

©Copyright 2025

Peter Thoreau

Advanced Techniques for Hall Thruster Research and Development

Peter Thoreau

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:
Justin M. Little, Chair
Michael R. Holmes
Uri Shumlak

Program Authorized to Offer Degree:
Aeronautics and Astronautics

University of Washington

Abstract

Advanced Techniques for Hall Thruster Research and Development

Peter Thoreau

Chair of the Supervisory Committee:

Justin M. Little

Department of Aeronautics and Astronautics

Hall thrusters are the most common form of electric propulsion on spacecraft currently in Earth orbit, they offer excellent efficiency, high thrust to power, and have relatively simple power processing units. They have seen significant development since their first commercial flight in 1972, however, methods for characterizing their operation for both development and flight have remained relatively consistent. Therefore, advanced techniques have been developed to more effectively, efficiently, and expeditiously characterize thruster operation. These methods have been advanced through four primary approaches.

The first approach was the development of a magnetically shielded Hall thruster with a movable inner pole to investigate the influence of magnetic field shape on thruster operation. The use of magnetically shielded geometries has been demonstrated to drastically reduce the rate of erosion on Hall thruster channel walls. However, the switch to a magnetically shielded configuration can increase plume divergence contributing to a decrease in efficiency and thrust. Varying the magnetic field shape near the exit of the thruster has been demonstrated to decrease plume divergence, however, the impact on thruster behavior is unknown. Three test campaigns on iterations of the ACME thruster, operating on both xenon and krypton, investigate the plume divergence, performance and efficiency changes, and map the plume of the thruster across a range of pole positions. The divergence, voltage, and current utilization efficiencies have strong dependencies on the relative pole position across all operating points.

The dependence of anode efficiency on pole position can be mapped to the other efficiencies within the uncertainty of the experiment. The change in inner pole position of -1 mm on xenon and - 4 mm on krypton has been shown to significantly increase the anode efficiency of the thruster, primarily through divergence and voltage utilization efficiency gains.

The second approach assesses the standard methods for characterizing thruster efficiency using plume diagnostics. The standard practice of measuring a centerline Ion Energy Distribution Function (IEDF) is insufficient for electric propulsion systems with unknown or highly divergent plume structures. This method is compared to spatially resolved IEDF measurements using a swept retarding potential analyzer to generate an angularly resolved IEDF. Using an adjustable Hall thruster, plumes ranging from highly divergent to over-focused were fully characterized and compared to the anode efficiency calculated using thrust stand measurements. Both centerline and spatially resolved measurements of the IEDF were sufficient to accurately measure the voltage utilization efficiency on well focused plumes. As the plume diverged, more complex plume structures were observed, and only the spatially resolved measurement maintained agreement with the thrust stand based efficiency. .

The third approach to expedite Hall thruster development created a method to rapidly and autonomously optimize thruster performance. This method combines rapid thrust measurement, thruster-in-the-loop control, and derivative-free optimization schemes to automate thruster optimization and rapidly find operational areas of interest, thus reducing the reliance on current-voltage-magnetic field (IVB) maps. Rapid thrust measurements are achieved by comparing test points to a known operating point, allowing fast and accurate thrust measurement while minimizing the effects of long-term thermal drift. Fast mapping of a Hall thruster's operation was demonstrated at 72 test points per hour, with an average thrust measurement error of <1% compared to conventional thrust measurements. Two-dimensional (thruster discharge voltage and magnetic field strength) Nelder-Mead and Powell optimization schemes are shown to converge rapidly to maxima in total efficiency or specific impulse

in fewer than 15 test points. The Powell optimization scheme remained effective in five dimensions, further increasing the peak thruster efficiency while adjusting three additional thruster dimensions (keeper current, cathode flow fraction, and magnetic field skew).

The fourth approach developed a method for autonomous optimization and modeling of an electric thruster that applies Bayesian optimization on a Gaussian Process Regression model generated in real time from experimental telemetry. The method can be combined with a prescribed objective function and optimization scheme to optimize the thruster for different mission objectives. A notional extrasolar probe mission powered by a Hall effect thruster is considered as an example where the goal of the optimization is to find the propellant gas mixture (argon:krypton:xenon) that minimizes overall mission cost. The results show that a thruster running on a mixture ratio of 11:87:2 benefits from a 7% reduction in total cost compared to the same thruster running on pure xenon. Analysis of the model reveals how the optimal propellant mixture depends strongly on propellant storage technologies, fluctuations in propellant price, and launch costs. Results from this analysis match trends seen in the commercial market with the move to cheaper propellants.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	x
Chapter 1: Introduction	1
1.1 Motivation and Research Direction	2
1.2 Thesis Structure	10
Chapter 2: Test Facilities and Diagnostics	14
2.1 The Space Test Facility (STF)	14
2.2 The Space Test Facility 2 (STF-2)	14
2.3 Thrust Stand	15
2.4 E×B Probe	18
2.5 Faraday Probe	20
2.6 Langmuir Probe	25
2.7 Retarding Potential Analyzer	25
Chapter 3: The ACME Hall Thruster	28
3.1 ACME Design and Construction	28
3.2 ACME Testing and Design Updates	34
Chapter 4: Plume Characterization with Spatially Resolved Ion Energy Distribution Measurements	56
4.1 Methodology	57
4.2 Results and Discussion	60
4.3 Chapter Conclusion	70

Chapter 5: Rapid Thruster-in-the-Loop Optimization for Hall Thrusters	72
5.1 Methodology	73
5.2 Results and Discussion	84
5.3 Chapter Conclusions	97
Chapter 6: Modeling and Optimization of Propellant Mixtures for Hall Thrusters .	98
6.1 Methods	100
6.2 Results and Discussion	111
6.3 Chapter Conclusions	133
Chapter 7: Conclusion	134
7.1 Peer-Reviewed Publications	138
Bibliography	139
Appendix A: Hall Thruster Diagrams	153

LIST OF FIGURES

Figure Number	Page
1.1 A cross section of a Hall thruster with a centrally mounted cathode showing the standard layout of features.	2
1.2 The magnetic field shape and plasma potential within the channel in unshielded and shielded Hall thruster configurations.	3
1.3 Changes to the plasma structure in the channel of the H6 Hall thruster in unshielded (US), and magnetically shielded (MS) configurations showing plasma potential (top), electron temperature (middle), and electron density with vectors of singly charged ion density overlaid. From [23].	5
1.4 Hall2De simulations of a) MaSMi-60 Hall thruster showing the electron temperature around the exit plane of the channel and b) MaSMi-DM with the updated field showing reduced curvature and shift in the electron population. [33, 34]	7
1.5 The range of grazing lines tested with their center-line field strength (a), and channel location, (b) [35].	8
1.6 The magnetic field inclinations used by Liu et al.. a) inwardly inclined, b) non-inclined, and c) outwardly inclined. From [36].	9
2.1 A Schematic of the STF showing the position of the diagnostics, chamber sections, and test equipment.	15
2.2 The STF-2 showing the position of thruster, thrust stand, and chamber sections.	16
2.3 A review of thrust stands (taken at the time of thrust stand design), comparing their thrust measurement range to their size. Data is from published test results, actual ranges are likely larger. Source: A [60], B [54], C [61], D [62], E [63], F [59], G [64], H [65], I [66], J [67].	17
2.4 a) ACME mounted to the thrust stand outside the STF. b) A sample thrust measurement taken during the V2 test campaign.	18
2.5 a) A schematic showing the electrical setup of the $E \times B$ probe. b) $E \times B$ data with Gaussian fits for three charge states of Krypton measured.	19
2.6 a) The Faraday probe mounted to the stainless steel arm. b) The measurement circuit used for the Faraday probe.	21

2.7	a) The electrical setup of the Langmuir probe. b) A sample Langmuir probe trace showing the natural log of the current, and ‘knee’ at the plasma potential.	25
2.8	a) The electrical setup of the RPA probe. b) A single point centerline RPA measurement showing the raw current, derived current, Gaussian fit, and most probably voltage.	27
3.1	The extremes of pole positions tested on ACME to date compared at 400 V, 1.5 A on krypton. The +4 mm pole position and associated photo, top, shows a highly divergent plume. The -12 mm pole position, bottom, shows an over-focused plume. The inner pole is located stationary in the frame of the photos, while the outer pole moves.	29
3.2	The magnetic field of ACME V1 with the under-shielded inner pole on the left and over-shielded outer pole on the right.	32
3.3	The full range of pole positions ACME can operate in: 0 mm inner position, left; maximum negative inner pole position of -35 mm, center; and the maximum positive inner pole position of +5 mm, right.	33
3.4	Conduction paths for plasma and magnet heating on version 1 of the ACME thruster.	34
3.5	The V1 ACME thruster without the centrally mounted hollow cathode . The heat sink shown was not used during V1 250 W tests.	35
3.6	The floating electrical configuration used for ACME V1 testing with inductive loads on both circuits.	36
3.7	ACME running on xenon. Left: DSLR image from outside the chamber showing dark regions on the edge of the plasma. Right: Image from the chamber internal camera showing a higher contrast view of the plasma.	37
3.8	A comparison of the plume divergence across the range of pole positions tested, showing: DSLR photos of the plume were taken at 70 mm f/4.5 1/20” ISO200; Field simulations with the average normal vector to the field; and Normalized swept Faraday probe measurements comparing current density at each position.	38
3.9	ACME version 2 mounted on the thrust stand outside the STF with new BN walls and alumina pole caps.	41
3.10	FEMM simulations of the magnetic field for ACME V2 with the new wall designs overlaid.	42
3.11	The floating electrical configuration with hybrid filter used for all tests in the V2 test campaign	43
3.12	The two distinct plume modes seen during thruster testing at the +1 mm position. The Jet mode, left, and Diffuse mode, right.	44

3.13	a) Plasma at the exit plane showing significant offset from the walls. b) Carbon-coated BN walls of the thruster and edge of the pole piece. The clean area, indicating net erosion, is seen on the inner section of the pole piece.	45
3.14	A comparison of the plume divergence across the range of pole positions tested. a) DSLR photos of the plume taken at 105mm f/16 1/10" ISO200. b) Field simulations with the average normal vector to the field c) Normalized swept Faraday probe measurements comparing current density at each position.	47
3.15	a) The momentum weighted divergence utilization efficiency of the thruster at each inner pole position. b) The current utilization efficiency at each position.	48
3.16	The maxima of the Gaussian fit of the IEDF at each of the pole positions tested.	49
3.17	a) The thrust measured at each of the pole positions. b) The relative anode efficiency at each pole position comparing the total influence of the efficiencies calculated. The points are normalized to the 0 mm pole position measurement. The reference anode efficiency at 0 mm is 51%.	49
3.18	Frequency spectra at each pole position showing the change in breathing mode frequency and amplitude.	50
3.19	Frequency spectra of the two distinct Hall thruster operating modes. Measured at the -1 mm position with otherwise identical operating parameters.	51
3.20	Frequency spectra showing the increasing breathing mode amplitude as the temperature increases from a) to c). Data taken at the -2 mm position.	51
3.21	a) The water-cooled block and new mounting system being test fit on the thrust stand. b) the new parts of the ACME v2.1 setup showing the block with cylindrical cathode heat shield, large heat sink, and modified rear casing of the thruster.	52
3.22	The cathode-tied electrical configuration with hybrid filter used for all tests with ACME V2W and V3.	53
3.23	a) A cross section of the A3 anode. b) A cross section of the A4 anode. c) The printed A3 anode. d) The printed A4 anode. Figures from [88].	55
4.1	The breakdown of the measured efficiencies across the range of pole positions showing: a) The momentum weighted divergence utilization efficiency, b) the current utilization efficiency, c) the charge utilization efficiency, d) the mass utilization efficiency, e) the three methods used for the voltage utilization efficiency, and f) The anode efficiency comparison between the product of the probe based efficiencies and the direct measurement from the thrust stand.	62

4.2	A comparison of the plume divergence across a subset of pole positions tested. Peak first derivative of current at each pole position is: a) $0.14 \mu A/V$ at +4 mm, b) $1.3 \mu A/V$ at 0 mm, c) $0.58 \mu A/V$ at -4 mm, and d) $0.09 \mu A/V$ at -8 mm. DSLR photos and contour plots of the swept IEDF measurements are shown for each position. Photos of the plume taken at 70 mm f/22 1/5" ISO200. Note that the inner pole is consistently positioned with each axis while the outer pole moves.	64
4.3	a) The field at the +4 pole position, b) the field at the -4 mm pole position, and c) the field at -8 mm showing the stretch in shielding on the outer pole. Note that the field lines are rendered to the same arbitrary limit to demonstrate the change.	66
4.4	a) A sample of the swept RPA from a 400 V, 1.5 A discharge at the +4 mm pole position. CEX dominated signal is seen below the red line. b) A comparison of the Faraday data, RPA current, and CEX correction for both signals.	67
4.5	a) Thrust measurements taken at each pole position at 400 V 1.5 A. b) The required flow rate through the anode to achieve a current of 1.5 A at thermal steady state.	68
4.6	A FFT of the discharge current at each pole position. Scale on each axis is identical.	69
5.1	a) A standard shutdown thrust measurement with a single calibration step. Note the under-damped oscillation before tuning the damping system. b) Rapid thrust measurement with the Control Point - Test Point methodology.	75
5.2	A flow chart showing the operation of the thruster using the Control Point - Test Point rapid thrust measurement technique.	78
5.3	The simplex formed in a 2D Nelder-Mead optimization showing the three current test points (P_0, P_1, P_n), the centroid (\bar{P}), and the reflected (P_R), expanded (P_E), contracted (P_C), and shrunk (P_{0S}, P_{1S}) points.	80
5.4	The evaluation of the first dimension along the vector, ξ_1 , in Powell's method. Data points taken at position \mathbf{x} with their evaluation $F(\mathbf{x})$ are shown. The estimated minima of the function at \mathbf{p}_1 is then the starting point of the normal vector, ξ_2 , used to evaluate that dimension.	83
5.5	a) Iterations between the control point, and test points with a decreasing discharge voltage, a fixed 200 G magnetic field strength, and dwell time of 15 s. Dashed vertical lines indicate the points where thrust is measured. b) Mean error and standard deviation of the rapid measurement methodology with increasing control and test point dwell time.	85

5.6	a) A contour map of the total efficiency in the two-dimensional test area. The map is generated using the rapid measurement technique and 117 test points. Eight comparison points showing the efficiency measured using the shutdown thrust measurement technique are overlaid for comparison. Dashed lines showing the region of peak efficiency are shown for comparison to subsequent optimizations. The approximate area where mode transitions occurred is shown as the broad dashed line. b) A map of the anode efficiency from the same tests.	88
5.7	a) A map of the two-dimensional total efficiency map with optimization paths shown for the Nelder-Mead optimization scheme. Starting points for the optimization are in blue with the resulting optimums in red. Lines showing the region of peak efficiency are shown for comparison other schemes. b) The total efficiency at each iteration for the same optimization paths.	89
5.8	a) A map of the two-dimensional total efficiency map with optimization paths shown for the Powell optimization scheme. Starting points for the optimization are in blue with the resulting optimums in red. Lines showing the region of peak efficiency are shown for comparison other schemes. b) The total efficiency at each iteration for the same optimization paths.	91
5.9	a) The two-dimensional specific impulse map with optimization paths shown for both the Nelder-Mead and Powell optimization schemes. Starting points for the optimization are in blue with the resulting optimums in red. b) The specific impulse at each iteration for the same optimization paths. The dashed lines show the Powell scheme optimizations.	92
5.10	Convergence for five-dimensional optimization showing: Total Efficiency, η ; Voltage (V), V; Field Strength (G), B; Keeper Current (A), KC; Cathode Flow Fraction, CF; and Field Skew, S. a) shows Nelder-Mead, and b) shows Powell. The iteration with peak total efficiency is marked with the dashed vertical line. The dimensions being optimized at each step are indicated by the shaded regions.	94
6.1	A flow chart showing steps to initialize a surrogate model with random data, and run BO of the thruster using experimental measurements. Thruster operations are shown in blue, optimizer functions in red, and start/stop points in green.	103
6.2	a) Performance of Bayesian minimization on the total cost after initialization on 77 randomly sampled points showing a trend in improvement of $-7.7 \times 10^4 \mathbf{x} + 1.56$. b) Convergence of towards the objective shown by the slope of linear fits to optimization data.	113

6.3	Evolution of the surrogate model for thrust with an increasing number of samples used to train the model. At each number of samples, the mean 95% confidence interval of the surrogate model, and the mean model error is evaluated for 1000 comparison points using randomized subsets of the data. .	116
6.4	a) Optimized thrust for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized thrust at each point. c) The discharge voltage for optimized thrust at each point. d) The magnetic field strength for optimized thrust at each point. e) The keeper current for optimized thrust at each point.	118
6.5	a) Optimized specific impulse for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized specific impulse at each point. c) The discharge voltage for optimized specific impulse at each point. d) The magnetic field strength for optimized specific impulse at each point. e) The keeper current for optimized specific impulse at each point. The discontinuities in the gradient show the locations of mode shifts in the Hall thruster discharge, indicating where the optimum discharge parameters change significantly. . .	120
6.6	a) Optimized total efficiency for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized total efficiency at each point. c) The discharge voltage for optimized total efficiency at each point. d) The magnetic field strength for optimized total efficiency at each point. e) The keeper current for optimized total efficiency at each point. The discontinuities in the gradient show the locations of mode shifts in the Hall thruster discharge, indicating where the optimum discharge parameters change significantly. . .	122
6.7	a) Optimized total cost for mixtures of argon, krypton, and xenon for mission 2 (VLEO satellite). b) The standard deviation of the optimization for mission 2. c) A repeat optimization of mission 2 with cryogenic propellant tanks. d) The standard deviation of the optimization for with cryogenic tanks.	125
6.8	Evolution of the optimal operating parameters for a minimized total cost for mission 2's specifications as a function of the increasing propellant cost factor of krypton and xenon. Data is from the evaluation of the surrogate model of thruster total cost.	128
6.9	Evolution of the optimal operating parameters for a minimized total cost for mission 2's specifications as a function of the increasing specific Launch Cost. Data is from the evaluation of the surrogate model of thruster total cost. SpaceX's ride-share price is indicated by the vertical dotted line.	131
A.1	The three sub-assemblies that make up ACME. Version 3 is shown.	153
A.2	The parts of the hollow cathode and keeper assembly.	154

A.3	Numbered parts of the ACME Hall thruster showing the separate inner and outer poles.	155
A.4	Numbered parts of the cooling block showing the separate inner and outer poles.	157

LIST OF TABLES

Table Number	Page
3.1 Thruster Performance at 0 mm on Xenon	37
3.2 Representative Thruster Performance at 0 mm on Xenon	46
5.1 ACME thruster dimensions. Fast dimensions, left, can be adjusted during operation. Slow dimensions, right, require removing the thruster from vacuum to change.	73
6.1 Tank specifications used for cost optimization [140, 141].	107
6.2 Propellant costs used for optimization set at the 2015 price, a time of relative stability, and comparable to current prices [142].	107
6.3 ACME thruster fast dimensions, $\mathbf{x} = (\gamma_{\mathbf{Ar}}, \gamma_{\mathbf{Xe}}, \mathbf{V}_d, \mathbf{B}, \mathbf{I}_k)$, that are computer controlled during operation.	109
6.4 Specifications for example mission optimizations (missions 1 & 2) and analysis of cost factor and specific launch cost (Missions 2 & 3).	112
6.5 Comparison of the minimized total cost for surrogate models of the full data set of 575 observations and the 36 BO observations.	114
A.1 Hollow cathode parts list.	154
A.2 ACME parts list.	156
A.3 Cooling block parts list.	156

ACKNOWLEDGMENTS

An immeasurable number of people are in some way responsible for getting me to where I am today, a few even helped. To all these people, even if you are not named here, Thanks.

To my advisor, Professor Justin Little, thank you for letting me loose in the lab and giving my a significant degree of autonomy in my research. I appreciate all that you have done for me, and apologize profusely for some of the drafts I've sent you (I can't be good at everything, right?). To the rest of my committee, thank you, I value all the advice and input I have received through the years.

To my parents, I couldn't have done it without you. Justin may have let me loose on the lab, but you let me loose on the world. To Jen, Michael, Steve, and Rachel, thanks for all the fun over the years. (Can you tell I dislike writing acknowledgments yet)

To Aria, thank you for the idea of applying machine learning which led to the full automation of thruster testing and eventually allowed me to run test campaigns on the couch.

To Sven: bork bork bork, good boy. While I'm at it, to Eclair and Bonita: rruuuuufff, good girlies.

To Mike Holmes, thanks for being chill, we still need to go skiing while the thruster tests itself.

And to all my friends in lab, you all know me well enough that you know I hate writing this. So thanks for being understanding about the single sentence that covers you all. Thank you Anna, Charlie, Nadiah, Curtis, Danny, Landon, Arvinhd, Patrick, Evan, Sari, Bennett, Jared, Gordon, Cameron, Sungyoung, and Diego.... oh no I've been in this lab too long.

xx Peter

Chapter 1

INTRODUCTION

Electric propulsion (EP) is a subset of rocket propulsion, used to produce thrust for spacecraft attitude adjustment, orbital maneuvering, and drag compensation [1, 2]. EP is advantageous, primarily due to the high specific impulse of electric thrusters, allowing for significantly lower propellant mass, or significantly higher total impulse [3]. Hall thrusters, a type of EP system, are the most common form of electric propulsion on spacecraft currently in orbit [4]. The key advantages of Hall thrusters over other EP systems are their high thrust to power ratio, and relative simplicity and mass of the power processing unit (PPU). A Hall thruster consists of an annular channel with a combined anode and gas supply/distribution system at the rear, see Fig. 1.1. A magnetic circuit, constructed of high magnetic permeability material and electromagnets, creates a radial magnetic field near the exit plane of the channel. The field strength is such that electrons in the field are magnetized with a Larmor radius significantly shorter than the channel width. These magnetized electrons drift in the azimuthal direction due to the $E \times B$ field, creating a Hall current, but are restricted in the axial direction. The electrons of the Hall current, and anode biased to a high potential, drive the axial electric field in the channel. As the field lines are thermalized equipotential lines, the magnetic field shape defines the electric field within the thruster, see Fig. 1.2. Plasma is isolated from the magnetic material with insulating ceramic walls and pole caps, although conducting wall concepts have been successfully demonstrated [5, 6, 7, 8]. Propellant injected into the rear of the channel, is ionized by the channel electrons and accelerated by the electric field, producing thrust. These high energy ions, combined with electrons from the cathode, form the thruster plume.




Figure 1.1: A cross section of a Hall thruster with a centrally mounted cathode showing the standard layout of features.

1.1 Motivation and Research Direction

Hall thrusters are a relatively mature technology with over 50 years of flight heritage. Development started in the Soviet Union in 1962, with the first launch and successful operation on the Meteor satellite in 1972 [9]. In the 1990's, the fall of the Soviet Union caused Russian design bureaus and agencies, including Fasel, Keldysh, and Tsniimash, to export thrusters and technology, greatly increasing the number of parties involved with research, development, and application [10]. Hall thruster efficiency, thrust to power ratio, and specific impulse all reached desirable levels, with lifetimes reaching 10,000 hours [9, 10]. Geostationary satellites have historically been the primary use of EP in spacecraft, and initially favored resistojets on western satellites until ~ 1994 , when arcjets and gridded ion thrusters took significant market share [4]. Since 2012, Hall thrusters have surpassed the other thruster types to be the most commonly flown EP system.

In recent years, Hall thruster design and mission use has seen a dramatic shift. There are two primary drivers to this shift: the advent of magnetic shielding in Hall thrusters providing significantly longer lifespans [11, 12]; and the rise of ‘New Space’ with the accompanying shift in thruster development and mission architectures [13, 14, 15]. This section will detail the effect of these two shifts and describe how they have guided the direction of research presented in this dissertation. Firstly, the influence of magnetic shielding on thruster design will be described, along with the research aimed at understanding and mitigating its drawbacks. Then, the influence of ‘New Space’ is described, with specific focus on the desire to cheaply and rapidly develop new thrusters, leveraging automation, machine learning, and other computationally intensive methods.

1.1.1 Magnetic Shielding of Hall Thrusters

Prior to the advent of magnetic shielding, the erosion of ceramic channel insulators was the most dominant cause of Hall thruster end-of-life. To minimize this erosion in unshielded thrusters, magnetic field lines were designed to be convex towards the anode with increasing strength towards the wall [16]. The field shape focused ions towards the peak electron density at the center of the channel, reducing the erosion and thermal loading, and focusing the ion beam, see Fig. 1.2.




Figure 1.2: The magnetic field shape and plasma potential within the channel in unshielded and shielded Hall thruster configurations.

Magnetic shielding developed from this concept, extending the convexity in the channel so the field lines are tangential to the insulator at the front of the thruster. Magnetic shielding was discovered after a 10,000 hour test of a BPT-4000 thruster effectively stopped eroding after 6,000 hours of operation [11, 12]. The erosion was found to be tangential to the field lines along the outer boron nitride rings, and when stopped, was the first extended lifetime test with no measurable channel erosion. Since this test, magnetic shielding of Hall thrusters has been demonstrated to drastically reduce the rate of erosion on the boron-nitride channel walls to less than 1/1000th of its unshielded counterpart [12, 17]. The degree of this reduction is sufficient to remove channel erosion as the primary cause of end-of-life, with the most significant point of erosion now occurring on the front poles of the thruster, where the magnetic field lines from the channel intercept the surface [18, 19, 20, 21].

However, the switch to a magnetically shielded configuration moves the Hall current, and acceleration and ionization regions downstream, increasing plume divergence. Tests on both magnetically shielded and unshielded versions of the H6 Hall thruster showed an increase in divergence half angle from 15 to 20 degrees and a decrease in divergence utilization of 5.2% [17]. Increased plume divergence, and the increased ion current at high plume angles associated with it, has the additional problem of energetic ions impinging on spacecraft surfaces. This can lead to erosion or damage, particularly changing the electrical, thermal, and optical properties of sensitive equipment [22]. Figure 1.3 shows Hall2De simulations of the H6 Hall thruster in both shielded and unshielded configurations. The differences in plasma structure are stark, with the potential drop pushed beyond the exit plane of the thruster and a significant increase in electron temperature in the shielded version. Of most interest to this dissertation, are the vectors of singly charged ion density overlaid on the electron density simulation. While the unshielded thruster has a relatively collimated plume, the shielded counterpart has a significantly more divergent plume, with the curved acceleration region beyond the channel causing a far wider range of ion trajectories.

The optimization of plume divergence in non-shielded thrusters has seen a range of methods. Altering the pole piece shape and position [24]; the use of trim coils to adjust exit plane

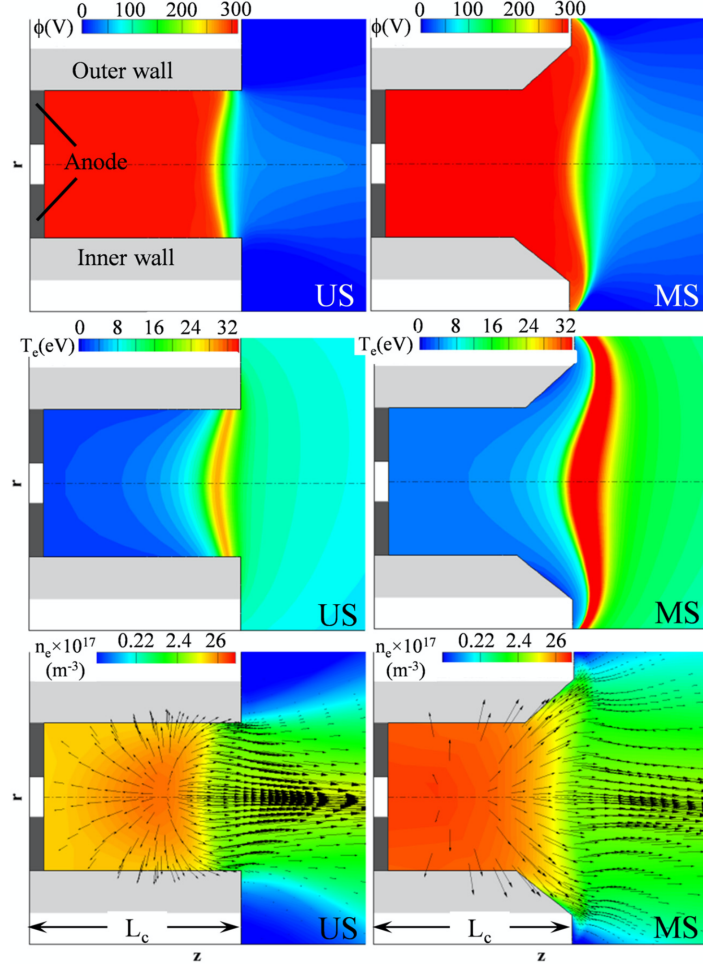


Figure 1.3: Changes to the plasma structure in the channel of the H6 Hall thruster in unshielded (US), and magnetically shielded (MS) configurations showing plasma potential (top), electron temperature (middle), and electron density with vectors of singly charged ion density overlaid. From [23].

field angle and channel plasma lensing [25, 26, 27]; and adjusting coil currents to affect both field shape and strength [28] have all been successfully demonstrated.

Mitigating the plume divergence of magnetically shielded thrusters has seen less investigation, particularly in low power Hall thrusters. A 200 W thruster, built to investigate plasma-wall interactions, has been operated in both shielded and non-shielded configurations [29]. It demonstrated reduction of erosion on the channel walls and continued ero-

sion of the front pole consistent with observations of high power magnetically shielded Hall thrusters. The erosion rates on the wall and inner pole were assessed by observing the areas of net erosion and sputtered carbon deposition. Ion velocity distributions in the channel near the exit plane were assessed with non-invasive Laser Induced Fluorescence spectroscopy (LIF) to determine the ion velocity distribution functions (IVDF). The magnetically shielded configuration showed significantly lower ion density, flux and energy close to the wall. LIF measurements also showed areas of high electric field strength further from the center-line of the channel in the shielded configuration. The location of this high field region stems from the increase curvature and strength of the magnetic field lines close to the exit plane seen in Figures 1.2 and 1.3.

The conclusions drawn from the LIF data collected agrees with a first principles analysis published on combined simulation and experimental data from existing literature [30]. The study concluded that two base mechanisms are primarily responsible for the reduction in erosion: reduction of the ion flux to the wall, and a reduction of energy of ions colliding with the wall. Both are due to the shift in ionization and acceleration regions downstream towards and past the exit plane, and the contoured field modifying the potential gradient in the near-wall region. Continued analysis suggested that plasma lensing was a key area of interest for reducing the plume divergence of the magnetically shielded geometry.

Iterations on the MaSMi thruster have significantly reduced the curvature of the field from the magnetic poles into the thruster, resulting in a significant increase in divergence utilization efficiency [31, 32, 33, 34]. Figure, 1.4 shows Hall2De simulations of electron temperature in the MaSMi-60 and MaSMi-DM thrusters. The updates to MaSMi for MaSMi-LM2 and MaSMi-DM reduced the over-shielding at the walls of the thruster, and show significantly less curvature in the high electron temperature region at either side of the channel. Total efficiency in the DM test improved to 41%, a $> 25\%$ improvement over the previous magnetic field configuration.

The curvature of the magnetic field over the pole and into the channel is used to define the shape of the walls in shielded thrusters. The field line that is closest to the contoured

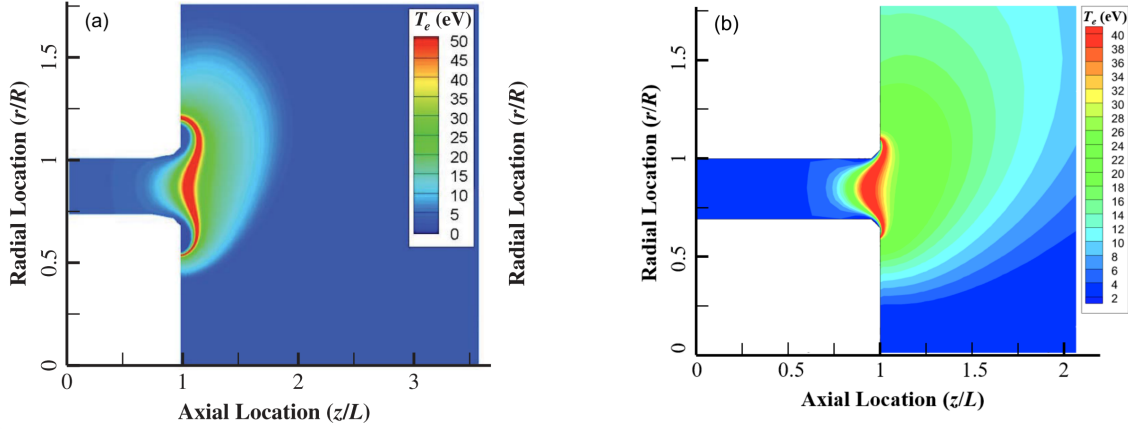


Figure 1.4: Hall2De simulations of a) MaSMi-60 Hall thruster showing the electron temperature around the exit plane of the channel and b) MaSMi-DM with the updated field showing reduced curvature and shift in the electron population. [33, 34]

wall is the grazing line. The grazing line can influence the plasma interaction with the wall and has been demonstrated to significantly affect shielding efficacy and thruster performance [35]. Performance measurements of a thruster with a consistent magnetic field shape were compared with a series of channel wall profiles. Each wall was contoured to match grazing lines, defined by different magnetic field strengths on channel center-line. Figure 1.5 shows the grazing lines used in the experiment, and their field strength at the midpoint of the channel. The lines tested ranged from 5 to 36% of the peak field strength, each with a matching chamfer on the wall. The deepest grazing lines, 5 and 12%, minimized erosion but suffered a performance penalty compared to the 20% experiment. Grazing lines higher than 20% of peak had significantly more erosion and dropped anode efficiency by $> 10\%$.

