays can complicate the interpretation of uncoupled variables X and Y .

5.3 Experimental Configuration

Extended cross-convergent mapping (eCCM) is applied to HET telemetry data collected in a previous facility-pressure-effect test on thruster impedance, as presented in [141]. This test was conducted with an H9 magnetically shielded HET developed by the University of Michigan, Air Force Research Laboratory, and Jet Propulsion Laboratory [142] and operated on krypton in VTF-2 at the Georgia Tech High-power Electric Propulsion Laboratory (HPEPL). The lowest operational pressure for the thruster was 6.8×10^{-6} Torr at the 4.5-kW operating condition and 8.5×10^{-6} Torr at the 6-kW operating condition. Telemetry data were captured with a Teledyne Lecroy Digital Oscilloscope HDO6104, using Powertek DP-25 differential voltage probes (bandwidth of 25 MHz) for both discharge and cathode-to-ground voltage measurements. Cathode current was measured with a Teledyne LeCroy CP150 current clamp (bandwidth of 10 MHz). Data from the oscilloscope was collected at 10 MHz. The thruster was operated with an inline RC filter to protect the power supply from unstable thruster operation (if it were to occur), with a resistance of 0.533Ω and a capacitance of $100 \mu\text{F}$ [14]. The thruster was operated with the thruster body tied to the chamber-ground. To adjust chamber pressure, krypton gas was flown through a $\frac{1}{4}$ inch Swagelok connection located 1.8 meters axially downstream of the thruster face and 1 meter below the thruster centerline. The flow rate was adjusted until the two and three-times baseline chamber pressures were achieved. Figure 25 shows a notional circuit diagram with depictions of the telemetry signals and

gas inlet for chamber pressure adjustment. The pressure gauge, a Varian 571 Bayard-Alpert hot filament with a neutralizing inlet [143], was located 1 m away from the thruster exit plane.




Figure 25: Simplified H9 testing schematic with telemetry measurements and notional pressure adjustment location.

5.4 Embedding Parameter Analysis

To perform cross-mapping of the telemetry signals, embedding parameters must be selected. Due to uncertainty in the relationship between the different dynamic signals, a parametric sweep of the embedding dimension m and time delay τ was conducted first to determine the optimal embedding parameters to use for cross-mapping. The time-delay range was 2 to 80 (or 0.2-8 μs), and the embedding dimension was varied from 2 to 100. The parametric sweep is performed without time shifts between the signals, and the embedding parameters that yield the highest cross-mapping skill for the reconstructed signal are then used in extended cross-convergent mapping, in which time shifts are applied to the signals. Although these signals are all measured from the same circuit operating a dynamical plasma system, there is no guarantee that a singular set of embedding parameters will yield maximum reconstruction quality for all pairs of signals used for cross-mapping.

5.4.1 4.5-kW Operation

The lowest chamber pressure of $6.8 \mu\text{Torr}$ was achieved with the 4.5 kW operating condition. After operating for several hours to achieve thermal steady state, the telemetry data were collected, and the chamber pressure was increased to twice the nominal background pressure, approximately $13 \mu\text{Torr}$, where telemetry was collected again. This process was repeated one more time to achieve three times the nominal background at approximately $20 \mu\text{Torr}$ where telemetry was collected one more time. Under all conditions, the thruster was operated at a steady-state discharge voltage of 300 V and a discharge current of 15 A, the discharge current being maintained by adjusting the thruster flow rate. The 4.5-kW operating condition had the lowest peak-to-peak current oscillation relative to the mean current in the telemetry waveforms collected.

The power spectral density plots for each signal under the three background pressure conditions are shown in Figure 26(a), Figure 27(a), and Figure 28(a). As shown in Figure 26(a1), Figure 27(a1), and Figure 28(a1), the cathode current exhibits a peak oscillation at the breathing mode around 16.4 kHz, with lower signal content at higher frequencies. Figure 26(a2,a3), Figure 27(a2,a3), and Figure 28(a2,a3) show that the discharge and cathode-to-ground voltage oscillations occur near the same breathing mode frequency at 17.6 kHz but with considerably more signal content in the 100 kHz–1 MHz range. In all cases, the breathing mode peak decreased with increasing pressure, accompanied by broader peaks relative to those at the baseline operating pressure. Cathode current showed the largest change in the oscillation amplitude at the breathing mode frequency, followed by cathode-to-ground voltage, with the smallest change in breathing mode peak height observed in discharge voltage.

For each signal combination, we perform two cross-mappings in each direction, *i.e.*, I_{cath} cross-maps V_{dis} and V_{dis} cross-maps I_{cath} . The combination of an unknown relationship between the signals, measurement uncertainty, and the associated uncertainty in properly embedding each signal is why one cannot assume that a single set of embedding parameters will perform well for all signals. Indeed, this can

be seen in the embedding parameters that result in the highest cross-mapping skill in Figure 26. The cross-mapping skill of I_{cath} and V_{dis} in Figure 26(b1) and (b2) show the significant change in embedding dimensions that result in optimal reconstruction. Beyond differences observed in the directionality of cross-mapping, we observe changes in the best embedding parameters for a single signal used to cross-map different signals, as shown in the reconstructions of V_{dis} and V_{CG} by I_{cath} in Figure 26(b1) and (b3), respectively. The same input signal, I_{cath} , is used for both reconstructions; however, significantly different embedding parameters are required to achieve the best cross-mapping. In general, the best embedding parameters to obtain the highest reconstruction skill vary widely across Figure 26(b1)-(b6).

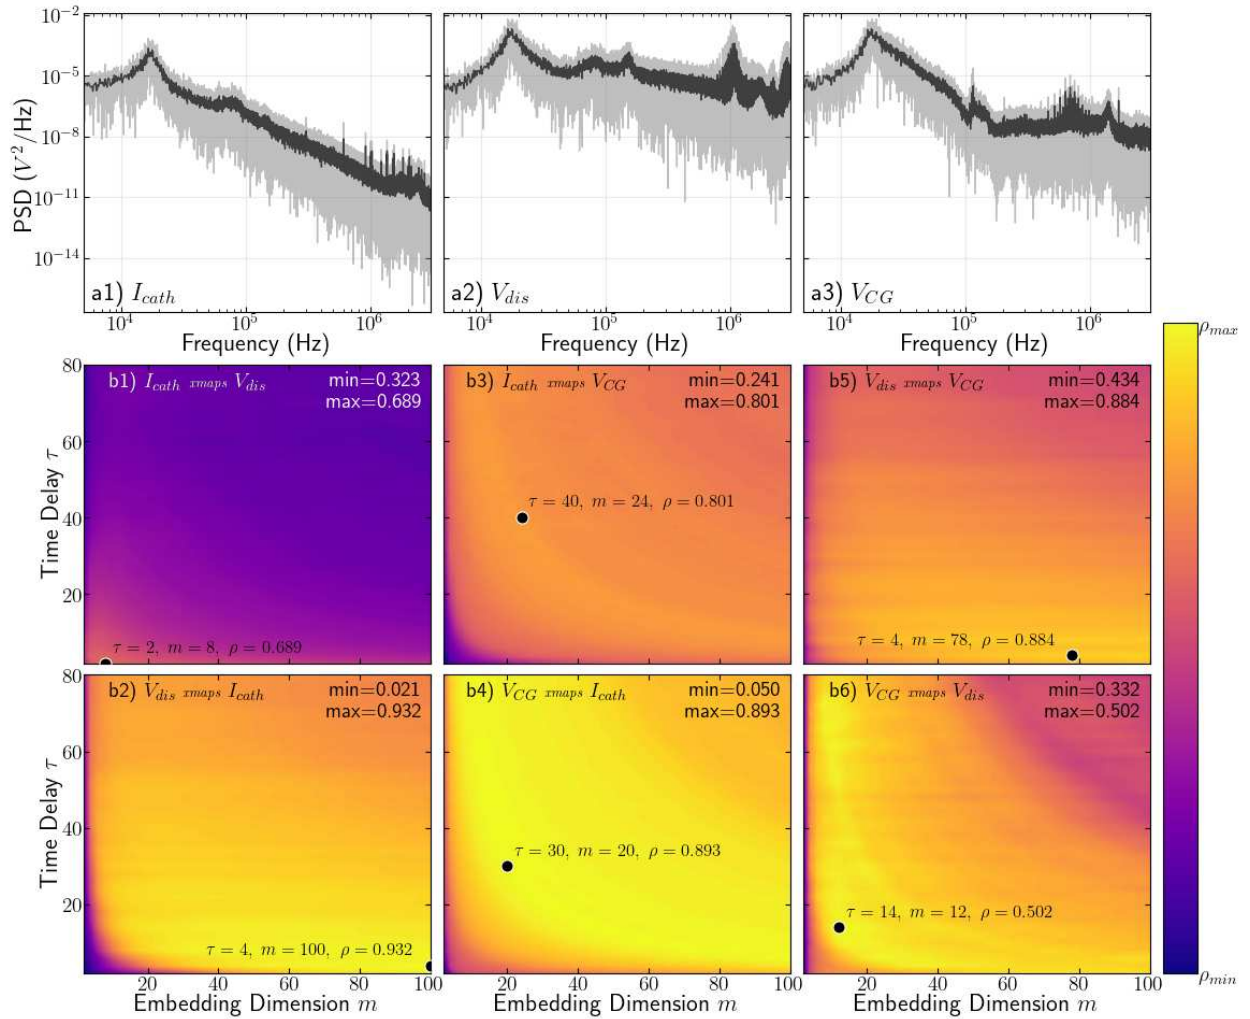


Figure 26: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 4.5 kW at baseline pressure, i.e., a chamber pressure of 6.8 μ Torr.

The power spectral density plots and parametric sweep of embedding parameters for the signals collected at two times baseline chamber pressure are shown in Figure 27. The embedding parameters that yield the best cross-mapping skill continue to depend on the signal being cross-mapped and the directionality of the cross-mapping between the two signals. A significant increase in the optimum τ value for V_{cg} cross-mapping V_{dis} in Figure 27(b6) is observed at the higher pressure, although the lower embedding τ in Figure 26(b6) remains as a relatively high cross-mapping skill region in both plots. This could be due to periodicity in the underlying dynamics, or to a change in the dynamics of the input or test signals, best captured by increasing the time delay for the embeddings by about $6 \mu s$. There is also a slight decrease in the peak cross-mapping skill achieved in each sweep at the higher pressure.

In cross-mappings V_{dis} shown in Figure 27(b2) and (b5), the highest cross-mapping skill was observed at $m = 100$, *i.e.*, at the limit of the dimensional range allowed for the sweep. Furthermore, additional sweeps to higher τ and m values were conducted and yielded only minor increases in cross-mapping skill. As a result, the eCCM analysis for all signals used embedding parameters that yielded the highest cross-mapping skill from within the parameter sweep conducted for a pair of signals.

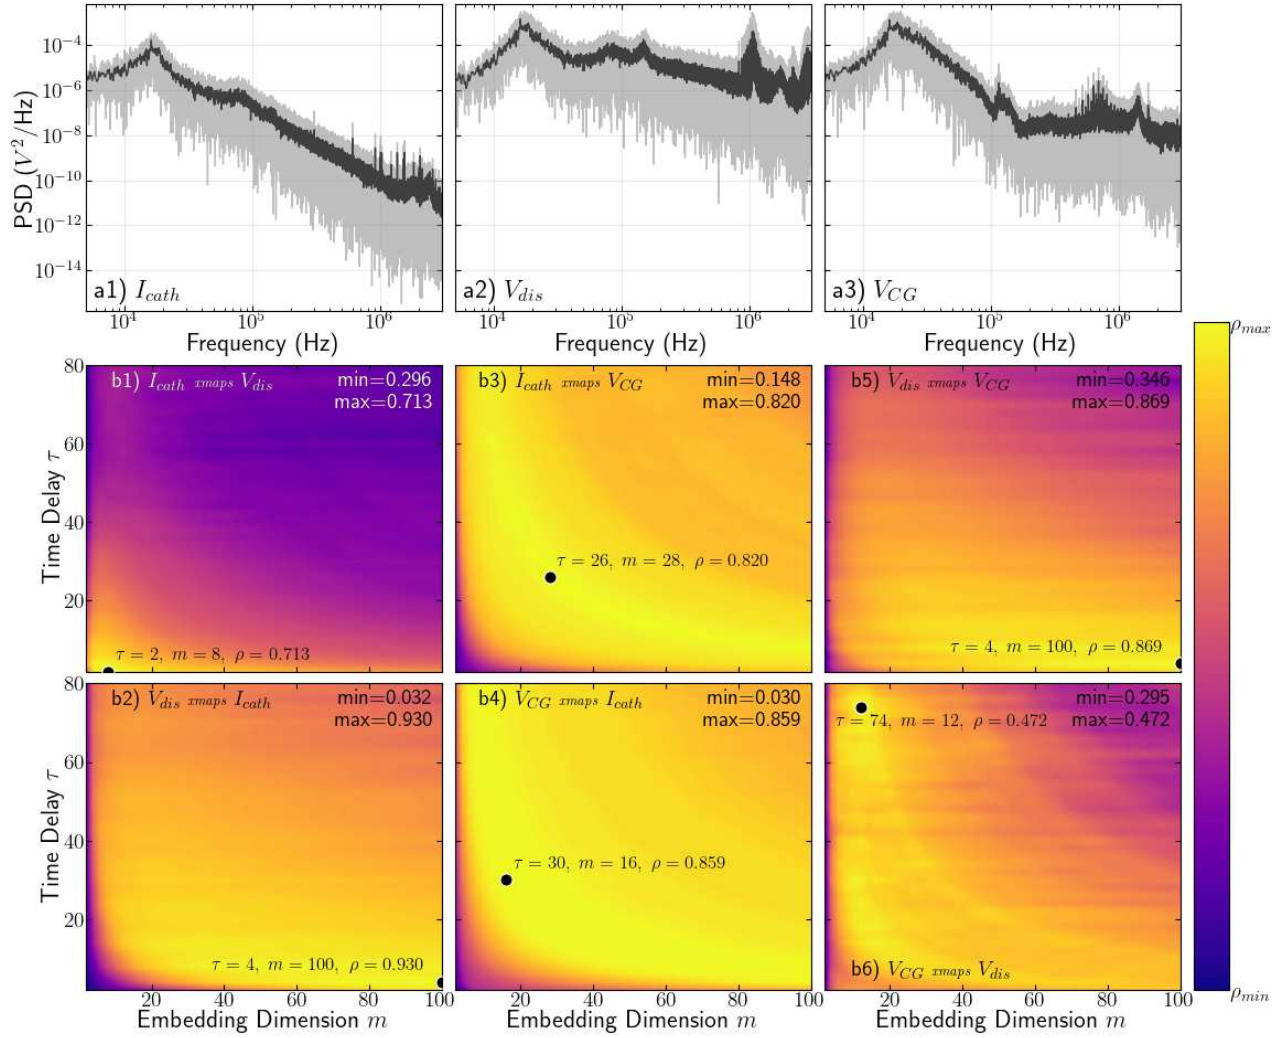


Figure 27: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 4.5 kW at twice the baseline pressure, i.e., a chamber pressure of $13 \mu\text{Torr}$.