Trim coils have been used to alter the field angle through the acceleration region of the plasma with a significant effect on the plume divergence [36]. Inclining the field towards the inner pole in the acceleration region produced the highest efficiency and minimized divergence when compared to flat and outwardly inclined fields, see Figure 1.6. The study did not extend the test beyond the single inwards inclination, and therefore does not give further information on where the peak in performance would be achieved. It concluded that

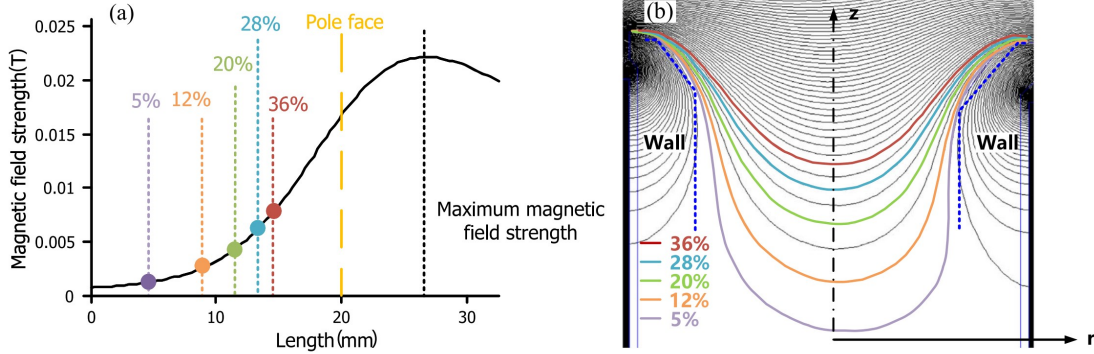


Figure 1.5: The range of grazing lines tested with their center-line field strength (a), and channel location, (b) [35].

within the range of field angles tested, the change in plume divergence dominated the thrust and efficiency changes. However, the study did not analyze the other efficiencies within the thruster.

This result is in contrast to an earlier study from the same institution, which studied the effects of field inclination on a non-shielded thruster. The non-shielded thruster was operated in three field configurations of varying inclination, adjusted by varying the current in the magnets. It found a minima in the plume divergence in the configuration with no inclination [37]. Comparing the photos of the plume between the two studies, the non-shielded thruster focused the plume to a similar distance downstream to the inwardly-inclined configuration of the shielded thruster.

The additional plume divergence observed in these magnetically shielded thrusters motivated the first research direction of this dissertation: the development of an adaptable Hall thruster to study the effect of plume structure on magnetically shielded thruster performance, and to develop methods to mitigate the increased divergence.

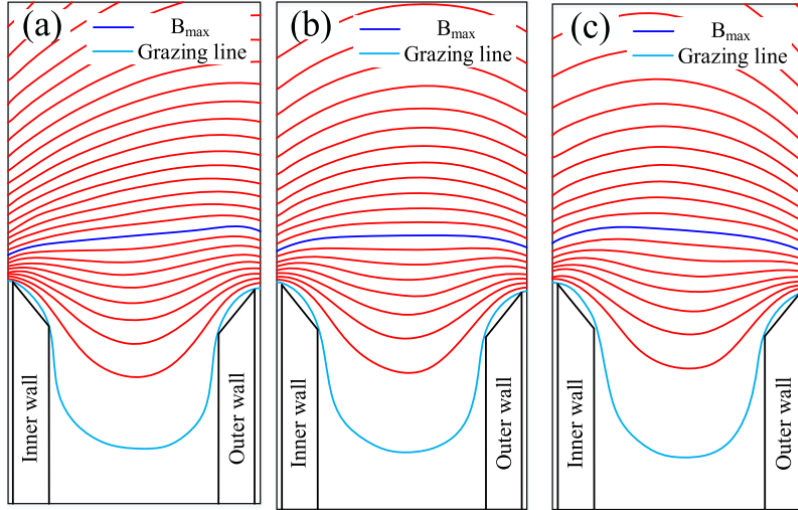


Figure 1.6: The magnetic field inclinations used by Liu et al.. a) inwardly inclined, b) non-inclined, and c) outwardly inclined. From [36].

1.1.2 Machine Learning in EP

The second main direction of research in this dissertation is motivated by the rise of ‘New Space’ which has seen decreasing costs for access-to-space and a rise in small satellite constellations, driving the need for a greater number of systems that are cheaper to design, test, and manufacture. The advancement of the machine learning field and parallel increase of available computational power provides our direction to explore this motivation, equipping us with new tools to tackle these problems, developing thrusters more efficiently and quickly.

Computationally driven development processes, including machine learning and algorithmic optimization, are well suited to the challenges of electric propulsion and have been demonstrated across a wide range of applications, including: developing Hall thruster scaling laws [38], reducing particle-in-cell (PIC) simulation time [39], characterizing physical process within a Hall thruster [40], predicting Hall thruster performance to enable increased autonomous control [41, 42], predicting Hall thruster erosion [43], aiding in the structural design process of Hall thrusters [44], automating advanced magnetic circuit design [45], and optimizing multi-frequency heating in an Electron Cyclotron Resonance (ECR) thruster [46].

However, there are key areas where these methods can be applied that remain unexplored. Specifically the characterization and optimization of electric propulsion systems, which has traditionally been conducted over long test campaigns. A significant component of that is the performance and stability mapping across a range of current-voltage-magnetic field (IVB) configurations [47, 48, 49]. IVB maps are useful for characterizing the global performance and more thoroughly analyzing changes in design [50]. However, mapping performance across the many dimensions of a thruster is costly in propellant and facility use, and development time. As Hall thrusters exhibit a range of physics that are difficult to model, development relies heavily on these experimental results to inform models and design iterations [51]. The results of these campaigns are typically very large data sets with methodical analysis of single variables, often relying on human interpretation to relate variables with complex dependencies [52, 53]. This approach is expensive and time consuming, and while successful, there is a need for more advanced methodologies.

This leads to us in two directions. While both methods leverage automation and computationally driven optimization, they differ greatly in their use. The first direction aims to expedite data collection, and rapidly find operational areas of interest. The second method shifts the trade-off from speed to understanding, aiming to capture the complex inter-dependencies of Hall thruster operation, and to improve analysis of these high dimensionality systems.

1.2 Thesis Structure

Having discussed some of the motivations and research directions of interest, we must define the primary research question that guides the work presented in this dissertation and outline the four primary studies presented in chapters 3-6. The primary research question of this dissertation is:

How can we advance the techniques used in electric propulsion research to more effectively, efficiently, and expeditiously characterize and develop thruster systems?

1. **The ACME Hall Thruster.** A magnetically shielded Hall thruster with a movable inner pole has been constructed to investigate the influence of magnetic field shape on thruster operation. The use of magnetically shielded geometries has been demonstrated to drastically reduce the rate of erosion on Hall thruster channel walls. However, the switch to a magnetically shielded configuration can increase plume divergence contributing to a decrease in efficiency and thrust. Varying the magnetic field shape near the exit of the thruster has been demonstrated to decrease plume divergence, however, the impact on thruster behavior is unknown. The influence of the inner magnetic pole position, and therefore topology of the magnetic field, on thruster performance is investigated with Faraday probe, retarding potential analyzer, ExB probe, and thrust stand measurements. This chapter will cover the design, initial test campaigns, and thruster updates of the ACME Hall thruster.
2. **Plume Characterization with Spatially Resolved Ion Energy Distribution Measurements.** Standard methods for characterizing thruster efficiency using plume diagnostics are insufficient for electric propulsion systems with unknown or highly divergent plume structures. The standard practice of measuring a center-line Ion Energy Distribution Function (IEDF) is compared to spatially resolved IEDF measurements using a swept Retarding Potential Analyzer (RPA) to generate an angularly resolved IEDF. Using an adjustable Hall thruster, plumes ranging from highly divergent to over-focused were fully characterized. This chapter compares to the anode efficiency calculated using thrust stand measurements to the full plume characterization with

both traditional and angularly resolved IEDFs to investigate methods for more complete plume characterization.

- 3. Rapid Thruster-in-the-Loop Optimization for Hall Thrusters.** In order to expedite Hall thruster development, a method to rapidly and autonomously optimize thruster performance was created. This method combines rapid thrust measurement, thruster-in-the-loop control, and derivative-free optimization schemes to automate thruster optimization and rapidly find operational areas of interest, thus reducing the reliance on current-voltage-magnetic field (IVB) maps. Rapid thrust measurements are achieved by comparing test points to a known operating point, allowing fast and accurate thrust measurement while minimizing the effects of long-term thermal drift. Fast mapping of a Hall thruster's operation was demonstrated at 72 test points per hour, with an average thrust measurement error of $<1\%$ compared to conventional thrust measurements. This chapter assesses two-dimensional (thruster discharge voltage and magnetic field strength) Nelder-Mead and Powell optimization schemes for their ability to rapidly converge to maxima in total efficiency or specific impulse. These schemes are then applied to five-dimensional optimization of Hall thruster performance to evaluate their usage to higher dimension systems.
- 4. Modeling and Optimization of Propellant Mixtures for Hall Thrusters.** A method is proposed for autonomous modeling and optimization of an electric thruster that applies Bayesian optimization on a Gaussian Process Regression model generated in real time from experimental telemetry. The method can be combined with a prescribed objective function and optimization scheme to optimize the thruster for different mission objectives. A notional extrasolar probe mission powered by a Hall effect thruster is considered as an example where the goal of the optimization is to find the propellant gas mixture (argon:krypton:xenon) that minimizes overall mission cost. This chapter investigates the use of Gaussian Process Regression based surrogate models and of Bayesian optimization to model and optimize for thruster performance

and mission cost for a set of sample missions. Further analysis of the model is used to relate the optimal propellant mixture to different propellant storage technologies, fluctuations in propellant price, and launch costs, including comparisons to trends seen in the commercial market.

Chapter 2

TEST FACILITIES AND DIAGNOSTICS

As electric propulsion systems are built to operate in space, the testing of EP systems is dependent on vacuum facilities to simulate the space environment, and diagnostics to characterize thruster operation and plasma properties. This chapter will detail the common test equipment used across the test campaigns presented in this dissertation. Where a test campaign diverges from the methods presented here, the updated and methods will be presented in the appropriate chapter.

2.1 The Space Test Facility (STF)

The SPACE Lab Space Test Facility (STF) is a 3 m³ chamber, pumped by up to three CP-20 and one CP-12 Sumitomo cryopumps and a Shimadzu TMP-V2304 turbopump giving a total pumping speed of > 26 kL/s for krypton and a base pressure of < 1×10^{-7} Torr, see Fig. 2.1. An all-metal gas feed system with two Alicat mass flow controllers deliver research grade 99.9999% krypton to the anode and cathode. Pressure is monitored by a nude ion gauge behind the thruster exit plane. Down stream of the thruster, a water-cooled graphite beam dump takes the majority of the beam energy while other downstream areas of the chamber are coated with a layer of graphite to reduce sputtering from the thruster plume and therefore reduce deposition of material onto the thruster. STF will be used for the test campaigns detailed in chapters 3-5.

2.2 The Space Test Facility 2 (STF-2)

The SPACE Lab Space Test Facility 2 (STF-2) is a 5 m³ stainless steel chamber, pumped by two custom cryosails and a Varian TV-551 turbopump totaling a pumping speed of > 75 kL/s




Figure 2.1: A Schematic of the STF showing the position of the diagnostics, chamber sections, and test equipment.

for argon, > 60 kL/s for krypton, and > 50 kL/s for xenon, with a base pressure of $< 2 \times 10^{-8}$ Torr. The inner walls of the chamber are lined with an ethanol-cooled cryo shrouds at -55 °C which is thermally isolated from the walls of the chamber by MLI blankets. Downstream of the thruster, an ethanol-cooled graphite beam dump takes most of the plume energy, while other downstream areas of the chamber are coated with a layer of graphite to reduce sputtering from the thruster plume and reduce deposition of material onto the thruster. An all-metal gas feed system with four MC-500 Alicat mass flow controllers delivers research grade 99.9999% argon, krypton, and xenon to the ACME anode and cathode. A baffled nude ion gauge monitors the pressure at the midpoint of the chamber. STF-2 will be used for the test campaign detailed in chapter 6.

2.3 Thrust Stand

The characterization of electric propulsion system thrust is challenging due to limitations of the test environment and thrust/weight ratios as low as 1/1000 [54]. While chemical rockets

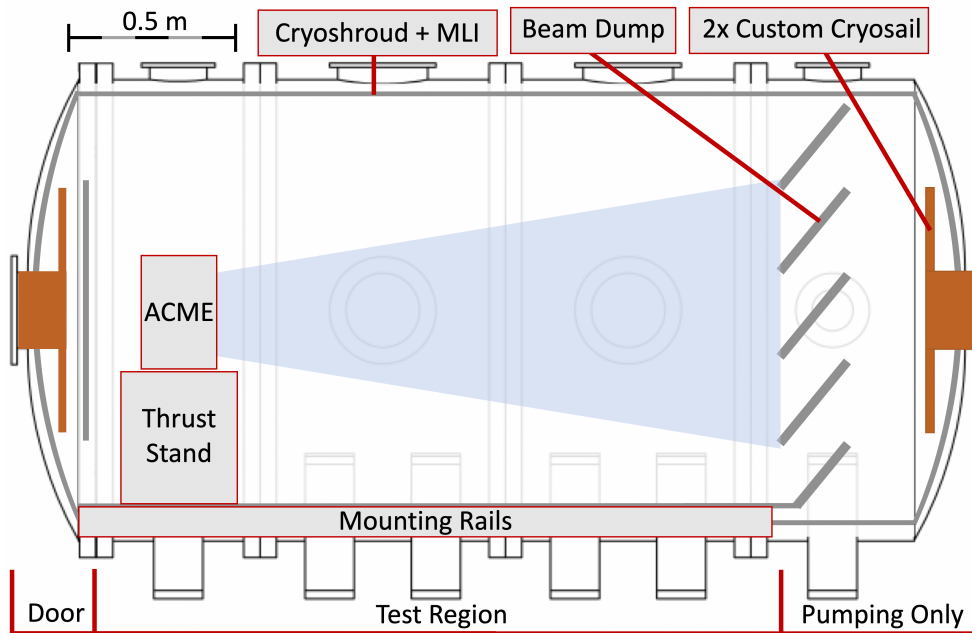


Figure 2.2: The STF-2 showing the position of thruster, thrust stand, and chamber sections.

are able to utilize load cells for direct measurement, electric propulsion systems typically rely on the motion of a pendulum to separate the effects of gravity and thrust [55]. However, the sensitivity of the thrust stand is directly proportional to the deflection, and therefore the length of the pendulum arm. As a result, their size is now a primary restriction to their sensitivity. If we also consider facility effects, which are minimized by centering the plume within the chamber, the available length for a thrust stand pendulum is further reduced [56].

The lowest thrust levels (nN- μ N) use torsional pendulums which provide the highest sensitivity levels and restrict motion to a plane perpendicular to gravity. They can be suspended by torsion flexures, thin filament, or be even magnetically levitated to minimize external forces [57]. Inverted and hanging pendulums are more typically used for measurements above 1 mN [55]. Greater sensitivity can be derived from an inverted design as gravity acts to amplify the deflection [58]. Hanging pendulums are more commonly used for simpler designs, although they have been used with the addition of mechanical amplification across

a broad range of thrust levels [59].

Figure 2.3, shows a review of thrust stands with published results in EP testing. This includes stands with additional factors such as double and counter-weighted pendulums [54, 60]. SPACE Lab required a thrust stand capable of measuring low thrust 10 W systems to high thrust multi-kW systems, all while keeping the thruster on the chamber center-line of STF. This presents a significant challenge to capture such a wide range of thrust measurements which span the two regions dominated by vertical and torsional configurations. To achieve this, a novel concept was created where the thruster was ‘nested’ in an inverted pendulum which was itself mounted in a water-cooled shroud. The thrust stand uses interchangeable flexures and adjustable eddy current damping, and is capable of a wide range of steady thrust measurements from $100 \mu\text{N}$ - 1 N [68]. A Keyence IL-100 laser with IL-1000 amplifier samples the position of the thruster at rate of 1 kHz from within the water-cooled shroud, where it is shielded from the plume and thermal radiation. The thruster mount is electrically isolating and provides additional thermal standoff. A waterfall feeds electrical, gas, water, and thermocouple connections through the shroud and to the thruster. Figure 2.4a shows

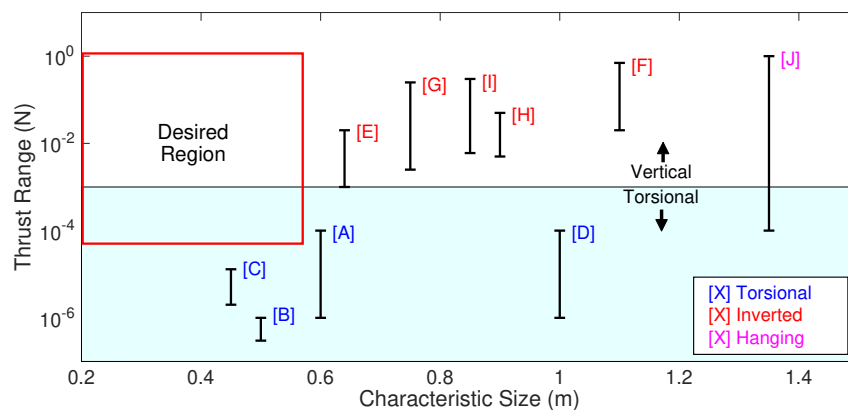


Figure 2.3: A review of thrust stands (taken at the time of thrust stand design), comparing their thrust measurement range to their size. Data is from published test results, actual ranges are likely larger. Source: A [60], B [54], C [61], D [62], E [63], F [59], G [64], H [65], I [66], J [67].

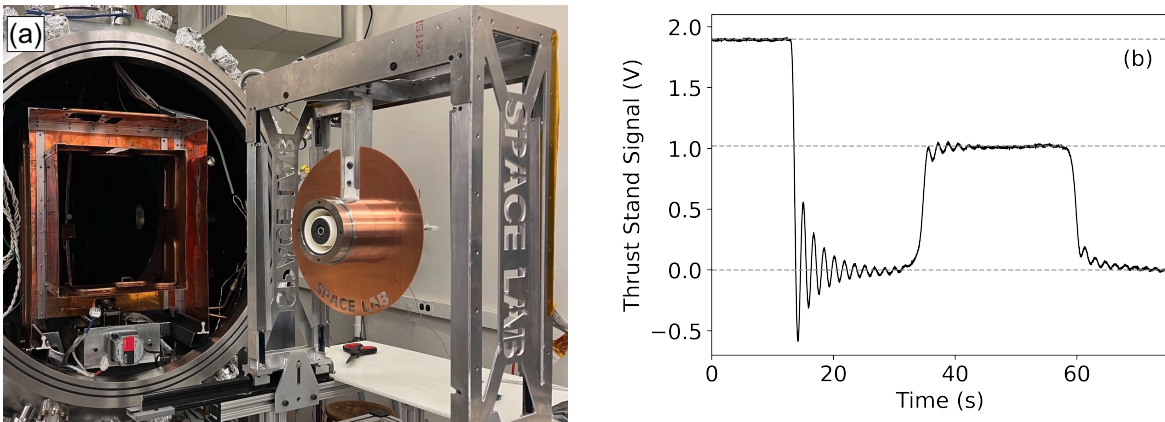


Figure 2.4: a) ACME mounted to the thrust stand outside the STF. b) A sample thrust measurement taken during the V2 test campaign.

ACME mounted to the thrust stand which sits on linear bearings and a rail system. This allows the anode and pole position to be reconfigured without removing the thruster from the thrust stand.

The process of measuring thrust consists of ignition, adjusting the thruster controller to the desired parameters, and waiting for thermal steady state. On reaching steady state, the thruster is shutdown and the thrust delta is measured. A single calibration mass is added and removed to confirm the calibration taken before testing. A plot of the thrust signal during shutdown and calibration is shown in Fig. 2.4b.

An additional thrust stand was built for STF-2, derived from the same design but without the use of the ‘nested’ geometry. As with other SPACE Lab thrust stands, interchangeable flexures and adjustable eddy current damping gives a wide range of steady thrust measurement.

2.4 $E \times B$ Probe

Calculating the mass and charge utilization efficiencies requires the charge state of the beam ions, measured using an $E \times B$ probe. Ions enter the probe through a collimated aperture and traverses a region with perpendicular electric and magnetic fields. Ions in this region

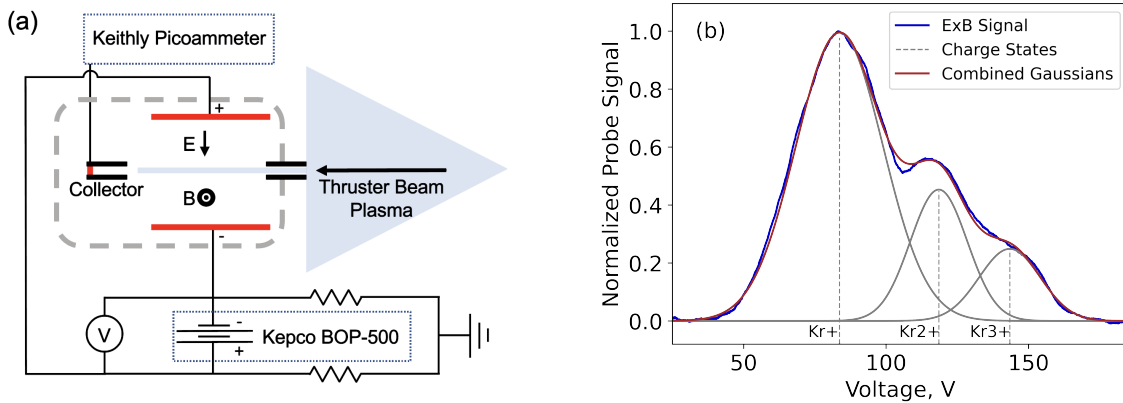


Figure 2.5: a) A schematic showing the electrical setup of the E×B probe. b) E×B data with Gaussian fits for three charge states of Krypton measured.

experience a Lorentz force of $F = q(E + v \times B)$ as they travel through the fields. If the ion velocity is such that the net force is zero, it passes through the region and is measured by a collector at the rear of the probe. As different ion charge states are accelerated through the same discharge potential in the thruster, they will have different velocities, and will therefore require different electric field strengths to counteract the $v \times B$ drift. Measuring the collected current as the electric field strength is varied, gives distinct peaks in the signal corresponding to different charge states.

The E×B probe is positioned 10 thruster diameters downstream (0.67 m) of the exit plane, offset from the centerline of the chamber to align with the thruster channel. Figure 2.5a shows the electrical setup of the E×B probe. Plasma from the thruster moves through a 35 mm long, 0.5 mm diameter collimator, into the probe housing. A Kepco BOP-500 driven at 1 Hz applies a 0 - 7000 V/m electric field, while permanent magnets apply a 1300 G field. The collector current is measured by a Keithly picoammeter.

Analysis of the peaks in the signal for each charge state can be done in several ways: peak height, triangle fitting, Gaussian and double Gaussian fitting, and variable exponential fitting are all accepted methods [69, 70]. Figure 2.5b shows a Gaussian fit of three Krypton charge states to the E×B probe signal. The gaussian method was used for this data analysis

as it was most effective at dealing with the overlap due to the broad peaks for each charge state.

Gaussian fitting starts with the first peak location, V_1 , taken at the point of highest current, while subsequent peak locations are set at $V_1\sqrt{Z_i}$ V for the i th charge state due to the ion velocity, v , scaling with $v \propto \sqrt{Z_i V/m}$ where Z_i is the ion charge state, V is the potential the ion is accelerated through, and m is the ion mass. Iterative Gaussian fits were calculated to minimize both the total and absolute residuals of the summed Gaussians. The integrated area under each gaussian is the total signal collected for each charge state. This is used to calculate the current fraction of a charge state, Ω_i as

$$\Omega_i = \frac{I_i}{I_d}, \quad (2.1)$$

where I_i is the current due to ions of charge state i being collected, and I_d is the total current collected. The charge utilization efficiency, η_q , can then be calculated as

$$\eta_q = \frac{\left(\sum \frac{\Omega_i}{\sqrt{Z_i}} \right)^2}{\sum_i \frac{\Omega_i}{Z_i}}, \quad (2.2)$$

using the sum of current fractions for each charge state. The charge utilization efficiency quantifies the losses due to multiply charged ions. As ionizing the propellant requires significant energy, doubly and triply charged ions have used additional energy that could have ionized additional propellant, while also using more power as they are accelerated to higher velocities.

2.5 Faraday Probe

Calculating the divergence and current utilization efficiencies requires an angularly resolved ion current density of the thruster plume. This is measured using a nude Faraday probe consisting of a collector and a guard ring, biased to a negative potential. Incoming ions

are collected by the probe and the current is measured, giving the ion current density over the known collection area. The negative potential repels electrons and must be sufficient to maintain ion saturation to ensure all ions are collected. The guard ring and the thin gap to the collector are designed to maintain a flat uniform sheath across the entire collector surface.

Far field measurements of the Hall thruster plume current density are taken with a Faraday probe mounted to a swept arm, see Fig. 2.6a. The probe sweep, driven by stepper motor, is limited to $\pm 90^\circ$ by the thrust stand shroud. Although experiments on spacecraft interaction have shown measurable plume as far as 140° from thruster centerline [71], the sweep to 90° covers $> 99\%$ of the plume current. A hemispherical coordinate geometry is used with the axis of rotation on the thruster exit plane, and the collector positioned at 0.54 m or 8 TDD. The thrust stand can deflect by up to 0.5 mm during operation, which amounts to 0.09% of the arc radius, with insignificant effect on the measurement.

The probe has a planar collector and guard made from 0.5 mm thick 99.9% molybdenum and is similar in design to probes used in literature [72]. Alterations to the construction give a clean front face and a tight gap between the collector and guard ring. The molybdenum is mounted to an alumina silicate backing with only four narrow tabs connecting inner and outer sections. This allows the plume impinging on the gap to be sufficiently vented.




Figure 2.6: a) The Faraday probe mounted to the stainless steel arm. b) The measurement circuit used for the Faraday probe.

Probe dimensions were chosen based on recommended practices [73] as a balance between angular resolution and signal for the wide range of thrusters tested in the STF. The collector has a diameter of 12.5 mm which gives an angular resolution of 1.25° at 8 TDD. The gap between the collector and guard is 0.4 mm and is sized to maintain a flat sheath on the probe while being reliably machinable. The probe is mounted 50 mm upstream of the stainless steel arm to reduce the influence of the conducting structure which is electrically floating, thus not providing a low impedance path to ground for the beam electrons.

Figure 2.6b shows the circuit used to take the current density measurements. The current is measured across a 100Ω shunt resistor using a floating differential measurement. A TLP7920F optically-coupled isolated amplifier conditions the floating measurement for the DAQ and has excellent common mode noise rejection. For all tests, the probes are biased to -30 V with a DC power supply to maintain the potential in the ion saturation region. The shunt resistor was sized to limit the the potential drop on the collector and maintain ion saturation. The circuit was calibrated and the signal checked for linearity using a series of applied resistances to ground, generating known currents.

The probe current is dominated by beam ions towards the centerline of the chamber. However, as the probe moves away from the centerline, charge exchange (CEX) ions become significant. This effect is exacerbated with increasing background pressure. Correcting for CEX ions in the Faraday signal is a major source of uncertainty in the total beam current and divergence measurements due to the relatively large area the signal is integrated over. In the far-field, recommended practice is to characterize the facility effects on the measurement with a series of tests at a range of background pressures [73]. A linear regression of current density as a function of pressure is used to achieve an expected current density in total vacuum. Alternatively the linear or extrapolated-wing method for correcting Faraday probe data can be used [74, 75]. This method takes the derivative of the current density in the plume dominated region and extrapolates it with an exponential fit to 90° . The uncertainty in the extrapolated-wing method increases significantly with more divergent plumes where the shape of the central plume does not fit the exponential decay. Tests of these methods

and their comparison to a novel method proposed in this chapter are detailed in Section 4.2.

The current from the probe must be corrected for the effective probe collection area, and secondary electron emission (SEE) [73]. This gives an effective current density, J , as

$$J = \frac{I_C}{A_C + \kappa_G} \kappa_{SEE}, \quad (2.3)$$

where J is the current density, I_C is the collector current, A_C is the collector area, κ_G is the effective probe collection area, and κ_{SEE} is the secondary electron emission correction. The effective probe area correction accounts for the thickness of the molybdenum probe surface, which gives additional surface area for ion collection in the gap between the collector and guard ring. The effective probe collection area, κ_G , is

$$\kappa_G = \pi(R_{GR}^2 - R_C^2) \left(\frac{2\pi R_C h}{2\pi R_C h + 2\pi R_{GR} h} \right), \quad (2.4)$$

where R_C is the outer radius of the collector, R_{GR} is the inner radius of the guard, and h is the thickness of the molybdenum. The thin molybdenum collector of the probe used has a thickness of 0.5 mm, resulting in a negligible value of κ_G . Although molybdenum has a low SEE yield, the number of sputtered electrons that accelerate away from the negatively biased surface is enough to increase the measured current, and must therefore be accounted for. The secondary electron emission correction, κ_{SEE} , is

$$\kappa_{SEE} = \frac{1}{1 + \sum_i \frac{\Omega_i \gamma_i}{Z_i}}, \quad (2.5)$$

where Ω_i is the current fraction of the ion charge state i as measured with the E×B probe, Z_i the ion charge state, and γ_i is the charge state dependent SEE yield of molybdenum. The values used are $\gamma_1 = 0.022$, $\gamma_2 = 0.20$, and $\gamma_3 = 0.70$ [73]. The correction is the fraction of measured current due to beam ions, out of the total measured current.

The corrected current density is used to calculate both the current and divergence utilization efficiencies. The divergence utilization efficiency, η_d , measures the plume divergence

and shows the efficacy in the ions at producing thrust. Ions that travel directly downstream contribute all their momentum to the thruster, while highly divergent ions contribute very little to the desired thrust vector. It is defined as

$$\eta_d = \left(\frac{I_{\text{axial}}}{I_{\text{beam}}} \right)^2, \quad (2.6)$$

where the total current in the beam is

$$I_{\text{beam}} = 2\pi R^2 \int_0^{\frac{\pi}{2}} J(\theta) \sin(\theta) d\theta, \quad (2.7)$$

and the current in the axial direction is

$$I_{\text{axial}} = 2\pi R^2 \int_0^{\frac{\pi}{2}} J(\theta) \cos(\theta) \sin(\theta) d\theta. \quad (2.8)$$

Here, the angularly dependent current density is integrated over the hemispherical coordinate system, at the radius of the probe sweep, R [73].

The current utilization efficiency, η_b , measures what fraction of the discharge current that is converted to beam current. It is the flux of ions leaving the thruster, and is reduced by electron flux to the anode. It is defined as

$$\eta_b = \left(\frac{I_{\text{beam}}}{I_d} \right), \quad (2.9)$$

where I_d is the nominal discharge current of the thruster.

The mass utilization efficiency, $\eta_m = \dot{m}_b / \dot{m}_a$, is defined as the ratio of the mass flow of beam ions, \dot{m}_b , to the total mass flow rate through the anode, \dot{m}_a . η_m represents the fraction of the propellant that is ionized and accelerated. Neutrals that exit the thruster are wasted as they produce very little thrust. Using the quantities measured above, we can calculate η_m as

$$\eta_m = \frac{\eta_b I_d M}{\dot{m}_a e} \sum_i \left(\frac{\Omega_i}{Z_i} \right), \quad (2.10)$$

where M is the propellant atomic mass, and e is the electron charge.

2.6 Langmuir Probe

A cylindrical Langmuir probe with a diameter of 0.8 mm and length of 10 mm was used to characterize the plasma potential of the plume. The probe was mounted the same rotating arm as the Faraday probe and RPA with a lateral offset of 10 cm to minimize interference and with the electrode pointing directly at the center of the thruster. A series of measurements were taken between 0 and 90° in 15° increments to characterize the plasma potential needed to correct the RPA measurements. The probe is swept between -150 and 150 V at 1 Hz by a Kepco BOP-500M and the current measured across a shunt resistor. Figure 2.7 shows a typical Langmuir probe signal. A fit of the log of the current from the Langmuir probe is taken, with the plasma potential ϕ_p taken as the point where the exponential current growth transitions to linear growth, resulting in a ‘knee’ in the log trace.

2.7 Retarding Potential Analyzer

Calculating the voltage utilization efficiency, and comparing the efficacy of the center-line and angularly resolved IEDF methodologies, requires RPA measurements of the plume. The RPA

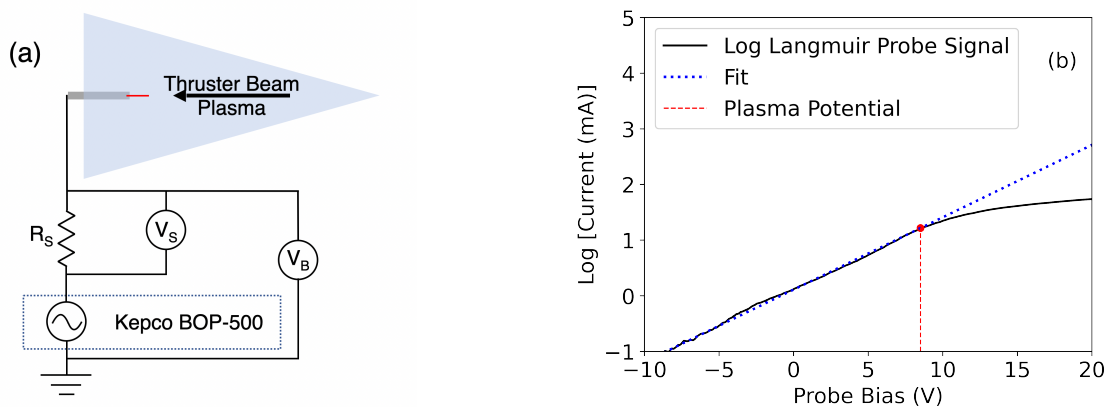


Figure 2.7: a) The electrical setup of the Langmuir probe. b) A sample Langmuir probe trace showing the natural log of the current, and ‘knee’ at the plasma potential.

consists of a conducting aluminum body which houses four parallel grids and a collector. The first grid is electrically floating to minimize perturbation to the plume plasma. The second grid in the RPA is negatively biased and repels electrons. Plume ions passing through these grids reach the third grid, which is biased to a positive potential. Ions with a higher energy than the grid potential will pass through it, while ions with lower energies will be retarded. As the retarding grid potential increases, so too does the fraction of the beam ions retarded. We therefore take the first derivative of the collector current as the IEDF, as it represents the current from ions with energy equal to the retarding potential. The grid potential is increased until all ions have been retarded and the collected current is 0. The fourth grid is negatively biased to repel electrons caused by secondary electron emission from ion impacts beyond the second grid.

The RPA is mounted to the swept arm at 8 TDD (0.54 m) of the exit plane. It is mounted level with the Faraday probe with a 10 cm lateral offset to avoid interference. The RPA is pointed directly at the center of the thruster with the axis of rotation of the swept arm passing through the exit plane.

The electrical setup of the RPA is shown in Fig. 2.8a. Grids 2 and 4 are biased at -30 V . Grid 3 is swept from 0 V to the Discharge voltage +75 V at 1 Hz by a Kepco BOP-500M, and retards incoming ions across the full range of ion energies. The collector current is measured using a Keithly 6485 picoammeter.

The center-line IEDF is the standard for EP and has been demonstrated to give accurate data with the well focused plumes seen in most Hall and gridded ion thrusters [76]. With the RPA positioned on thruster centerline in the far field, the voltage of the retarding grid is swept and the collected current is measured. Figure 2.8b shows a typical RPA signal and its dependence on the retarding voltage. As the retarding voltage is measured relative to ground, the signal must be corrected for the plasma potential to get the acceleration voltage, V_a , as

$$V_a = V_{RPA} - \phi_p \quad (2.11)$$

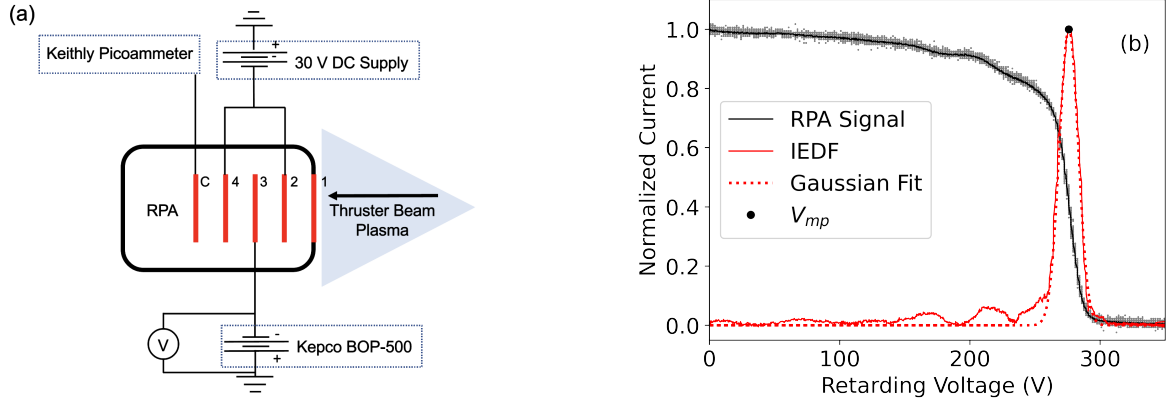


Figure 2.8: a) The electrical setup of the RPA probe. b) A single point centerline RPA measurement showing the raw current, derived current, Gaussian fit, and most probable voltage.

where V_{RPA} is the retarding voltage, and ϕ_p is the local plasma potential measured by the Langmuir probe. The acceleration voltage is the potential a plume ion was accelerated over.

The voltage utilization efficiency, $\eta_v = V_b/V_d$, is defined as the ratio of the beam voltage V_b , to the thruster discharge voltage, V_d . It measures how effective the anode to cathode potential is at accelerating ions. Propellant that is accelerated through a lower potential is less energetic and therefore less effective at producing thrust. The value of V_a at the peak of a Gaussian fit to IEDF, see Fig. 2.8b, as the most probable voltage, V_{mp} , and is assumed to represent the beam voltage across the entire plume. This gives η_{v1} , as

$$\eta_{v1} = \left(\frac{V_{mp}}{V_d} \right). \quad (2.12)$$

Chapter 3

THE ACME HALL THRUSTER

The Adaptive-field Central-cathode Magnetically-shielded Electric (ACME) Hall thruster is a magnetically-shielded adaptive Hall thruster, designed to study the influence of magnetic field structure on thruster behavior and performance. ACME is both a versatile research device, built to enable the other research presented in this dissertation, and in itself a key component of the research. The adaptive design allows it to be reconfigured to a wide range of experimental setups without requiring disassembly or rebuilding. The adaptability exhibits a continuum of plume structures from highly divergent to over-focused, see Fig. 3.1. This chapter details the design of ACME, the test campaigns used to initially characterize the thruster, and the design updates used to improve performance throughout its life.

3.1 ACME Design and Construction

To achieve the goals of studying the influence of magnetic field structure on Hall thruster performance, we require a thruster capable of a range of field shapes while still being representative of flight Hall thrusters. Addressing these specifics presents a challenge in design, resulting in the ACME thruster. The thruster has three physical degrees of freedom, all in the axial dimension: the relative position of the poles, the position of the anode, and the position of the centrally mounted cathode. While external and centrally mounted cathodes are both common, the changes in pole position may cause potential issues with cathode coupling, leading to the central cathode as more favorable option.

The design of ACME follows Hall thruster scaling laws from Refs. [77, 78], only diverging where the adjustability of the poles restricts the design envelope. The ACME design begins with the discharge power, which is then used to define channel dimensions and the required




Figure 3.1: The extremes of pole positions tested on ACME to date compared at 400 V, 1.5 A on krypton. The +4 mm pole position and associated photo, top, shows a highly divergent plume. The -12 mm pole position, bottom, shows an over-focused plume. The inner pole is located stationary in the frame of the photos, while the outer pole moves.

magnetic field strength. Combining these with the desired magnetic field shape, forms the core of the thruster design, which drives the mechanical design of the thruster.

The thruster power is primarily limited by the pumping speed of STF, and desired background pressure based on recommended practices for electric propulsion testing [79, 80]. The design power of ACME is 500 - 750 W is based on a discharge voltage of 300 V commonly used in thrusters at this power level [81, 78, 82], the use of xenon and krypton propellants, and scaling laws of thruster specific impulse and efficiency [78].

With the power range defined, the channel cross section of 16.5 cm^2 was calculated based on an ideal current density range of 0.1 to 0.15 A/cm² [2]. The channel dimensions are influenced by the magnetization of the electrons and ions, the ratio of the channel surface

area to volume, and the inner pole diameter required to house the cathode, inner magnet, magnetic circuit, and walls. Scaling laws of low power Hall thrusters (< 1.35 kW) give a channel width of

$$h = 0.242d, \quad (3.1)$$

where h is the channel width, and d is the channel mean diameter [78]. This would give ACME a mean channel diameter of 46.6 mm and channel width of 11.3 mm, however, designs of the inner pole required a minimum of 49 mm to fit the required structure and magnets. Therefore, the mean channel diameter was increased to 57 mm with a channel width of 8 mm. This gives a channel width to diameter ratio of 0.14, significantly lower than the unshielded thrusters used for the scaling laws, but closer to the comparably shielded MaSMi-60 Hall thruster which has a ratio of 0.157 [83].

The channel's width determines the required electron Larmor radius, r_{Le} , for its confinement through the magnetic field strength. This relationship is

$$r_{Le} = \frac{m_e v_e(T_e)}{eB}, \quad (3.2)$$

where m_e mass of an electron, $v_e(T_e)$ is the temperature dependent electron velocity, e is the elemental charge, and B is the magnetic field strength. Here, we differ from the scaling laws for unshielded thrusters due to the location of the Hall current and peak magnetic field. Simulations of Hall thruster discharges with Hall2De, have seen significantly hotter electrons in shielded thrusters than their unshielded counterparts[33], however the larger Larmor radius is mitigated by the position of the Hall current, which is at or beyond the exit plane of the channel. Using an electron temperature of 40 eV, taken from a representative thruster [34], we calculate that a 250 G field is required to confine the electrons to one tenth of the channel diameter.

The minimum channel depth is defined by the Melikov–Morozov criterion, such that

$$\lambda_i = \frac{v_n(T_n)}{n_n \sigma_i(T_e) v_e(T_e)} \ll L \quad (3.3)$$

where λ_i is the ionization mean free path, L is the channel depth, $v_n(T_n)$ is the temperature dependent neutral velocity, n_n is the neutral density, and $\sigma_i(T_e)$ is the ionization cross section. Here, the ionization mean free path is < 1 mm for an ideal channel neutral density of $1.2 \times 10^{19} \text{m}^{-3}$ with an expected neutral temperature of 800 K [77]. The channel length is therefore derived from similar thrusters, with an additional allowance for adjustment, particularly for using alternate propellants.

The magnetic shielding profile is designed to be similar to that of JPL’s MaSMi and shielded versions of the H6 thruster [81][23]. The magnetic field design used Finite Element Method Magnetics (FEMM) simulations to determine the layout of the Hiperco 50A magnetic circuit and electromagnets. After machining the Hiperco components, they were annealed to achieve the desired magnetic properties. Figure 3.2, shows the field of ACME version 1 (V1) after the simulation was corrected for the measured magnetic properties of the material, and maps of the resulting field taken with axial and radial Hall effect probes.

The resulting field shape is slightly over-shielded on the outer pole, similar to earlier versions of MaSMi [32] and under-shielded inner pole. Here, over-shielded refers to an excessive curvature of the magnetic field around the pole caps, leading to higher plume divergence. The field is under-shielded when the curvature is insufficient to clear the pole cap and instead intercepts the channel wall. At all pole positions, the center-line field monotonically increases from the anode to a peak beyond the exit plane. Improvements to the field will be detailed in section 3.2.2.

To achieve the desired magnetic field inclination, the thruster’s magnetic circuits are adjusted by physically moving the inner and outer poles assemblies relative to each other. This allows changes to the field structure beyond changes to the inner and outer electromagnet currents. An interface between the magnetic circuits for the inner and outer poles at the rear




Figure 3.2: The magnetic field of ACME V1 with the under-shielded inner pole on the left and over-shielded outer pole on the right.

of the thruster allows each pole to move as a single unit. Because the channel wall, magnet, insulators, and Hiperco magnetic circuit all move together axially, the shape of the shielding on each of the poles remains consistent. The magnetic circuit at the rear of the thruster allows the magnetic flux to flow consistently across the range of pole positions. While this enables the thruster to be a highly configurable test platform, the design comes at a cost, requiring significantly higher magnet power than similarly sized thrusters. The thruster is capable of 40 mm of total change in relative pole position. A +5 mm pole position, is the most positive value possible with the thruster and corresponds to an inner pole positioned 5 mm in front of the outer pole, see Fig. 3.3. 0 mm corresponds to a flat exit plane, while negative pole positions see the inner pole inset into the thruster.

Both inner and outer magnets are single pole designs of 15 awg wire insulated by a high temperature silicone fiberglass wrap in a ceramic binder, limiting the thruster body temperature to 400C. Finite element analysis of the thermal performance indicates that localized heating of the Hiperco pole caps from the combined thermal loading of the plasma,




Figure 3.3: The full range of pole positions ACME can operate in: 0 mm inner position, left; maximum negative inner pole position of -35 mm, center; and the maximum positive inner pole position of +5 mm, right.

cathode, and magnets, sets the temperature limit due to field saturation in the inner Hipercop pole cap. As the inner pole has a lower radiative view factor, and higher expected plasma heating due to the larger values of inset, it is mounted as the stationary pole to improve conductive heat transfer. Figure 3.4 shows the conduction paths for the plasma and magnet thermal loads. As the outer pole slides further out, its contact with the casing and therefore thermal conductivity to the heat sink decreases, however, its ability to radiate heat improves. Thermocouples to track the internal temperature were attached to the backside of the inner and outer poles with a ceramic adhesive. A copper heat sink is mounted to the rear of the stainless steel rear case. It provides a secondary function of ensuring no plasma interaction with the electrical connections at the rear of the thruster.

The hollow cathode is a heaterless design, manufactured by Plasma Controls using barium-based porous tungsten emitter [84]. The emitter is housed in a 1/8" tantalum tube with a 0.51 mm orifice. The tube is electrically and thermally insulated from the stainless steel keeper assembly by a 1/4" alumina tube. An interchangeable graphite keeper electrode with 0.6 mm orifice is mounted at the end of the keeper assembly, 3 mm from the cathode orifice.

The channel walls in V1 are machined from graphite to be thin enough to work with




Figure 3.4: Conduction paths for plasma and magnet heating on version 1 of the ACME thruster.

the movement of the adjustable magnetic circuit, while remaining within the budget of the project. A thin coating of alumina was applied to the channel walls for electrical insulation. Figure 3.5 shows ACME V1 before its initial test campaign. An exploded view of the thruster and list of parts can be found in appendix A, Figs. A.1, A.2, A.3, A.4, and tables A.1, A.2, A.3.

3.2 ACME Testing and Design Updates

After the manufacturing of ACME was complete, a series of experimental campaigns were used to characterize the thruster operation. These campaigns were used to assess the thruster for any design shortfalls, and begin studying the effects of changing pole position. This section details two test campaigns and the resulting thruster design updates.

3.2.1 ACME V1 Testing

ACME, being the first prototype of an adaptable Hall thruster, required an initial check-out focused on developing the ignition methodology, finding stable operating regions, and verifying the magnetic shielding.




Figure 3.5: The V1 ACME thruster without the centrally mounted hollow cathode . The heat sink shown was not used during V1 250 W tests.

Figure 3.6 shows the electrical configuration used for V1 testing. The main and keeper discharges were driven by a TDK-Lambda 600-2.6 and 600-1.3 DC power supplies respectively, with EMS 100-10 and EMS 20-50 DC power supplies running the magnets. Flyback diodes and ballast are used to protect the power supply during ignition. The discharge current was measured by an Elditest CP6990 current probe on the cathode electrical connection between the filter and cathode external to the chamber. The thruster is electrically floating in this configuration.

Early visual indications showed the plasma offset from the both channel walls for field

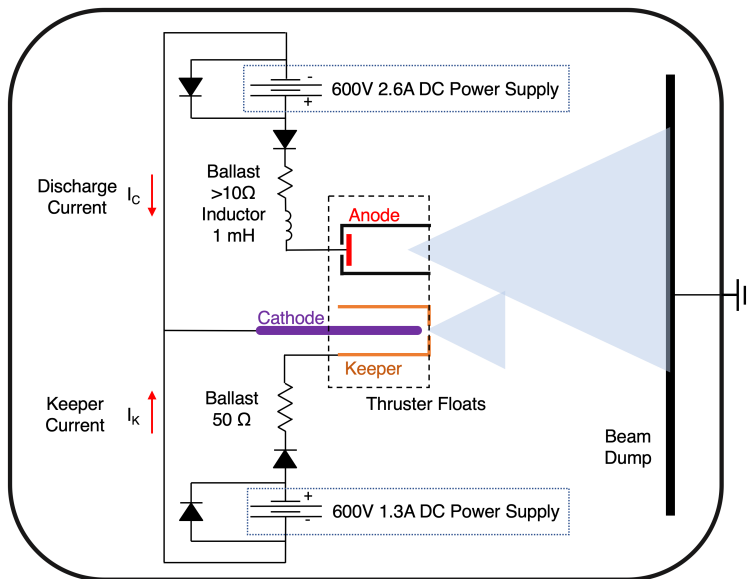


Figure 3.6: The floating electrical configuration used for ACME V1 testing with inductive loads on both circuits.

strengths greater than 150 G in the 0 mm pole position. Figure 3.7 shows both external DSLR and internal camera images of the thruster running on xenon. The DSLR image shows darker areas near the channel walls, similar to the shielding seen on MaSMi [81]. The sweep of the internal camera gives a clear view of the anode from either side of the channel plasma, further indicating that the thruster may be suitably shielded.

Continued testing covered a range of voltages from 200 to 350 V, currents from 0.8 A to 1.3 A, and magnetic fields from 140 to 280 G. This aimed to find a stable operating point to be used in the first tests of the adjustable field concept. A stable operating range with a well defined plume shape was found centered around 250 V, and 1 A, at the 0 mm pole position. Table 3.1 shows the performance measured at this point. This campaign neglected optimization, instead simply aiming to verify proper thruster operation at the given pole positions. Measured anode efficiency and specific impulse running on xenon matched performance figures from comparable thrusters running at similar power levels [34]. Uncertainties in performance metrics include estimated error from entrained neutrals at the




Figure 3.7: ACME running on xenon. Left: DSLR image from outside the chamber showing dark regions on the edge of the plasma. Right: Image from the chamber internal camera showing a higher contrast view of the plasma.

2×10^{-5} Torr operating pressure.

To investigate the changes in plume structure, ACME was run at five pole positions (+3, +2, +1, 0, and -1 mm). In each pole position, the plume was qualitatively assessed using DSLR images, and the plume divergence measured using a swept Faraday probe using the method described in section 2.5. Figure 3.8 shows the DSLR images, simulated magnetic field, and Faraday probe trace for each pole position. The plume images clearly show a wide plume at the +3 mm inner pole position, with an increase in focus at each step towards

Table 3.1: Thruster Performance at 0 mm on Xenon

Power	250 V, 1 A	
Propellant Flow	1.1 mg/s	
Thrust	16 mN	+1 -3 mN
Specific Impulse	1450 s	+50 -200 s
Anode Efficiency	45%	+2 -8%
Total Efficiency	29%	+1 -5%

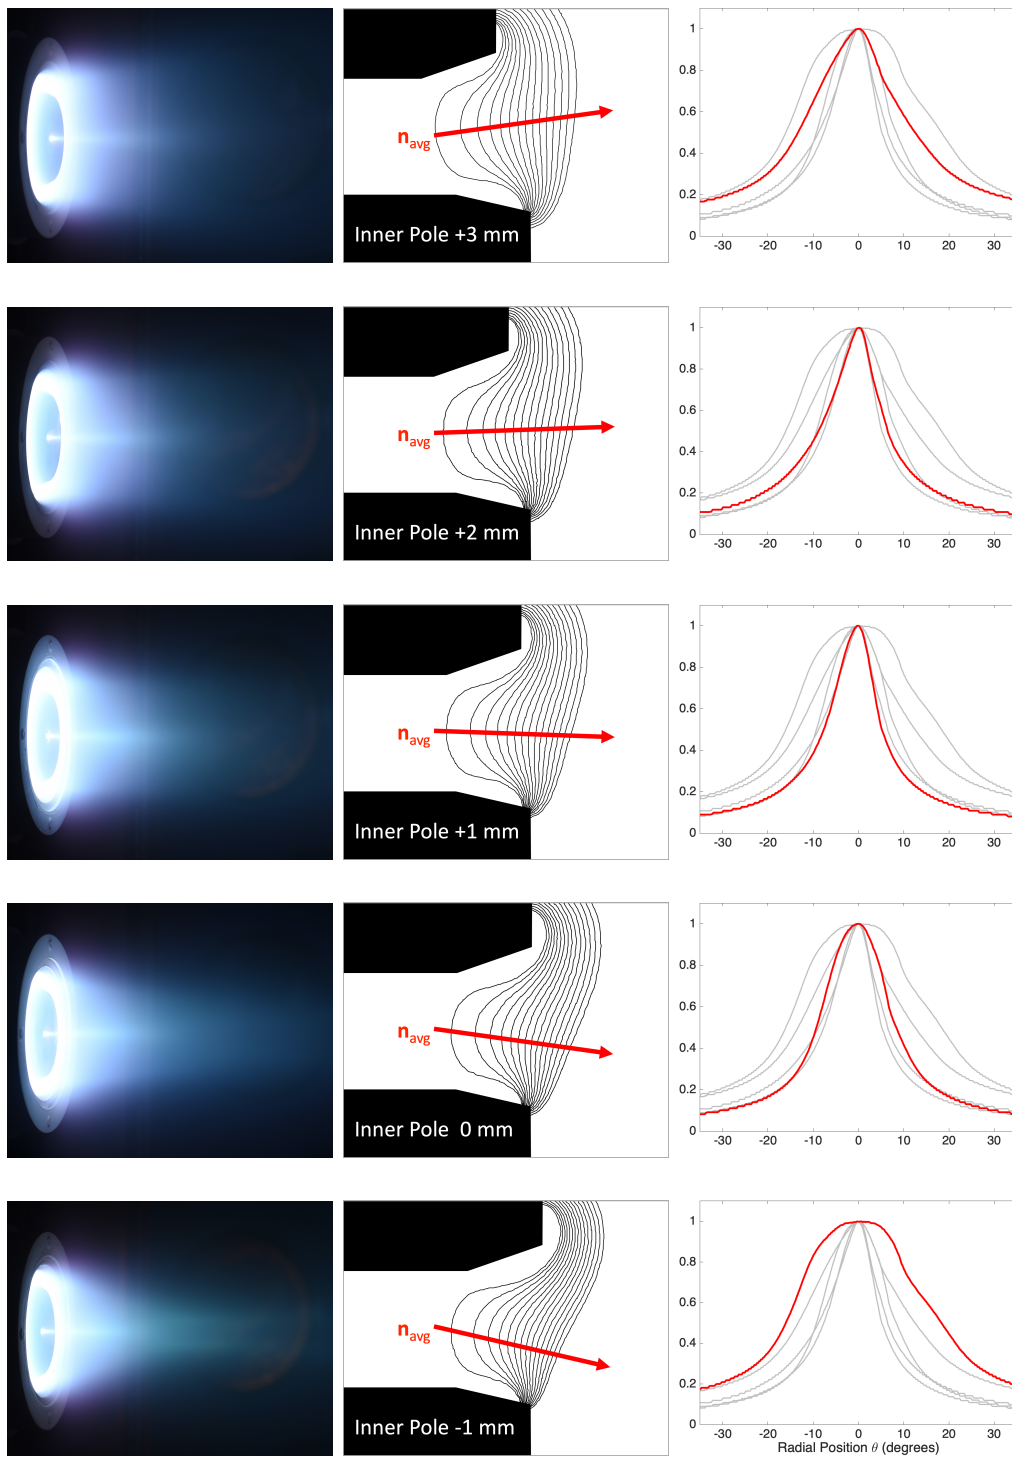


Figure 3.8: A comparison of the plume divergence across the range of pole positions tested, showing: DSLR photos of the plume were taken at 70 mm f/4.5 1/20" ISO200; Field simulations with the average normal vector to the field; and Normalized swept Faraday probe measurements comparing current density at each position.

the slightly over-focused plume seen at the -1 mm position. The change in field angle due to the pole position was quantified by integrating the radial and axial magnetic fields along the centerline of the thruster in the ionization and acceleration regions. Taking the vector normal to the magnetic field roughly indicate the direction of the electric field, and was used as a predictor for the plume angle. The average normal vector, n_{avg} was calculated as

$$n_{avg} = \frac{1}{n} \sum_{i=1}^n \tan^{-1} \left(\frac{B_Z}{B_R} \right) \quad (3.4)$$

where B_R and B_Z are the radial and axial magnetic fields strengths along the channel centerline. The vector and plume angle at the exit of the channel show good qualitative agreement and are linearly proportional for the range of positions tested.

Faraday probe sweeps show plume divergence decreasing as the pole position reduces from +3 to +1 mm, and over-focusing when at 0 and -1 mm causing an increase in measured divergence. The Faraday signal has been filtered for noise but has not been corrected for facility effects which cause a broadening of the measured plume due to interaction with background neutrals.

The gas feeds at the rear of the anode are positioned in the plane of the Faraday probe sweep, creating regions of slightly higher density due to uneven distribution from the single layer anode design. This was observed with the internal camera which showed brighter plasma in two regions coinciding with the injection points. This was observed across all pole positions and thus do not introduce significant error when comparing the plume divergence. The range of plume divergence seen with different pole positions, and the continued magnetic shielding of the poles is a sufficient validation of the adaptive field capabilities and meets the goals of the V1 test campaign. The results showed that the thruster performance is similar to other published research thrusters and that the adjustable poles were able to change the field angle while maintaining the shielded geometry. The predicted plume angles based on the averaged normal vectors were in good agreement with the measured plume, ranging from wide to over-focused. This relationship agrees with the channel field inclination tests from Ref. [36], and will be used to guide field profile updates to produce a magnetic field parallel

to the exit plane for ACME V2.

3.2.2 ACME Update to Version 2

The design updates for V2 aim to improve the reliability of thruster operation, and enhance the magnetic field shape through the channel and around the exit plane for more even magnetic shielding.

During testing of V1, the alumina coating on the channel walls of the thruster showed significant charring, limiting the duration of each test and required rebuilding between each pole position test. Once significant charring was observed, the performance and stability of the thruster was altered. Boron nitride (BN) walls, the most common wall material [2], needs to be thicker than the coated graphite. Therefore, in order to use BN walls without narrowing the channel, the Hiperco magnetic shields on either side of the channel require machining. Characterization of the Hiperco 50A and analysis of the shield design showed that the magnetic screens are thicker than required and can therefore be machined. A thicker section of the shield towards the rear of the thruster is used to complete the magnetic circuit as the outer pole moves. Therefore only a portion of the magnetic shield closest to the exit plane can be machined to allow more room for a BN wall without affecting the magnetic flux between the poles.

During the redesign for the walls, the shape of the magnetic shielding was updated. The field shown in 3.2 was unbalanced due to differences between the ideal and actual magnetic properties of the annealed Hiperco 50A used to manufacture the thruster. The shielding on the inner pole was insufficient, with a grazing field line from the outer pole impacting the wall on the inner pole. Figures 3.9 and 3.10 show the new BN walls and alumina pole caps designed for ACME V2. The inner Hiperco pole cap, and both magnetic screens were machined to create the updated field, while the walls and pole caps were chamfered to match the desired grazing line. Using the ideal grazing line strength described in Ref. [35], the field was contoured so the grazing line had a strength of 40 G on center-line for the 180 G field. Shielding on the inner and outer poles is even with a consistent grazing line across

each pole. The field at the exit plane has no axial component on centerline for the 0 mm pole position in V2, in contrast with the inclined field seen in V1. Simulations of the field at pole positions from -4 to +4 mm, were used to generate second order polynomial fits of the magnet currents between 140 and 250 G for each pole position. Each pole position shows a consistent field structure with similar curvature around the poles and channel walls, and a monotonically increasing field with a peak past the exit plane.

The electrical configuration was updated to add a hybrid inductive-capacitive filter to the discharge side of the circuit, see Fig. 3.11. The thruster is ignited in an inductive configuration on both the cathode-keeper and cathode-anode circuits which prevents arcing of the thruster by increasing the current rise time. Once operating, the inductive filter and




Figure 3.9: ACME version 2 mounted on the thrust stand outside the STF with new BN walls and alumina pole caps.




Figure 3.10: FEMM simulations of the magnetic field for ACME V2 with the new wall designs overlaid.

ballast is switched off on the cathode-keeper, and the main discharge is switched to the capacitive filter. This electrical setup is used for V2 thruster testing.

3.2.3 Version 2 Testing

The goal of the ACME V2 test campaign was to assess updates to the magnetic field and channel walls, and to characterize the performance of the thruster across the range of pole positions. The tested range of pole positions, -2 to $+2$ mm, was similar in scope to the V1 campaign, but centered around 0 mm due to the updated magnetic field shape. At each pole position the field simulation was compared to a channel center-line field profile, measured using a radial hall probe. At the 0 mm pole position, 200 points mapped the magnetic field confirming the machined magnetic circuit matched simulation. To standardize the channel depth across changing pole positions, the anode was positioned at a set distance from the midpoint of the exit plane (i.e. the average distance from the inner and outer pole caps). The anode position for xenon was based channel depths of other thrusters at similar power levels [78]. A field strength of 180 G was chosen based on the maxima of a current to mass flow rate analysis at 0 mm.

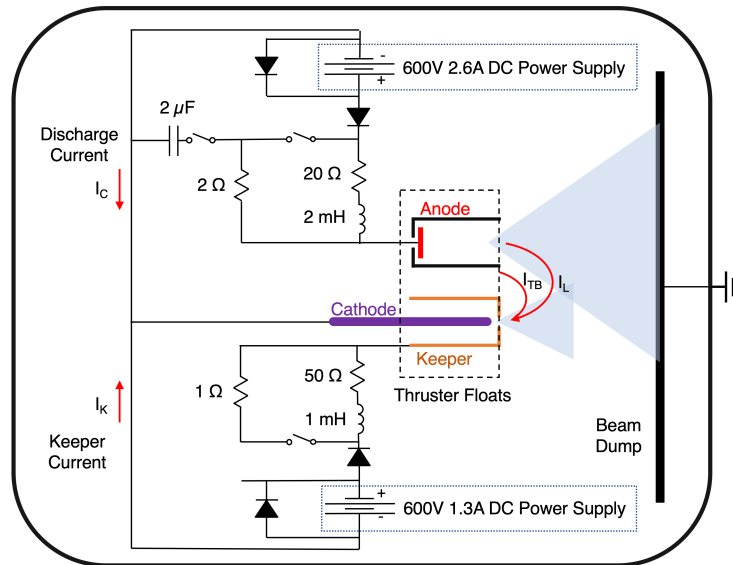


Figure 3.11: The floating electrical configuration with hybrid filter used for all tests in the V2 test campaign

As in the V1 campaign, a single operating point is required for this test. Two distinct modes were observed during the checkout and initial testing of ACME V2. Mode shifts are changes between two regions of operation where small changes to the thruster operating point cause large changes to thruster operation. These include: the anode gas flow required to maintain the discharge current, thruster oscillation frequency and magnitude, stability, plume shape, and efficiency. Figure 3.12 shows the Jet mode, left, and Diffuse mode, right. The Jet mode was found to be dominant at higher voltages and field strengths and had a significantly lower breathing mode amplitude.

Although under-shielded on the inner pole, the ACME V1 gave some visual indicators of magnetic shielding with significant plasma stand-off from the channel walls. Figure 3.13a shows a photo of the ACME V2 operating at the 0 mm position. The dense plasma in the Hall current appears close to or at the exit plane of the thruster and has a notable offset from the channel walls. More significantly, after testing the thruster, the BN channel walls were uniformly carbon coated, see Fig. 3.13b. The uniform deposition of carbon on the




Figure 3.12: The two distinct plume modes seen during thruster testing at the +1 mm position. The Jet mode, left, and Diffuse mode, right.

channel walls is an indicator of working magnetic shielding [85]. As a point of comparison, the thruster was run with an unbalanced field, resulting in a clean, carbon deposition-free, band on the chamfered section of the inner wall. This indicates that a field line intercepted that section of the BN, subjecting it to higher ion flux at higher temperatures, and creating a region with net erosion [86, 35, 85]. Figure 3.13b also shows clean and carbon-coated regions on the inner pole cap, indicating that the inner area has net erosion, while the edges of the cap have net deposition. Comparing this with the field simulations shown in Fig. 3.10, shows that this is where the field lines in the hall current region intercept the pole cap, resulting in net deposition within the shielded region and net erosion beyond it.

A thermocouple attached to the rear of the inner Hiperco pole piece was used to monitor temperature during testing. Balancing the field at different inner pole positions requires changes to the currents in the inner and outer magnets, resulting in changes to resistive heating. The inner pole experiences significant thermal loading from the plasma at negative




Figure 3.13: a) Plasma at the exit plane showing significant offset from the walls. b) Carbon-coated BN walls of the thruster and edge of the pole piece. The clean area, indicating net erosion, is seen on the inner section of the pole piece.

pole positions, making temperature monitoring necessary during operation and resulting in steady state temperatures of 380 - 420 °C. At the -1 and -2 mm positions, the operating point was close to a mode shift and temperature dependent mode shifts were observed. This was most likely the inner pole piece, which is under the highest thermal load and is closest to its magnetic saturation limit in normal operation. This is similar to observations of MaSMI from Ref. [85], which reported thermally dependent mode shifts which were linked to inner pole saturation. In ACME, simulations of the thermally induced saturation of the inner pole reduces the shielding strength near the channel wall. Once the field strength decreases sufficiently, the plasma interacts with the inner wall and the discharge shifts to a Diffuse mode.

Table 3.2 shows representative thruster performance taken in Jet mode at the chosen operating point of 300 V, 1.33 A, at the 0 mm pole position. These performance numbers show give an indication that the level of thruster performance is similar to other published research thrusters [81][32]. It shows a significant performance increase over the V1 test campaign with anode efficiency increasing from 45 to 51%. This improvement is due to

the increased discharge voltage, increased current density, improved shielding, and Jet mode operation.

Figure 3.14 shows a comparison of the plume divergence across the range of pole positions tested. DSLR photos were taken of the thruster at $1/5''$, $f/16$, ISO 200, with a 105 mm lens. The plume shapes are visually distinct and show a wide plume at +2 mm pole position with increasing focus with each negative step in pole position. Comparing the photos to those in the V1 test campaign (Fig. 3.8), the V2 Jet mode plume has a distinct high density plume directed from the channels, while the Diffuse mode plume of V1 appears less directed. The simulated fields with the average normal vector once again show good agreement with the visual plume shape.

Faraday probe traces taken at each position are used to compare the angularly dependent current density. A set of measurements were taken with two pumping configurations, the standard configuration operating at 2×10^{-5} Torr, and with one cryopump gated at 4.5×10^{-5} Torr. These tests were used to correct for background charge exchange collisions in the wider plume with a linear extrapolation of the current density to perfect vacuum. As with the DSLR photos, it is clear that the focus is increasing with each step. As with the V1 test, the divergence reaches a minima, then becomes over-focused and begins to diverge.