For a chamber pressure of $20 \mu\text{Torr}$, corresponding to three times the base operating pressure for the thruster at 4.5 kW, the power spectral density plots and the parametric-sweep contours for cross-mapping are shown in Figure 28. The continued reduction in peak height with increased broadening is observed in Figure 28(a1)-(a3), while the continued reduction in peak and high cross-map skill regions is observed in the contour plots in Figure 28(b1)-(b6). A reversion to the smaller τ value for V_{CG} cross-mapping V_{dis} was observed in Fig. 4(b6), along with a notable increase in dimension m for I_{cath} cross-mapping V_{CG} in Fig. 4(b3).

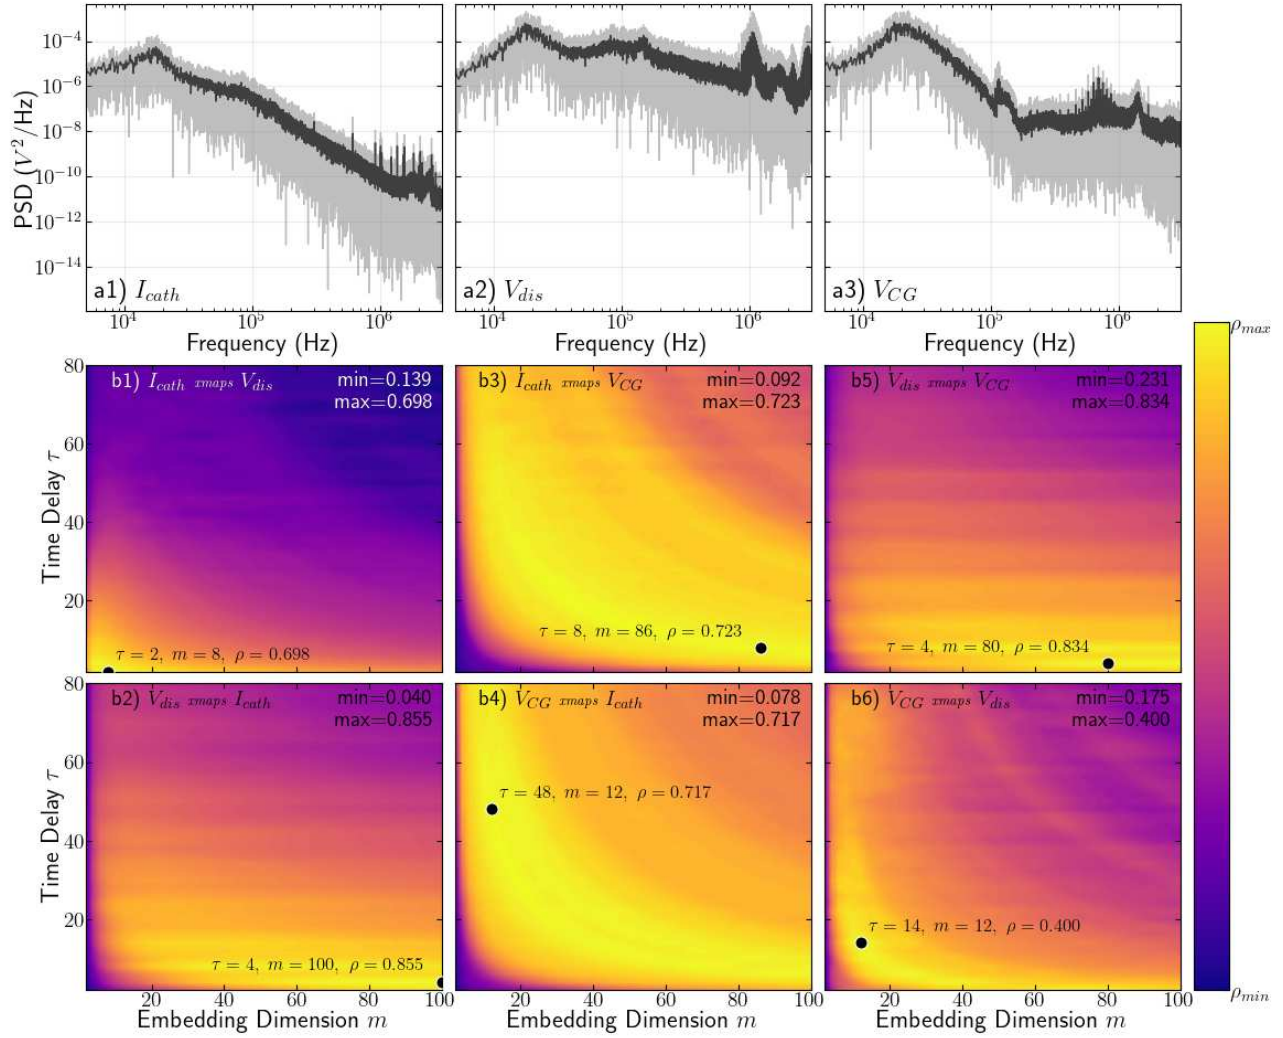


Figure 28: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 4.5 kW at three times the baseline pressure, i.e., a chamber pressure of 20 μ Torr.

In general, the parametric sweep of embedding parameters appears justified, as the reconstructed signals and the direction of cross-mapping influenced the embedding parameters that yield the highest cross-mapping skill. In some instances, particularly for I_{cath} and V_{CG} , the relationship of the $m\tau$ product appears to produce regions of higher cross-mapping skill, as may be expected for signals of the same dynamical system, as this scalar product is related to the effective time period captured by a single embedding vector at a given time step [75]. Specifically, for I_{cath} and V_{CG} , it seems to be loosely grouped around $m\tau \approx 500$, or 50 microseconds, which is near the breathing-mode oscillation period.

5.4.2 6-kW Operation

As a result of the strong breathing-mode oscillation present in all telemetry signals when the thruster is operated at 6 kW, a general baseline of high cross-mapping skill across many embedding parameters is observed, driven by the common 14.52 kHz oscillation that dominates the power spectra of each time series. As a result, the information about changes in cross-mapping skills should be evaluated on a relative basis, and therefore so should the dependence on embedding parameters. To aid visualization of relative trends in the effect of embedding parameters on cross-mapping skill at each chamber pressure, an arbitrary minimum cross-mapping skill of 0.8 was applied to the embedding parameter sweeps in Figure 29 through Figure 31.

Despite the significant increase in cross-mapping skill amongst all signals in the 6 kW condition operated at 8.5 μ Torr, the embedding parameter regions that yielded the highest cross-mapping skill in Figure 29(b3)-(b6) remain similar to those of 4.5 kW shown in Figure 26(b3)-(b6) with slight differences in specific peak embedding parameters shown in Figure 29(b3)-(b6). The smaller τ values for I_{cath} and V_{dis} reconstructions in Figs. 6(b1) and (b2) are similar, with a preference for larger dimensions for cross-mapping performed by V_{dis} . The regions of high cross-mapping skill between I_{cath} and V_{cg} , in Figure 29(b3) and (b4) appear to follow a similar $m\tau$ product relationship, as seen with the 4.5 kW condition.

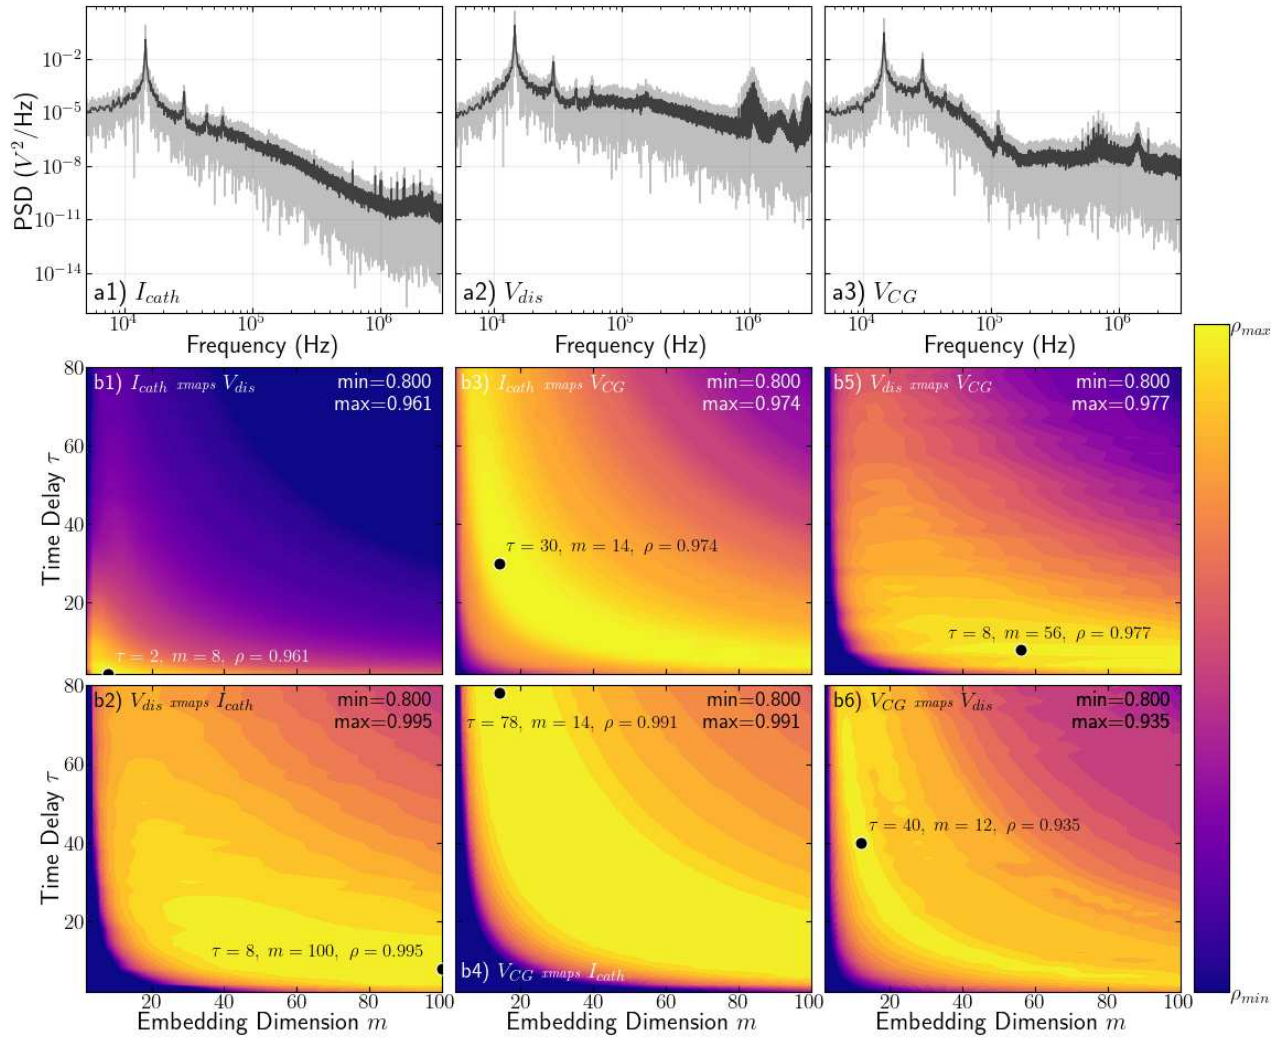


Figure 29: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 6 kW at baseline pressure, i.e., chamber pressure of $8.5 \mu\text{Torr}$.

The background pressure in the chamber was increased to $16 \mu\text{Torr}$, roughly twice the base operating pressure at 6 kW, and telemetry was collected as the thruster reached steady-state operation. The resulting embedding parameter sweep, shown in Figure 30, indicates a reduction in the size of the regions producing the highest cross-mapping skill, similar to the effect observed in the 4.5-kW operation. This trend continues, as shown in Figure 31, as the background pressure was further increased to $25 \mu\text{Torr}$, roughly three times the base operating pressure during 6-kW thruster operation. There is also a decrease in the peak cross-mapping skill achieved by each sweep with increasing pressure.