Using the method described in Chapter 2, the momentum weighted divergence utilization efficiency is calculated from the Faraday probe data, see Fig. 3.15a. The 95% momentum

Table 3.2: Representative Thruster Performance at 0 mm on Xenon

Power	300 V, 1.33 A	
Peak Center-line Field	181 G	
Propellant Flow	1.5 mg/s	
Thrust	24 mN	+1 / -3 mN
Specific Impulse	1615 s	+50 / -200 s
Anode Efficiency	51%	+2 / -8%
Total Efficiency	40%	+1 / -5%

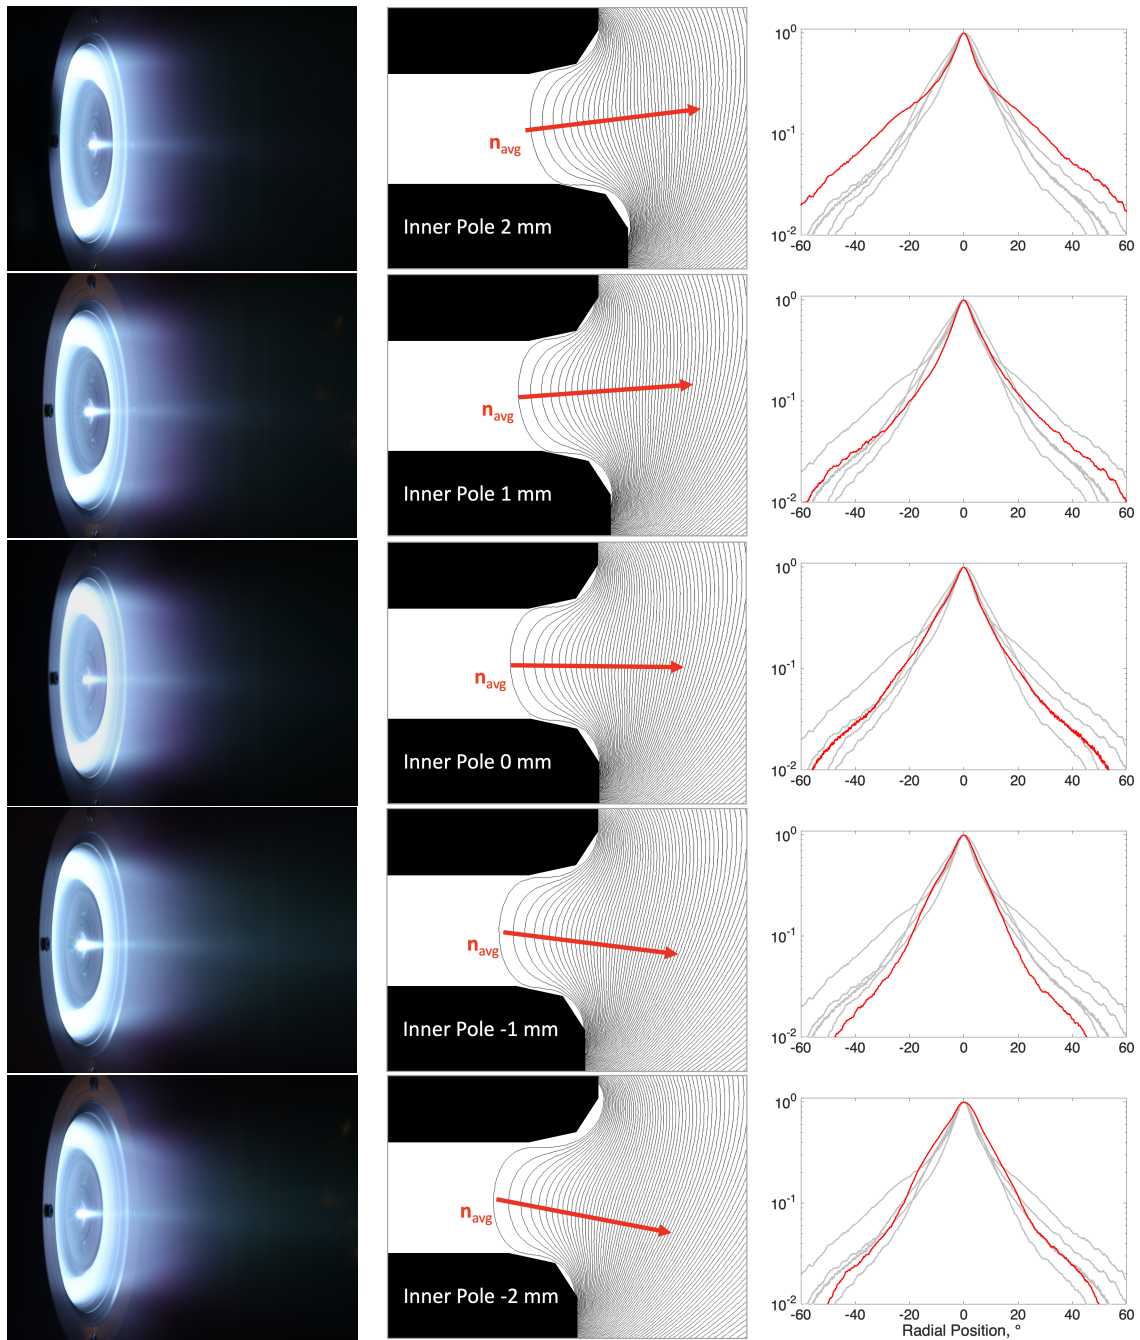


Figure 3.14: A comparison of the plume divergence across the range of pole positions tested. a) DSLR photos of the plume taken at 105mm f/16 1/10" ISO200. b) Field simulations with the average normal vector to the field c) Normalized swept Faraday probe measurements comparing current density at each position.

weighted ion current ranges from a 45.7° half angle for the -1 mm position to 65.0° for +2 mm. This corresponds to a peak divergence utilization efficiency of $0.715 + 0.016 / - 0.009$ at the -1 mm position. The uncertainty shown in Fig. 3.15b is dominated by the uncertainty due to the background charge exchange correction. The current utilization efficiency in Fig. 3.15b shows a similar trend, with a maxima of 0.63 ± 0.01 at 1 mm and minimum of 0.49 ± 0.01 at -2 mm. Although the divergence utilization efficiency was not calculated for V1, comparison of the plume traces show a significant improvement to the plume divergence in V2.

The voltage utilization efficiency was calculated using the method described in chapter 2. The peak of the IEDF, shown in Fig. 3.16, is highest at the -2 mm pole position, with a plasma potential corrected peak of 287 V, and decreases significantly with more positive pole positions.

Figure 3.17a shows the recorded thrust at each pole position. Measurements for -2 mm, and 0 to 2 mm fit show a trend similar to the current utilization efficiency with a maxima at 0 mm. However the thrust measured at -1 mm, the position with the highest divergence utilization efficiency, is the lowest thrust recorded. Figure 3.17b shows the the anode efficiency calculated from the thrust measurement, and from the combined efficiencies. The combined efficiency are calculated as a function of the current, voltage, and divergence utilization




Figure 3.15: a) The momentum weighted divergence utilization efficiency of the thruster at each inner pole position. b) The current utilization efficiency at each position.

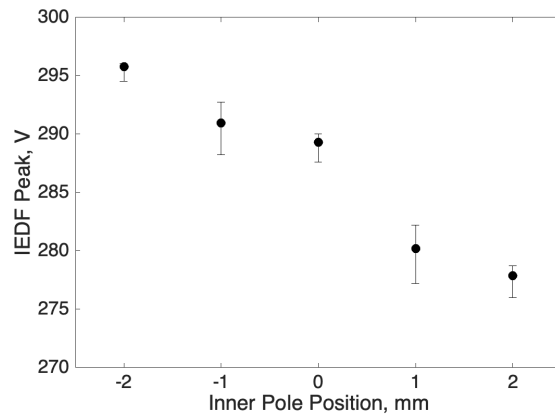


Figure 3.16: The maxima of the Gaussian fit of the IEDF at each of the pole positions tested.

compared to a normalized value at the 0 mm pole position. The normalized value is used as no $E \times B$ probe was available, and therefore the charge and mass utilization efficiencies cannot be calculated. We see agreement in the relative values at the -2 mm pole position, and significant deviation at the -1, 1 and 2 mm.

While all measurements were taken in the Jet mode, some operating points were closer to




Figure 3.17: a) The thrust measured at each of the pole positions. b) The relative anode efficiency at each pole position comparing the total influence of the efficiencies calculated. The points are normalized to the 0 mm pole position measurement. The reference anode efficiency at 0 mm is 51%.




Figure 3.18: Frequency spectra at each pole position showing the change in breathing mode frequency and amplitude.

mode transitions than others, and thermal mode transitions were observed. Figure 4.6 shows a fast Fourier transform (FFT) of the discharge current oscillations. From the 0 to +2 mm positions, there is a clear breathing mode frequency which ranges from 19.5 kHz at +2 mm to 25 kHz at 0 mm. The -2 mm position shows a small second peak, while at the -1 mm pole position, there are multiple distinct harmonics in the spectrum and the amplitude is 13.5 dB higher than at -2 mm and > 15 dB higher than the other three positions. Comparing these to FFTs of the Jet and Diffuse mode, see Fig. 3.19, we see that the the -1 mm peak is close to the mode transition to Diffuse mode. This could be due to the field shape moving the mode transition closer to the operating point, or due to the previously noted thermal dependence due to higher thermal loading on the inner pole at negative pole positions.

To relate this to the thermal dependence, a series of current measurements were taken at increasing temperature in the -2 mm position. Figure 3.20 shows the FFTs of the current signals. From a) to c) there is a clear increase in the amplitude of the breathing mode as




Figure 3.19: Frequency spectra of the two distinct Hall thruster operating modes. Measured at the -1 mm position with otherwise identical operating parameters.

the temperature increases, with multiple frequency peaks becoming prominent. After the mode transition to the Diffuse mode, the multiple peaks become more distinct with a 29.3 dB increase in amplitude. Asynchronous measurements of by the Faraday probe sweep and thrust likely resulted in differences at the -1 mm pole position, as the thrust measurement was taken after the Faraday probe sweeps were completed, resulting in a higher thruster operating temperature, and a thermal mode shift.

The difference in the +1 and +2 mm pole positions is less easily explained. While this is




Figure 3.20: Frequency spectra showing the increasing breathing mode amplitude as the temperature increases from a) to c). Data taken at the -2 mm position.

only a relative comparison, the differences in these points is significant enough that it merits further investigation and will be the focus of chapter 4 where the plume is fully characterized.

3.2.4 ACME Update to Version 2W

The next incremental update for ACME is to (V2W), with the addition of watercooling and a new keeper design. Figure 3.21 shows the new water-cooled block and heat sink before assembly. The concept for cooling the inner pole remains the same; heat is conducted through the Hiperco and stainless steel to the rear of the thruster. However, the previous heat sink has been replaced with a water-cooled 101 copper block with three channels flowing deionized water at 20°C. The block retains the external heat sink to separate the rear connections from plasma, and adds a copper heat shield between the keeper and the body of the thruster to reduce the heating effects of the cathode on the inner pole.

A hollow cathode redesign extends keeper assembly, allowing for a greater range of axial cathode locations. The new keeper electrode extends the cathode 10 mm, 5 mm of which compensates for the watercooling block. Additional alumina insulators in 1 mm increments can be added to position the keeper orifice exit +5 to -5 mm from the exit plane of the




Figure 3.21: a) The water-cooled block and new mounting system being test fit on the thrust stand. b) the new parts of the ACME v2.1 setup showing the block with cylindrical cathode heat shield, large heat sink, and modified rear casing of the thruster.

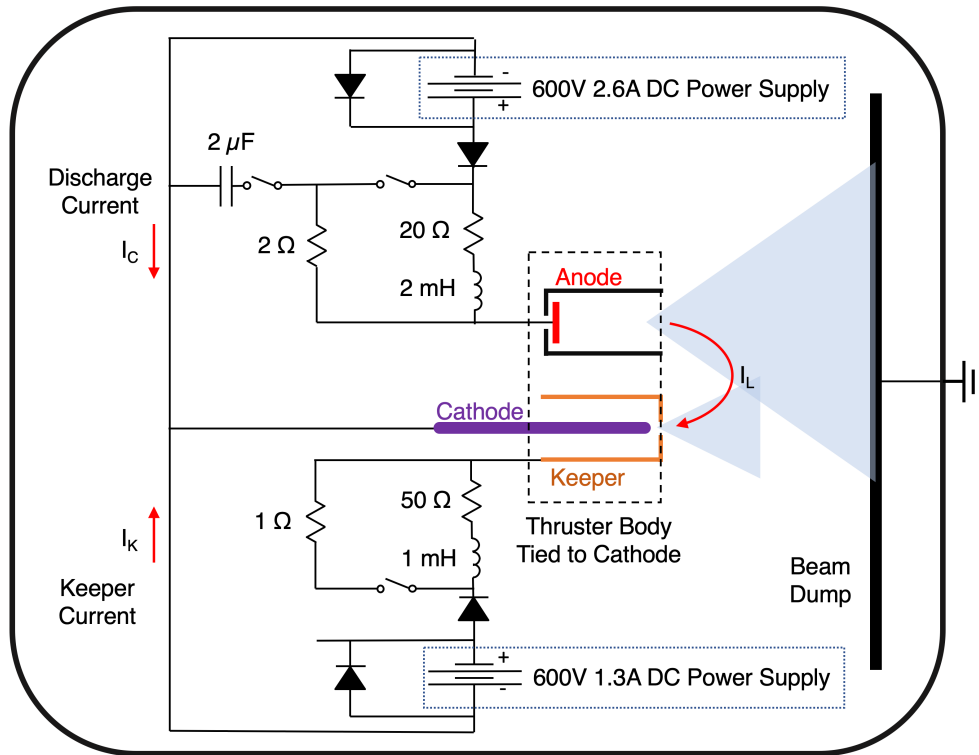


Figure 3.22: The cathode-tied electrical configuration with hybrid filter used for all tests with ACME V2W and V3.

thruster. The aim of this update is to isolate the effects of the inner pole movement and cathode positioning. This is particularly important for studying the cathode coupling across pole positions and attempting to decouple it from the voltage utilization efficiency. A study of this effect is not part of the work presented in this dissertation, but the upgrade enables increased capability in the thruster for future studies.

The final update is a change in the electrical configuration to a cathode tied setup. In this configuration, the thruster is electrically isolated from the facility, and the thruster body is tied to the cathode potential. Additional electrons that are collected by the thruster body can be emitted by the cathode. This configuration most accurately models in space operation of the thruster system [87]. Version 2W is used for the test campaigns described in chapters 3 and 4.

3.2.5 ACME Update to Version 3

The initial ACME anode is the most significant shortcoming of the thruster. To address this, a final update to ACME, version 3 (V3), tests two new anode designs with new BN walls. Improvements to anode design for more even gas distribution have shown significant improvements to measured anode efficiency on comparable thrusters [34]. Two new anode designs were manufactured in 316 stainless steel using direct metal laser sintering. Cross sections of the design, and the printed anode can be seen in Fig. 3.23. In both anodes, the gas flow moves between stacked and concentric mixing chambers with a series of holes that double in number for each chamber, resulting in 64 holes at the exit. Characterization of the designs showed a significant improvement to gas uniformity, with a peak-to-peak variation in neutral density of 0.7 and 1.3% for the A3 and A4 anodes respectively, compared to the original variation of 3.3% [88]. Anode A4 was used for all ACME V3 test campaigns. The new BN walls do not change the shape of the channel, but were updated to with a tighter interface with the pole caps.

3.2.6 Chapter Conclusions

This chapter detailed the design and initial test campaigns of the ACME thruster. The adaptive design has allowed it to be reconfigured to a wide range of experimental setups, with observed plume shapes ranging from highly divergent to over-focused. Updates to the thruster from the initial design have: improved the magnetic shielding of the poles; improved the channel wall material and structure; updated the electrical configuration; added watercooling; increased the adaptability of the thruster with more variable cathode positions; and redesigned the anode to improve gas distribution.

Testing of the thruster showed performance was strongly influenced by the pole position, with significant changes to plume divergence, and current and voltage utilization. While the V2 test campaign did not fully characterize the thruster at each pole position, the diagnostics that were used suggest that the accepted methods for plume characterization may

not be sufficient for the more divergent plume seen at +2 mm. Therefore, the full plume characterization will be investigated further, with methods to more thoroughly measure divergent plumes presented in chapter 4.




Figure 3.23: a) A cross section of the A3 anode. b) A cross section of the A4 anode. c) The printed A3 anode. d) The printed A4 anode. Figures from [88].

Chapter 4

PLUME CHARACTERIZATION WITH SPATIALLY RESOLVED ION ENERGY DISTRIBUTION MEASUREMENTS

Electric propulsion systems exhibit a range of physics that are difficult to model. Therefore, understanding the underlying drivers of efficiency in the system, beyond measurements of pure performance, is critical.

The measurement of thruster performance falls into two distinct categories: direct measurement of the performance, and characterization of the thruster plume [2]. Directly measuring the thrust gives an accurate measurement of the anode and total efficiency, specific impulse (ISP), and thrust-to-power ratio. However, this gives limited insight into the drivers of the performance. Using a suite of diagnostics in the plume, thruster operation can be more thoroughly characterized, allowing the divergence, current, charge, mass, and voltage utilization efficiencies to be calculated [76]. The product of each of these efficiencies is the anode efficiency of the thruster, which generally shows good agreement with the anode efficiency calculated from the thrust stand measurement [74, 75, 86]. These diagnostics are well understood, and have long proven and accepted recommended practices, particularly for highly developed systems such as Hall and gridded ion thrusters [55, 79, 69, 89, 73, 90, 91, 92].

However, applying these methods to less developed thrusters, where the form of the plume is unknown or highly divergent, may not be sufficient to fully characterize operation. This is primarily due to the positioning of the $E \times B$ probe and Retarding Potential Analyzer (RPA) which are stationary and take measurements on or near the centerline of the thruster. A study of angularly-resolved $E \times B$ probe spectra indicates that the measurement of ion species fractions at the channel centerline is sufficient to describe the charge utilization efficiency of

the full plume [93]. However, the voltage utilization efficiency, which is calculated from the Ion Energy Distribution Function (IEDF) measured by a RPA, can exhibit a strong angular dependence [87, 71] which may lead to erroneous results on more divergent plumes.

The goal of this study is to investigate the efficacy of centerline and angularly resolved IEDFs in deconvolving the efficiencies of electric propulsion systems with varying plume structures. The relative efficacy of the methods can be used to guide plume characterization in divergent systems or novel thruster concepts. In this chapter we compare a full plume characterization, with both centerline and angularly resolved IEDFs, to thrust stand measurements of an adjustable Hall thruster with a range of plume shapes from over-focused to highly divergent.

4.1 Methodology

In order to compare the efficacy of centerline and angularly resolved IEDFs, an experimental setup was developed to characterize thruster operation using both plume diagnostics and direct thrust measurements across a range of plume shapes, ranging from highly divergent, to over-focused. The setup includes the Hall thruster, vacuum chamber, thrust stand, and suite of diagnostics required to fully characterize the plume. Each of these components are described in detail below where they differ from the standard methods presented in chapter 2.

4.1.1 Criteria for Comparison

Comparing the centerline and angularly resolved IEDFs requires the thruster to be fully characterized using both thrust measurements and plume diagnostics. Thrust measurements will be used as the ground truth for the comparison. The anode efficiency, η_a , used as the comparison metric is

$$\eta_a = \frac{T^2}{2\dot{m}_a P_a}, \quad (4.1)$$

where T is the thrust, \dot{m}_a is the anode mass flow rate, $P_a = V_d I_d$ is the power delivered to the anode, V_d is the discharge voltage, and I_d is the discharge current. The anode efficiency

represents how well the power and mass flow being delivered to the anode are being converted into thrust, giving a good assessment of the overall performance.

The anode efficiency calculated from the plume characterization is the product of five efficiencies,

$$\eta_a = \eta_d \eta_b \eta_v \eta_m \eta_q \quad (4.2)$$

where η_d is the momentum weighted divergence utilization efficiency, η_b is the current utilization efficiency, η_v is the voltage utilization efficiency, η_m is the mass utilization efficiency, and η_q is the charge utilization efficiency [2]. Here, the voltage utilization efficiency will be calculated using both the centerline and angularly resolved IEDFs, and the resulting anode efficiencies will be compared to the anode efficiency calculated from thrust measurements. These component efficiencies are calculated using an RPA, and Faraday, E×B, and Langmuir probes which will be described later in this section. They are measures of the physical processes that occur within the thruster and each one allows greater physical insight into how losses in efficiency occur.

4.1.2 The ACME Thruster, Vacuum Setup, and Standard Plume Diagnostics

The wide range of plume shapes required for this experiment are produced by ACME V2W. While the bounds of pole position are +5 to -35 mm, this experimental campaign reduces the range to a set of pole positions of -12, -8, -6, -4, -2, -1, 0, 1, 2, and 4 mm. Instability of the plasma discharge at the extreme setpoints prohibited full characterization of the +4 and -12 mm positions, reducing the full characterization to 8 pole positions. The thruster was run in a cathode tied configuration in the STF using the thrust stand described in chapter 2. The thruster was run at a single operating point with a 400 V, 1.5 A main discharge, a peak centerline magnetic field strength of 180 G, a fixed keeper current of 0.5 A, and cathode flow rate of 0.25 mg/s of krypton for all test points.

During operation of ACME, the pressure ranged from $1 - 2 \times 10^{-5}$ Torr across the flow rates tested. The thruster was mounted on the nested inverted-pendulum thrust stand on the

centerline of the chamber. The process of measuring thrust consisted of ignition, adjusting the thruster controller to the desired parameters, and waiting for thermal steady state. On reaching steady state, the thruster is shutdown and the thrust delta is measured. A single calibration mass is added and removed to confirm the calibration taken before testing.

The Langmuir and Faraday probes are mounted to a swept arm positioned at 0.54 m or 8 TDD, while the $E \times B$ probe is positioned 0.67 m or 10 TDD and offset from the centerline of the chamber to align with the thruster channel. The Faraday, $E \times B$, and Langmuir probes use recommended practices as described in chapter 2. A series of additional Langmuir probe measurements were taken between 0 and 90° in 15° increments to characterize the angularly dependent plasma potential needed to correct the RPA measurements.

4.1.3 Retarding Potential Analyzer

Calculating the voltage utilization efficiency, and comparing the efficacy of the center-line and angularly resolved IEDF methodologies, requires RPA measurements of the plume. The RPA is pointed directly at the center of the thruster with the axis of rotation of the swept arm passing through the exit plane.

The voltage utilization efficiency will be calculated in three ways to get the three voltage utilization efficiencies, $(\eta_{V1}, \eta_{V2}, \eta_{V3})$, used for the comparison.

The center-line IEDF will be used to calculate η_{V1} and η_{V2} and is measured with the standard method described in chapter 2. The first voltage utilization efficiency, η_{V1} , is calculated using the peak of a Gaussian fit to IEDF. This is defined as the most probable voltage, V_{mp} , and is assumed to represent the beam voltage across the entire plume. This gives η_{V1} , as

$$\eta_{V1} = \left(\frac{V_{mp}}{V_d} \right). \quad (4.3)$$

The second voltage utilization efficiency, η_{V2} , calculates the beam voltage by integrating

over the centerline IEDF, with

$$\eta_{V_2} = \frac{\int_0^{V_{max}} V_a J(V_a) dV}{V_d J}, \quad (4.4)$$

where $J(V_a)$ is the collected current at voltage V_a , and J is the total current collected in the sweep. The wide acceptance angle and noise sensitivity of RPAs can result in broadening of the IEDF distribution leading to underestimation of the voltage utilization efficiency when integrated [74]. This centerline integrated method is therefore included in the analysis to investigate the broadening and determine its effect on the full plume integration.

The angularly resolved IEDF is generated in the same way as the center-line IEDF, but is repeated at 136 positions equally spaced between 0° , at thruster centerline, and 90° , thus giving an angularly resolved IEDF. The angularly resolved IEDF is integrated across the hemispherical coordinate system to calculate the third voltage utilization efficiency, η_{V_3} , as

$$\eta_{V_3} = \frac{\int_0^{\frac{\pi}{2}} \left(\int_0^{V_{max}} V_a J(V_a, \theta) dV \right) \cos(\theta) \sin(\theta) d\theta}{\int_0^{\frac{\pi}{2}} V_d J(\theta) \cos(\theta) \sin(\theta) d\theta} \quad (4.5)$$

where the collected current $J(V_a, \theta)$ now has angular dependence. The integration uses the axial current, see Eq. 2.8, and therefore weights the influence of ions by their contribution to thrust. For example, an ion measured at 90° from centerline is counted in the total current integration, but does not provide axial thrust and is therefore not counted in the axial current. When the voltage utilization efficiency is calculated, only the fraction of current producing thrust, and the ion energy distribution of that population, are integrated.

4.2 Results and Discussion

The aim of this chapter is to assess methods for characterizing divergent plumes in electric propulsion systems with a specific focus on the efficacy of centerline and angularly resolved IEDFs in calculating the voltage utilization efficiency. To compare the three methods of cal-

culating the voltage utilization efficiency described in section 4.1.3, the ACME Hall thruster was run at a single operating point of 400 V, 1.5 A, at 10 pole positions from +4 mm to -12 mm, (+4, +2, +1, 0, -1, -2, -4, -6, -8, -12). Due to instability in operating the thruster at the extrema of pole position, the points at +4 and -12 mm could not be fully characterized. This section will initially analyze the plume characterization, comparing the efficacy of the three voltage utilization efficiencies in describing the plume. Then, methods for CEX correction will be assessed including a novel method utilizing the angularly dependent IEDFs. Finally, the performance of the thruster across the pole positions will be discussed.

4.2.1 Plume Characterization

At each pole position, a suite of diagnostics was used to calculate the divergence, current, charge, mass, and voltage utilization efficiencies. The anode efficiency calculated as a product of these efficiencies, Eq. 4.2, was then compared to the anode efficiency calculated from thrust stand measurements, Eq. 4.1. Figure 4.1 shows the breakdown of these efficiencies, along with the anode efficiency at each of the pole locations tested.

The divergence utilization efficiency, see Fig. 4.1a, showed a clear trend across the pole positions with a maxima at the -2 mm pole position. This agrees with visual observations of the plume and previous investigations of ACME [94, 95], that showed significant changes in plume shape with changing pole position. The plume at the +4 mm position is highly divergent, becoming more focused as the pole position moves negatively. Beyond the -4 mm pole position, the plume narrows significantly and becomes over-focused.

The current utilization efficiency in Fig. 4.1b has a clear peak at the -4 mm pole position, decreasing in both directions. The primary hypothesis for this result is the change to cathode coupling at the different pole positions due to changes in field shape and relative cathode-channel position. Investigations of cathode placement have observed significant changes to current utilization with changing cathode position [96]. Another study of cathode coupling found that a thruster with an extended outer pole, similar to the -2 mm pole position, changed the magnetic field structure between the channel and the cathode, improving coupling [97].

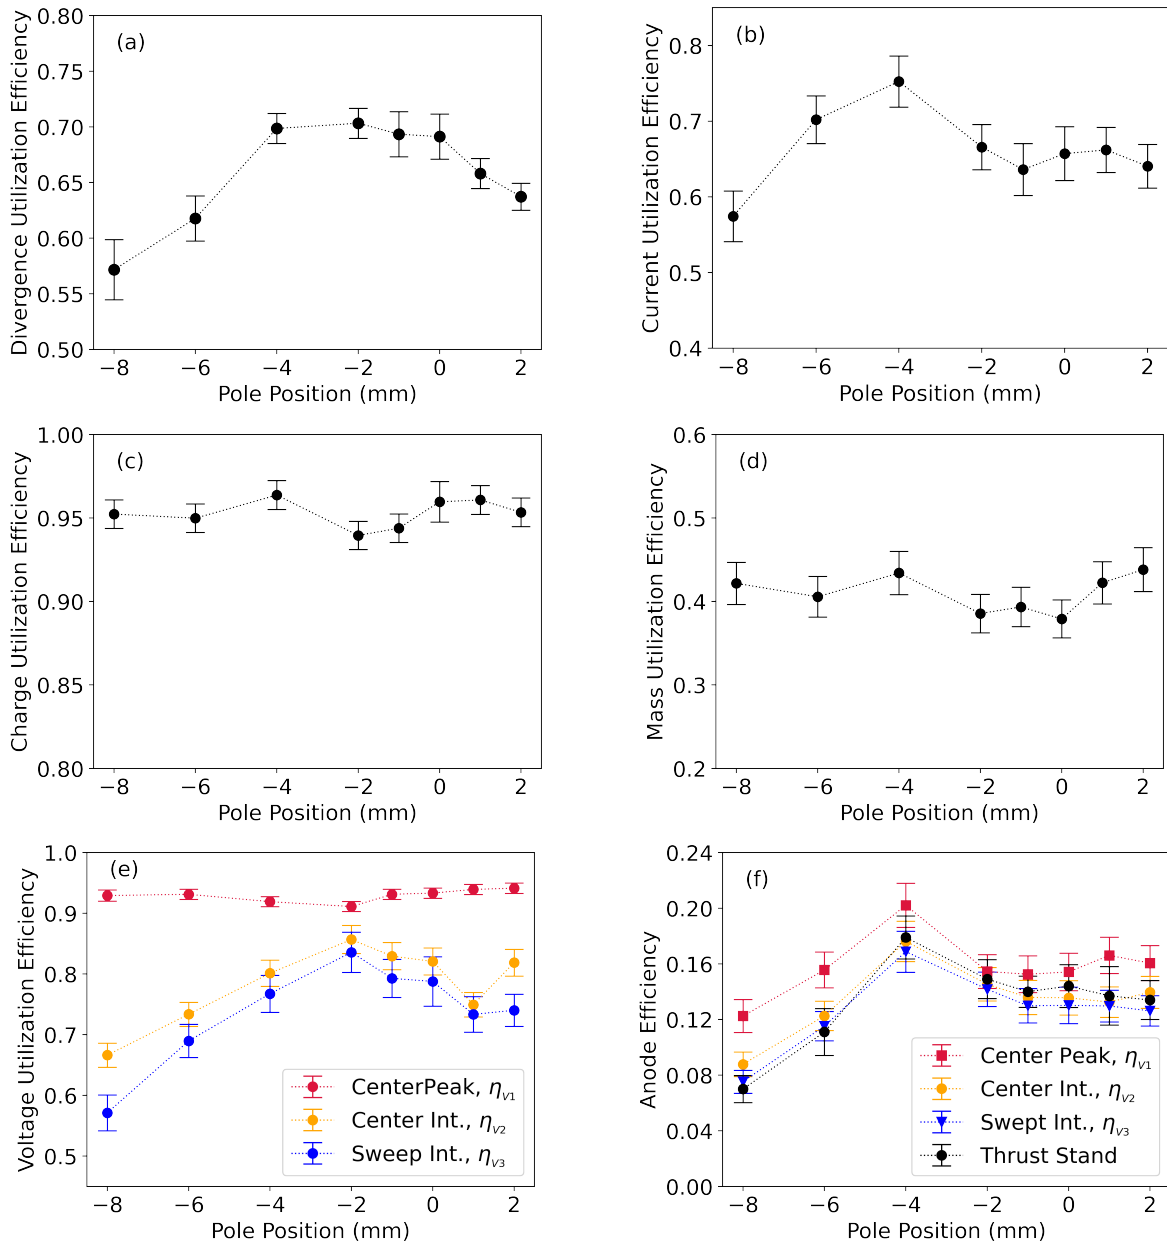


Figure 4.1: The breakdown of the measured efficiencies across the range of pole positions showing: a) The momentum weighted divergence utilization efficiency, b) the current utilization efficiency, c) the charge utilization efficiency, d) the mass utilization efficiency, e) the three methods used for the voltage utilization efficiency, and f) The anode efficiency comparison between the product of the probe based efficiencies and the direct measurement from the thrust stand.

Further evidence of this hypothesis is discussed in section 4.2.3.

The charge utilization efficiency, see Fig. 4.1c and mass utilization efficiency, Fig. 4.1d, both remain relatively consistent across the pole positions, with a slight decrease around the -1 mm position. The symmetry around the -1 or -2 mm pole position is evident across each of the component efficiencies, and is likely due to the field stretching symmetrically as the poles are positioned either side of the optimum.

The voltage utilization, see Fig. 4.1e, shows significant variation between the methods. The centerline peak value, η_{V1} , remains relatively consistent across the range of pole positions, indicating that a similar peak in the IEDF on centerline exists at every pole position. The other voltage utilization efficiencies, η_{V1} and η_{V2} , are closest to η_{V1} at the -2 mm pole position where the plume is most focused, but diverge significantly in the other pole positions. The effect of integrating the centerline for η_{V2} appears similar to the integration of the full plume in η_{V3} , with the values only significantly diverging at the extremes of pole position.

Figure 4.1f shows a comparison of the anode efficiencies calculated using the combined efficiencies with $(\eta_{V1}, \eta_{V2}, \eta_{V3})$, and from thrust stand measurements. At the -2 mm pole position, the three methods show closest agreement among themselves and with the thrust stand. Notably, this pole position also corresponds to the highest divergence utilization efficiency. The reason for good agreement among all the cases is due to a significant, high energy, ion population on the centerline of the plume. As we move to more divergent plumes, the anode efficiency from η_{V2} , and η_{V3} match the thrust stand data within the uncertainty of each measurement. While the centerline integrated measurement of η_{V2} provides an improved match compared to η_{V1} , it only provides additional information on center-line and is therefore of minimal additional utility. The agreement between thrust stand and η_{V3} based anode efficiencies across the full range of pole positions indicates that the angularly resolved IEDF is beneficial in the study of divergent plumes. The method could be further improved with the use of a collimated RPA to minimize the broadening effect [98, 71].

To visualize the plume shapes, Fig. 4.2 shows angularly resolved IEDFs in conjunction with plume photos for a subset of the pole positions tested. The plume images show a




Figure 4.2: A comparison of the plume divergence across a subset of pole positions tested. Peak first derivative of current at each pole position is: a) $0.14 \mu A/V$ at +4 mm, b) $1.3 \mu A/V$ at 0 mm, c) $0.58 \mu A/V$ at -4 mm, and d) $0.09 \mu A/V$ at -8 mm. DSLR photos and contour plots of the swept IEDF measurements are shown for each position. Photos of the plume taken at 70 mm f/22 1/5" ISO200. Note that the inner pole is consistently positioned with each axis while the outer pole moves.

clear trend from a highly divergent plume at +4 mm to an over-focused plume at -8 mm. The angularly resolved IEDF in Fig. 4.2b shows a single well focused peak at 0 mm. At this pole position, the field is flat across the exit plane of the thruster channel and the majority of ions are accelerated downstream in a relatively tight plume. At the +4 mm pole position, the peaks in the IEDF are off center-line due to the inclination of the magnetic field resulting in a significant radial component to the electric field in the acceleration region. The ions are therefore accelerated at a significant angle, resulting in η_{V1} and η_{V2} missing crucial information in the off axis plume. Similarly, as the pole position moves to -4 and -8 mm, the magnetic field is angled towards the center-line, and the plume becomes over-focused with ions passing through the centerline and continuing to diverge before reaching the RPA at 8 TDD.