Figure 30: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 6 kW at twice the baseline pressure, i.e., a chamber pressure of 16 μ Torr.




Figure 31: Power spectral density plots (a1-a3) and embedding parameter sweep (b1-b6) of I_{cath} , V_{dis} , and V_{CG} waveforms captured from an H9 thruster operated at 6 kW at three times the baseline pressure, i.e., a chamber pressure of 25 μ Torr.

5.5 Extended Cross-Convergent Mapping Results

After selecting the appropriate embedding parameters for each individual signal-specific reconstruction, time shifts can be applied to provide context around the flow of information between the two signals. These relationships are evaluated at two operating power levels (4.5-kW and 6-kW, both operated at 300V) across three background pressures.

5.5.1 Extended cross-convergent mapping results of the 4.5-kW condition

As discussed in Section 5.2.1, one of the primary advantages of eCCM over previous techniques is the identification of time shifts between the causal responses of one variable to another, reflecting the

information-transit time between the two signals and the direction of information flow. The difference between the resulting peaks in cross-mapping skill relative to one another, as well as the no-shift case ($l = 0 \mu s$), is used to determine whether there is a lag in the information flow between the signals. Peaks in the negative time-shift direction indicate that the reconstructed signal is causal with respect to the signal used for the reconstruction, with a delay on the order of the observed time shift. Peaks at a positive time shift indicate that the reconstructed signal is not causal with respect to the signal used for reconstruction, but it does share the signal's dynamics. This can occur when similar external dynamics drive non-interacting variables, or when the signal performing the cross-mapping is unidirectionally driving the signal being cross-mapped. In the case of a unidirectional relationship, there should be both a positive and a negative time-shift peak in the respective causal direction. This is the classic case covered by Ye et al [78]. Finally, as discussed in 5.2.4 and by Ye *et al.* [78], there are limits to the accuracy of the interpreted time shift if the time shift is less than that of the $m\tau$ product, as the shift value can be influenced by centering of the delay vector and $m\tau$ product size. We found in our example using negative-lag embeddings, following $[x(t), x(t - \tau), x(t - 2\tau), \dots, x(t - (m - 1)\tau)]$, that the negative shift consistently identified the correct lag value, while overpredicting in the forward-shifted peak direction, as shown in the plots in 5.2.4.

Figure 32 shows the results of eCCM of the H9 telemetry signals at baseline, *i.e.*, $6.8 \mu\text{Torr}$ (a1-a3), $2x$ baseline, *i.e.*, $13\mu\text{Torr}$ (b1-b3), and $3x$ baseline, *i.e.*, $20 \mu\text{Torr}$ (c1-c3). As shown in the first column, Figure 32 (a1), V_{dis} has the highest cross-mapping skill of I_{cath} with a nearly instantaneous response. The slight negative shift would indicate that V_{dis} is causally driven by I_{cath} , with a little or no delay, and an observed reduction in cross-mapping skill with increasing chamber pressure, as shown in Figure 32(b1) and (c1). In Figure 32(a2), we can see that the negative peak of V_{cg} relative I_{cath} indicates that I_{cath} is causal on V_{cg} at the baseline pressure. However, the reduction in cross-mapping skill of V_{cg} with increased pressure below that of the positive-shifted peak shown in Figure 32(b2) and (c2), indicates that this may deteriorate into an indirect relationship that is still reliant on a common set of dynamics with reduced unique information about I_{cath} in the time history of V_{cg} . Finally, it is shown in Figure 32(a3)-

(c3) that a strong indirect relationship exists where common dynamic information in V_{dis} is also present in V_{cg} ; however V_{cg} cannot reconstruct the dynamics of V_{dis} well. Again, the reconstructions follow the trend of the other two signal pairs, with decreasing reconstruction skill observed at higher pressures, indicating eroded mutual dynamic information.

As discussed in the parametric sweeps of embedding parameters, there are cases when the maximum embedding dimension (*e.g.*, $m = 100$) was selected as the highest cross-mapping skill. In these cases, the signal doing the cross-mapping exceeded the cross-mapping skill in the reverse direction, leading to no change in the identified directionality or time shift l in the subsequent causality analysis. Furthermore, minor gains in the highest cross-mapping skill from expanding the parameter sweep ranges also reduce the likelihood of differences arising from sweep limitations.

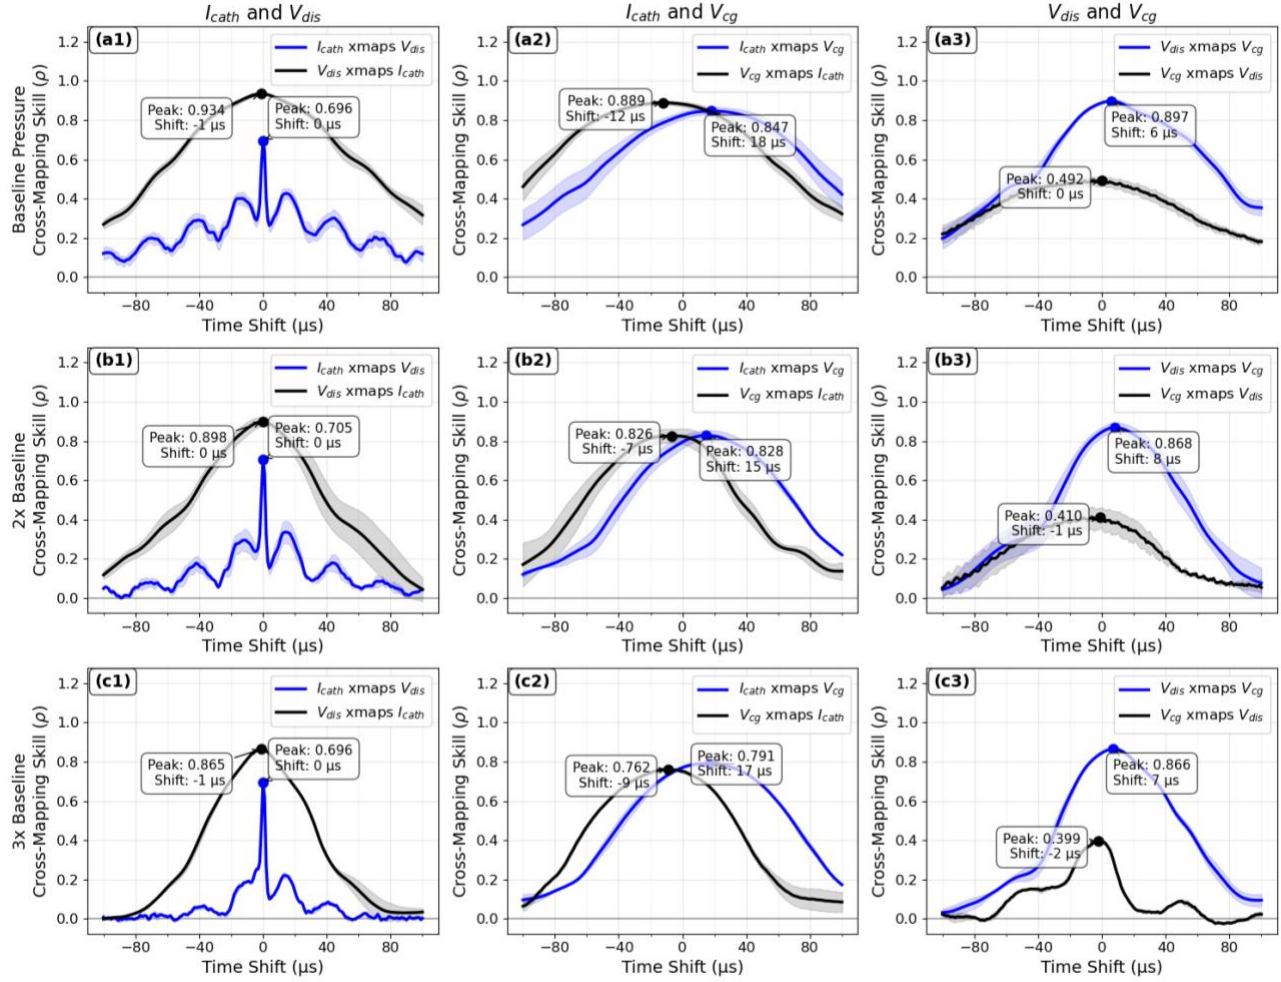


Figure 32: Extended cross-convergent mapping results of telemetry collected from an H9 thruster operated at 4.5 kW and at three background pressure conditions.

Discharge voltage V_{dis} appears to reconstruct the other signals most effectively, even if the positive peak with cross-mapping V_{cg} indicates a form of an indirect relationship, V_{dis} and is still able to capture unique states best. This high cross-mapping skill underscores the importance of including V_{dis} in dynamic, data-driven models. Cathode-to-ground voltage V_{cg} originally appears to be driven by cathode current I_{cath} , as might be expected; however, it continues to lose relevant dynamic information of I_{cath} with higher chamber pressures. This could indicate changes in the current paths near the cathode plume, driven by changes in chamber pressure.

5.5.2 Extended cross-convergent results of the 6-kW operation

The results of the eCCM analysis applied to the H9 telemetry collected during operation at 6 kW are shown in Figure 33. The results of cross-mapping skill and time shifts at the lowest achievable pressure, referred to as baseline, *i.e.*, 8.5 μ Torr (a1-a3), 2x baseline, *i.e.*, 16 μ Torr (b1-b3), and 3x baseline, *i.e.*, 25 μ Torr (c1-c3). Similar to the results from telemetry collected with the thruster operating at 4.5 kW, the discharge voltage appears to reconstruct all other signals best, but continues to show a decrease in peak cross-mapping skill with increasing pressure. Again V_{dis} appears to respond nearly instantaneously to changes in cathode current. The causal relationship between I_{cath} driving V_{cg} in Figure 33(a2), (b2), and (c2) appears to be similar in the direction and time shift values to those at the 4.5-kW baseline condition in Figure 32(a2), but the cathode to ground voltage V_{cg} is able to maintain a higher cross-mapping skill than I_{cath} through all increases in pressure. It could be that due to more coherent and larger amplitude breathing-mode oscillations V_{cg} remains casually driven by the primary dynamics in I_{cath} . The results of cross-mapping between V_{dis} and V_{cg} appear to indicate again an indirect relationship, with the peak cross-mapping skill resulting from V_{dis} cross-mapping V_{cg} at a positive time shift. This indicates that the two signals may share common dynamics, but that cathode-to-ground voltage V_{cg} does not capture all the dynamics of the discharge voltage V_{dis} .

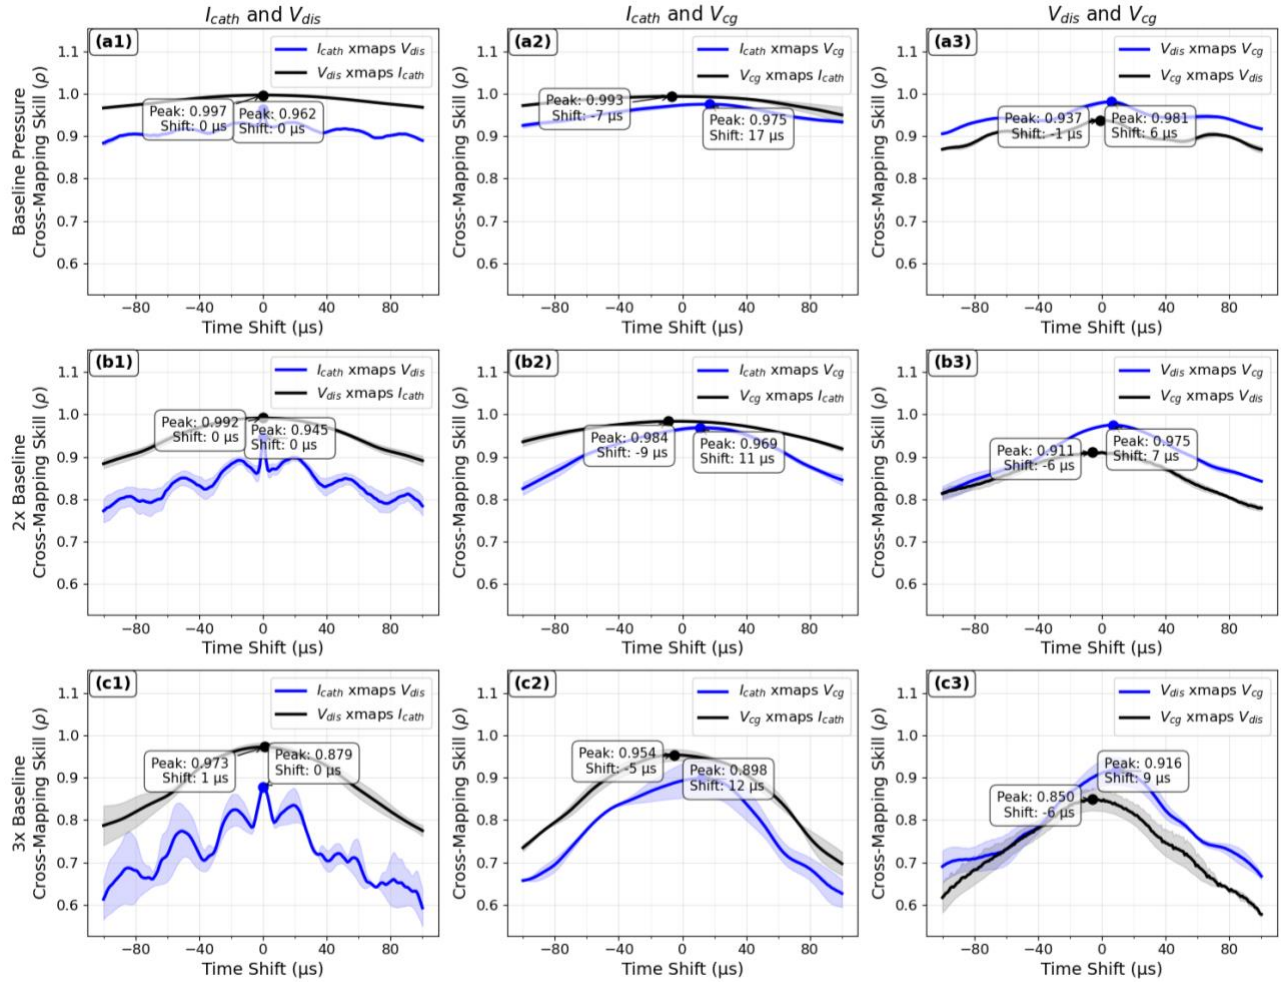


Figure 33: Extended cross-convergent mapping results of telemetry collected from an H9 thruster operated at 6 kW and at three background pressure conditions.

5.5.3 Pressure effects on cross-mapping

Understanding the causal relations in HET telemetry can help better select which data to include in data-driven models, but understanding how environmental factors, such as background chamber pressure, affect these relationships is crucial. If a given signal is selected for its causal relationships with many dynamic data, but those relationships are sensitive to environmental factors such as background pressure, then care must be taken in using the data, as it may no longer contain the relevant information needed to improve state estimation. To this end, we explore how the dynamic relationships among thruster telemetry signals are altered by pressure across the two thruster operating conditions, as shown in Figure 32 and Figure 33. During thruster operation at 4.5 kW, the peak cross-mapping skill monotonically

decreased with increasing chamber pressure for nearly all signals. This trend is illustrated directly in Figure 34, which plots the maximum cross-mapping skill against background pressure for all pressure conditions. Asymmetric changes in the effect of signals are present in all three, but are most clear in the largest cross-mapping skill change with pressure observed for V_{CG} cross-mapping I_{cath} , in Figure 34(b). Here, the change in peak cross-mapping is significant enough to drop the reconstruction quality below that of the I_{cath} cross-mapping V_{CG} , indicating that there may not be a clear unidirectional causal relationship between the two, especially at higher pressures. This may be due to changes in the coupling of cathode current to the near-field plasma region, which are influenced by additional charge-exchange current paths to the thruster plume and the chamber wall. If the dynamics of I_{cath} are altered by elevated charge-exchange ions, the time history of cathode current may not be captured well by the cathode-to-ground voltage. If this effect of chamber pressure is due to the coupling of cathode current to the vacuum chamber wall, it also underscores the importance of the grounded chamber surfaces near the cathode and thruster plume in shaping the relationship between I_{cath} and V_{CG} . Therefore, if data-driven models rely on the causal relationship of V_{CG} for ground testing, it is important to remember that dynamic information of cathode current could be lost with increases in chamber pressure.




Figure 34: Maximum cross-mapping skill of each telemetry signal pair with increasing chamber pressure of an H9 operated at 4.5 kW.

The effect of pressure on the peak cross-mapping of telemetry signals collected as the H9 was operated at 6 kW is shown in Figure 35. As discussed in 5.4.1, the base cross-mapping skill is significantly higher for this condition due to the dominance of the breathing mode in the system dynamics. As a result, changes in the causal relationships are evaluated at higher levels of relative

reconstruction skill. Yet again, a monotonically declining cross-mapping skill with increased chamber pressure was observed for all signals. The largest change in cross-mapping skill is observed between V_{dis} and V_{cg} in Figure 35(c), however, V_{cg} cross-mappings were strongly influenced by I_{cath} in Figure 35(b), as well. There appears to be a larger change between the two times baseline (16 μ Torr) and the three times baseline (25 μ Torr) condition; however, it is unclear whether this change would remain nonlinear with further increases in pressure.




Figure 35: Maximum cross-mapping skill of each telemetry signal pair with increasing chamber pressure of an H9 operated at 6.0 kW.

Finally, it is important to note that these signals are an incomplete set of the dynamic variables that make up the dynamic plasma system. As such, the causal relationships identified in eCCM are likely a subset of signals related to specific aspects of the underlying dynamic system. An example of this may be seen in the relationship between V_{dis} and V_{cg} in Figure 32 and Figure 33. In both operating conditions, V_{dis} can better reconstruct V_{cg} but with a forward time shift, meaning that V_{dis} can correctly capture dynamics common to both signals, although it may also be driven by additional dynamics not shared with V_{cg} , as indicated by the inability of V_{cg} to accurately cross-map V_{dis} to the same degree. In addition, it appears that V_{dis} responds much quicker, $< 1 \mu s$ relative to 7-9 μs to the common dynamics of I_{cath} , than V_{cg} does, producing results similar to those discussed in when two indirect variables respond to a common driver with different response delays. One might propose a common driver could be I_{cath} however, the negative time-shifts identified in V_{cg} cross-mapping I_{cath} exceeds the delays between V_{dis} and V_{cg} . The apparent indirect relationship with V_{cg} and I_{cath} in Figure 32(c2) indicates a more complicated dynamic

network. A more comprehensive set of dynamic measurements would be needed to further clarify the relationship between the common driver and each signal.

5.6 Discussion

As shown in the analysis of both 4.5-kW and 6-kW thruster operation, there is a reduction in the peak cross-mapping skill of the telemetry signals with elevated chamber pressure. It can be hypothesized that this degradation in causality stems from pressure-dependent changes in the cathode's current pathways. Under normal operation of the thruster in an on-orbit environment, circuit continuity would ensure that cathode current and voltage would remain linked to the anode behavior. However, the presence of background pressure could result in a higher near-field plasma density, enabling electron current flow to the electrically conducting vacuum facility walls or other components, *e.g.*, the thruster body. In this scenario, discharge ions can be neutralized by electrons pulled from the target and chamber wall surfaces with the cathode-supplied electrons arriving to the facility walls elsewhere and flowing within the walls to the location where ions are arriving. This type of cathode-wall coupling behavior could partially decouple the direct neutralizing behavior from the cathode current. To better understand if this is a reasonable proposal for the observed behavior, a larger set of system variables would need to be included in the analysis.

Expanding the set of available signals to better capture the dynamical network could include both high-speed probes and optical measurements coupled with expanded telemetry to improve our ability to capture changes in the underlying dynamics. This would help provide a clearer understanding of how the environmental factors, such as chamber pressure, alter the dynamic relationships that exist between the discharge, cathode, and nearfield plasmas. Another point of interest is to explore how to incorporate environmental conditions that affect causal relations into multi-modal data-driven frameworks that allow the use of signals whose dynamic relationships are sensitive to such factors (*e.g.*, gain adjustments to Kalman filters under specific environmental conditions).

5.7 Summary

This chapter discusses and applies a dynamic causality determination technique, *i.e.*, extended cross-convergent mapping, to a subset of Hall effect thruster telemetry signals to evaluate the causal dynamic relationships among the signals and their sensitivity to an environmental factor: chamber pressure, at two different operating conditions of 4.5 kW and 6 kW. As part of this process, the impact of embedding parameters used for reconstructions in cross-convergent mapping is investigated. It is found that optimal embeddings, in the ranges of τ and m we investigated, were sensitive to all three, although there are entire regions of embedding parameter combinations that yielded relatively high cross-mapping skill. The extended cross-convergent mapping technique is then applied, using the best embedding parameters from the sweeps conducted for each signal set, condition, and pressure combination, to evaluate the cross-mapping skill of the cathode current (I_{cath}), discharge voltage (V_{dis}), and cathode-to-ground voltage (V_{cg}) collected from the H9 thruster tested in Georgia Tech's VTF-2 vacuum facility. Across all pressures and operating conditions that V_{dis} produced the best cross-map of the other telemetry signals, although it appears to be indirectly related V_{cg} . The cathode current I_{cath} initially appears to act unidirectionally on the cathode-to-ground voltage V_{cg} , although increases in chamber pressure in the 4.5-kW operation indicate that this is likely an indirect relationship to an underlying common driver or set of drivers that are not fully captured in the subset of signals. It is observed that chamber pressure negatively affects cross-mapping skill across nearly all cross-mappings, suggesting that facility pressure can alter the dynamic relationships between these signals.

The selection of dynamically relevant signals for data-driven models can be aided by extended cross-convergent mapping, which identifies causal relationships among multiple dynamic signals. Our results indicate that signal selection in one environment may not retain the same information in another, *e.g.*, chamber pressure may affect the quality of the causal relationship of the selected signal. Therefore, models that rely on signals sensitive to such environmental changes should be evaluated to confirm that they retain causal relationships with the variables of interest when implemented with data from multiple

facilities or pressures. Furthermore, our results indicate that discharge voltage measurements retain the most dynamically relevant information among the other signals, even when there are indirect relationships, making it a strong candidate for inclusion in dynamic data-driven models. Finally, the three telemetry signals used here account for less than half of the available circuit telemetry and include no in-vacuum measurements, such as light emission, ion current, or ion energy. However, in situations where the available measurement sets are limited, the present study shows that maximal information can be obtained by selecting signals with the strongest causal relationship with the other signal dynamics.

6 BROADER IMPLICATIONS

The methods for improving measurement and information specificity demonstrated in Chapters 3-5 are useful on their own, but when used within a unifying framework, they become a more directed diagnostic tool to feed the next generation of data-dependent modeling efforts. However, the methods presented have practical hurdles and computational limitations that need to be examined, and this activity may provide key insights into the next steps to further improve them. This chapter synthesizes how the techniques can be used together to further improve the value of diagnostic measurements made with existing diagnostics, thereby achieving greater measurement specificity.

6.1 Measurement Specificity as a Unifying Framework

Difficulty in estimating plasma properties and capturing anomalous mobility effects in traditional Hall thruster modeling has led to reliance on a select few diagnostics to provide the calibration information needed to match experimental data to model results [20]. Data-driven models, such as those demonstrated with EKF frameworks [22], [26], [27], provide a unique avenue for Hall thruster modeling efforts that capture important effects and provide important insights with reduced computational burden by using experimental data to estimate unknown parameters. However, these models can be limited in their usefulness by the level of information contained in their input measurements, not just by the model physics included. The degree to which a diagnostic measurement provides unambiguous, targeted information about a particular plasma property or model parameter influences the efficacy of the data-driven methods in capturing the unknown parameters that define the thruster's operation. Improving measurement specificity in existing diagnostics will not only enable a wider deployment of data-driven frameworks but also advance fundamental studies needed to generate true predictive modeling tools.

Chapters 1-5 presented three ways to improve diagnostic specificity. The first was through species specificity, achieved through serial operation of filtering diagnostics with the combined EVADER probe. The second was through temporal specificity, whereby the temporal resolution of sweep-based

measurements, such as ion-current density profiles, was spatiotemporally reconstructed using a nonlinear reconstruction technique, shadow manifold interpolation. The final form of specificity was causal signal selection, which used a causality analysis tool (eCCM) to guide and justify signal selection for data-driven models and was demonstrated to aid in selecting the most dynamically relevant signals for temporal data-driven modeling frameworks. These improvements in specificity are largely independent, suggesting that further gains in information for modeling efforts could be achieved by combining these techniques.

6.2 Toward Integrated Diagnostic Frameworks

The greatest gains in model information do not come from any single technique in isolation but from composing them into an integrated measurement pipeline. Apart from the use of performance metrics and IVDF profiles used to create specificity in existing non-temporal models [20], [51], it is not often that a combination of measurements with partial information of a system state is integrated in a unified manner to improve state estimation. However, the use of species-resolved energy distributions, temporally resolved ion currents, and a causally linked high-speed telemetry signal could be combined to maximize the information available to a multimodal data-driven model.

A major source of uncertainty in the thruster ion current density profile is the high-angle ion current density, which can strongly influence the estimated beam currents and has been shown to be affected by the facility pressure [39]. Approaches to better identify the true ion current density profile have tried to weight ions by the current density assumed to be part of the beam ions, relative to that of charge exchange. This is done by weighting the amounts of each ion current measurement using angularly swept RPA measurements and selecting a threshold ion energy that can be used to subtract low-energy ion current from the total ion current measured by the Faraday probe [144]. However, using a discriminator grid in front of an ion collector of a Faraday probe, similar to an RPA probe, but with limited high-pass conditions, could quickly evaluate the shape of these plumes. If combined with temporal reconstruction techniques, a unique picture of the temporally resolved beam current with filtered and unfiltered CEX

ions could be evaluated with respect to thruster oscillations. This provides specificity in both the selected ion current profile and the time-dependent spatial profile. Thereby limiting the spatial location in which ions are likely to be generated and accelerated. The fusion of serial filtering for ion specificity and temporal measurement data further specifies the beam ions of interest for models and performance metrics, and also reveals how the facility background pressure affects the thrust plume properties.