At both the +4 and -8 mm pole positions, several populations of ions become visible. Significant tails of lower energy ions are present on and near the centerline, which may be due to the stretched field shape causing a wider range of ionization locations, or elastic collisions occurring outside the thruster [99, 71]. Examining the magnetic field lines for these pole positions we can see that at the extreme values, the curvature around the poles begins to extend and the assumption of consistent shielding on the poles begins to break down, see Fig. 4.3. The asymmetry of the grazing line and increased curvature on the poles significantly pushes the potential drop to one side of the channel. The increased curvature and skew in the field shape may be the driver of the secondary populations seen in the angularly resolved IEDF. As the magnetic field lines are isopotential lines, this increased curvature creates electric fields that are at large angles from the thruster center-line. Neutrals near the outer wall in Fig. 4.3c may reach beyond the channel exit while remaining in the high potential region. When they are ionized, they could be accelerated at a wide range of angles including perpendicularly to the center-line.




Figure 4.3: a) The field at the +4 pole position, b) the field at the -4 mm pole position, and c) the field at -8 mm showing the stretch in shielding on the outer pole. Note that the field lines are rendered to the same arbitrary limit to demonstrate the change.

4.2.2 Charge Exchange Ion Characterization

The characterization of facility effects is a key factor in understanding the performance of EP systems in vacuum facilities [56, 100]. Here, the angularly dependent IEDF can be further utilized to characterize CEX ions and provide corrections for facility effects. As the plumes become more divergent, CEX ions become significantly more prevalent as the high energy ions diverge over a larger area, scaling with $\sin(\theta)$, and therefore interacting with a greater number of neutrals. This has been previously demonstrated while characterizing the plume of a BPT-4000 Hall thruster [71]. The CEX population is visible in the RPA data as a distinct low energy population present across the whole sweep, particularly for divergent plumes. Figure 4.4a shows a section of the angularly resolved IEDF at the +4 mm pole position, scaled to emphasize the lower current density areas.

As the RPA is able to measure the CEX populations across the plume, a method to calculate the current density distribution was developed. Figure 4.4b shows a comparison of Faraday probe and RPA sweeps for a 400 V, 1.5 A discharge at the -2 mm pole position. A scalar correction factor was applied to RPA current data to account for the grid conductance and collector area differences of the RPA and Faraday probe. This correction was calculated from the -2 mm pole position to match the centerline current density and then applied to all RPA sweeps. The Faraday probe data was corrected for CEX ions using




Figure 4.4: a) A sample of the swept RPA from a 400 V, 1.5 A discharge at the +4 mm pole position. CEX dominated signal is seen below the red line. b) A comparison of the Faraday data, RPA current, and CEX correction for both signals.

both the extrapolated wing and pressure characterization methods described in section 2.5. The extended-wing method significantly underestimated the current in the wider plume, a problem exacerbated by more divergent plumes.

The RPA data was corrected for CEX by removing all current measured below a threshold voltage which is assumed to be dominated by CEX ions [101]. The threshold voltage was iteratively increased until the RPA data most closely matched the pressure corrected Faraday probe trace. Using a threshold voltage of 40 V showed the closest match to the Faraday correction, see Fig. 4.4b. This value also matches the distinct low energy population seen below the line showing the 40 V threshold in Fig. 4.4a. Correcting for CEX within the RPA sweep allows for CEX ions to be removed from the voltage utilization efficiency no matter the plume shape, removing a significant source of uncertainty.

4.2.3 Influence of Pole Position on Thruster Performance

Thruster performance changed significantly across the range of pole positions tested. Thrust measurements, see Fig. 4.5a, showed a maxima at -4 mm with a significant drop off at more negative pole positions. The plume characterizations indicates that the peak at -4 mm can




Figure 4.5: a) Thrust measurements taken at each pole position at 400 V 1.5 A. b) The required flow rate through the anode to achieve a current of 1.5 A at thermal steady state.

be attributed to the peak in current utilization at this pole position, see Fig. 4.1a, combined with the high divergence utilization efficiency shown in Fig. 4.1b. A similar investigation of unshielded Hall thruster performance compared two thruster configurations, one with a flat exit plane, and one with the outer pole extended 2 mm. The extended outer pole, equivalent to a -2 mm pole position, showed significant performance improvements with an increase in total efficiency of 10% [102]. This peak in thrust and anode efficiency, and the agreement with a similar comparison in literature, suggests that negative pole positions could be advantageous in magnetically shielded thrusters and may warrant continued investigation.

Figure 4.5b shows the anode flow required to achieve the 1.5 A discharge, which significantly increases with positive pole positions. As pumping speed is consistent across tests, the >30% change in flow required between the extreme pole positions resulted in higher background pressures, and consequently, higher CEX populations were observed in the positive pole positions.

Measurements of the discharge oscillations, see Fig. 4.6 show increased breathing mode amplitude and an increase in the number of harmonics at more positive pole positions. There are at least two likely drivers of this shift. Firstly, analysis of facility effects has previously




Figure 4.6: A FFT of the discharge current at each pole position. Scale on each axis is identical.

demonstrated that higher facility pressures result in changes to breathing mode frequency and amplitude which may be a factor in this trend [79]. Secondly, increased breathing mode amplitude due poor cathode coupling has previously observed in Hall thrusters by varying the cathode flow [103]. As the cathode position moves with inner pole, and more positive pole positions require a higher inner magnet current to balance the fields. As a result there is a significantly higher field between the cathode and channel in the +4 mm position, which when combined with the lower plasma density due to the divergent plume, likely results in poorer cathode coupling [102, 104]. With more negative pole positions, the cathode is positioned in a region of higher plasma density. The combination of a higher plasma density and lower magnetic fields on the inner pole improve the cathode coupling, decreasing the required mass flow rate.

4.3 Chapter Conclusion

The aim of this chapter is to investigate the efficacy of centerline and angularly resolved IEDFs in describing the operation of electric propulsion systems with varying plume structures. To assess these methods, the ACME Hall thruster was operated at a range of pole positions ranging from highly divergent to over-focused. At each of these pole positions the performance of the thruster was characterized with both plume diagnostics and a thrust stand. The standard method of taking the peak of the IEDF on the centerline of the plume was sufficient to describe the voltage utilization while the plume was well focused. However, as the plume becomes more divergent at the +1 and +2 mm pole positions, and becomes over-focused at the -4, -6, and -8 mm positions, the agreement with the thrust stand is lost. At these positions, only the anode efficiency calculated using the integrated IEDF matches the thrust stand data within the uncertainty of each measurement. Integrating the angularly resolved IEDF is therefore clearly beneficial in the study of divergent plumes or novel thruster concepts where the structure of the plume is unknown. Improving the method with the use of a collimated RPA is suggested as it would minimize the broadening effect and decrease the underestimation of the voltage utilization efficiency leading to a more accurate characterization the plume [98, 71].

We anticipate the proposed method for calculating voltage utilization efficiency will be particularly important for thrusters that utilize alternative propellants due to increased plume divergence from shifts in the ionization and acceleration regions of lightweight or multi-species propellants [105, 106, 107, 108, 109, 110, 111]. These propellants also require increased flow and may not be pumped as effectively, and therefore may benefit from the improved CEX correction over the full plume, without having to take measurements at multiple background pressures.

The operation of ACME changed significantly over the range of pole positions tested. Angularly resolved IEDFs showed more complex plume structures at the extrema of pole position with multiple distinct ion populations. It is hypothesized that these may be due

to the increased curvature of the magnetic shielding at the poles, causing multiple distinct populations to be ionized and accelerated. Peak efficiency was not found at 0 mm in the most standard 'flat' field configuration seen in most Hall thrusters, instead the peak was at -4 mm, driven by improved current utilization efficiency and good divergence utilization efficiency. Continued testing of ACME beyond the scope of this test has shown vast improvement to the performance at higher current densities, increased channel depth due to anode positioning, and better gas distribution due to printed stainless anodes. While these results only represent a single operating point at each pole position, trends in the performance could be useful for future developments. The significant decrease in flow required at negative pole positions is of particular interest for future investigations. It is likely due to the cathode positioned closer to the higher density plasma combined with the lower magnetic fields required on the inner pole improving cathode coupling. Combined with the improvements to the divergence and current utilization efficiency, there may be gains that can be achieved with other Hall thrusters with centrally mounted cathodes.

The investigation of plume characterization methods has been submitted for publication in the Journal of Propulsion and Power.

Chapter 5

RAPID THRUSTER-IN-THE-LOOP OPTIMIZATION FOR HALL THRUSTERS

The characterization and optimization of Hall thrusters has traditionally been conducted over long test campaigns. A significant component of that is the performance and stability mapping across a range of current-voltage-magnetic field (IVB) configurations [47, 48, 49]. IVB maps are useful for characterizing the global performance and more thoroughly analyzing changes in design [50]. As discussed in chapter 1, this is a significant limitation as the mapping performance across the many dimensions of a thruster is costly in propellant and facility use, and development time. Decreasing costs for access-to-space and the rise in small satellite constellations has driven the need for a greater number of systems that are cheaper to design, test, and manufacture. Methods for rapid and efficient thruster development are therefore of interest.

Previous applications of computationally driven development processes including machine learning and algorithmic optimization have been outlined in chapter 1. They have proved to be well suited to the challenges of electric propulsion research and development. However, a method to expedite the development of high Technology Readiness Level (TRL) thrusters, decreasing test time and associated costs for new systems, is missing.

In this chapter we present results that describe an active control system and methodology that is capable of rapidly mapping and optimizing the performance of Hall thrusters. A rapid thrust measurement method is developed that significantly reduces the time required to characterize a thruster setpoint. Full end-to-end computer control and real time thruster-in-the-loop optimization is applied to the ACME Hall thruster [95]. Nelder-Mead and Powell optimization schemes are assessed for their efficacy in rapidly optimizing ACME, a medium

Table 5.1: ACME thruster dimensions. Fast dimensions, left, can be adjusted during operation. Slow dimensions, right, require removing the thruster from vacuum to change.

Fast Dimension	Minimum	Maximum	Slow Dimension	Minimum	Maximum
Discharge Voltage	100 V	400 V	Pole Position	-12 mm	+4 mm
Field Strength	140 G	220 G	Anode Position	10 mm	35 mm
Keeper Current	0 A	1.5 A	Cathode Position	-2 mm	+2 mm
Cathode Flow Fraction	10%	40%	Anode Design	3 Available	
Field Skew	-0.2	0.2			

(5-20) dimension thruster system.

5.1 Methodology

The following section describes the ACME Hall thruster, vacuum facilities, thrust stand, and optimization methods used in this research.

5.1.1 ACME Hall Thruster

ACME V2W, with a cathode-tied electrical configuration was used during this test campaign. ACME has 10 dimensions, the thruster parameters that the optimizer can adjust. These thruster parameters can be categorized into the ‘fast’ dimensions which are computer controlled in real time, and the ‘slow’ dimensions which require mechanical changes to the thruster. The fast dimensions are: applied anode voltage, anode mass flow, cathode mass flow, keeper current or voltage, inner magnet current, and outer magnet current. For optimizing the thruster at a constant power, these dimensions are reduced to the five fast dimensions listed in Table 6.3. The minimum and maximum values listed are the bounds of the optimization search space.

To maintain a balanced magnetic field, the control parameters of the inner and outer magnet currents are replaced with “Field Strength” and “Field Skew”, a measure of how far the magnetic field is moved out of balance from its evenly shielded shape. Changing the field skew changes the current in the inner and outer magnets to compensate for differences

between the true field and simulations, along with changes due to field saturation at high operating temperatures. A field skew of 0.1 would increase the inner magnet current by 10% and decrease the outer current by the same fraction. Mapping the control parameters to the input parameters of the magnet currents is done through 2nd order polynomial fits to data from Finite Element Method Magnetics (FEMM) [112] simulations of the magnetic field. These fits, confirmed with a grid of axial and radial Hall probe measurements, act as input functions to the physical control parameters translating strength and skew into inner and outer currents.

Mode shifts between two distinct thruster modes, Jet and Diffuse, have been observed during operation of ACME. Mode shifts are changes between two regions of operation where small changes to the thruster operating point cause large changes to thruster operation. These include: the anode gas flow required to maintain the discharge current, thruster oscillation frequency and magnitude, stability, plume shape, and efficiency. In ACME, the Diffuse mode is prevalent at lower operating voltages. It has a lower breathing mode frequency compared to Jet mode, but with a far greater amplitude and well defined harmonics. The efficiency of the system can change significantly either side of a mode shift, however, it is not always the case that one is better. At 500 W on xenon, the Jet mode has better anode efficiency with a peak of 50% around 300 V, while on krypton, the Diffuse mode has better anode efficiency with a peak of 30% under 200 V. Regions near the mode shift can flicker between the two modes, presenting a challenge to the automation of the system, particularly the control of the anode gas flow required to maintain the discharge current.

Optimization over the slow dimensions of the ACME thruster is beyond the scope of this chapter, optimization methods using the slow dimensions will be discussed in Section 5.2.5. Throughout the testing detailed in this chapter, a single configuration of slow dimensions was used with a pole position of 0 mm, anode position of 25 mm from the exit plane, and cathode 2 mm beyond the center pole. The configuration had not been used before and was chosen to simulate working on a new thruster design.

5.1.2 Vacuum Chamber and Thrust Stand

All testing of the control and optimization system on ACME was completed in STF using the nested inverted pendulum thrust stand. During operation of ACME, the pressure ranged from $1 - 5 \times 10^{-5}$ Torr across the flow rates tested.

The process of measuring thrust normally consists of ignition, adjusting the thruster controller to the desired parameters, and waiting for thermal steady state. On reaching steady state, the thruster is shutdown and the thrust delta is measured. A single calibration mass is added and removed to confirm the calibration taken before testing. A plot of the thrust signal during shutdown and calibration is shown in Fig. 5.1a. Taking another measurement requires significant time to reignite the thruster and return to thermal steady state due to the system cooling while the thruster is off. At 250 V, 2.0 A, ACME requires > 15 minutes to achieve quasi-steady state after a full shutdown measurement and calibration.

To improve the measurement rate, a novel rapid thrust measurement methodology was developed to avoid thruster shutdown and minimize thermal drift. The method is shown in Fig. 5.1b. Thrust measurements of “Test Points” are taken while the thruster is running and compared to a well characterized “control point”. The thrust delta between the test

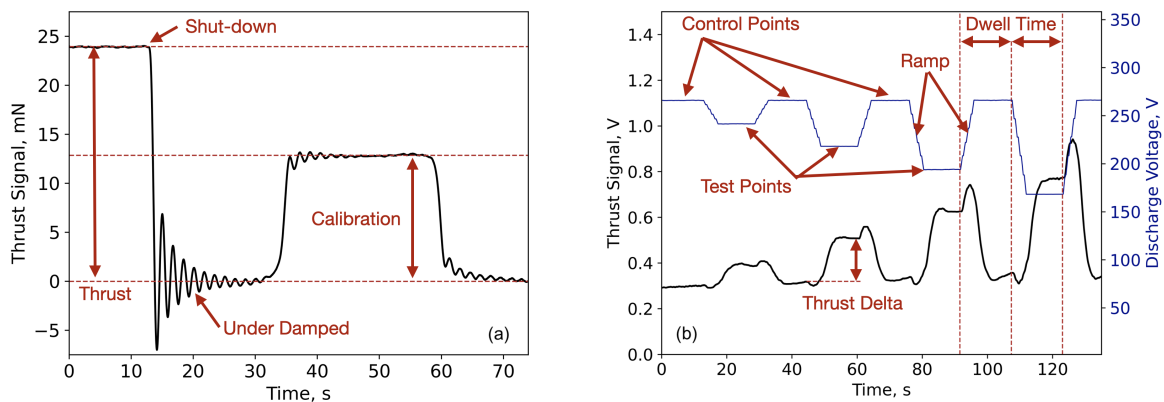


Figure 5.1: a) A standard shutdown thrust measurement with a single calibration step. Note the under-damped oscillation before tuning the damping system. b) Rapid thrust measurement with the Control Point - Test Point methodology.

and control points is then used to calculate the test point thrust, minimizing the error due to long term thermal drift. While there is still thermal drift, the error introduced to the measurement from thermal drift is restricted to drift occurring during the dwell time, t_d , at the test and control points. With this rapid measurement methodology, the new limiting factors to the measurement frequency become a combination of the anode thermalization rate, thrust stand oscillation damping time, and the PID control of the anode gas flow as it reaches quasi-steady state. Tuning the thrust stand flexures and dampers to achieve critical damping significantly reduces the settling time to well below the anode thermalization time. Investigations into achievable measurement frequency of the rapid methodology are presented in Section 5.2.1.

5.1.3 Control System

Thruster-in-the-loop optimization requires full computer control of the power supplies, flow controllers, and ignition ballast switching, along with all data acquisition. This is handled by a custom thruster controller, developed by EOI, which uses a Raspberry Pi computer running Python code to control Pi-Plate interface modules. These modules then connect to and control the thruster electrical supplies and can measure analog telemetry. Digital signals from the Pi-Plate interfaces control high-voltage, high-speed switching through fiber optics for isolation. The propellant mass flow controllers are controlled via USB connection to the Raspberry Pi. Only low-speed telemetry is obtained via the PI-Plate DAQ boards. For high-speed data collection, oscilloscopes connected via an ethernet LAN are used. The control algorithm can use data from any of the above sources to control thruster operation, including direct prompting of human control if necessary. This approach gives a portable system that can be used in any lab to optimize and map thruster performance.

Testing with the control system follows the flow chart shown in Fig. 6.1. Before testing, the thruster is warmed up to thermal steady state at the control point. Steady state is achieved when both the thruster temperature and propellant flow required to achieve the discharge current are constant. Once the thruster is ready, the optimizer is initialized and

generates the first test point. The thrust is measured using the control/test point method, and all control inputs, telemetry, and additional diagnostics are recorded. Using feedback from the system, the objective function is evaluated at that point and recorded. The algorithm then checks for convergence to the optimum. If the convergence criterion is met, the system is returned to the control point where it awaits user input. If it is not reached, the algorithm proceeds to the next test point. In the event of thruster shutdown, the control system detects the drop in anode current, restarts the thruster, and achieves steady state before resuming the test.

5.1.4 Optimization

The characteristics of a search space determine which class of optimization scheme is appropriate. ACME has one discrete and nine continuous dimensions making it a medium-dimension system. The optimization objectives for this test, total efficiency and specific impulse (I_{sp}), do not have a known functional form as no existing data is used for each optimization, and the convexity of the search space is unknown. To efficiently and rapidly determine these optima, the class of methods appropriate are derivative-free approaches as they do not require knowledge of the local search space. Among the suitable options given these considerations are Nelder-Mead [113] and Powell's [114] methods which are described more thoroughly below. These methods cannot deal with large measurement noise, and although there are uncertainties in the thruster measurements, these are small enough to not cause significant issues.

Optimization Objectives

The objective functions are the mathematically defined goals given to the optimizer. They can be any quantifiable measure of the thruster. In this chapter, the two optimization goals are: total efficiency and specific impulse. As the optimization scheme acts to minimize the objective function, the inverse of these values are used.

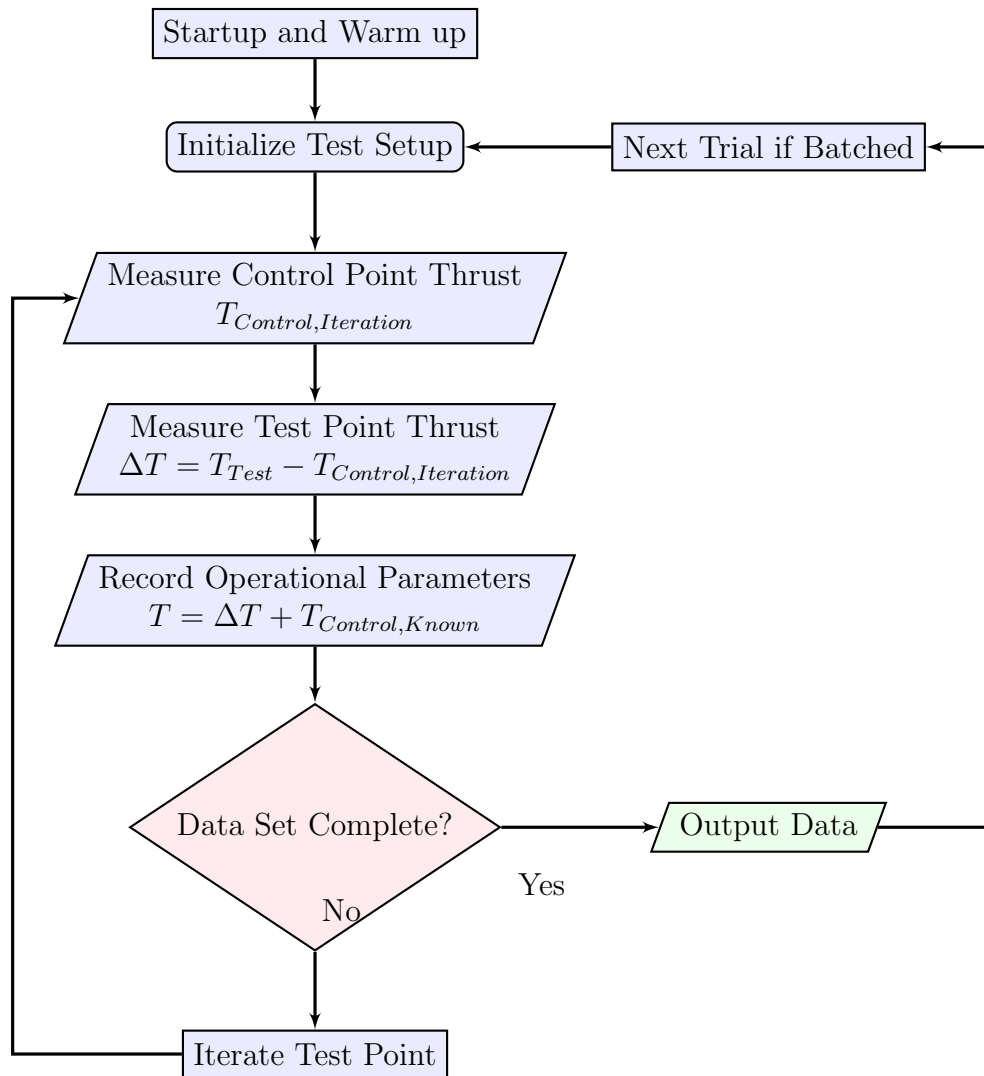


Figure 5.2: A flow chart showing the operation of the thruster using the Control Point - Test Point rapid thrust measurement technique.

The objective function for determining the best total efficiency, η_T , is defined as:

$$F_\eta(\mathbf{x}) = \frac{1}{\eta_T} = \frac{2\dot{m}_T P_T}{T^2} \quad (5.1)$$

where \mathbf{x} are the set of input parameters to the system, P_T is the total thruster power including

cathode and magnets, \dot{m}_T total mass flow rate from the anode and cathode, and T is the thrust. Similarly, the objective function for maximizing I_{sp} is defined as:

$$F_{I_{sp}}(\mathbf{x}) = \frac{1}{I_{sp}} = \frac{\dot{m}_T g_0}{T} \quad (5.2)$$

where the specific impulse, I_{sp} , is calculated from the thrust T , total mass flow rate \dot{m}_T , and standard gravity g_0 .

These functions are calculated from thrust measurements and the automated control system inputs. The bounds of the search space, Table 6.3, are the limits given to the optimizer for each dimension to prevent damage to the system and limit operation to the area of interest.

Nelder-Mead Method

Nelder-Mead is a simplex optimization method using thruster test points P_0, P_1, \dots, P_n to form $n+1$ vertices of a polytope, where n is the number of dimensions of the optimization. The simplex is the simplest n -dimensional polytope of a given dimension. In 0D it is a point, 1D is a line segment, 2D is a triangle, 3D is a tetrahedron, etc. The vertices of the polytope represent the thruster test points where the objective function is evaluated. Once these test points are sampled, they are used to extrapolate the position of the expected minimum value of the objective function using a calculated polytope centroid and three new test points using the reflection, contraction and expansion functions described below [113]. Figure 5.3 shows these points and how they relate to the simplex formed in a 2D optimization.

The initial simplex of the Nelder-Mead method is created with points separated by 5% of the span of \mathbf{x} in each dimension from the starting point. The $n+1$ points of the simplex are ordered by the evaluation of the objective function at each location, so that

$$F(P_0) \geq F(P_1) \geq \dots \geq F(P_n), \quad (5.3)$$




Figure 5.3: The simplex formed in a 2D Nelder-Mead optimization showing the three current test points (P_0 , P_1 , P_n), the centroid (\bar{P}), and the reflected (P_R), expanded (P_E), contracted (P_C), and shrunk (P_{0s} , P_{1s}) points.

where P_0 is the worst point and P_n is the best point.

The centroid of the polytope, \bar{P} , is calculated using all test points except the test point with the worst evaluation of the objective function, P_0 . The reflected test point, P_R , is then found by creating a vector between the worst test point and the polytope centroid, using a positive reflection coefficient α , through the relation

$$P_R = (1 + \alpha)\bar{P} - \alpha P_0 \quad (5.4)$$

where α is the reflection coefficient. Taking the reflected point with an α value of 1 maintains the current polytope size and therefore the rate at which it can search the parameter space. $F(P_R)$ is then evaluated by the automated control system. If $F(P_1) \geq F(P_R) \geq F(P_n)$ then a new simplex is formed with P_R and the other n best points, and the optimization moves to the next iteration.

If the reflection test point has produced a better value than P_n , then the value P_R is expanded through

$$P_E = \gamma P_R + (1 - \gamma)\bar{P} \quad (5.5)$$

to get the expanded point P_E where γ is the expansion coefficient. A γ value of 2 accelerates the exploration of the parameter space by stretching the simplex without being so large that it misses regions of interest. If the evaluation $F(P_E)$ is better than $F(P_R)$ then a new simplex is formed with the n previous best points and P_E . If not, the simplex is formed with P_R and the optimization moves to the next iteration.

However, if the evaluation of $F(P_R) \geq F(P_1)$, then a contracted point P_C is calculated in one of two ways. If P_R is the worst point then

$$P_C = \beta P_0 + (1 - \beta)\bar{P} \quad (5.6)$$

where β is the reflection coefficient. This creates the contracted point within the simplex. Else $F(P_0) \geq F(P_R)$ and

$$P_C = \beta P_R + (1 - \beta)\bar{P} \quad (5.7)$$

creating the contracted point between the reflected point and the centroid. The β value of 0.5 mirrors the value of γ allowing the search to decelerate at a similar rate once the simplex is in an operational area of interest. If the evaluation of P_C is better than P_0 then a new simplex is formed with the n best points and P_C , and the optimization moves to the next iteration. If the contracted point is the worst point, then the simplex is shrunk by moving every point half way to the best point. For point 1 this would be

$$P_{1S} = (P_1 + P_n). \quad (5.8)$$

The algorithm continues, with the polytope evolving through n -dimensional space, until one of two convergence criteria are met. The standard deviation of $F(\mathbf{x})$ and \mathbf{x} for the $n + 1$ points on the polytope and the next selected test point are compared to defined tolerances. When either drops below their tolerance value, the convergence criteria is satisfied and the iterative search is ceased. The tolerance chosen is a trade off between the accuracy of the solution and speed of convergence. For the presented optimizations the tolerances are 1% of $F(\mathbf{P}_n)$ for $F(\mathbf{x})$ and 5% of the normalized dimensions in \mathbf{x} . The shape of the search space also plays a role. Flat objective functions require a smaller tolerance value for $F(\mathbf{x})$ to achieve the same accuracy in \mathbf{x} , while more defined minima can be found with larger tolerance values in $F(\mathbf{x})$ but require smaller values for \mathbf{x} . The thruster parameters are normalized for \mathbf{x} such that the bounds become $[0, 1]$, ensuring that the differing scales of each parameter do not affect the optimization. If a calculated point has a system parameter beyond the bounds of the optimization, that parameter is set to the limit of the bounds. The values $\alpha = 1.0, \beta = 0.5, \gamma = 2.0$ used for all tests are those recommended by Nelder as they maximize convergence speed while minimizing false optima [113].

Powell's Method

Powell's method is an efficient way to use one-dimensional optimization to solve a multidimensional problem when the derivative of the objective function is unknown or expensive to calculate. Effective up to 20 dimensions [114], this method initially uses non-parallel vectors, ξ_r , which are the unit vectors of the native search space dimensions (discharge voltage, field strength, etc.). These vectors map the search space of dimension n , for $r = 1, \dots, n$.

Figure 5.4 shows an example evaluation using Powell's method. The algorithm runs a one-dimensional optimization by sampling along the first vector, ξ_1 . Each one-dimensional optimization uses Brent's method, a hybrid root-finding algorithm, to select test points and then determine the minima $F(\mathbf{p}_1)$, and location of the minima, \mathbf{p}_1 [115]. Brent's method maintains the convergence reliability of the bisection search method while increasing conver-




Figure 5.4: The evaluation of the first dimension along the vector, ξ_1 , in Powell's method. Data points taken at position \mathbf{x} with their evaluation $F(\mathbf{x})$ are shown. The estimated minima of the function at \mathbf{p}_1 is then the starting point of the normal vector, ξ_2 , used to evaluate that dimension.

gence speed with inverse quadratic interpolation of the test points when appropriate.

Similarly to the implementation of Nelder-Mead, the thruster parameters of the search space are normalized such that the bounds become $[0, 1]$. If a test point is selected beyond the bounds of the optimization, that point is taken at the bound. Once Brent's method meets the convergence criteria, \mathbf{p}_1 is chosen as the value in ξ_1 and a new one-dimensional optimization is performed along ξ_2 . This is continued in subsequent dimensions, until the minima on ξ_n is found, where a conjugate vector is determined from the locations of \mathbf{p}_1 and \mathbf{p}_n . The algorithm continues along this new conjugate vector, again solving a one-dimensional optimization. As the algorithm continues to iterate, new conjugate directions are determined by the optima determined in previous vectors until the convergence criteria are met. The convergence criteria for each dimensional optimization include tolerances for $F(\mathbf{x})$ and \mathbf{x} . After each measured point the convergence criteria for ξ_r are checked by comparing the delta between sequential test points to the tolerances. For the presented optimizations the tolerances are 1% of $F(\mathbf{p})$ for $F(\mathbf{x})$ and 2% of the normalized dimensions in \mathbf{x} . The global convergence criteria compares optima in each sequential ξ_r to the tolerances.

5.2 Results and Discussion

Applying the rapid thrust measurement method effectively requires understanding how the thruster reacts under autonomous control. Firstly, the frequency at which measurements can be taken is determined, along with step sizes in each of the dimensions. Once the limits of these have been assessed, rapid measurements of varying dwell times are compared to shutdown thrust measurements. The Nelder-Mead and Powell optimization schemes are then assessed in two, then five, dimensions. Finally, limitations to the rapid thrust measurement methodology and selected optimization schemes are discussed.

5.2.1 Measurement Frequency

The first step in setting up the system is to determine how the thruster and thrust stand respond to changes in each dimension, and how quickly these dimensions can be ramped. Each change in thruster operating point can involve a simultaneous change in all five dimensions. Large changes in some dimensions can cause thruster shutdown, causing significant delay while the control system restarts the thruster and returns to thermal steady state. For example, a significant increase in discharge voltage will create a nearly instantaneous response, far faster than the PID control can adjust the gas flow to achieve the desired current. This results in thruster shutdown if the voltage is reduced, or extreme overshoot of the test point if the voltage is increased. These problems can be mitigated by controlling the maximum ramp rate in each dimension. In addition to the thruster, the response of the PID controller and propellant feed line length affect the maximum ramp rate. The required anode gas flow in a Hall thruster is primarily dependent on the desired current. As the optimizations are run at a constant anode power, changes to the gas flow are therefore dominated by changes to the anode voltage. Tuning the PID values and ramp rate together gave optimal results under 40 V/s for the anode potential, minimizing overshoot and eliminating thruster shutdown. This was simplified to a 5 s ramp time for all dimensions and a maximum discharge voltage step size of 150 V.

The dwell time at each test and control point is the limiting factor in how quickly the system can take data and optimize the thruster. It depends on how quickly the thruster can reach the desired operating point, and then how quickly it can reach quasi-steady state where the anode is thermalized. The ramp rate and PID tuning have minimized the time to switch operating points, allowing the remainder of the dwell time for the thruster to reach quasi-steady state.

Figure 5.5a illustrates how the response of the Hall thruster depends on the step size between the control and test points. Thrust measurements are taken at each of the dashed lines before the controller ramps to the next point. Three test points are measured in Fig. 5.5a at 340, 315, and 290 V, with a 265 V control point between each. At the 340 and 315 V test points, the thrust signal is dynamic when the measurement is taken as the gas flow and anode temperature have not had sufficient time to equilibrate. As the delta between the test and control points decreases, the gas flow and thermal response is smaller, allowing the system to equilibrate faster. At the 290 V test point, the system settles rapidly and the thrust measurement is taken at quasi-steady state.

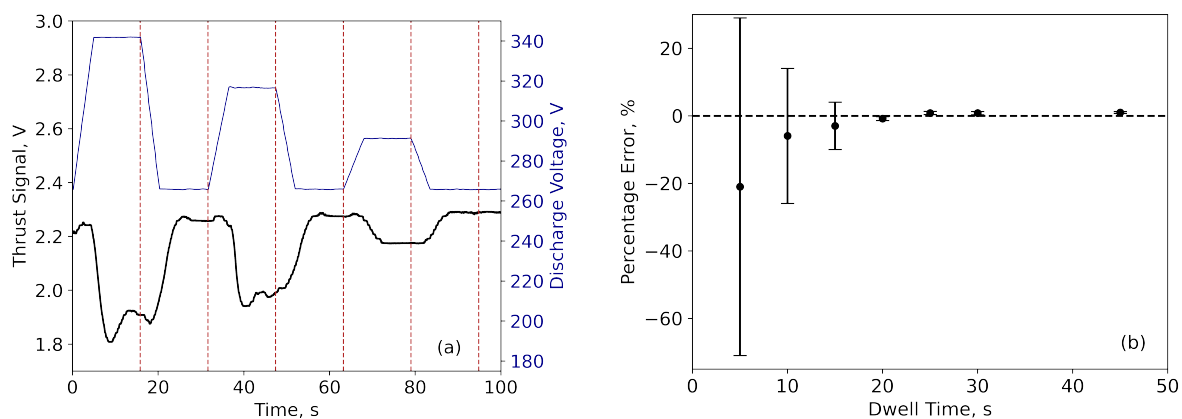


Figure 5.5: a) Iterations between the control point, and test points with a decreasing discharge voltage, a fixed 200 G magnetic field strength, and dwell time of 15 s. Dashed vertical lines indicate the points where thrust is measured. b) Mean error and standard deviation of the rapid measurement methodology with increasing control and test point dwell time.

Measurements of the anode efficiency were used to determine the minimum t_d for optimization as they capture the effects of both the thrust measurement methodology, and of the changing potential, magnetic field, and gas flow on the thruster. A series of tests with t_d from 5 to 45 s were conducted with 30 repeated measurements between a test point with a discharge voltage of 200 V, and magnetic field strength of 210 G, (200 V, 210 G), and a control point of (300 V, 180 G). The mean and standard deviation of these thrust measurements, compared to a shutdown thrust measurement, are shown in Fig. 5.5b. At a t_d of 5 s, the control system has only just finished ramping the input when the measurement is taken. This results in both high variance and mean error as the measurement was taken while the system was highly dynamic. As the t_d increases, the mean error and variance both decrease until the 20 s test. Here the variance is low, but the system slightly underestimates the efficiency compared to the shutdown measurement. At 25 s and above, the system consistently overestimates the efficiency with an error of 0.8 to 1% while the variance remains small. In this case, the test point discharge voltage is lower than the control point, causing the bulk thruster temperature to decrease during the test point and increase at the control point. This effect may be the source of a slight change in mean value of the control point when using this method, leading to the slight overestimation of efficiency seen in Fig. 5.5b. Further characterization of this error, using a wider range of test points with discharge voltages both above and below the control point, is presented in Section 5.2.2.

All data shown in the following tests was taken with a t_d of 25 s, to minimize test time while maintaining accuracy. When comparing the shutdown measurements to the new method, the uncertainty in the measurement stated is only from the rapid method in comparison to the shutdown thrust stand measurements. It does not include uncertainty in the calibration, or other systematic errors as they are consistent for both measurements.

5.2.2 Mapping

With rapid thrust measurements now possible, and full automation of the control system complete, the performance of the thruster was mapped in two dimensions. The map is

a useful development tool and allows for easy validation of the optimization schemes in two dimensions. The mapping was run at 500 W constant power, varying the voltage and magnetic field strength. The system is capable of sampling 72 test points per hour. For comparison, the shutdown method can take four to six measurements of ACME in the same time. In under two hours, a contour map of the thruster operation comprised of 117 points was completed with the rapid method. The map, Fig. 5.6a, shows the total efficiency of the thruster, with areas where the thruster operation was not stable and resulted in shutdown shown as 0% efficient. For comparison sake, eight additional test points representing the results of the shutdown thrust measurement procedure are overlaid both Fig. 5.6a and b. Thrust measurements utilizing the rapid methodology overestimated the thrust by 0.9% and total efficiency by 2.7% on average. For example, the rapid measurement at (150 V, 170 G) gave a thrust of 27.7 mN and total efficiency of 20.9%, while the shutdown thrust measurement gives 27.5 mN and 20.5% respectively. This error is consistent with the thrust error measured in the repetition tests in Fig. 5.5b, including those points taken at a higher test voltage than control voltage where the slight overestimation continued. The efficiency error stems from the thrust measurement which accounts for 1.8% and the lower total mass flow rate which accounts for 0.9% as per the relationship in Eq. 5.1. While this is significant for final performance measurements, it is well within the accuracy required for rapid mapping and optimization.

The total efficiency map shows a region of peak efficiency from (130 V, 140 G) to (170 V, 220 G), varying by less than 0.7% along its length. This region is highlighted in Fig. 5.6a with twin dotted lines. The improved performance in this high current density region, and lack of stability at higher voltages, can be attributed to the relatively short channel, a result of the anode position chosen for these tests. The anode efficiency map seen in Fig. 5.6b shows a single region of peak efficiency of 29.0% around (170 V, 220 G). The difference between maps is primarily due to the adaptive design of ACME where significant power is needed to run the magnets. Although an increased magnetic field strength improves anode efficiency, the increased magnet power usage hinders the total efficiency to a greater extent




Figure 5.6: a) A contour map of the total efficiency in the two-dimensional test area. The map is generated using the rapid measurement technique and 117 test points. Eight comparison points showing the efficiency measured using the shutdown thrust measurement technique are overlaid for comparison. Dashed lines showing the region of peak efficiency are shown for comparison to subsequent optimizations. The approximate area where mode transitions occurred is shown as the broad dashed line. b) A map of the anode efficiency from the same tests.

than a conventional Hall thruster design. To demonstrate the five-dimensional optimization, the total efficiency will be used as it includes both these effects, and those due to cathode flow and keeper current.

Mode shifts from Diffuse to Jet mode were observed around 300 V, with Jet mode dominating as the voltage increased [95]. The mapping was not significantly affected by the mode shifts. Control continued as expected without issue on either side of the mode shift. In the transitional regime between the modes, the low pass filtering of the discharge current monitor was not affected by the flickering between modes, and the PID loop controlling gas to the anode acted on the average current of the two modes. As the flickering was faster than the thrust stand response time, the thrust measurement was also not affected. Additionally, optimization trajectories explored the transitional region, but did not dwell in the Jet mode due to the lower efficiency in this setup.




Figure 5.7: a) A map of the two-dimensional total efficiency map with optimization paths shown for the Nelder-Mead optimization scheme. Starting points for the optimization are in blue with the resulting optimums in red. Lines showing the region of peak efficiency are shown for comparison other schemes. b) The total efficiency at each iteration for the same optimization paths.

5.2.3 Optimization in 2D

Although rapid thrust measurements can map two dimensions quickly, optimization schemes become advantageous at higher dimensions. While a two-dimensional, likely convex, search space does not strictly require an optimization algorithm, it does provide a simple proof of concept test. The map previously created provides an excellent proving ground as the optimum is already known, giving an easy way to check the convergence of the method in lower dimensions before adding the complexity of higher dimensional spaces. Prior knowledge of the behavior of the anode efficiency in other configurations of ACME shows that the dominant dimensions are voltage and field strength. Here, the two derivative-free algorithms of Nelder-Mead and Powell are demonstrated in these two dimensions, showing how the schemes work and their functional limitations. A single control point of (Discharge Voltage: 250 V, Magnetic Field Strength: 180 G, Keeper Current: 0.5 A, Cathode Flow Fraction: 0.15, Magnetic Field Skew: -0.06) is used for all optimizations as it represents the center of the parameter space and therefore minimizes the maximum possible step size.

Figure 5.7a shows the same map of total efficiency with five optimization trajectories taken with the Nelder-Mead method. The optimizer has no prior knowledge of the test space; the trajectories are overlaid on the maps to aid in visualizing the process. The starting point of each optimization is a blue point, and the found optimum of that trajectory is a red point. Each optimization found its optimum within 15 iterations, and stopped searching when their tolerance criteria was met within 20 iterations, see Fig. 5.7b. The points shown on the trajectories are the new vertices of the polytope in each iteration. Any additional points sampled to select the next vertex are not shown. For example, if the expanded point P_E is selected, the algorithm must first have sampled the reflected point, P_R . As expanded points generally occur at the start of the trajectory, and contracted points generally occur near the minima, the number of additional points required is small. Of the trajectories shown, none required more than 6 additional samples.

In two of the demonstrated trajectories, the algorithm continues to search a significant number of steps after the reported minima was found. At each of these points, neither of the tolerances were met due to large deltas between vertices of the polytope. The returned optima is always the best value, but it is not necessarily in the final polytope.

Since the search operations in Nelder-Mead are displacements, dependent on the centroid formed by the best n points and the worst permitted point, the subsequently sampled point will be somewhere along the line that joins the centroid and worst evaluation. This tends to make this optimization scheme's trajectory relatively straight. Depending on the initial points chosen, this can result in the algorithm determining an 'optimum' that is not actually a global optima. As the dimensionality increases, this effect is exacerbated as will be seen later in Section 5.2.4.

As mentioned in Section 5.2.2, the thruster design and physical configuration resulted in a region of peak efficiency from (130 V, 140 G) to (170 V, 220 G), varying by less than 0.7% total efficiency along its length. While all the optimizations did end on this region, the slight gradient towards the high field region was not followed. However, rapid optimization is a trade off between optimization speed and accuracy. Adjusting the criteria for defining

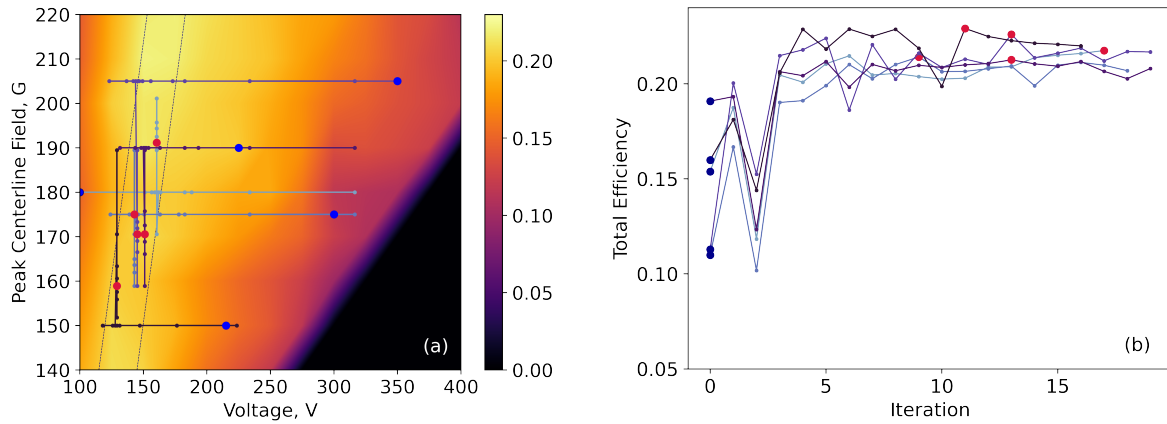


Figure 5.8: a) A map of the two-dimensional total efficiency map with optimization paths shown for the Powell optimization scheme. Starting points for the optimization are in blue with the resulting optimums in red. Lines showing the region of peak efficiency are shown for comparison other schemes. b) The total efficiency at each iteration for the same optimization paths.

an optimum, along with the positive reflection coefficient α , the expansion coefficient γ , and the contraction coefficient β would allow for the trajectory to follow the smaller gradient. However, rather than extend the test time of a single optimization, multiple initial points are tested and their optimum points compared. This gives a better understanding of the thruster operation as the trajectories effectively search the parameter space allowing multiple local minima with different operating modes to be found. Multiple fast runs can be completed in the time taken for a fine tuned optimization and will be more likely to find a global optima rather than accurately finding a local optima.

Figure 5.8a shows optimization paths taken using Powell's method starting from the same five initial test points. Each trajectory converged, meeting their tolerance criteria within a similar number of iterations required as the Nelder-Mead method, see Fig. 5.8b. All optima found using Powell's method also fell in the region of high efficiency. In some cases, the reported minima was found while the step size, and therefore delta between subsequent measurements, was larger than the tolerance. These trajectories continued until the convergence criteria were satisfied. Similarly to Nelder-Mead, the criteria for defining an optima can be




Figure 5.9: a) The two-dimensional specific impulse map with optimization paths shown for both the Nelder-Mead and Powell optimization schemes. Starting points for the optimization are in blue with the resulting optimums in red. b) The specific impulse at each iteration for the same optimization paths. The dashed lines show the Powell scheme optimizations.

tightened, however, multiple optimizations from a range of starting points will yield better information on the system to be optimized and may find more than one optima.

Overall, the highest total efficiencies found in two dimensions were 22.5% in the map, 22.9% for Nelder-Mead, and 22.9% for Powell. The higher efficiency values found during the optimizations are due to the optimal points being between mapped points and are therefore not captured by the map.

To test the adaptability of the optimization methodology for different objectives, both schemes were applied to optimizing for I_{sp} . The thruster was in the same configuration, and no changes to the optimizer were made except to swap objective functions. Figure 5.9a shows both Nelder-Mead and Powell optimization paths, starting from the same five locations as previous tests, overlaid on a map of the specific impulse. The map shows a single, relatively flat region of peak I_{sp} , in contrast to the line seen in previous total efficiency optimizations. All five paths converged on the same global maxima within their tolerance in either $F(\mathbf{x})$ or \mathbf{x} . The paths converged within a similar number of iterations to the total efficiency tests and all identified the operational area of interest. This adds confidence in the optimizer when

changing objective functions. As with previous tests, a finer tolerance requirement would increase the grouping on the flat maxima, however the area of interest has been suitably identified and further accuracy isn't required.

5.2.4 Optimization in 5D

Now that the schemes have been tested and developed in the proving ground of 2D optimization, we can verify that these methodologies work effectively in five dimensions. Here, the search space is large enough that fully mapping the dimensions becomes untenable, even with rapid thrust measurement. Using the same resolution as the 2D maps, a 5D map would take 2500 hours of continuous running. Optimizing in five dimensions is therefore hugely advantageous as operational areas of interest can be rapidly identified. In more traditional thruster development, lower dimensional maps would be used to inform localized mapping, one-dimensional sweeps, and other more focused tests requiring human intervention. Automated optimization in these higher dimensional spaces is therefore of interest as it can expedite several steps by rapidly finding key areas of interest. The optimization uses all five fast dimensions of ACME shown in Table 6.3.

Figure 5.10a and b show the evolution of the thruster efficiency in parallel to each of the thruster dimensions for Nelder-Mead and Powell, respectively. The different approaches of the schemes can be seen in how each of their variables change.

The Nelder-Mead method evolves all dimensions simultaneously, as indicated by the shaded regions in Fig. 5.10a. The first 6 points sampled are to generate the initial polytope before the expansion, contraction, and reflection operations begin. As the polytope must be generated for each new optimization, the speed of batched or sequential optimizations is limited. The scheme reached a maximum efficiency of 21.35% after 44 iterations at (Discharge Voltage: 165.0 V, Magnetic Field Strength: 219.3 G, Keeper Current: 0.519 A, Cathode Flow Fraction: 0.166, Magnetic Field Skew: -0.079). The dominance of the first two dimensions (Discharge Voltage and Magnetic Field Strength) drives this optimization method as a 5% change in the normalized discharge voltage may have a larger effect than a 50% change in




Figure 5.10: Convergence for five-dimensional optimization showing: Total Efficiency, η ; Voltage (V), V; Field Strength (G), B; Keeper Current (A), KC; Cathode Flow Fraction, CF; and Field Skew, S. a) shows Nelder-Mead, and b) shows Powell. The iteration with peak total efficiency is marked with the dashed vertical line. The dimensions being optimized at each step are indicated by the shaded regions.

the keeper current, cathode flow fraction, or skew. This makes fine tuning of the less gross parameters difficult without significantly increasing the time to convergence.

The Powell method initially optimizes each dimension individually, see Fig. 5.10b. From the starting point, only the discharge voltage, ξ_1 , is adjusted with each iteration until it converges. Then, the optimizer moves on to optimize the magnetic field, ξ_2 , keeping the optimum voltage. This continues through each dimension until iteration 27, where the optimizer forms a conjugate vector from the initial minima of ξ_1 and the minima determined from the ξ_5 vector. The optimizer then iterates one step either side of the current optima, and

not finding improvement, confirms the optima meets convergence criteria. Five-dimensional Powell increased the total efficiency from the two-dimensional optimization, reaching 23.24% after three dimensions, 23.35% after four, and 23.41% in the fifth dimension. These efficiency gains are in the order of the $\pm 0.5\%$ uncertainty of the rapid measurement method, however the data is consistent relative to the points around it, and the optimizers' exploration in the area increases confidence in the optimum location. The optimum was achieved after 27 iterations at (Discharge Voltage: 164.9 V, Magnetic Field Strength: 217.6 G, Keeper Current: 0.392 A, Cathode Flow Fraction: 0.176, Magnetic Field Skew: -0.092).

The two algorithms agreed on the optimum voltage and magnetic field with a difference of 0.1 V and 1.7 G, or 0.03% and 1.9% of full span, respectively. They differed more in the less dominant dimensions with a keeper current delta of 0.127 A or 8.5% of full span, cathode flow fraction delta of 0.01 or 3.33%, and magnetic field skew delta of 0.013 or 3.25%. As previously discussed, the lower Nelder-Mead efficiency stems from the three less dominant dimensions, as the algorithm cannot easily separate their influence from the more dominant dimensions while maintaining the rate of convergence.

5.2.5 Discussion of Challenges and Limitations

It was found that the rapid thrust measurement method described in Section 5.2.2 achieved similar accuracy to shutdown thrust measurements, provided sufficient dwell time at each test point. However, the rapid measurement method has some limitations that must be understood. These limitations stem from the differences between quasi-steady state, where the anode has thermalized, and full steady state where the entire thruster has thermalized.

Before all tests the thruster was warmed up to steady state at the control point. However, each test point will have a different thermal equilibrium that would take significant time to reach, far longer than the dwell time. For instance, tests of a 4.5 kW Hall thruster with a distributed array of 35 thermocouples investigated changes to thruster thermal performance when adjusting operational inputs [116]. Doubling the anode mass flow rate required an hour to achieve thermal steady state while changing the magnetic field setup took > 1.5

hours. Further, the changes in the thruster temperature were non-homogeneous. Lower magnet current resulted in increased plasma thermal loading near the exit plane, while lower magnet power reduced temperature in the upstream parts of the channel. Increased thermal loading from plasma also increases magnet power as they become more resistive with increasing temperature. This could lead to magnetic circuit saturation which has been shown to significantly affect performance and can lead to thermal mode shifts [94, 85]. Anode temperature has been shown to have a significant effect on the neutral transit time of the propellant, affecting the required mass flow rate, beam divergence, breathing mode frequency and magnitude, and causing mode shifts [117, 118, 119]. The rapid measurement method assumes the change in performance due to continued anode thermalization after quasi-steady state is achieved is small compared to changes due to the different operating points. This has been confirmed for total efficiency and thrust with the comparison points taken in Fig. 5.6a showing good agreement between the thrust measurement methods. This holds true for test points across the 2D parameter space including those with a higher discharge voltage than the control point voltage. However, additional effects including changes to the discharge oscillations in the thruster have not been studied.

There is precedent for using high-speed mapping without full thermal steady state. Maps of the oscillations in Hall thrusters, including HERMeS and the XR-5, have been generated using high-speed sweeps [49, 120]. This gives confidence that the operating point achieved in the dwell time is representative of the final steady state operating point. While there is still added uncertainty when using the rapid thrust measurement, the applicability of the method as a tool to rapidly characterize and optimize a thruster is not diminished.

Limitations of the thruster-in-the-loop optimization vary by the algorithm used. While the Powell method is efficient for five-dimensional rapid optimization, and scalable to higher dimensions, this cannot be extended into the slow dimensions of ACME as iterations in any of these slow dimensions or conjugate vectors would require a setup change between every measurement. The methods used thus far do not have a complete ‘memory’ of the test points recorded which is key for the optimizer to understand the relationship between the continuous

and discrete variables, and to interpolate between the physical setup changes of the slow dimensions. Therefore, when including the slow dimensions for a global optimization, more complex methods are required. In this scenario, algorithms such as Bayesian Optimization or Data-based Online Nonlinear Extremumseeker (DONE) are more appropriate, as they use the results of all previous data to create a surrogate model of how the objective behaves in its domain[121]. While these methods are more complex to implement, the surrogate models they generate can provide greater detail on the performance of the thruster across higher dimensional search spaces.

5.3 Chapter Conclusions

Novel experimental methods to rapidly characterize and optimize a Hall thruster have been developed. The rapid thrust measurement method was demonstrated, mapping 72 test points per hour, a $> 6\times$ improvement over shutdown measurements with an average error of 0.9% compared to shutdown thrust measurements. Combining existing algorithms, rapid thrust measurements, and a thruster-in-the-loop controller facilitated fully automated optimization. Two-dimensional Nelder-Mead and Powell optimization schemes were able to converge rapidly to maxima in total efficiency and specific impulse in under 15 test points. Powell’s method remained effective in five dimensions with convergence in under 30 test points. Automated optimization in these higher dimensional spaces can expedite many stages of thruster development by rapidly finding key operational areas of interest. The demonstrated plug-and-play nature of the objective functions can further enhance the optimizer’s capabilities. More complex functions allow for real time thruster-in-the-loop optimization that incorporates system mass, mission cost, required performance for maneuvers, etc., vastly improving the potential functionality for mission specific optimization.

The investigation of rapid optimization was invited to, and published in, the Journal of Electric Propulsion’s “Select Contributions from the 38th International Electric Propulsion Conference”.

Chapter 6

MODELING AND OPTIMIZATION OF PROPELLANT MIXTURES FOR HALL THRUSTERS

The study, development and application of electric propulsion (EP) systems presents a significant engineering challenge on many fronts. EP systems exhibit a range of physics that are difficult to model [122, 123, 22], increasing the reliance on experimental results to study observed phenomena, inform models, improve thruster designs, and characterize performance for flight [50, 49, 124]. The experimental methods traditionally used involve long test campaigns, iterating through large test matrices or performance maps. The results of these campaigns are typically very large data sets with methodical analysis of single variables, often relying on human interpretation to relate variables with complex dependencies [52, 53]. This approach is expensive and time consuming, and while successful, there is a need for more advanced methodologies capable of capturing the complex inter-dependencies of these high dimensionality systems.

The application of automation and computationally driven development described in chapter 5 is continued in this chapter. However, where the focus was previously on rapid optimization, we move to consider the more complete characterization of Hall thruster operation. Here, we present a highly adaptable method, capable of creating surrogate models of any quantifiable aspect of thruster operation using Gaussian Process Regression (GPR), and of optimizing that aspect of thruster operation in real time using Bayesian Optimization (BO). To demonstrate these methods, they are applied to the problem of modeling and optimizing Hall thruster operation on argon, krypton, and xenon gas mixtures.

Xenon has been the propellant of choice for Hall thrusters over the majority of their fifty years of flight heritage. With a low ionization energy, high ionization cross section, high

atomic mass, and excellent storage density, xenon is ideal for EP. As a result, EP accounted for 30% of global xenon usage in 2022 [125]. The limited and variable supply of xenon drives cost fluctuations, causing many EP spacecraft, particularly large constellations, to use alternative propellants [126]. Using krypton comes with decreased performance due to the lower atomic mass, lower ionization cross section, and higher ionization energy compared to xenon [53]. Some satellites and constellations have chosen to use even cheaper noble gases by switching to argon, further exacerbating the performance drop [105].

The use of gas mixtures for EP aims to reduce costs while limiting the decrease in performance. Studies of xenon/krypton mixtures have focused on ratios with large amounts of both gases in step sizes of 20 - 25%, which found decreased efficiency with increasing krypton fraction [127, 128, 129, 130]. However, specific impulse was observed to be relatively consistent across the range of mixtures at set discharge powers when krypton was run at higher current densities [129]. This aligns with a study which showed vastly improved krypton performance with increasing current density [131]. In addition to the previous findings, xenon injection at 1-4% as a tracer for LIF in krypton Hall thrusters was observed to significantly increase the discharge current, suggesting that a small fraction of xenon aids in the ionization of krypton [132]. A study using a mixture of 26% argon and 74% krypton, found improvements in the specific impulse and anode efficiency [106]. Similarly, another study found using xenon as an ionization enhancer improved the performance in an argon Hall thruster [105]. The results of these studies suggest that mixtures of xenon, krypton, and argon may be beneficial in some operating regimes.

The study of gas mixtures adds additional dimensionality to the optimization of thruster performance, resulting in a prohibitively large search space for mapping using conventional approaches, thus creating an ideal test bed for these methods. The motivation for using lighter, more abundant propellants, must also be captured in the analysis. While the previously noted studies focused on direct performance comparisons, this study will use the total mission cost of a set of industry-relevant sample missions as an objective, therefore encompassing mission-relevant parameters including propellant costs, launch costs, and storage

considerations. This chapter presents real-time thruster-in-the-loop optimizations and surrogate models for mission cost using experimental data of a Hall thruster operating on mixtures of argon, krypton, and xenon.

6.1 Methods

Real-Time optimization, and data collection for surrogate models, requires the synthesis of algorithmic elements with the automated control of test facilities, diagnostics, and thruster. This section first details the algorithms used for the optimization and modeling of the Hall thruster, and the objective functions and cost model used to define the optimization goal. Then the experimental setup, including Hall thruster, vacuum chamber, thrust stand, and control system are described in detail.

6.1.1 Modeling and Optimization

The modeling and real-time optimization of mission cost requires a set of methods, which when combined, can generate a model of the total cost for a sample mission based on thruster performance, and autonomously make decisions to sample the search space and find the optima. This section will detail the GPR method used to generate surrogate models of the search space, the BO used for real-time optimization, and the cost model used to define the optimization objective for sample missions.

Gaussian Process Regression

Originally created for mining precious ores efficiently, the interpolation method known as GPR is a powerful tool for making predictions with estimated uncertainty on the behavior of an unknown function [133]. The interpolation is a type of surrogate model, which are used in lieu of a physics based model due to the difficulty in directly modeling Hall thruster operation. A surrogate model is a direct model of a system output (i.e. thrust) based on the the system input (i.e thruster operating parameters). It does not contain any additional

information of the system. Therefore to model another output (i.e. efficiency), a separate surrogate model is trained. They are very efficient both computationally, and in the required experimental data, however, their use comes at the cost of losing the deep understanding physics-based models provide.

Physical measurements of the system, called observations, separated in control parameter space can create a Gaussian kernel, or covariance function. Observations are interpreted as draws from a random process, which we assume to be Gaussian, despite identical measurements when an observation is repeated. The kernel predicts the surrogate model at a new location in the search space by assuming a prior about the behavior of the space. We assume as our prior the search space to be regular and continuous, and able to be approximated by a constant locally.

The evidence of the observed points then forms their own probability distribution and joint probability with the prior assumption. This gives the posterior probability distribution using Bayes' Theorem [134, 135]:

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(B)}, \quad (6.1)$$

where $\mathbb{P}(A|B)$ is the posterior probability, $\mathbb{P}(B|A)$ is the likelihood, $\mathbb{P}(A)$ is the prior, and $\mathbb{P}(B)$ is the evidence. Two or more spatially separated samples create a Gaussian kernel or covariance function, approximating any continuous function arbitrarily well. The kernel predicts the objective function at a new location using the Gaussian prior and a Gaussian likelihood in Bayes' Theorem. The resulting posterior distribution is also Gaussian. Due to the use of probability distribution functions, an estimate of the mean and variance of the surrogate model behavior is calculable using the Gaussian kernel:

$$K(\mathbf{x}_i, \mathbf{x}_j) = k \exp \left[-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2} \right], \quad (6.2)$$

where k is the autocorrelation, and σ is the standard deviation, where both are positive to

ensure a positive definite kernel. With this basis function, and $f(\mathbf{x}_i)$ being the observation (i.e. calculated total cost) at the point \mathbf{x}_i (i.e. the thruster operating point), the mean and the standard deviation can be estimated from a series of points in the search space, $\{\mathbf{x}_0, \dots, \mathbf{x}_n\}$, and their observations, $\{f(\mathbf{x}_0), \dots, f(\mathbf{x}_n)\}$. Here, the mean and standard deviation given by the surrogate model are the predicted value and uncertainty at that operating point. In this chapter we present surrogate models used to directly model thrust, specific impulse, efficiency, and total cost of sample missions. The surrogate model is also key part in the real-time Bayesian optimization of the thruster which will be detailed in the next section.

Bayesian Optimization

BO is a statistical reinforcement learning technique for discrete, multidimensional search spaces where the objective function is expensive to evaluate, and the evaluation itself is a black box [121]. Expense in evaluation here is the cost of taking data points by physically running the thruster, while ‘black box’ refers to the inability to analytically predict the performance. An objective function is defined as the goal of the optimization.

The process used in this application of BO for optimizing the thruster is shown in Fig. 6.1. From the test start, the model can be initialized by previous data, or new random samples in the search space can be generated and evaluated. A surrogate model of the search space is created using the selected prior, and GPR using the initialization observations of the objective process, $\{f(\mathbf{x}_0) \dots f(\mathbf{x}_i)\}$ at the set of coordinates $\{\mathbf{x}_0, \dots, \mathbf{x}_i\}$. From the surrogate model, an acquisition function is required to choose the next point for BO. This acquisition function balances exploration and exploitation through calculating the posterior’s mean and variance. Generally, acquisition functions are the expectation of observing a point which produces a better result in the objective function than all points previously observed. It does this through considering the mean and the standard deviation of the surrogate model, versus the best point yet observed. While there are a number of possible acquisition functions, “Expected Improvement (EI)” and, more recently, “log Expected Improvement (logEI)” have dominated use in BO as they are simple to understand, provide a balance of exploration and

exploitation, and are appropriate for scenarios where finding the maximum of the acquisition function is less expensive than sampling more data [136]. Avoiding any possible vanishing gradients in the acquisition function, “logEI” has recently shown consistent performance with higher dimensional, noiseless systems, which can otherwise cause convergence issues with an “EI” acquisition function [137]. As a consequence, “logEI” is used in our results.

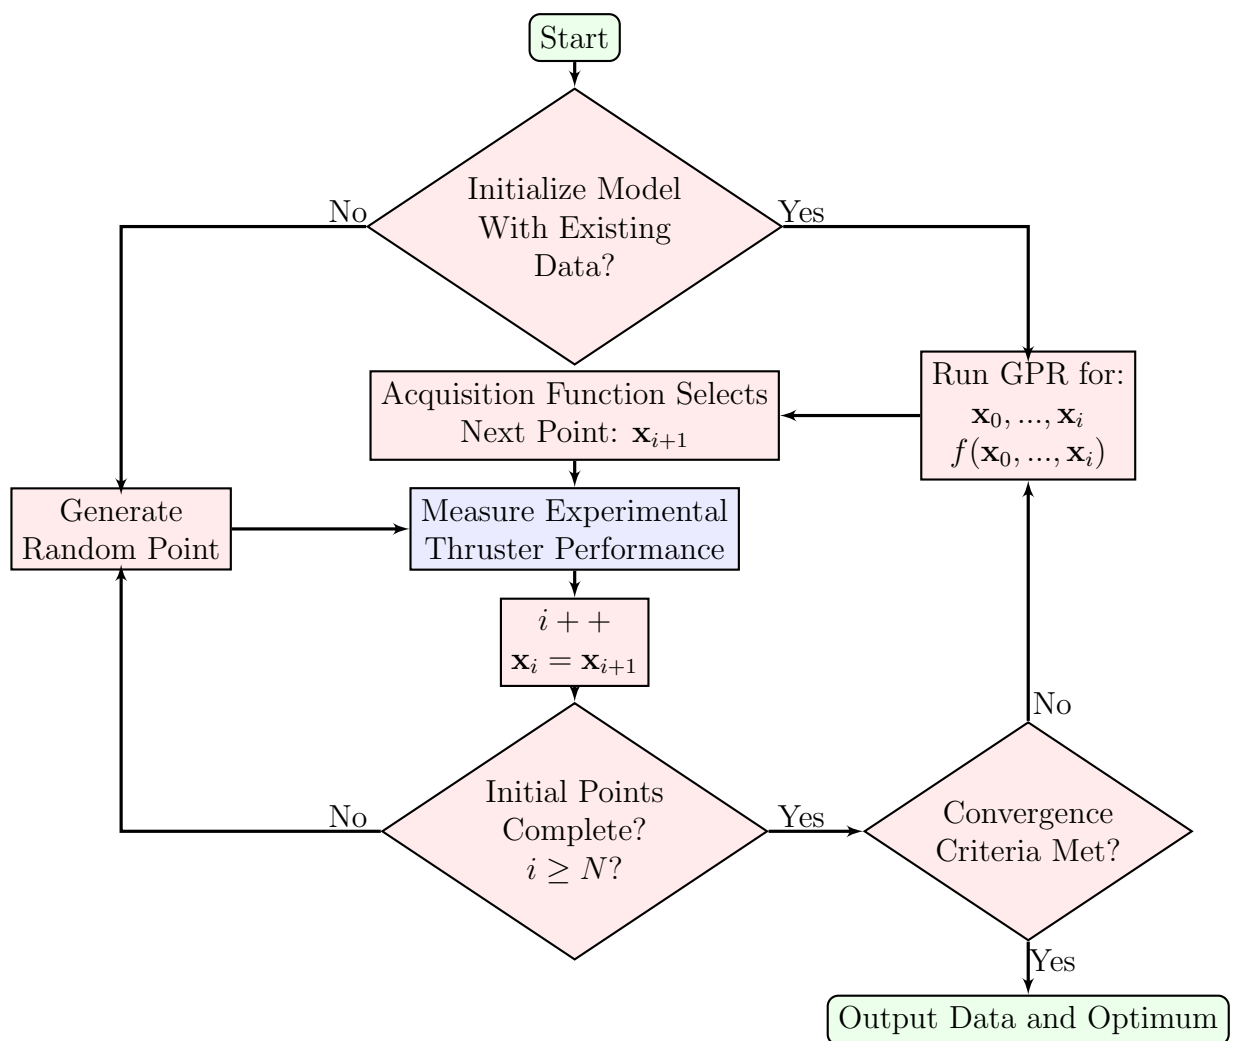


Figure 6.1: A flow chart showing steps to initialize a surrogate model with random data, and run BO of the thruster using experimental measurements. Thruster operations are shown in blue, optimizer functions in red, and start/stop points in green.