Another common challenge is determining the source of a given set of energetic ions that may be causing erosion of the inner pole piece of high-power HETs [145]. For example, if temporally resolved energetic ions exhibit dynamic behavior that is recovered only by the cathode keeper voltage, indicating a causal relationship, whereas the same causal relationship was not found for the discharge current or voltage signals, then the source of the ions can be traced to local plasma phenomena occurring near the keeper and cathode, and not the discharge. In this way, unique causal relationship mapping enables investigations into the origins of convoluted ion species that are not easily distinguishable. In other words, the unique dynamic nature and evaluation of causality would allow for greater specificity in identifying origins of measured plasma properties, which could be used to design experiments to understand the phenomena better and then make changes to the thruster to eliminate undesirable plasma conditions.

This framework can also be used to improve the information gained in these techniques in new ways. For example, if a causally relevant signal can be identified with a key plasma property of interest, such as neutral density, it is possible to use this relationship in a sparse nonlinear reconstruction framework. Light emission from a metastable related to a neutral population, found to be causally related to the discharge voltage or anode current signals, could then be used to reconstruct neutral fluctuations from TALIF signals, which are limited in collection rate to 10 Hz by the pump laser used in such schemes. This is because, although the TALIF signal fires only at 10 Hz, the measurement speed is significantly faster. As a result, the uniqueness of the location within the dynamic manifold where the measurement is made is captured by the other signals that retain a causal relationship to the underlying property (neutral density). In practice, this would mean that high-speed light measurement and causally

linked telemetry signals could be captured before and after a laser firing to provide a temporal reconstruction of the rate-limited measurement at much faster time scales than the 10 Hz repetition rate of the pump laser. The confidence in the manifold location can be confirmed by evaluating signals that include dynamic causal relationships with the signal of interest.

6.3 Implications for Hall Thruster Modeling

Given current computational capabilities and our understanding of Hall thruster operation, physics-based HET models (fluid, hybrid, PIC) require empirically driven anomalous transport closures or methods to speed up computation, which limit the usefulness of the model output. Species-specific IEDFs, provided by tools such as EVADER, for example, could be used to constrain ionization models more tightly by specifying the fraction of ions generated for each species and the potential profile required to produce the resulting energy distributions. Time-resolved measurements, such as those used to reconstruct ion current density profiles, could be used to constrain time-dependent ion production and the resulting temporal anomalous transport [21], [22]. In combination with species-specific profiles, these data can better improve existing physics-based models but may also be used in data-driven frameworks.

The current PC-EKF data-driven frameworks developed for Hall thruster modeling [22], [26], [27] have focused on using discharge current due to its ease of integration into the 0-D global models and its assumed causal relationship with the quantities of interest. Indeed, many other 0-D models use this current fluctuation to predict a given set of thruster effects [146], [147]. However, as shown in Chapter 5, the discharge-voltage signal was more causally relevant than the discharge-current signal. It is possible that the voltage signal is causally linked to more processes occurring in the thruster. As a result, if models can be modified to include discharge-voltage fluctuations by adding a circuit model, they gain access to a data-driven input that has been shown to be more strongly linked to changes in the overall system dynamics. There may also be situations where the signal of interest, such as the cathode current, is causally influenced by the environment in ways that make it less reliable for dynamic data integration.

Using changes in causal relationships over, say, vacuum test chamber conditions, for example, may also inform gain adjustments for a given data stream, as discussed below.

The sensitivity of EP devices to their operating environment further underscores the need to develop models that capture these effects. Environment-aware modeling could leverage metrics such as causal relationships among dynamic data to identify the extent to which facility effects may be affecting dynamic behavior. For example, strong sensitivity of the causal relationship in each pressure range or operating condition would be indicative of model-significant facility influences under such regimes. These facility-aware modes are the goal of institutes like JANUS, which aim to capture the effects of these facilities [148]. Although no dynamic models are currently used in the JANUS predictive engineering model framework, the ability to determine whether facility effect thresholds have been reached can identify key parameters of interest for evaluation.

6.4 Limitations and Scope

There are limitations to the applications of these techniques, both in their applicability and in the duration of data collection during testing and the computational time required to analyze the data. In terms of data collection, a drawback of double-filtered diagnostics such as the EVADER probe is that the total number of sweeps required increases with the energy resolution of the first stage. As discussed in Chapter 1, this is necessary to resolve each species' ion energy sufficiently. Depending on the program used to control the plate voltages and read the collector current, this can increase the time for a single sweep of one filter from 30 seconds up to a total measurement duration of 45 minutes. The data collection scheme has not been optimized and can likely be reduced by 50% or more if different power supplies are used to sweep the voltages over limited ranges where information is contained. In addition, collecting the data required for spatiotemporal fusion generally takes longer than a typical spatial or internal probe voltage sweep. This is due to the increased data volume collected for the SMI process. At each measurement state, the high-speed collector signal and reference signals must be recorded, resulting in many oscilloscope traces being collected for each state rather than just a current or light-intensity

measurement. Although SMI has shown good performance in data sufficiency needs relative to FFT methods [42], the data can still exceed gigabytes for a given fine-step sweep with more than 20 measurements, although compression can help limit the size of these datasets. The required data volume adds complexity to collecting data quickly. Since a full oscilloscope waveform of multiple signals is collected at each measurement state, delays in the sweep must be introduced to allow sufficient time to read and save the oscilloscope data before collecting the next measurement state. This can similarly increase the duration of a 30-second sweep to 10s of minutes. It is also important to evaluate the first and last signals to ensure there is no major change in the dynamic behavior from the beginning to the end of these measurements. This ensures system stationarity and allows data from the beginning or end of the measurement sweep to be reconstructed from the input reference signal.

Stationarity is a critical requirement for the SMI-generated spatiotemporal reconstructions and for eCCM causality evaluations. If the system that these signals are collected from is undergoing a change in the underlying deterministic relationships, then the cross-mapping process may not be valid. In addition to the stationarity requirement, another important aspect is the resolution of the signals. Low-bit-depth signals can limit the available data for reconstruction and make it more difficult to find appropriate, near neighbors in the shadow manifolds. When possible, 12-bit waveforms are preferred over 8-bit, along with high-resolution sensors and probes. In addition to high resolution, the probes should also have sufficient bandwidth to resolve the key time scales of interest. Differential probes are typically easy to find in the tens of MHz, but many low-cost current probes are limited to less than 1 MHz. If interested in resolving high-frequency current oscillations, the probe bandwidth should be greater than 10 times the bandwidth of the oscillation of interest. This enables proper identification of near-neighbor candidates needed to resolve these changes.

The next limitation is due to the time required to analyze the data; both SMI and eCCM must read in many large temporal files and perform piecewise computations at each time step. Shadow manifold generation and cross-mapping for a single signal pair generally take less than a minute or two, depending

on the amount of data collected. The increased computational cost comes with any parametric sweeps that are performed. In Chapter 5, the need to select an appropriate embedding parameter was shown to sometimes require parametric investigation. If the parameter ranges for each delay and dimension pair are used in reconstruction, this can result in a large number of cross-mappings. Furthermore, any additional windowing to generate multiple cross-mappings and evaluate statistical variation will further require additional time. The analysis presented in this dissertation used a Python script that was parallelized on a Linux server with 64 cores. The parallelized computations still took several hours per parametric-sweep, requiring cross-mapping of 3 signals across 2500 parametric signal combinations. We note that this level of calculation may only be required in the initial parametric sweep, and it can be significantly reduced later by changing the window to explore. Additionally, using selection criteria from MITC [125], in which sweeps over the maximal delay locations in both signals are explored, may further reduce the computational cost.

Finally, it is important to discuss what this work does not address. The developments here do not implement new physics-based models or propose a new implementation of extended Kalman filter frameworks. It provides methods to improve the input measurements for such data-driven models by increasing measurement specificity and improving signal selection. The empirical anomalous transport estimation is not improved over existing LIF methods, but rather measurement specificity that better constrains the uniqueness with these traditional downstream diagnostics is presented. The species-resolved, time-resolved, and causally selected measurements described here have not yet been fully integrated into an EKF framework as a closed-loop demonstration. That integration is beyond the scope of this dissertation, but is suggested for possible future work.

7 CONCLUSION

The work presented in this dissertation represents unique contributions to the field of diagnostic measurements in electric propulsion systems. Although these techniques are applied and discussed in the context of the diagnostic frameworks needed to improve understanding of HET physics and modeling, they are in no way limited to only these devices. Their use in plasma processing and fusion applications is directly translatable from the work presented in this dissertation, as the understanding of dynamic phenomena in complex systems with convoluted measurements is common to both fields. Furthermore, as discussed above, the use of these techniques in combination with one another to improve measurement specificity presents a major opportunity for further investigation.

7.1 Summary of Contributions

The intention of this work is to show how improvements can be made to measurement specificity across ion species, temporal, and dynamically causal dimensions to enhance the quality of inputs available to data-driven models for improving HET modeling and characterization. This was demonstrated through three independent investigations, with results published or submitted for peer-reviewed publication.

The combined EVADER probe diagnostic was demonstrated to improve species distributions, with overlapping, wide energy distributions. This probe was also further demonstrated to provide species-specific ion energy distributions. This can show a significant change in the species fractions assumed by a VDF fit with highly overlapping species distributions. Although this does not change the overall thruster efficiency much, it does provide unique information about the plasma in the discharge that would be required to determine such species fractions. The species-specific information of the energy distribution can be used to even further inform the available thruster state-space required to produce such distributions.

Shadow manifold interpolation with time-delay embeddings was used to reconstruct the time-resolved ion-current distribution of a highly oscillatory thruster. Evaluating the time between peaks in the current indicates that the doubly charged ions could be distinguished from singly charged ions by adding such time resolution in such a highly oscillatory regime. This method has been demonstrated for many state-based diagnostics and, in combination with eCCM methods, may provide a way to select appropriate reference signals for reconstructing sparsely measured signals.

Extended cross-convergent mapping was applied to Hall thruster telemetry signals collected while operated under several pressure environments. This was done to demonstrate a method for evaluating the causality of signals to further inform signal selection in data-driven models that require temporal data, and to show how these signals can be affected by environmental conditions such as pressure. A parametric sweep was conducted to select the appropriate embedding parameters, revealing a signal-to-signal dependence and slight pressure dependence. The eCCM results indicated that many relationships that may be thought to be causally related are likely to be indirectly related. Furthermore, the discharge voltage appeared to be the signal that could best reconstruct all other evaluated signals. The strength of the relationships was shown to be pressure-dependent, further highlighting the importance of understanding the environmental conditions under which the data were collected and analyzed.

These three improvements in measurement specificity are needed to bridge the gap between current measurement techniques and the information required for data-driven models to capture thruster physics with downstream probes. Although demonstrated independently in Chapters 3-5, these techniques can be combined to further improve the information obtained from these measurements. Time-resolved, species-specific, ion energy and current density distributions reconstructed from the most causally relevant signal selected for the data-driven framework would be an example of leveraging all techniques together to reduce the available discharge plasma property state space that would have produced that data.

7.2 Future Work

There are several natural next steps that follow from the concepts discussed in Chapter 6. Some of the most immediate would be to apply SMI to the species-specific measurements of an EVADER probe. This presents a challenge because it requires measuring low-current signals at high speed, though adjustments to the probe and measurement circuit could help address it. If a trans-impedance amplifier were integrated more directly near the probe and only had to send amplified signals to vacuum-chamber feedthroughs to measurement equipment located outside the chamber, the measurement speed and signal would likely be achieved. If the temporally resolved species-specific measurements were fused with spatiotemporally fused ion current density distributions, then unique spatial and species information would be redundantly present in the reconstructed signals, further improving the certainty in the temporal reconstruction.

Another near-term solution is to evaluate SMI using causally guided signal selection and estimated uncertainty bands for reconstruction quality. In the majority of reconstructions to date, the discharge current has been the primary signal used to reconstruct time-resolved state-based measurements with SMI. However, as the eCCM results revealed, the discharge current may contain slightly less information than other signals, such as discharge voltage. Additionally, in the SMI framework, there is often a single reconstruction presented for all states, but by windowing sufficiently long segments of the temporal signals and performing the reconstruction several times, any variation from the reconstructions can also be reported as a standard deviation, just as the shaded regions of the eCCM results were. Taken together, causality-related frameworks could enable better reconstruction signal selection, with variability in reconstruction reported to provide additional information about the uncertainty in the reconstruction that can be leveraged with the standard diagnostic analysis to propagate a better understanding of the uncertainty in the temporal measurement.

As mentioned earlier, some medium-term efforts could focus on incorporating the information into a closed-loop EKF framework, where the identified signal for data-driven modeling is used within a

multimodal framework to better inform model parameters, including anomalous electron transport. The combination of species-specific ion energy information with temporal data could demonstrate the degree to which the state space is constrained by these data. Also in the near term, extending the SMI and eCCM frameworks further independently would be useful. As discussed in Chapter 6, the view of the deterministic dynamical view of Hall thruster may allow the use of additional temporal signals to accurately estimate the manifold placement of sparse data collected by rate-limited systems. This would enable measurement systems, such as Thomson scattering and TALIF, that rely on rate-limited pump-laser firings to produce temporally resolved measurements.