With the acquisition defined, the BO algorithm proceeds by finding the location of the maximum of the acquisition function, giving the next point, \mathbf{x}_{i+1} , for observation. This is added to the data set of points and observations, and GPR is performed again on this updated data set, continuing the optimization loop. The convergence criteria for the experiments run was set to a number of iterations, allowing observation of the process for longer periods.

In order to evaluate each point, the thruster must be ignited and run at the operating point, \mathbf{x}_{i+1} , the performance measured, and objective function, $f(\mathbf{x}_{i+1})$, calculated. All these processes, including running the optimization algorithm, require a dedicated control system which will be detailed in section 6.1.2.

Objective Function and Cost Model

The objective function is the mathematical definition of the optimization goal. In this case, we are aiming to minimize the total mission cost and must therefore be able to model mission cost based on the thruster performance at an operating point, \mathbf{x} , and the desired mission specifications. To do this, a simplified cost model based on mass models in Ref. [138] is used. The objective function, $f(\mathbf{x})$, is taken as

$$f(\mathbf{x}) = C_{\text{total}} = C_{\text{launch}}[m_0(\mathbf{x})] + C_{\text{sat}}[m_p(\mathbf{x})] = \kappa_{\text{launch}}m_0(\mathbf{x}) + C_{\text{sat}}[m_p(\mathbf{x})], \quad (6.3)$$

where C_{total} is the total mission cost, calculated from the launch cost, C_{launch} , the cost of the fully fueled satellite, C_{sat} , the specific launch cost in \$/kg, κ_{launch} , the launch mass of the satellite, m_0 , and the set of input parameters, $\mathbf{x} = (\gamma_{Ar}, \gamma_{Xe}, V_d, B, I_k)$, defined by the gas fractions of argon, xenon, and the discharge voltage, magnetic field strength, and keeper current, respectively. The launch mass is constrained to $50 \leq m_0 \leq 831\text{kg}$ with a specific launch cost of $\kappa_{\text{launch}} = 6500\text{\$/kg}$ to represent a smallsat ride-share on a Falcon 9 to a SSO, LEO, or polar orbit [139].

With the specific launch cost defined, the dependence of the launch mass on the input parameters and system needs to be considered. The launch mass of the satellite is the sum

of its constituent masses,

$$m_0(\mathbf{x}) = m_{\text{payload}} + m_{\text{bus}} + m_{\text{propulsion}} + m_{\text{power}} + m_{\text{tank}}(\mathbf{x}) + m_p(\mathbf{x}) \quad (6.4)$$

where m_{payload} , m_{bus} , $m_{\text{propulsion}}$, m_{power} , $m_{\text{tank}}(\mathbf{x})$, $m_p(\mathbf{x})$ are the masses of the satellite payload, bus, propulsion system, power system, propellant tank, and propellant, respectively.

The tankless mass, m_{tankless} , is the sum of the payload, satellite bus, propulsion, and power system masses, and is set as a constant as part of the mission specifications. The tank mass scales linearly with propellant mass [140], and is defined as $m_{\text{tank}}(\mathbf{x}) = \chi_t(\mathbf{x})m_p(\mathbf{x})$, where $\chi_t(\mathbf{x})$ is referred to as the tankage fraction, and depends on the tank type, material, and the propellant mixture.

The launch mass may therefore be simplified to a linear function of the propellant mass

$$m_0(\mathbf{x}) = m_{\text{tankless}} + [1 + \chi_t(\mathbf{x})]m_p(\mathbf{x}). \quad (6.5)$$

The ideal rocket equation closes the above equation, giving the $m_p(\mathbf{x})$ as,

$$m_p(\mathbf{x}) = m_0(\mathbf{x}) \left[1 - e^{\left(\frac{-\Delta v}{g_0 I_{sp}(\mathbf{x})}\right)} \right], \quad (6.6)$$

where ΔV is a constant specified by the mission requirements, g_0 is the standard gravity, and $I_{sp}(\mathbf{x})$ is the measured specific impulse of the thruster calculated from the measured thrust and propellant mass flow at operating point \mathbf{x} . The launch mass can then be calculated using the mission specifications, the performance measured from the thrust stand, and the calculated tankage fraction as

$$m_0(\mathbf{x}) = \frac{m_{\text{tankless}}}{1 - [1 + \chi_t(\mathbf{x})] \left(1 - e^{\left[\frac{-\Delta v}{g_0 I_{sp}(\mathbf{x})}\right]} \right)}. \quad (6.7)$$

To examine the influence of propellant storage technology on mission cost, we consider

two different tank styles in our cost model: supercritical and cryogenic. Most noble gas used in satellite propulsion is stored as a supercritical fluid, which allows for relatively light tanks for xenon but incurs significant mass penalties when used for lighter gasses. With lighter propellants, cryogenic tanks offer significant mass savings, but their use is complicated by the proper balancing of the thermal load to maintain proper propellant flow [140]. The tankage fraction for a supercritical thin-walled spherical tank is

$$\chi_{t,sup}(\mathbf{x}) = \frac{3RT\beta\rho_{\text{tank}}}{2\sigma_y M(\mathbf{x})}, \quad (6.8)$$

where R is the universal gas constant, T is the temperature, β is the safety factor, ρ_{tank} is the density of the tank material, σ_y is the yield strength, and $M(\mathbf{x})$ is the mean atomic mass of the gas mixture [140].

The cryogenic tankage fraction is

$$\chi_{t,cryo}(\mathbf{x}) = \frac{3P_{\text{tank}}\beta\rho_{\text{tank}}}{2\sigma_y\rho_p(\mathbf{x})} \quad (6.9)$$

where P_{tank} is the tank pressure, and the density of the propellant mixture,

$$\rho_p(\mathbf{x}) = \sum \gamma_j(\mathbf{x})\rho_j, \quad (6.10)$$

is the sum of the constituent species' densities, ρ_j , scaled by their respective gas fractions, $\gamma_j(\mathbf{x})$ [140] with $\gamma_{Kr}(\mathbf{x}) = 1 - [\gamma_{Xe}(\mathbf{x}) + \gamma_{Ar}(\mathbf{x})]$. For all tank related calculations, the values in Table 6.1 are used.

The launch cost is now fully defined; the other half of the total cost in this model, the satellite cost, $C_{\text{sat}}(\mathbf{x})$ is

$$C_{\text{sat}}(\mathbf{x}) = C_{\text{payload}} + C_{\text{bus}} + C_{\text{propulsion}} + C_{\text{power}} + C_{\text{tank}}(\mathbf{x}) + C_p(\mathbf{x}) \quad (6.11)$$

where C_{payload} , C_{bus} , $C_{\text{propulsion}}$, C_{power} , $C_{\text{tank}}(\mathbf{x})$, $C_p(\mathbf{x})$ are the costs of the satellite payload,

Table 6.1: Tank specifications used for cost optimization [140, 141].

Specification	Value
Tank Temperature (K)	300
Tank Density (kg/m^3)	2700 (Aluminum)
Tank Yield Strength (Pa)	4.8×10^8 (Aluminum)
Supercritical Safety Factor	2.0
Cryogenic Safety Factor	3.0
Argon Cryo Density (kg/m^3)	1392
Krypton Cryo Density (kg/m^3)	2412
Xenon Cryo Density (kg/m^3)	3080
Cryogenic Tank Pressure (Pa)	101325

bus, propulsion system, power system, propellant tank, and propellant, respectively. For the purposes of this analysis, all but the propellant cost is reduced to a constant and are therefore excluded from the optimization costs. The propellant cost C_p is then

$$C_p(\mathbf{x}) = \sum_j \alpha_j \kappa_j m_j(\mathbf{x}), \quad (6.12)$$

where α_j , κ_j , and $m_j(\mathbf{x})$ are the propellant cost factor, specific cost and total mass of propellant, j . The propellant cost factor is a term used to scale the price of krypton and xenon in an optimization to represent changes in the market price of the propellants. As argon is more commonly available and has not seen significant historical fluctuations, α_{Ar} is set at 1.0 for all tests. The specific cost of the propellants used are shown in Table 6.2.

An important factor not considered in this study is the influence of propellant mixtures

Table 6.2: Propellant costs used for optimization set at the 2015 price, a time of relative stability, and comparable to current prices [142].

Propellant	Specific Cost (\$/kg)
Argon	5
Krypton	108
Xenon	3420

on thruster lifetime. Previous studies suggest that the use of lightweight propellants increases the rate of thruster erosion, reducing their life expectancy [143, 105]. This could be addressed in future studies by creating an outcome constraint based on the lifetime required by the mission specifications, and the life expectancy of the thruster given the propellants used. However the lifetime of thruster is not sufficiently understood for an arbitrary gas mixture and set of operating parameters, and therefore cannot be easily mathematically defined as a constraint. Additional mission constraints can be added, such as a minimum required thrust for maneuvers, maximum propellant mass, or even a maximum propellant volume. As long as they can be mathematically defined and calculated from the thruster telemetry and mission specification, they can be used as a constraint.

6.1.2 Experimental Setup

An experimental setup was developed to characterize the performance of GPR for surrogate model development and demonstrate our optimization method using the sample problem of propellant gas mixtures in a Hall thruster. The setup includes the Hall thruster, vacuum chamber, thrust stand, and control system used to run the thruster, collect test data, and iterate the optimization algorithm. Each of these components are described in detail below.

ACME Setup

All testing was conducted using ACME V3, with a cathode-tied electrical configuration. ACME was mounted in on the centerline of STF-2, on an inverted pendulum thrust stand. An all-metal gas feed system with four MC-500 Alicat mass flow controllers delivered research grade 99.9999% argon, krypton, and xenon to the ACME anode and cathode. During operation, the pressure did not exceed 2×10^{-5} Torr for the flow rates and gas mixtures tested.

ACME has more than 10 dimensional control parameters that can be adjusted. These parameters can be categorized into the real-time computer controllable ‘fast’ dimensions, and the ‘slow’ dimensions which require mechanical changes to the thruster. The fast dimensions

Table 6.3: ACME thruster fast dimensions, $\mathbf{x} = (\gamma_{Ar}, \gamma_{Xe}, V_d, B, I_k)$, that are computer controlled during operation.

Fast Dimension	Minimum	Maximum
Argon Fraction, γ_{Ar}	0	1
Xenon Fraction, γ_{Xe}	0	1
Discharge Voltage, V_d	200 V	350 V
Field Strength, B	140 G	230 G
Keeper Current, I_k	0 A	1.0 A

are: xenon mass flow, krypton mass flow, argon mass flow, cathode mass flow, discharge voltage, keeper current, inner magnet current, and outer magnet current. To reduce the size of the search space, the data presented in this chapter was collected for a constant power of 1 kW. Gas fractions for the three propellant species are constrained to sum to 1. Therefore, the gas mixture can be fully defined with the argon and xenon fractions.

The discharge current of a Hall thruster is controlled by the gas flow through the anode. A PID controller reads the discharge current, and sets the total mass flow, which is then split into the mass flow fractions for each of the gasses and sent to the three flow controllers that feed the anode. The cathode operates on pure krypton to reduce the complexity of the experimental setup with a fixed cathode mass flow fraction of 0.1 to reduce search space size and to limit the influence of the cathode’s krypton flow on the main discharge. To further reduce the dimensionality of the optimization, the field skew is fixed, resulting in the five fast dimensions, $\mathbf{x} = (\gamma_{Ar}, \gamma_{Xe}, V_d, B, I_k)$, listed in Table 6.3. The minimum and maximum values listed are the bounds of the optimization search space. The optimization over the slow dimensions of the ACME thruster is beyond the scope of this chapter, optimization methods using the slow dimensions will be discussed in Section 6.2.4. A single configuration of slow dimensions was used, with a pole position of 0 mm, anode position of 28 mm from the exit plane, and keeper orifice flush with the center pole.

A vacuum gapped magnetic circuit limits the heat conduction from the poles on ACME, resulting in higher thruster operating temperatures. The anode temperature significantly

affects the neutral transit time of the propellant, changing the required mass flow rate, beam divergence, and breathing mode frequency and magnitude [117, 118, 119]. Lighter propellants suffer more at higher thruster temperatures as the shorter neutral transit time exacerbates problems due to the small cross section and higher ionization energy resulting in low mass utilization efficiency [131]. With previous ACME setups, high operating temperatures also led to magnetic circuit saturation, showing significant performance loss, and thermal mode shifts [94, 85]. To reduce these thermal effects during optimization, all thrust measurements are taken at a thruster temperature of 250 °C, as measured with a platinum RTD in the inner pole of the thruster.

Control System and Real-Time Optimization

Thruster-in-the-loop optimization requires full computer control of power supplies, flow controllers, and ignition ballast switching, along with all data acquisition. This is handled by an updated version of the controller described in chapter 5.

Executing the optimization with the control system follows the flow chart shown in Fig. 6.1 and detailed in section 6.1.1. Before starting the test, the thruster is warmed up to a thermal steady state by operating at an standardized point. Steady state is achieved when both the thruster temperature and propellant flow required to achieve the discharge current are constant. Thermalization of the thruster to the target temperature before optimization is key to taking repeatable measurements. Once the thruster is ready, the thruster shuts off and cools to a start temperature of 220 °C. The optimization algorithm selects a test point, \mathbf{x}_{i+1} , and the process of measuring thrust begins, consisting of ignition at the set gas fractions, ramping the thruster setpoint to the desired parameters, and waiting for the test end temperature of 250 °C. Once at temperature, the thruster is again shut off and the thrust delta is measured. Taking another measurement requires cooling the thruster to the start temperature, and repeating the process. The full cycle is completed with the shutoff gas lines venting into the chamber during cool-down to eliminate cross contamination and ensure accurate mixtures. After the optimization is complete, a single calibration mass is

added and removed to confirm the calibration taken before testing. The thrust data is then combined with the operating point and mass flows to calculate thruster performance, mission parameters, and evaluate the objective function. Once the thruster has cooled, the cycle begins again, with the new data point updating the prior. In the event of unplanned thruster shutdown, the control system detects the drop in anode current, restarts the thruster, and resumes the test. The end of the data set is defined as a set number of test points; once this is achieved, the system shuts down and exports the data.

6.2 Results and Discussion

The aim of this chapter is to determine the efficacy of GPR in generating surrogate models, and of BO in determining optima for given objective functions using experimental data. To achieve this, these methods are applied to the operation of a Hall thruster running argon, krypton, and xenon gas mixtures. This section initially demonstrates the application of BO to a total cost optimization of an industry-relevant mission. Then, surrogate models are assessed for accuracy, and the required number of data points is evaluated. Performance trends of the gas mixture data are examined using these models to determine the behavior of the propellant mixtures across the search space. Once models of the performance have been assessed, the total cost surrogate model is used to optimize another sample mission, including an analysis of how propellant storage technology influences that optimization. Finally, the models are used to gauge the effects of real-world market forces with changing launch costs, and propellant pricing.

The mission optimizations within this section use the three sample missions specified in Table 6.4. They represent real-world EP spacecraft that have been designed for a diverse range of missions. Mission 1 is an exploration probe, designed for analyzing near-Earth asteroids for mineral content [144, 145]. It is a relatively light satellite with a very high mission Δv and will be used as the test case for assessing the efficacy of BO. Mission 2 is a Very Low Earth Orbit (VLEO) satellite with a large Δv requirement for orbit raising. This mission will be used to assess the surrogate model optimization for both supercritical

and cryogenic tanks, and to study the effect of propellant cost factor on the optimization. Mission 3 is typical of a Low Earth Orbit (LEO) communications satellite used in large constellations, it has a low Δv requirement, and will be used to assess the effect of specific launch cost changes on the gas mixture optimization.

6.2.1 Bayesian Optimization of Missions

The goal of using Bayesian optimization is to rapidly investigate operational areas of interest, in this case, we are interested in minimizing mission cost. As described in section 6.1.1, the BO uses a surrogate model generated from an initial set of test points in the search space. An acquisition function then selects new test points, \mathbf{x} , with the aim of lowering uncertainty in operational areas of interest, and of finding the optimum operating point. This rapidly generates high confidence in the area of interest, rather than the ‘big picture’ given by the surrogate models of the full search space presented in the next section, which require a significantly higher number of test points, but do not necessarily have the sampling density around the optimum.

To demonstrate the efficacy of Bayesian optimization, it was applied to the cost optimization of mission 1, the asteroid-rendezvous probe, see Table 6.4. An initial surrogate model was formed using 77 test points within the 5-dimensional parameter space in Table 6.3, and at a constant power of 1 kW. These points were a combination of randomly sampled points, and a previous optimization of a different mission. The surrogate model can use any existing

Table 6.4: Specifications for example mission optimizations (missions 1 & 2) and analysis of cost factor and specific launch cost (Missions 2 & 3).

	Mission 1	Mission 2	Mission 2 Cryo	Mission 3
Mission Type	Exploration Probe	VLEO Satellite	VLEO Satellite	LEO Com. Sat.
m_{tankless} (kg)	115	300	300	300
Δv (m/s)	5000	4000	4000	600
Tank Type	Supercritical	Supercritical	Cryogenic	Supercritical
Cost Factor, α_p	1.0	0.75	0.75	20.0

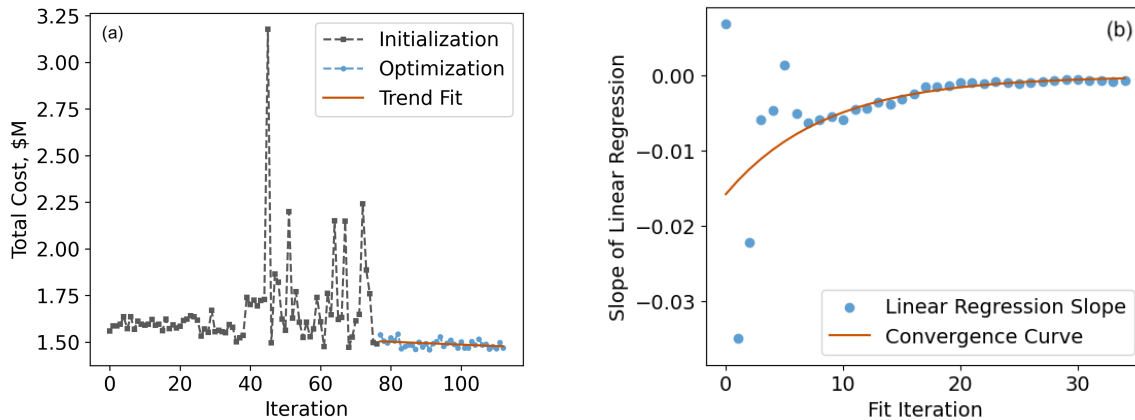


Figure 6.2: a) Performance of Bayesian minimization on the total cost after initialization on 77 randomly sampled points showing a trend in improvement of $-7.7 \times 10^4 \mathbf{x} + 1.56$. b) Convergence of towards the objective shown by the slope of linear fits to optimization data.

data within the search space as all data is converted to the desired mission specification.

Figure 6.2a shows the objective, total cost, for the 77 points used in the initialization, and the 36 iterations of the optimization. The initialization points show two clear regions, the first 38 points are from an optimization to a different mission, while points 39 to 77 are random samples from across the search space. The optimization of mission 1 shows a clear trend with an overall minimization of the total cost function. A linear fit to the optimization data gives a slope of $-\$77k$ per iteration as shown in Fig. 6.2a. Showing a consistent convergence towards the minimized total cost, the majority of the linear fits to subsets of the optimization trajectory data have negative slopes, shown in Fig. 6.2b. This shows the the BO very rapidly finds the area where the minima of total cost exists, and iterates in that region to reduce uncertainty in both the optimal operating parameters, \mathbf{x} , and total cost in that location, $f(\mathbf{x})$.

We can compare the results of this local optimization to the full data set of 575 points taken during this test campaign to determine the efficacy of BO in optimizing the mission, and ensure the correct operational area of interest was identified. As with the initialization set used for BO, the full data set is comprised of both random initialization points and

Table 6.5: Comparison of the minimized total cost for surrogate models of the full data set of 575 observations and the 36 BO observations.

Value at Optima	All Data	Bayesian Optimization
Total Cost (\$M)	1.421	1.419
Xenon Fraction	0.00	0.0141
Argon Fraction	0.00	0.114
Discharge Voltage (V)	320	323
Magnetic Field Strength (G)	141	151
Keeper Current (A)	0.276	0.617

optimizations to a range of mission specifications. This test can determine how well local, directed sampling can match the larger data set. Table 6.5 shows a comparison between the surrogate model from the BO data, to the model formed from the full data set. The difference between the reported minima for total cost is less than 0.2% at \$1.2k, well within the uncertainty of the full model. The operating points chosen by each model are close in xenon fraction, discharge voltage, and field strength. However they are significantly different in argon fraction and keeper current, for two different reasons. The keeper current varies between the points as the influence on the total cost is relatively low in this region. Querying the full model for this operating point with keeper currents of 0.276 and 0.617 A (the reported optima), gives only an \$800 difference between them. The argon fraction has a much larger influence on the optimization, however in the region of interest, the specific impulse remains relatively constant, and the additional launch cost due to the higher tankage fraction is balanced by the lower propellant cost. This can be visualized in Fig. 6.5, and the trends in performance will be discussed in section 6.2.2.

6.2.2 Surrogate Models

With the local accuracy of BO built GPR models demonstrated, the requirements for achieving good accuracy across the whole search space must be assessed. To calculate the dependence on sample density, a series of tests were run with an increasing number of samples,

n_s , used to train a surrogate model of thrust. At each value of n_s , 10 randomly selected comparison points, \mathbf{x}_c , with thrust values T_c , are removed from the pool of all test points \mathbf{x}_t . A model is then trained on n_s randomly selected from the remainder of the pool. As the comparison points are not used to train the model they are compared against, they can be considered independent. This model is then evaluated at each \mathbf{x}_c and the normalized absolute error $|(T_c - F(\mathbf{x}_t))/T_c|$, and the model's confidence interval from the Gaussian kernel in Eq. 6.2 are recorded. Bootstrapping is used with randomized resampling until 1000 \mathbf{x}_c evaluations have been completed. Figure 6.3 shows a comparison between the mean 95% confidence interval of the surrogate model and the normalized absolute error for models trained with 2 to 550 samples. As the number of samples increases, both the model's confidence interval and the mean error initially decrease before leveling off. The value they level off to is roughly equivalent to the uncertainty in thrust due to the accuracy and repeatability of the control system. The most significant source of error in the control system is the real time measurement of the anode current used to control the PID for anode gas flow. Error due to thrust measurement, and power supply control also contribute, albeit to a lesser extent. Increasing the number of samples beyond this point cannot improve accuracy beyond the uncertainty in the physical system. Therefore the 575 points measured during testing are sufficient to train a model which fully describes the performance of the system.

This is a relatively sparse data set for the utility it provides, particularly because new surrogate models with new objective functions can be calculated using the operating parameters, the measured thrust and propellant flow rates, and the mission specifications. This is the key difference between the more directed BO informed models, and the larger datasets used in this section. New objective functions or mission specifications can form a new model without the need to re-run the thruster, allowing a wide range of mission and operational analysis. To demonstrate the application of surrogate models to modeling and optimizing a range of objective functions, trends in the thrust, specific impulse, and total efficiency of gas mixtures will be analyzed across the fast, computer-controlled dimensions shown in Table 6.3. These models will give insight into the behavior of the thruster across the search




Figure 6.3: Evolution of the surrogate model for thrust with an increasing number of samples used to train the model. At each number of samples, the mean 95% confidence interval of the surrogate model, and the mean model error is evaluated for 1000 comparison points using randomized subsets of the data.

space and can be compared to previous investigations of gas mixtures. Additionally, these performance optimizations will be used to assess how the thruster is operating for the mission optimizations in section 6.2.3.

Results from this analysis are shown in Figures 6.4-6.7. For each of the objectives presented, the model is evaluated on a triangular grid of gas mixtures. This is not required for an optimum to be found from the surrogate model, but is presented to visualize the behavior of the thruster in a human-interpretable way. Each of the 1035 points on the grid represents a fixed gas mixture, the optimal free parameters (discharge voltage, magnetic field strength, and keeper current) are determined at each point using the Nelder-Mead optimization scheme on the surrogate model posterior [113]. The five ternary plots for each objective show the best optimization for each gas mixture point, the associated uncertainty in the model, and the discharge voltage, magnetic field strength, and keeper current.

Thrust Optimization

Figure 6.4 shows five ternary plots of the surrogate model trained to maximize thrust. A peak thrust of 71.2 mN is seen at 100% Xe and a discharge voltage of 241 V, field strength of 216 G, and keeper current of 0.59 A. The optimum thrust drops to 54.5 mN on pure Kr, and 36.7 mN on pure Ar. The optimal thrust on each of the pure propellants is towards the upper end of the expected ranges of thrust to power ratio observed in other Hall thrusters [146, 77, 147, 129, 105] which is the expected result of the optimization. The plot shows a smooth monotonic decrease in thrust away from the optima with no significant drops, indicating that there are no mode shifts with changing gas fractions at these operating points. The only exception to this is the noise around the high krypton, low argon region where two modes or operating points may have similar thrust causing the optimization to converge to different points. The standard deviation in the model, Fig. 6.4b, shows high confidence in the thrust model, particularly in the xenon rich region. The uncertainty increases in the region with high xenon and argon fractions as no optimizations converged in this space leaving only the initialization samples.

The operating conditions corresponding to the optimal thrust at each gas fraction are shown in Fig. 6.4c-e. Looking at the discharge voltage trend, we see that the discharge voltage is relatively low across the full range of gas mixtures, indicating a strong preference for high current densities in the thrust optimization. This trend agrees with previous studies of alternative propellants for Hall thruster operation which found that high current densities were key to improving mass utilization efficiency and global performance on krypton [131]. The optimal discharge voltage is highest in the region from 100% xenon to a mix of 60% xenon and 40% krypton, dropping off significantly with small fractions of argon, or larger fractions of krypton. This indicates that larger xenon fractions enhance the ionization of the krypton, reducing the reliance on higher current densities, and allowing for higher ion energies. Increasing xenon in a xenon/krypton mixture has previously shown enhancement in krypton ionization rates as xenon dominates upstream ionization, increasing the ionization

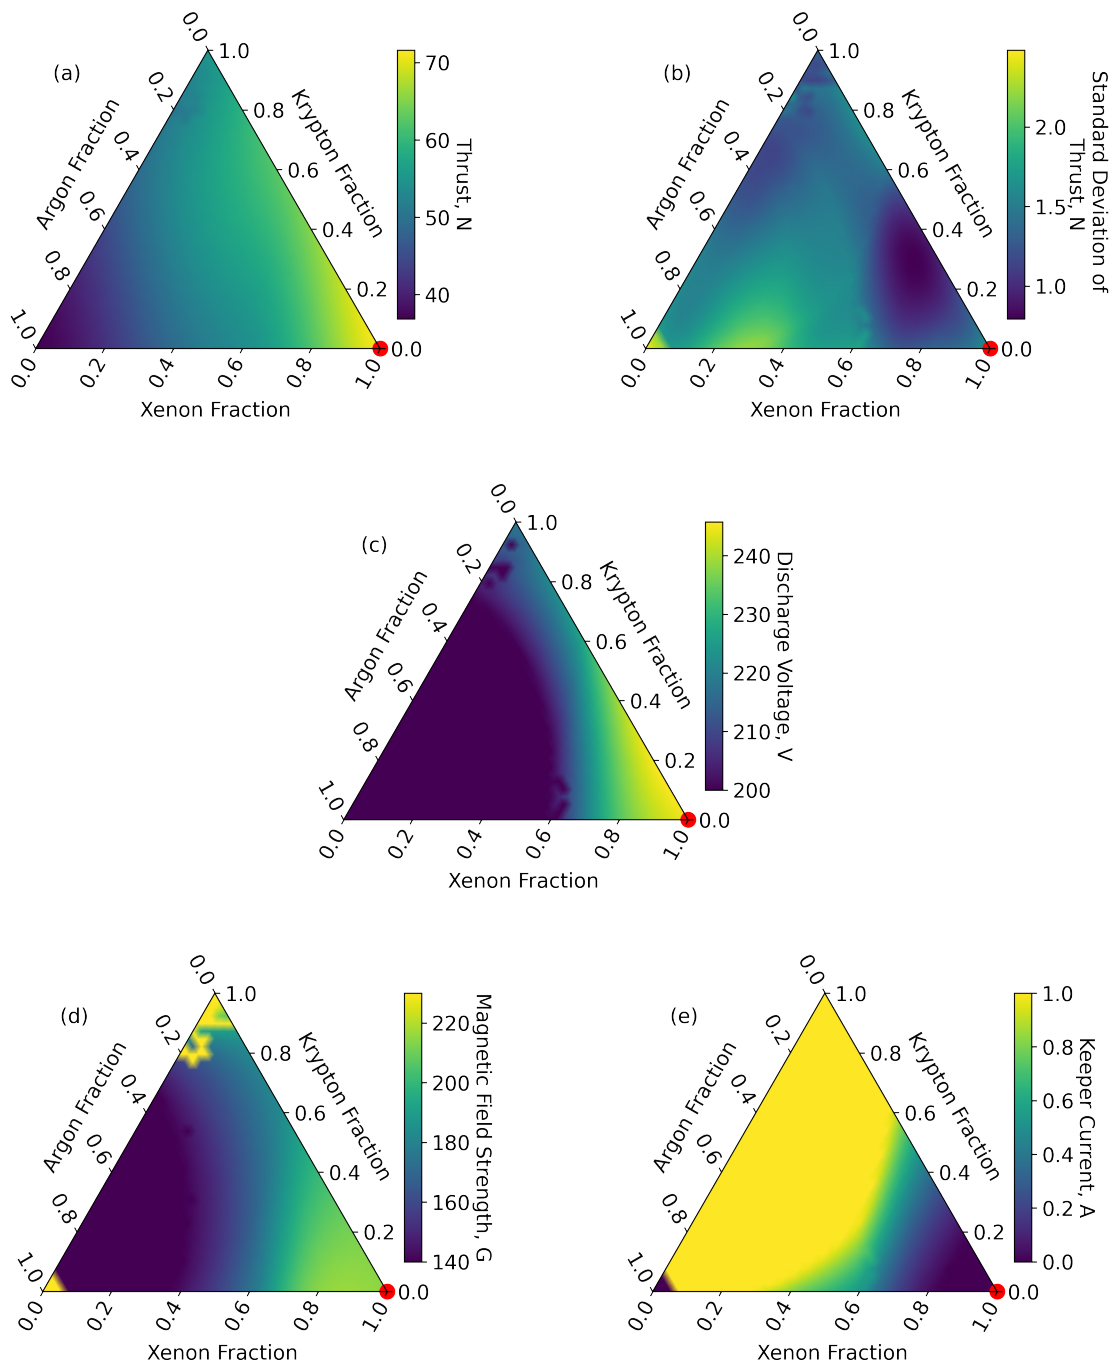


Figure 6.4: a) Optimized thrust for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized thrust at each point. c) The discharge voltage for optimized thrust at each point. d) The magnetic field strength for optimized thrust at each point. e) The keeper current for optimized thrust at each point.

rate of krypton in the downstream region [107].

The optimal magnetic field strength shown in Fig. 6.4d shows a dominant trend of increasing field with increasing average atomic mass, matching the increasing discharge voltage. The keeper current in Fig. 6.4e decreases with increasing xenon. However, in comparison to the discharge voltage and magnetic field strength, the influence of the keeper current is relatively small. The reason for the invariance with respect to krypton/argon balance could be attributed to the relatively consistent discharge voltage and magnetic field strength in the low xenon region.