Another extension of existing techniques would be a more comprehensive investigation of Hall thruster telemetry tested in imperfect facilities, which could help to better fill in the missing dynamic links identified in Chapter 5. Specifically, the inclusion of both anode and cathode current would be helpful in understanding the degree to which the two plasma devices are affected by the facility in which they are tested. Another direct extension would be to apply the eCCM framework to the telemetry of thrusters tested in clusters. Many high-power nuclear electric propulsion architectures plan to leverage Hall thruster clusters to provide sustained thrust, but how these thrusters will interact dynamically is still an outstanding question [149]. The eCCM framework could be applied to assess how the cathode-discharge causal relationship is affected when another thruster is operated in close proximity and to quantify the degree to which thrusters interact with each other. For example, if the distinct dynamic data of thruster A's discharge is reproducible in the current or cathode-to-ground voltage of thruster B's cathode, the degree to which they are causally linked could be identified and examined more closely. This would have implications for understanding the conditions under which the thrusters may be better suited to operate together as well as the degree to which these metrics change with increased thruster spacing.

As the community moves toward predictive engineering models that integrate experiments and simulations with improved uncertainty quantification, the specificity of diagnostic measurements will become the rate-limiting factor in reducing model uncertainty. The techniques for specificity in species

selection, temporal resolution, and causal signal selection present a framework to address that bottleneck with existing diagnostics. As Lord Kelvin would describe it, these measurements allow us to “know” so that we may “improve”.

REFERENCES

- [1] D. M. Goebel and I. Katz, *Fundamentals of electric propulsion: ion and Hall thrusters*. John Wiley & Sons, 2008.
- [2] K. A. Polzin, A. K. Martin, F. M. Curran, R. M. Myers, and M. A. Rodriguez, “Strategy for developing technologies for megawatt-class nuclear electric propulsion systems,” in *37th International Electric Propulsion Conference*, Boston, MA, 2022.
- [3] J. Brophy *et al.*, “Non-Nuclear Exploration of the Solar System Study,” May 2022. doi: 10.7907/h62p-6328.
- [4] J. S. Snyder *et al.*, “Initial checkout of the Psyche electric propulsion system,” *Journal of Electric Propulsion*, vol. 4, no. 1, p. 53, 2025, doi: 10.1007/s44205-025-00148-8.
- [5] D. Y. Oh *et al.*, “Design and Development of the Psyche Spacecraft for NASA’s Discovery Program,” *Space Sci. Rev.*, vol. 221, no. 7, p. 96, 2025, doi: 10.1007/s11214-025-01222-1.
- [6] R. R. Hofer, J. B. Simmonds, D. M. Goebel, J. M. Steinkraus, and A. R. Payman, “The H10 high power density hall thruster,” *Journal of Electric Propulsion*, vol. 4, no. 1, p. 35, 2025, doi: 10.1007/s44205-025-00129-x.
- [7] R. Hofer *et al.*, “Development and qualification of a 12-kW hall thruster for deep-space missions,” in *73rd International Astronautical Congress (IAC)*, Sep. 2022.
- [8] K. Diamant, R. Spektor, E. Beiting, J. Young, and T. Curtiss, “The Effects of Background Pressure on Hall Thruster Operation,” in *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, in Joint Propulsion Conferences. , American Institute of Aeronautics and Astronautics, 2012. doi: doi:10.2514/6.2012-3735.
- [9] K. D. Diamant, R. Liang, and R. L. Corey, “The Effect of Background Pressure on SPT-100 Hall Thruster Performance,” in *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, in AIAA Propulsion and Energy Forum. , American Institute of Aeronautics and Astronautics, 2014. doi: doi:10.2514/6.2014-3710.
- [10] J. S. Snyder, G. Lenguito, J. D. Frieman, T. W. Haag, and J. A. Mackey, “Effects of Background Pressure on SPT-140 Hall Thruster Performance,” *J. Propuls. Power*, vol. 36, no. 5, pp. 668–676, May 2020, doi: 10.2514/1.B37702.
- [11] W. Huang, H. Kamhawi, and P. Y. Peterson, “Effects of Background Pressure and Electrical Configuration on the Velocity Field of the HERMeS Hall Thruster,” in *AIAA Propulsion and Energy 2020 Forum*, in AIAA Propulsion and Energy Forum. American Institute of Aeronautics and Astronautics, Aug. 2020. doi: doi:10.2514/6.2020-3616.

- [12] J. A. Walker, S. J. Langendorf, M. L. R. Walker, V. Khayms, D. King, and P. Pertson, “Electrical Facility Effects on Hall Current Thrusters: Electron Termination Pathway Manipulation,” *J. Propuls. Power*, vol. 32, no. 6, pp. 1365–1377, Aug. 2016, doi: 10.2514/1.B35904.
- [13] J. Walker, D. Lev, M. L. R. Walker, V. Khayms, and D. King, “Electrical characteristics of a Hall effect thruster body in a vacuum facility testing environment,” *Journal of Electric Propulsion*, vol. 1, no. 1, p. 18, 2022, doi: 10.1007/s44205-022-00016-9.
- [14] D. R. Jovel, “Impedance Characterization of a Hall Effect Thruster Discharge in a Ground-Based Vacuum Test Facility,” Georgia Institute of Technology, 2024. [Online]. Available: <https://www.proquest.com/dissertations-theses/impedance-characterization-hall-effect-thruster/docview/3275478504/se-2?accountid=10223>
- [15] R. B. Lobbia, J. E. Polk, R. R. Hofer, V. H. Chaplin, and B. Jorns, “Accelerating 23,000 hours of ground test backspattered carbon on a magnetically shielded Hall thruster,” in *AIAA Propulsion and Energy 2019 Forum*, 2019, p. 3898.
- [16] J. H. Gilland, G. Williams, J. M. Burt, and J. Yim, “Carbon Back Sputter Modeling for Hall Thruster Testing,” in *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, 2016, p. 4941.
- [17] K. Hara, “An overview of discharge plasma modeling for Hall effect thrusters,” *Plasma Sources Sci. Technol.*, vol. 28, no. 4, p. 044001, 2019, doi: 10.1088/1361-6595/ab0f70.
- [18] J.-P. Boeuf, “Tutorial: Physics and modeling of Hall thrusters,” *J. Appl. Phys.*, vol. 121, no. 1, p. 011101, Jan. 2017, doi: 10.1063/1.4972269.
- [19] I. D. Kaganovich *et al.*, “Physics of $E \times B$ discharges relevant to plasma propulsion and similar technologies,” *Phys. Plasmas*, vol. 27, no. 12, p. 120601, Dec. 2020, doi: 10.1063/5.0010135.
- [20] I. G. Mikellides and A. L. 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 Sci. Technol.*, vol. 28, no. 1, p. 014003, 2019.
- [21] E. T. Dale and B. A. Jorns, “Non-invasive time-resolved measurements of anomalous collision frequency in a Hall thruster,” *Phys. Plasmas*, vol. 26, no. 1, p. 013516, Jan. 2019, doi: 10.1063/1.5077008.
- [22] D. E. Troyetsky, C. M. Greve, S. Tsikata, and K. Hara, “State estimation of the dynamic behavior of plasma properties in a Hall effect thruster discharge,” *J. Phys. D Appl. Phys.*, vol. 56, no. 44, p. 444001, 2023, doi: 10.1088/1361-6463/ace7da.
- [23] E. T. Dale and B. A. Jorns, “Experimental characterization of Hall thruster breathing mode dynamics,” *J. Appl. Phys.*, vol. 130, no. 13, p. 133302, 2021.
- [24] T. Lafleur, P. Chabert, and A. Bourdon, “The origin of the breathing mode in Hall thrusters and its stabilization,” *J. Appl. Phys.*, vol. 130, no. 5, p. 053305, Aug. 2021, doi: 10.1063/5.0057095.
- [25] J. Gottfried, S. Antozzi, J. Stienike, S. Thompson, J. Williams, and A. P. Yalin, “Temporally Resolved Neutral Density Measurements of Hall Effect Thruster Breathing Mode by Two-Photon Absorption Laser Induced Fluorescence (TALIF),” in *AIAA SCITECH 2024 Forum*, 2024, p. 0802.

- [26] C. M. Greve, M. Majji, and K. Hara, “Real-time state estimation of low-frequency plasma oscillations in Hall effect thrusters,” *Phys. Plasmas*, vol. 28, no. 9, p. 093509, Sep. 2021, doi: 10.1063/5.0057751.
- [27] T. Zivre and K. Hara, “Estimation of in-space thruster performance of a Hall effect thruster through extrapolation of plasma states using extended Kalman filter,” *J. Appl. Phys.*, vol. 138, no. 16, p. 163305, Oct. 2025, doi: 10.1063/5.0284774.
- [28] C. M. Greve, “Real-time state estimation in plasma modeling applications,” Texas A&M University, 2022.
- [29] C. C. Farnell, C. C. Farnell, S. C. Farnell, and J. D. Williams, “Recommended Practice for Use of Electrostatic Analyzers in Electric Propulsion Testing,” *J. Propuls. Power*, vol. 33, no. 3, pp. 638–658, Nov. 2016, doi: 10.2514/1.B35413.
- [30] K. M. Lemmer, A. J. Thomas, W. Huang, D. M. Goebel, and R. B. Lobbia, “Recommended Practices for Design and Use of Retarding Potential Analyzers in Electric Propulsion,” in *International Electric Propulsion Conference*, 2025.
- [31] J. L. Rovey *et al.*, “Recommended Practice for Use of ExB Probes in Electric Propulsion Testing,” in *39th International Electric Propulsion Conference (IEPC)*, 2025.
- [32] R. B. Lobbia and B. E. Beal, “Recommended Practice for Use of Langmuir Probes in Electric Propulsion Testing,” *J. Propuls. Power*, vol. 33, no. 3, pp. 566–581, Apr. 2017, doi: 10.2514/1.B35531.
- [33] B. Jorns, D. M. Goebel, and R. R. Hofer, “Plasma Perturbations in High-Speed Probing of Hall Thruster Discharge Chambers: Quantification and Mitigation,” in *51st AIAA/SAE/ASEE Joint Propulsion Conference*, in AIAA Propulsion and Energy Forum. , American Institute of Aeronautics and Astronautics, 2015. doi: doi:10.2514/6.2015-4006.
- [34] R. R. Hofer and A. D. Gallimore, “High-specific impulse Hall thrusters, part 2: efficiency analysis,” *J. Propuls. Power*, vol. 22, no. 4, pp. 732–740, 2006.
- [35] V. H. Chaplin *et al.*, “Recommended Practices in Laser-Induced Fluorescence (LIF) Diagnostics for Electric Propulsion,” in *International Electric Propulsion Conference 2025*, 2025, p. IEPC-2025.
- [36] K. Hara, V. H. Chaplin, L. K. Johnson, G. Xu, and D. Levin, “Recommended Practices in Electric Propulsion for Optical Emission Spectroscopy,” in *39th International Electric Propulsion Conference*, London, UK, Sep. 2025.
- [37] D. M. Goebel, I. Katz, and I. G. Mikellides, *Fundamentals of Electric Propulsion*. John Wiley & Sons, 2023.
- [38] W. Huang and R. Shastry, “Analysis of Wien filter spectra from Hall thruster plumes,” *Review of Scientific Instruments*, vol. 86, no. 7, p. 073502, Jul. 2015, doi: 10.1063/1.4923282.

- [39] D. L. Brown, M. L. R. Walker, J. Szabo, W. Huang, and J. E. Foster, “Recommended Practice for Use of Faraday Probes in Electric Propulsion Testing,” *J. Propuls. Power*, vol. 33, no. 3, pp. 582–613, Sep. 2016, doi: 10.2514/1.B35696.
- [40] R. B. Lobbia, “A time-resolved investigation of the Hall thruster breathing mode,” 2010.
- [41] C. J. Durot, “Development of a Time-Resolved Laser-Induced Fluorescence Technique for Nonperiodic Oscillations,” University of Michigan, 2016.
- [42] D. Eckhardt, J. Koo, R. Martin, M. Holmes, and K. Hara, “Spatiotemporal data fusion and manifold reconstruction in Hall thrusters,” *Plasma Sources Sci. Technol.*, vol. 28, no. 4, p. 045005, 2019, doi: 10.1088/1361-6595/ab0b1f.
- [43] R. Hofer *et al.*, “Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations,” in *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, in Joint Propulsion Conferences. , American Institute of Aeronautics and Astronautics, 2008. doi: doi:10.2514/6.2008-4924.
- [44] E. Choueiri, “An overview of plasma oscillations in Hall thrusters,” *Spacecraft Propulsion*, vol. 465, p. 287, 2000.
- [45] M. S. McDonald and A. D. Gallimore, “Rotating spoke instabilities in Hall thrusters,” *IEEE Transactions on Plasma Science*, vol. 39, no. 11, pp. 2952–2953, 2011.
- [46] S. Tsikata, “Small-Scale Electron Density Fluctuations in the Hall Thruster, Investigated by Collective Light Scattering,” Ecole Polytechnique X, 2009. [Online]. Available: <https://pastel.hal.science/tel-00484027>
- [47] M. L. R. Walker *et al.*, “Overview of the joint advanced propulsion institute (JANUS),” in *37th International Electric Propulsion Conference*, 2022.
- [48] J. P. Sheehan, Y. Raitses, N. Hershkowitz, and M. McDonald, “Recommended Practice for Use of Emissive Probes in Electric Propulsion Testing,” *J. Propuls. Power*, vol. 33, no. 3, pp. 614–637, Nov. 2016, doi: 10.2514/1.B35697.
- [49] R. Hofer, D. Goebel, I. Mikellides, and I. Katz, “Design of a laboratory Hall thruster with magnetically shielded channel walls, phase II: Experiments,” in *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2012, p. 3788.
- [50] W. Huang, R. Shastry, G. C. Soulas, and H. Kamhawi, “Farfield Plume Measurement and Analysis on the NASA-300M,” in *60th JANNAF Propulsion Meeting*, 2013.
- [51] L. L. Su, T. A. Marks, and B. A. Jorns, “Trends in mass utilization of a magnetically shielded Hall thruster operating on xenon and krypton,” *Plasma Sources Sci. Technol.*, vol. 33, no. 6, p. 065008, 2024, doi: 10.1088/1361-6595/ad52be.
- [52] T. Wegner, S. J. Thompson, J. Williams, and A. P. Yalin, “Two-Photon Absorption Laser Induced Fluorescence (TALIF) Of Neutral Xenon In A Hall Effect Thruster Plasma,” in *AIAA Propulsion and Energy 2021 Forum*, in AIAA Propulsion and Energy Forum. , American Institute of Aeronautics and Astronautics, 2021. doi: doi:10.2514/6.2021-3391.

- [53] J. A. Gottfried, S. Antozzi, J. Stienike, S. J. Thompson, J. D. Williams, and A. P. Yalin, “Temporally resolved relative krypton neutral density during breathing mode of a hall effect thruster recorded by TALIF,” *Journal of Electric Propulsion*, vol. 3, no. 1, p. 9, 2024, doi: 10.1007/s44205-024-00070-5.
- [54] S. Tsikata, N. Lemoine, V. Pisarev, and D. M. Gresillon, “Dispersion relations of electron density fluctuations in a Hall thruster plasma, observed by collective light scattering,” *Phys. Plasmas*, vol. 16, no. 3, p. 33506, 2009.
- [55] S. Tsikata, K. Hara, and S. Mazouffre, “Characterization of hollow cathode plasma turbulence using coherent Thomson scattering,” *J. Appl. Phys.*, vol. 130, no. 24, p. 243304, Dec. 2021, doi: 10.1063/5.0071650.
- [56] P. J. Roberts and B. A. Jorns, “Laser Measurement of Anomalous Electron Diffusion in a Crossed-Field Plasma,” *Phys. Rev. Lett.*, vol. 132, no. 13, p. 135301, Mar. 2024, doi: 10.1103/PhysRevLett.132.135301.
- [57] J. L. S. Betancourt, N. Butler-Craig, J. Lopez-Uricoechea, A. M. Steinberg, and M. L. R. Walker, “Laser Thomson scattering measurements indicate non-isothermal magnetic field lines in magnetically shielded Hall effect thrusters,” *Phys. Plasmas*, vol. 31, no. 11, p. 113106, Nov. 2024, doi: 10.1063/5.0229827.
- [58] C. Eichhorn *et al.*, “Single- and two-photon absorption laser-induced fluorescence spectroscopy in rare gases for gridded ion thruster diagnostics,” *EPJ Techn Instrum*, vol. 9, no. 1, 2022, [Online]. Available: <https://doi.org/10.1140/epjti/s40485-022-00077-y>
- [59] J. Lopez-Uricoechea, J. L. Suazo Betancourt, N. Butler-Craig, and M. L. R. Walker, “Spatially resolved Thomson scattering measurements of electron properties across the acceleration region of a high-power magnetically shielded Hall effect thruster,” *J. Appl. Phys.*, vol. 136, no. 11, p. 113303, Sep. 2024, doi: 10.1063/5.0213928.
- [60] B. Vincent, S. Tsikata, and S. Mazouffre, “Incoherent Thomson scattering measurements of electron properties in a conventional and magnetically-shielded Hall thruster,” *Plasma Sources Sci. Technol.*, vol. 29, no. 3, p. 035015, 2020, doi: 10.1088/1361-6595/ab6c42.
- [61] A. J. Friss and A. P. Yalin, “Cavity-enhanced Thomson scattering measurements of electron density and temperature in a hollow cathode discharge,” *Opt. Lett.*, vol. 43, no. 21, pp. 5343–5346, 2018, doi: 10.1364/OL.43.005343.
- [62] B. A. Jorns and R. R. Hofer, “Plasma oscillations in a 6-kW magnetically shielded Hall thruster,” *Phys. Plasmas*, vol. 21, no. 5, p. 53512, 2014.
- [63] N. Kuwabara, M. Chono, N. Yamamoto, and D. Kuwahara, “Electron Density Measurement Inside a Hall Thruster Using Microwave Interferometry,” *J. Propuls. Power*, vol. 37, no. 3, pp. 491–494, Mar. 2021, doi: 10.2514/1.B38163.
- [64] C. R. Mullins *et al.*, “Non-invasive Hall current distribution measurement in a Hall effect thruster,” *Review of Scientific Instruments*, vol. 88, no. 1, p. 13507, 2017.

- [65] K. Hara, M. J. Sekerak, I. D. Boyd, and A. D. Gallimore, “Perturbation analysis of ionization oscillations in Hall effect thrusters,” *Phys. Plasmas*, vol. 21, no. 12, p. 122103, Dec. 2014, doi: 10.1063/1.4903843.
- [66] T. Marks, P. Schedler, and B. Jorns, “HallThruster.jl: a Julia package for 1D Hall thruster discharge simulation,” *J. Open Source Softw.*, vol. 8, no. 86, p. 4672, 2023, doi: 10.21105/joss.04672.
- [67] I. G. Mikellides and I. Katz, “Numerical simulations of Hall-effect plasma accelerators on a magnetic-field-aligned mesh,” *Phys. Rev. E*, vol. 86, no. 4, p. 46703, Oct. 2012, doi: 10.1103/PhysRevE.86.046703.
- [68] S. Cho, K. Komurasaki, and Y. Arakawa, “Kinetic particle simulation of discharge and wall erosion of a Hall thruster,” *Phys. Plasmas*, vol. 20, no. 6, p. 063501, Jun. 2013, doi: 10.1063/1.4810798.
- [69] J. C. Adam, A. Héron, and G. Laval, “Study of stationary plasma thrusters using two-dimensional fully kinetic simulations,” *Phys. Plasmas*, vol. 11, no. 1, pp. 295–305, Jan. 2004, doi: 10.1063/1.1632904.
- [70] M. D. Campanell, A. V Khrabrov, and I. D. Kaganovich, “General Cause of Sheath Instability Identified for Low Collisionality Plasmas in Devices with Secondary Electron Emission,” *Phys. Rev. Lett.*, vol. 108, no. 23, p. 235001, Jun. 2012, doi: 10.1103/PhysRevLett.108.235001.
- [71] K. Hara and S. Tsikata, “Cross-field electron diffusion due to the coupling of drift-driven microinstabilities,” *Phys. Rev. E*, vol. 102, no. 2, p. 23202, 2020.
- [72] L. A. McGee and S. F. Schmidt, “Discovery of the Kalman filter as a practical tool for aerospace and industry,” 1985.
- [73] C. M. Greve and K. Hara, “Estimation of plasma properties using an extended Kalman filter with plasma global models,” *J. Phys. D Appl. Phys.*, vol. 55, no. 25, p. 255201, 2022, doi: 10.1088/1361-6463/ac5c1c.
- [74] I. G. Mikellides, A. L. Ortega, V. H. Chaplin, and J. S. Snyder, “Facility pressure effects on a Hall thruster with an external cathode, II: theoretical model of the thrust and the significance of azimuthal asymmetries in the cathode plasma,” *Plasma Sources Sci. Technol.*, vol. 29, no. 3, p. 035010, 2020, doi: 10.1088/1361-6595/ab6c7f.
- [75] H. Kantz and T. Schreiber, *Nonlinear time series analysis*. Cambridge University Press, 2003.
- [76] F. Takens, “Detecting strange attractors in turbulence,” in *Dynamical Systems and Turbulence, Warwick 1980*, D. Rand and L.-S. Young, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 1981, pp. 366–381.
- [77] G. Sugihara *et al.*, “Detecting Causality in Complex Ecosystems,” *Science (1979)*, vol. 338, no. 6106, pp. 496–500, Oct. 2012, doi: 10.1126/science.1227079.
- [78] H. Ye, E. R. Deyle, L. J. Gilarranz, and G. Sugihara, “Distinguishing time-delayed causal interactions using convergent cross mapping,” *Sci. Rep.*, vol. 5, no. 1, p. 14750, 2015, doi: 10.1038/srep14750.

- [79] C. W. J. Granger, “Investigating Causal Relations by Econometric Models and Cross-spectral Methods,” *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969, doi: 10.2307/1912791.
- [80] R. R. Hofer and A. D. Gallimore, “Development and characterization of high-efficiency, high-specific impulse xenon Hall thrusters,” Ph.D. Dissertation, University of Michigan, Ann Arbor, 2004. [Online]. Available: <https://www.proquest.com/dissertations-theses/development-characterization-high-efficiency/docview/305180370/se-2?accountid=10223>
- [81] R. Shastry, R. R. Hofer, B. M. Reid, and A. D. Gallimore, “Method for analyzing E×B probe spectra from Hall thruster plumes,” *Review of Scientific Instruments*, vol. 80, no. 6, p. 063502, Jun. 2009, doi: 10.1063/1.3152218.
- [82] J. Szabo and M. Robin, “Plasma Species Measurements in the Plume of an Iodine Fueled Hall Thruster,” *J. Propuls. Power*, vol. 30, no. 5, pp. 1357–1367, Aug. 2014, doi: 10.2514/1.B35075.
- [83] A. C. Hillier, “Revolutionizing space propulsion through the characterization of iodine as fuel for Hall-effect thrusters,” M.S. Thesis, Air Force Institute of Technology, 2011.
- [84] A. Gurciullo, A. L. Fabris, and M. A. Cappelli, “Ion plume investigation of a Hall effect thruster operating with Xe/N₂ and Xe/air mixtures,” *J. Phys. D Appl. Phys.*, vol. 52, no. 46, p. 464003, 2019, doi: 10.1088/1361-6463/ab36c5.
- [85] T. M. Andreano, “Performance and Plume Characterization of a Laboratory Krypton Hall Thruster,” MS Thesis, 2020.
- [86] T. Andreano, J. Williams, M. Clement, and C. Farnell, “Performance Comparison of a 1.5 kW Hall Thruster with Center-Mounted and Outer-Pole-Mounted Heaterless Cathodes,” in *AIAA Propulsion and Energy 2019 Forum*, in AIAA Propulsion and Energy Forum. American Institute of Aeronautics and Astronautics, Aug. 2019. doi: doi:10.2514/6.2019-4083.
- [87] C. Bundesmann *et al.*, “Note: An advanced in situ diagnostic system for characterization of electric propulsion thrusters and ion beam sources,” *Review of Scientific Instruments*, vol. 81, no. 4, p. 046106, Apr. 2010, doi: 10.1063/1.3386585.
- [88] D. L. Brown, “Investigation of low discharge voltage Hall thruster characteristics and evaluation of loss mechanisms,” Ph.D. Dissertation, 2009.
- [89] C. C. Farnell, D. L. Brown, G. Willis, R. Branam, and J. D. Williams, “Remote diagnostic measurements of Hall thruster plumes,” Air Force Research Lab Edwards AFB CA Propulsion Directorate, 2009.
- [90] SpaceX, “Among other enhancements, V2 minis are equipped with new argon Hall thrusters for on orbit maneuvering,” X (formerly Twitter).
- [91] J. L. Rovey *et al.*, “Review of multimode space propulsion,” *Progress in Aerospace Sciences*, vol. 118, p. 100627, 2020, doi: <https://doi.org/10.1016/j.paerosci.2020.100627>.
- [92] M. Nakles, L. Brieda, G. Reed, W. Hargus, and R. Spicer, “Experimental and Numerical Examination of the BHT-200 Hall Thruster Plume,” in *43rd AIAA/ASME/SAE/ASEE Joint*

- Propulsion Conference & Exhibit*, in Joint Propulsion Conferences. , American Institute of Aeronautics and Astronautics, 2007. doi: doi:10.2514/6.2007-5305.
- [93] B. E. Beal, A. D. Gallimore, J. M. Haas, and W. A. Hargus, “Plasma Properties in the Plume of a Hall Thruster Cluster,” *J. Propuls. Power*, vol. 20, no. 6, pp. 985–991, Nov. 2004, doi: 10.2514/1.3765.
- [94] G. Reed and W. Hargus, “Complementary Density Measurements for the 200-X3 Hall Thruster Plume,” in *42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, in Joint Propulsion Conferences. , American Institute of Aeronautics and Astronautics, 2006. doi: doi:10.2514/6.2006-4992.
- [95] W. Hargus, “Laser-Induced Fluorescence of Neutral Xenon in the Near Field of a 200 W Hall Thruster,” in *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, in Joint Propulsion Conferences. , American Institute of Aeronautics and Astronautics, 2005. doi: doi:10.2514/6.2005-4400.
- [96] W. A. Hargus and C. S. Charles, “Near Exit Plane Velocity Field of a 200-Watt Hall Thruster,” *J. Propuls. Power*, vol. 24, no. 1, pp. 127–133, Jan. 2008, doi: 10.2514/1.29949.
- [97] J. Ekholm and W. Hargus, “E x B measurements of a 200 W xenon Hall thruster,” in *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2005, p. 4405.
- [98] D. L. Brown and A. D. Gallimore, “Evaluation of facility effects on ion migration in a Hall thruster plume,” *J. Propuls. Power*, vol. 27, no. 3, pp. 573–585, 2011.
- [99] J. D. Frieman, J. A. Walker, M. L. R. Walker, V. Khayms, and D. Q. King, “Electrical Facility Effects on Hall Thruster Cathode Coupling: Performance and Plume Properties,” *J. Propuls. Power*, vol. 32, no. 1, pp. 251–264, Aug. 2015, doi: 10.2514/1.B35683.
- [100] W. Huang, H. Kamhawi, and T. Haag, “Plasma oscillation characterization of NASA’s HERMeS Hall thruster via high speed imaging,” in *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, 2016, p. 4829.
- [101] M. L. Hause, B. D. Prince, and R. J. Bemish, “Krypton charge exchange cross sections for Hall effect thruster models,” *J. Appl. Phys.*, vol. 113, no. 16, p. 163301, Apr. 2013, doi: 10.1063/1.4802432.
- [102] I. Katz, R. R. Hofer, and D. M. Goebel, “Ion Current in Hall Thrusters,” *IEEE Transactions on Plasma Science*, vol. 36, no. 5, pp. 2015–2024, 2008, doi: 10.1109/TPS.2008.2004219.
- [103] V. V Zhurin, H. R. Kaufman, and R. S. Robinson, “Physics of closed drift thrusters,” *Plasma Sources Sci. Technol.*, vol. 8, no. 1, p. R1, 1999, doi: 10.1088/0963-0252/8/1/021.
- [104] H. Price, S. Thompson, C. Farnell, C. Farnell, and J. Williams, “A Model of a Combined ESA and ExB Probe,” in *39th International Electric Propulsion Conference*, London, UK, Sep. 2025.
- [105] N. Gascon, M. Dudeck, and S. Barral, “Wall material effects in stationary plasma thrusters. I. Parametric studies of an SPT-100,” *Phys. Plasmas*, vol. 10, no. 10, pp. 4123–4136, Oct. 2003, doi: 10.1063/1.1611880.

- [106] B. A. Jorns, I. G. Mikellides, and D. M. Goebel, “Ion acoustic turbulence in a 100-A LaB6 hollow cathode,” *Phys. Rev. E*, vol. 90, no. 6, p. 63106, Dec. 2014, doi: 10.1103/PhysRevE.90.063106.
- [107] K. Hara and C. Treece, “Ion kinetics and nonlinear saturation of current-driven instabilities relevant to hollow cathode plasmas,” *Plasma Sources Sci. Technol.*, vol. 28, no. 5, p. 055013, 2019, doi: 10.1088/1361-6595/ab18e4.
- [108] S. Tsikata, K. Hara, and S. Mazouffre, “Characterization of hollow cathode plasma turbulence using coherent Thomson scattering,” *J. Appl. Phys.*, vol. 130, no. 24, p. 243304, Dec. 2021, doi: 10.1063/5.0071650.
- [109] A. P. Scott and D. M. Goebel, “Hollow cathode discharge instability onset in electric thrusters,” *J. Appl. Phys.*, vol. 135, no. 12, p. 123301, Mar. 2024, doi: 10.1063/5.0188988.
- [110] C. V Young, A. L. Fabris, N. A. MacDonald-Tenenbaum, W. A. Hargus, and M. A. Cappelli, “Time-resolved laser-induced fluorescence diagnostics for electric propulsion and their application to breathing mode dynamics,” *Plasma Sources Sci. Technol.*, vol. 27, no. 9, p. 094004, 2018, doi: 10.1088/1361-6595/aade42.
- [111] N. A. MacDonald, M. A. Cappelli, and W. A. Hargus Jr., “Time-synchronized continuous wave laser-induced fluorescence on an oscillatory xenon discharge,” *Review of Scientific Instruments*, vol. 83, no. 11, p. 113506, Nov. 2012, doi: 10.1063/1.4766958.
- [112] A. Diallo, S. Keller, Y. Shi, Y. Raitses, and S. Mazouffre, “Time-resolved ion velocity distribution in a cylindrical Hall thruster: Heterodyne-based experiment and modeling,” *Review of Scientific Instruments*, vol. 86, no. 3, p. 033506, Mar. 2015, doi: 10.1063/1.4914829.
- [113] V. H. Chaplin *et al.*, “Time-resolved ion velocity measurements in a high-power Hall thruster using laser-induced fluorescence with transfer function averaging,” *Appl. Phys. Lett.*, vol. 116, no. 23, p. 234107, Jun. 2020, doi: 10.1063/5.0007161.
- [114] A. Jardin and P. Lubiński, “Searching for chaotic behavior in the experimental ion current and discharge current waveforms of a Hall effect thruster,” *Phys. Plasmas*, vol. 31, no. 4, p. 042114, Apr. 2024, doi: 10.1063/5.0197676.
- [115] M. Baird, R. McGee-Sinclair, K. Lemmer, and W. Huang, “Time-resolved ion energy measurements using a retarding potential analyzer,” *Review of Scientific Instruments*, vol. 92, no. 7, p. 073306, 2021.
- [116] A. Thomas and K. Lemmer, “Time-resolved ion energy measurements using a retarding potential analyzer for electric propulsion applications,” *Review of Scientific Instruments*, vol. 95, no. 2, p. 023505, Feb. 2024, doi: 10.1063/5.0176167.
- [117] M. J. Baird, T. A. Kerber, K. M. Lemmer, and W. Huang, “Hall Thruster Plume Measurements of Time Resolved Ion Energy,” in *36th International Electric Propulsion Conference*, Vienna, Austria, 2019.
- [118] A. Thomas, K. Lemmer, J. Cabrera, and M. L. R. Walker, “High-Speed Far-Field Energy Measurements in a 9-kW Hall Effect Thruster for Model Inference and Calibration,” in *38th International Electric Propulsion Conference*, Toulouse, France, 2024.

- [119] P. J. Roberts, V. H. Chaplin, and B. A. Jorns, “Time-Resolved Laser Diagnostics in Hall Thrusters with Time-Delay Embedding Data Fusion,” in *38th International Electric Propulsion Conference*, Toulouse, France, 2024.
- [120] Z. K. Robertson *et al.*, “Single Power Supply Operation of a Hall Thruster,” in *38th International Electric Propulsion Conference*, Toulouse, France, 2024.
- [121] C. M. Greve, K. Hara, R. S. Martin, D. Q. Eckhardt, and J. W. Koo, “A data-driven approach to model calibration for nonlinear dynamical systems,” *J. Appl. Phys.*, vol. 125, no. 24, p. 244901, Jun. 2019, doi: 10.1063/1.5085780.
- [122] G. Sugihara and R. M. May, “Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series,” *Nature*, vol. 344, no. 6268, pp. 734–741, 1990, doi: 10.1038/344734a0.
- [123] E. George *et al.*, “Decomposing Signals from Dynamical Systems Using Shadow Manifold Interpolation,” *SIAM J. Appl. Dyn. Syst.*, vol. 20, no. 4, pp. 2236–2260, Jan. 2021, doi: 10.1137/20M1350923.
- [124] A. M. Fraser and H. L. Swinney, “Independent coordinates for strange attractors from mutual information,” *Phys. Rev. A (Coll Park)*, vol. 33, no. 2, p. 1134, 1986.
- [125] R. S. Martin, C. M. Greve, C. E. Huerta, A. S. Wong, J. W. Koo, and D. Q. Eckhardt, “A robust time-delay selection criterion applied to convergent cross mapping,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 34, no. 9, p. 093110, Sep. 2024, doi: 10.1063/5.0209028.
- [126] L. Cao, “Practical method for determining the minimum embedding dimension of a scalar time series,” *Physica D*, vol. 110, no. 1–2, pp. 43–50, 1997.
- [127] C. Huerta, C. Greve, and A. Wong, “Comparison of causality determination techniques in studying Hall-effect thrusters,” *Journal of Electric Propulsion*, vol. 3, no. 1, p. 23, 2024, doi: 10.1007/s44205-024-00084-z.
- [128] B. A. Jorns, M. Byrne, P. Roberts, L. Su, E. Dale, and R. R. Hofer, “Prediction and Mitigation of the Mode Transition in a Magnetically Shielded Hall Thruster at High-Specific Impulse and Low Power,” 2022.
- [129] S. J. Thompson, S. C. Farnell, C. C. Farnell, C. C. Farnell, T. M. Andreano, and J. D. Williams, “Combined electrostatic analyzer—Wien filter probe for characterization of species distributions in Hall thrusters,” *J. Appl. Phys.*, vol. 130, no. 23, p. 233302, Dec. 2021, doi: 10.1063/5.0071656.
- [130] E. T. Dale and B. A. Jorns, “Experimental characterization of Hall thruster breathing mode dynamics,” *J. Appl. Phys.*, vol. 130, no. 13, p. 133302, Oct. 2021, doi: 10.1063/5.0046048.
- [131] K. Hara and S. Tsikata, “Cross-field electron diffusion due to the coupling of drift-driven microinstabilities,” *Phys. Rev. E*, vol. 102, no. 2, p. 23202, Aug. 2020, doi: 10.1103/PhysRevE.102.023202.

- [132] J. P. Boeuf and L. Garrigues, “E × B electron drift instability in Hall thrusters: Particle-in-cell simulations vs. theory,” *Phys. Plasmas*, vol. 25, no. 6, p. 061204, May 2018, doi: 10.1063/1.5017033.
- [133] W. Hargus and B. Pote, “Examination of a Hall thruster start transient,” in *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2002, p. 3956.
- [134] L. Brieda, “Molecular Outgassing and Deposition in EP Applications,” in *34th International Electric Propulsion Conference*, Hyogo-Kobe, Japan, Jul. 2015.
- [135] C. M. Greve and C. Marsh, “Evaluating Hall thruster startup transients using dynamics-based embeddings,” *Journal of Electric Propulsion*, vol. 4, no. 1, p. 36, 2025, doi: 10.1007/s44205-025-00136-y.
- [136] M. G. Allen, P. J. Roberts, and B. A. Jorns, “Direct measurements of the pressure dependence of the electron temperature in a magnetically shielded Hall thruster,” in *Proceedings of the 38th International Electric Propulsion Conference, Toulouse, France, IEPC-2024-821 (Electric Rocket Propulsion Society, 2024)*, 2024.
- [137] S. J. Thompson, Z. K. Robertson, S. C. Farnell, C. C. Farnell, C. C. Farnell, and J. D. Williams, “Time-Resolved Measurements of Plasma Properties in the Beam Plasma of a Hall Thruster,” in *38th International Electric Propulsion Conference*, Toulouse, France, 2024.
- [138] P. A. P. Moran, “The statistical analysis of the Canadian Lynx cycle.,” *Aust. J. Zool.*, vol. 1, no. 3, pp. 291–298, Dec. 1953, doi: 10.1071/ZO9530291.
- [139] C. E. Huerta, R. S. Martin, D. Q. Eckhardt, and J. W. Koo, “Determining causality in Hall effect thrusters using extended convergent cross mapping, part I**,” *Plasma Sources Sci. Technol.*, vol. 30, no. 7, p. 075004, 2021, doi: 10.1088/1361-6595/ac0a95.
- [140] C. E. Huerta, R. S. Martin, D. Q. Eckhardt, and J. W. Koo, “Determining causality in Hall effect thrusters using extended convergent cross mapping: II*,” *Plasma Sources Sci. Technol.*, vol. 31, no. 3, p. 035015, 2022, doi: 10.1088/1361-6595/ac511e.
- [141] J. Cabrera, D. Jovel, and M. L. R. Walker, “Pressure Facility Effects on the Impedance Characteristics of a 9-kW HET,” in *38th International Electric Propulsion Conference*, Toulouse, France, 2024.
- [142] R. R. Hofer, S. E. Cusson, R. B. Lobbia, and A. D. Gallimore, “The H9 magnetically shielded Hall thruster,” in *35th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2017, pp. 2017–2232.
- [143] J. W. Dankanich, M. Walker, M. W. Swiatek, and J. T. Yim, “Recommended Practice for Pressure Measurement and Calculation of Effective Pumping Speed in Electric Propulsion Testing,” *J. Propuls. Power*, vol. 33, no. 3, pp. 668–680, Aug. 2016, doi: 10.2514/1.B35478.
- [144] M. G. Allen and B. Jorns, “A Method Comparison for Charge Exchange Corrections for Hall Thrusters,” in *AIAA SCITECH 2026 Forum*, in AIAA SciTech Forum. , American Institute of Aeronautics and Astronautics, 2026. doi: doi:10.2514/6.2026-2133.

- [145] W. Huang, H. Kamhawi, and D. A. Herman, "Evidence of Counter-Streaming Ions near the Inner Pole of the HERMeS Hall Thruster," in *AIAA Propulsion and Energy 2019 Forum*, in AIAA Propulsion and Energy Forum. , American Institute of Aeronautics and Astronautics, 2019. doi: doi:10.2514/6.2019-3897.
- [146] B. A. Jorns, E. Dale, and R. R. Hofer, "Mode transitions in a magnetically shielded Hall thruster. I. Experimentally informed model," *J. Appl. Phys.*, vol. 136, no. 5, p. 053301, Aug. 2024, doi: 10.1063/5.0205983.
- [147] B. A. Jorns, M. Byrne, P. Roberts, L. Su, E. Dale, and R. R. Hofer, "Mode transitions in a magnetically shielded Hall thruster. II. Stability criterion," *J. Appl. Phys.*, vol. 136, no. 5, p. 053302, Aug. 2024, doi: 10.1063/5.0205985.
- [148] M. L. R. Walker *et al.*, "Overview of the Joint Advanced Propulsion Institute (JANUS)," 2022.
- [149] K. A. Polzin *et al.*, "Technology Maturation Plan for High Power Nuclear Electric Propulsion," National Aeronautics and Space Administration, 2026.