Specific Impulse Optimization

The specific impulse optimization, see Fig. 6.5, will give a strong indication of the results of the mission optimizations as it determines the required propellant mass, see eqn. 6.6. As the specific launch cost is significantly higher than even the cost of xenon, the propellant mass will be a strong driver of the mission optimizations. As with the thrust optimization, the five ternary plots show the optimized specific impulse, the standard deviation in the model, and the discharge voltage, magnetic field strength, and keeper current at that optima. The specific impulse optimized to a mixture of 24.1% xenon and 75.9% krypton with a specific impulse of 1420 s at a discharge voltage of 348 V, a field strength of 230 G, and a keeper current of 0.38 A. For the pure propellants, the optimal specific impulses are: 1411 s for Xe, 1357 s for Kr, and 1123 s for Ar. As expected, the specific impulse optimized to high discharge voltages, as the increased potential accelerates ions to a higher velocity. Unlike the thrust optimization, three distinct regions are visible. These regions are indicative of different discharge modes of the plasma, and are made apparent by the sharp transitions consistent across the plots. The xenon rich region where the optima was found was observed to be in a Jet mode and is operating at a higher discharge voltage and field strength. As the argon and krypton fractions increase, there is a discontinuity in the specific impulse, and the optimal operating points. In this middle region, observed to be a Diffuse mode, the specific

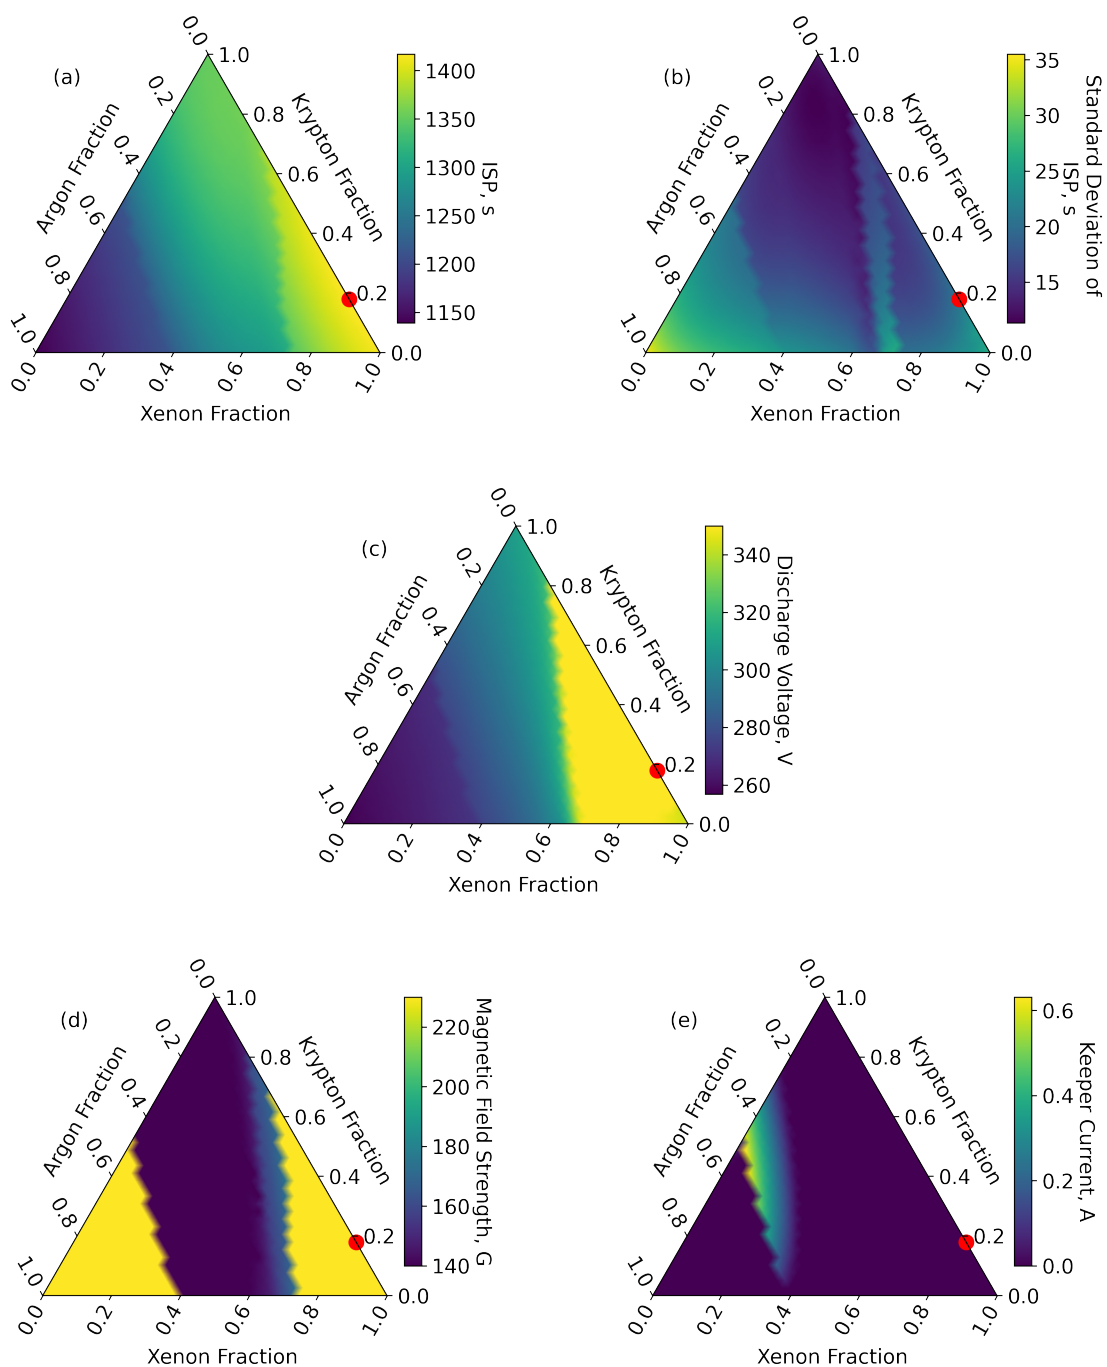


Figure 6.5: a) Optimized specific impulse for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized specific impulse at each point. c) The discharge voltage for optimized specific impulse at each point. d) The magnetic field strength for optimized specific impulse at each point. e) The keeper current for optimized specific impulse at each point. The discontinuities in the gradient show the locations of mode shifts in the Hall thruster discharge, indicating where the optimum discharge parameters change significantly.

impulse has optimized to a low field, but still relatively high discharge voltage. As the argon fraction increases further, there is another discontinuity where the potential drops and the field strength jumps to the upper bound. Argon, being more difficult to ionize, required a larger magnetic field strength but a smaller discharge voltage and therefore higher current density to maximize the specific impulse, which is reflected in the plot for higher argon gas fractions.

The optimization of specific impulse seen here with a significant fraction of krypton is similar to the results of Ref. [106], where a mixture of 26% argon and 74% krypton improved specific impulse over pure krypton propellant. While these are mixtures of different propellants, they indicate that the configuration of the thruster, particularly channel length, and power level used may favor propellant mixtures. The results of Ref. [129] showed a nearly constant specific impulse across the full range of tested xenon/krypton mixtures. This also agrees with the xenon/krypton mixtures demonstrated, where the specific impulse varies by only ~ 10 s in the regions either side of the mode transition. As previously observed on ACME, the Jet mode is dominant at higher powers, indicating that the high specific impulse region would expand with increasing power, resulting in a single mode covering the full xenon/krypton axis and therefore consistent specific impulse across the full length.

Total Efficiency Optimization

The total efficiency optimization gives a third comparison point which will be used to assess the mission optimizations and understand why they converge to particular operating points. The optimization seen in Fig. 6.6a shows a peak total efficiency of 0.442 while running on 100% xenon at a discharge voltage of 270 V, a field strength of 207 G, and a keeper current of 0.00 A. The optimum total efficiency drops to 0.329 on pure Kr, and 0.188 on pure Ar. The optimized efficiency with xenon, and the drop in efficiency to krypton, then argon, due to lower ionization cross section and higher ionization energies, are in line with other comparisons in literature [106, 148, 129, 105]. The uncertainty in the model is low across the high krypton and xenon regions with increasing uncertainty closer to argon. This increased

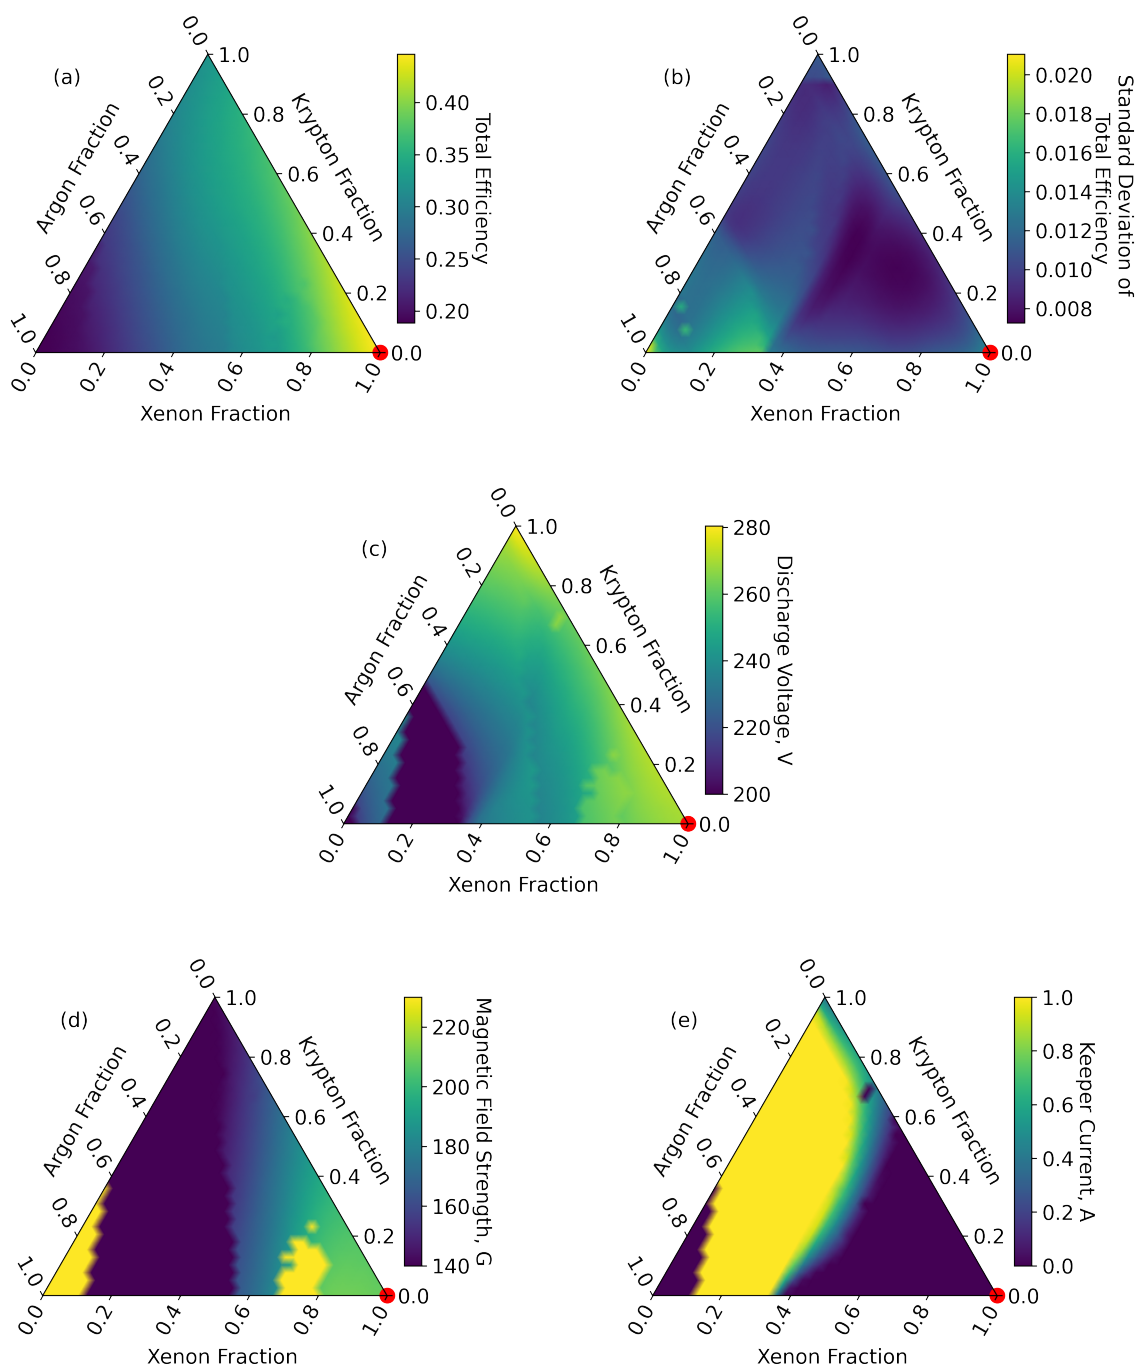


Figure 6.6: a) Optimized total efficiency for mixtures of argon, krypton, and xenon. b) The standard deviation of the optimized total efficiency at each point. c) The discharge voltage for optimized total efficiency at each point. d) The magnetic field strength for optimized total efficiency at each point. e) The keeper current for optimized total efficiency at each point. The discontinuities in the gradient show the locations of mode shifts in the Hall thruster discharge, indicating where the optimum discharge parameters change significantly.

uncertainty can be attributed to the operating points in this region optimizing to the upper or lower bounds of the discharge voltage, magnetic field strength, and keeper current. At the bounds, the surrogate model only has data on one side of the point for these dimensions, thus reducing the ability of the model to interpolate and increasing the uncertainty.

The efficiency drops significantly towards regions with high argon fractions, but has a relatively low gradient in the region near pure krypton. This indicates that mixtures of krypton with low fractions of argon may reduce propellant costs while mitigating performance drop. The plots show three mode regions which are particularly visible on the magnetic field plot, see Fig. 6.6d. The majority of gas mixtures show a single trend of decreasing magnetic field strength away from pure xenon, and discharge voltage decreasing away from the krypton/xenon boundary. Similar to the thrust optimization, the current density is high across this region; as previously discussed, this improves mass utilization on lighter propellants, improving total efficiency. Close to pure argon, the magnetic field strength jumps to the upper bounds, likely to reduce electron diffusion to the anode from the high neutral density. An additional, distinct region appears around a xenon fraction of 0.75 and argon fraction of 0.25. This region has a field increase of around 20 to 30 G, and slightly higher discharge voltage relative to the nearby mode. This may be indicative of a change in the convergence of the optimization at these points due to the sparse sampling in that region as it does not resemble the defined lines from other observed mode transitions.

With the exception of boundaries between different modes, small fractions of propellants were not observed to have significant impact on thrust, specific impulse, or total efficiency. The significant increase in current, reported in Ref. [132], and attributed to increased ionization due to small xenon fractions, was not observed. Trends in all three performance models showed relatively smooth monotonic improvements in performance, and good agreement with previous investigations of gas mixtures.

This indicates that there is no special advantage to the use of pure propellants, as they do not offer any measurable advantage when compared to the mixtures around them. The relatively high specific impulse seen with large krypton fractions, and the significantly lower

cost, suggest that this area will be favored by mission optimizations, as would be expected with current trends in real-world spacecraft propellant.

6.2.3 *Surrogate Model Mission Optimization*

With the performance across the gas mixtures having been assessed, the surrogate model will now be used with the more complex objective function of mission cost. This will demonstrate the surrogate model describing and optimizing this five-dimension system with complex interdependencies to a more convoluted objective function. The objective of total mission cost depends on the performance of the thruster, the mission specifications, and the operating point. Initially the system is optimized to the given mission specifications, then output constraints are added to assess the effect of minimum thrust requirements on the operating point. Finally the dependence on mission specifications will be investigated, demonstrating how the models can be used to explore the relationship between model inputs and resulting optima.

Mission Cost Optimization

The surrogate model for total cost is trained with the same data used for the performance optimizations, the objective is simply changed to total cost, and a set of mission specifications are provided. Two optimizations are run using mission 2, a small VLEO observation satellite, with m_{tankless} of 300 kg, that experiences significant atmospheric drag requiring the large Δv of 4000 m/s to maintain orbit. The optimization will be run with both supercritical or cryogenic tanks, see Table 6.4. The propellant price used reflects prices in 2016-2018 when propellant prices were relatively constant before the price fluctuations seen since 2020 [142]. The comparison of cryogenic tanks is included to see how advances in propellant storage technologies influence optimal propellant mixtures.

Figure 6.7a and b show the total cost optimization and associated uncertainty for mission 2 with supercritical propellant storage. It optimized to a mixture of 72.7% krypton and 27.3% xenon, a discharge voltage of 350 V, a field strength of 230 G, and a keeper current of 0.368 A,




Figure 6.7: a) Optimized total cost for mixtures of argon, krypton, and xenon for mission 2 (VLEO satellite). b) The standard deviation of the optimization for mission 2. c) A repeat optimization of mission 2 with cryogenic propellant tanks. d) The standard deviation of the optimization for with cryogenic tanks.

giving a tankage fraction of 0.45. The operating points are very close to the specific impulse optimization at the same gas fraction shown in Fig. 6.5. The significant xenon fraction is due to the specific launch cost being higher than the price of xenon, penalizing high propellant mass and tankage fraction in the optimization. The increased performance and decreased tankage fraction of xenon is therefore enough to counteract the increased propellant price. The uncertainty in total cost is low in the region of interest, giving good confidence in the model, and increases with the total cost further from the optima.

Figure 6.7c and d show the optimization and uncertainty for mission 2 with cryogenic propellant storage. Here the tankage fraction is very small, $< 1\%$ of the supercritical tankage fraction, and therefore has far less effect on the optimization. The cryogenic storage shifted the optimum significantly, with a mixture of 93.2% krypton and 6.8% argon, a discharge voltage of 311 V, a field strength of 140 G, and a keeper current of 0.427 A, giving a tankage fraction of 0.001. As with the supercritical storage, this operating point is very close to the same voltage and field strength as the specific impulse optimization for this gas fraction. As a demonstration of surrogate models, this shows how different mission profiles can be compared without taking additional data.

Output constraints can be used to add additional mission requirements into the optimization. For example, it is possible that a spacecraft could require a minimum level of thrust to complete required orbital maneuvers, therefore setting a constraint on the thrust. The results of our thrust optimization show that lighter propellants generally operate at smaller thrust levels (Fig. 6.4a). Therefore, it is possible that some mixture ratios or other operating parameters may not be capable of satisfying a minimum thrust requirement. We find that adding a minimum thrust to mission 2 moves the optimum towards heavier propellants, as they produce higher thrust at equivalent operating parameters. For example, a minimum thrust of 45 mN moves the optimum to a gas fraction of 63% krypton and 37% xenon for a supercritical propellant tank. A 50 mN results in a fraction of 54% krypton and 46% xenon. In each case, the discharge voltage, magnetic field strength, and keeper current remained constant. The invariance of these operating parameters when adding the thrust constraints

can be explained by the specific impulse optimization in Fig. 6.5a. These optima are all within the jet mode region, where the ISP optimizes to the upper limit of voltage and field strength, and the lower limit of keeper current. Moving to heavier propellants provides additional thrust while maintaining the other operating parameters to maximize specific impulse and therefore minimize launch mass.

Propellant Cost Sensitivity

The sample mission demonstrated the use of surrogate models to optimize thruster operation to a more complex objective function than pure performance. However, the surrogate models can further breakdown the different drivers behind these optimizations and allow us to study how external factors can affect the mission optimizations. The most significant driver behind the move to lighter propellants is the limited and variable supply of xenon. The optimizations in section 6.2.3 were run with propellant prices prior to the fluctuation seen in 2021-2023, where krypton and xenon prices increased by a factor of ~ 20 . We will now investigate how changing propellant prices influence mission optimization.

Figure 6.8 shows the optimization for total cost of mission 2 with a variable propellant cost factor ranging from 0.01 to 1000. Recall that the cost factor represents the increase in xenon and krypton price to model changes in supply and demand (see Eq. 6.12). Both the supercritical and cryogenic tanks are shown with the solid, and dashed lines, respectively. While the propellant cost factor is very low, propellant is essentially free, and the optimal solution is found to correspond to thruster operating points that maximize specific impulse, and therefore minimize launch mass. This gives an operating point for both tank types close to the optima for specific impulse seen in section 6.2.2. As the cost of krypton and xenon increase, the optimal point for the supercritical tank slowly increases in krypton fraction before fully switching at a propellant cost factor of 0.45. At the switchover point, the propellant cost decreases significantly and the launch cost increases by the same amount.

The operating point maintains a high discharge voltage until a propellant cost factor of 40 where the argon fraction begins increasing and the optimal operating point moves to a

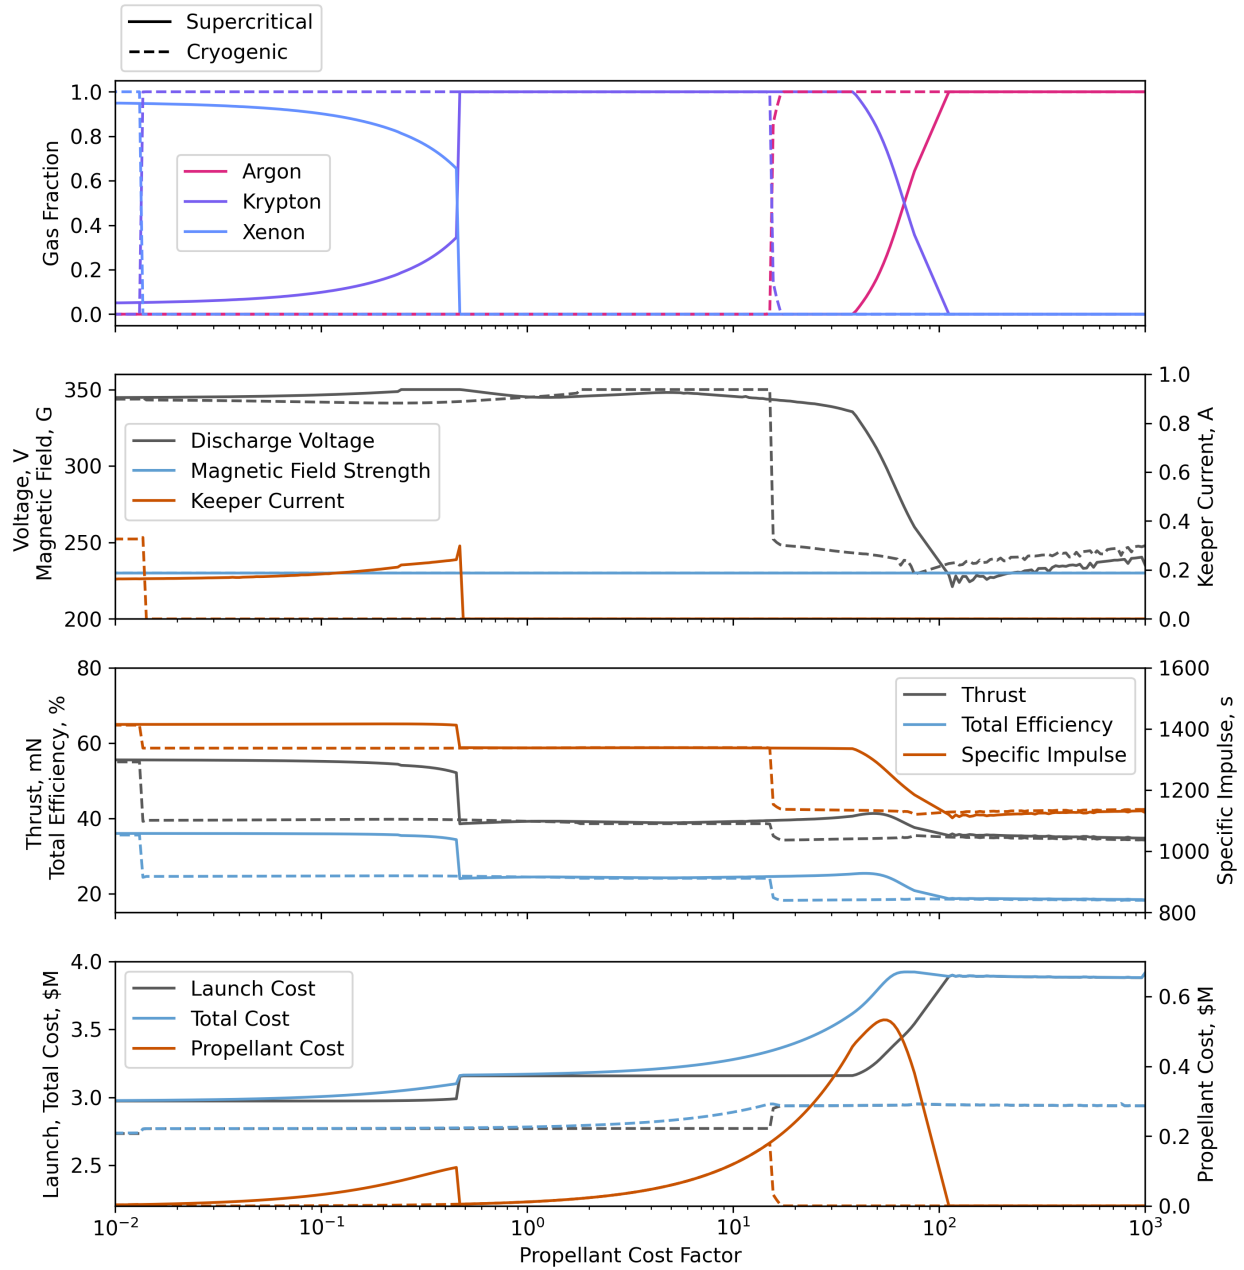


Figure 6.8: Evolution of the optimal operating parameters for a minimized total cost for mission 2’s specifications as a function of the increasing propellant cost factor of krypton and xenon. Data is from the evaluation of the surrogate model of thruster total cost.

higher current density as seen in section 6.2.2. The change in optimal mix from pure krypton to pure argon is a balance of launch cost and propellant cost. Once the optimal mix is pure argon, the total cost remains consistent as the cost factor only affects krypton and xenon. The total cost slightly overshoots during the transition to argon due to the Nelder-Mead optimization at these points converging with a very slight error which is amplified by the high propellant prices.

The cryogenic tankage fraction is very small, causing the optimal propellant mix to switch to pure krypton much earlier than with the supercritical tank, at a propellant cost factor of 0.014. The increase in total mission cost at this point is around 1% due to the small increase in propellant mass being matched by the cheaper propellant. As the propellant cost factor continues to increase past 15, the optimal propellant abruptly switches to pure argon, and the discharge voltage decreases from the high specific impulse operating point into the lower voltage, higher current density mode preferred by argon. Once again, the total cost remains constant past this point.

The thrust, efficiency, and specific impulse, along with the operating parameters, show that both tank types run at the same three operating points across the full range. The performance drops significantly as the optimization moves to each successively lighter gas. The higher tankage fractions of the supercritical propellant storage results in larger regions of cost factor optimizing to mixtures of gasses. This is because the increasing propellant cost is balanced with the decreasing cost to launch both the tank and the propellant. For the cryogenic tank, the switch to pure argon occurs before the historical peak of ~ 20 , suggesting that, within the limits of both the thruster surrogate and mission cost models, the cheapest mission in 2021 would have been a fully argon satellite with a cryogenic tank.

Launch Cost Sensitivity

While cost factor is the main driver of propellant choice, the strongest driver of total cost is the specific launch cost. Therefore, we now examine the influence of specific launch cost on propellant mixture optimization. The optimization will be applied to mission 3, a LEO

communication satellite, similar to some current and planned satellite constellations. A propellant cost factor of 20 is used to highlight the state of the propellant market in 2021 that led to the decision to use lighter propellants. In all previous models, the specific launch cost has been set at SpaceX’s advertised 6500\$/kg ride-share price [139]. Here, we extend the range to a maximum specific launch cost of 10,000 \$/kg, corresponding to values found in the early 2000s [149].

Starting at the maximum launch cost shown in Fig. 6.9, we see that the propellant cost factor causes the optimum propellant mix to be pure krypton, with an operating point to maximize the specific impulse. As the price decreases to the current ride-share price (indicated by the vertical dotted line), the optimal operating point remains constant. As the price decreases below 6000\$/kg, the optimal discharge voltage begins to decrease slightly, and at 5000\$/kg, the optimal argon fraction begins to rise. As the argon fraction increases, we see the voltage continue to decrease, moving to the higher current density mode seen in the specific impulse optimizations in section 6.2.2. At a specific launch cost of 2800\$/kg, the optimal gas mix is an even mixture of krypton and argon, and at 2100\$/kg the optimal propellant is pure argon.

We can compare these thresholds to reported launch costs for similar real-world missions. The public price of a full Falcon 9 was 2700\$/kg in 2010 [149], with internal pricing estimated to be as low as 1000\$/kg in 2024 [150]. At these price ranges, the results of the model align with the observed decisions made across several satellites and constellations with moves to lighter, cheaper propellants due to the potential for future market fluctuations.

6.2.4 Adaptability and Limitations of Gas Mixture Optimization

The results seen in the previous sections indicate that gas mixtures can be beneficial to mission cost optimization and, at the right operating points, provide a relatively continuous range of performance. Due to the constraints of the laboratory power supplies, the bounds of the optimization were set to a discharge voltage of 200 to 350 V, resulting in a channel current density between 0.17 and 0.3 A/cm². It is clear from both the thrust and efficiency




Figure 6.9: Evolution of the optimal operating parameters for a minimized total cost for mission 2's specifications as a function of the increasing specific Launch Cost. Data is from the evaluation of the surrogate model of thruster total cost. SpaceX's ride-share price is indicated by the vertical dotted line.

models, which optimized to the upper current density limit for much of the test area, that even higher current densities would further improve performance. Demonstrations of the H9 Hall thruster on krypton at high current densities showed improvements to both efficiency and specific impulse with increasing current density, a trend which continued to the tested limit of 0.45 A/cm^2 [131]. This result suggests that there is increased performance that could be extracted for this setup of ACME.

The performance comparisons presented were demonstrated with, and therefore representative of, a single physical configuration, of a particular thruster (ACME), at a single power level. They therefore do not account for changes in thruster design that would improve performance in some regions, particularly on lighter propellants. This could be investigated with the adaptability of ACME, where the slow dimensions (pole position, anode position, and cathode position) could be added as dimensions into the surrogate models, allowing for trends to be observed across physical thruster changes. Adding this would incorporate thruster design into the modeling of gas mixtures and give insight into the optimization of lightweight propellants, and the relationship between the channel depth, field strength, discharge voltage, and field inclination. Setups of ACME with a range of inclined fields have demonstrated reduced plume divergence and improved performance on krypton, counteracting the increased divergence seen from the ionization region moving towards the channel exit [107, 95]. As demonstrated in section 6.2.3, decreasing launch costs give justification for focused development of an argon, or predominantly argon, specific Hall thruster. As ACME was not designed, or specifically setup, to run on argon, there are significant performance gains that could be achieved in changes in the design, higher current densities, and higher power [151]. Even small performance improvements could yield significant cost savings, particularly in a volatile and capacity constrained propellant market.

While the cost model is simplified, it demonstrates the technique's ability for finding optimal operating parameters, and describing the effects of gas mixtures on performance and total cost. More complex cost and mass models, adding scaling for thermal, power and propulsion systems, or even full satellite scaling can be all be applied to the existing models

of performance, however, such additions require detailed subsystem scaling are therefore beyond the scope of this work.

6.3 Chapter Conclusions

In order to more efficiently characterize and develop electric propulsion systems, a method capable of creating generalizable models with Gaussian Process Regression and running real-time Bayesian optimization has been developed. To demonstrate these methods, they were applied to the problem of modeling and optimizing Hall thruster operation on argon, krypton, and xenon gas mixtures with a series of industry-relevant sample missions. Real-time Bayesian optimization of an asteroid-rendezvous mission converged rapidly on a global minima, with a mission cost within 0.2% of the cost found on a larger surrogate model. This demonstrated both the ability of the optimization methodology to quickly find the optimum in five dimensions, and the localized accuracy of the surrogate models created from optimization data. Surrogate models of performance and sample mission cost accurately captured the operation of a Hall thruster in five dimensions, revealing the ability to capture mode changes across gas mixtures. For the sample missions explored, surrogate models demonstrated the capability of finding optimal gas mixtures and operational parameters to minimize total mission cost for a given cost model. Investigating the effects of the volatile propellant market and decreasing launch costs showed trends consistent with real world changes to spacecraft propellants, particularly those in large constellations. Notably, the methods developed here were able to analyze an electric thruster in five dimensions using fewer experimental data points than would be required using traditional experimental methods in just three dimensions. Such a tool has the potential to drastically increase the scope of experimental test campaigns by efficiently modeling complex inter-dependencies in higher dimensional spaces.

The investigation was invited, and has been submitted, to be published in the Journal of Electric Propulsion's "Special Issue of the 39th International Electric Propulsion Conference".

Chapter 7

CONCLUSION

Experimental characterization of Hall thrusters and other electric propulsion systems is a fundamental cornerstone in their research, development, and testing. Thus the advancement in experimental tools and methods is of the utmost importance. This dissertation covers the four primary approaches aimed to advance techniques in electric propulsion research, to more effectively, efficiently, and expeditiously characterize and develop electric propulsion systems. The methods and key findings are summarized in this section.

The first approach used to advance electric propulsion testing was the development of the The Adaptive-field Central-cathode Magnetically-shielded Electric (ACME) Hall thruster. The design of ACME is unique in the Hall thruster world. It was designed as a dedicated research device, to be rapidly reconfigured to a wide range of experimental setups without disassembly or rebuilding. The variable pole positions of the thruster were designed to study the influence of magnetic field structure on thruster behavior and performance. To achieve this, the thruster's magnetic circuits are adjustable, not just through current in the inner and outer electromagnets, but by physically moving the inner and outer poles relative to each other. The design of the poles and magnetic circuit ensured the shape of the shielding on each of the poles remains consistent. The initial test campaigns presented were used to inform design iterations, and demonstrated a range of plume shapes from highly divergent to over-focused. The evolution of the thruster through three design iterations has expanded its capabilities and remedied its most significant deficiencies, resulting in a highly capable research thruster for SPACE Lab. The ACME Hall thruster has played a key role in the research presented in this dissertation, providing a versatile platform that could be adapted for each of the test campaigns.

The ACME V2 test campaign revealed that the combined efficiencies of plume diagnostics likely weren't sufficient to fully characterize the most divergent plumes. While this was not conclusive, methods for improved plume characterization were investigated. The focus of this investigation was the Retarding Potential Analyzer which is traditionally located on the centerline of the thruster. Previous studies have shown the Ion Energy Distribution Function measured by a RPA can exhibit a strong angular dependence [87, 71]. Therefore this study aimed to investigate the efficacy of centerline and angularly resolved IEDFs in describing the operation of electric propulsion systems with varying plume structures. To assess these methods, the ACME Hall thruster was operated at a range of pole positions ranging from highly divergent to over-focused. At each of these pole positions the performance of the thruster was characterized with both plume diagnostics and a thrust stand. While the standard method of taking the peak of the IEDF on the centerline of the plume was sufficient to describe the voltage utilization while the plume was well focused. Only the integrated IEDF was sufficient to fully characterize the plume at all pole positions. Integrating the angularly resolved IEDF is therefore clearly beneficial in the study of divergent plumes or novel thruster concepts where the structure of the plume is unknown. Improving the method with the use of a collimated RPA is suggested as it would minimize the broadening effect and decrease the underestimation of the voltage utilization efficiency [98, 71]. While not necessary in most operational Hall or Gridded ion thrusters, current research into alternative lightweight or molecular propellants, or the recent push to air-breathing EP, may require such methods due to increased plume divergence and shifts in the ionization and acceleration regions of multi-species propellants [105, 106, 107, 108, 109, 110, 111]. These propellants also require increased flow and may not be pumped as effectively, and therefore may benefit from the improved CEX correction over the full plume, without having to take measurements at multiple background pressures.

This test campaign found a peak efficiency at the -4 mm pole position, driven by improved current utilization efficiency and good divergence utilization efficiency. The significant decrease in flow required at negative pole positions is of particular interest for future inves-

tigations. It was hypothesized that this was due to the cathode positioned closer to the higher density plasma combined with the lower magnetic fields required on the inner pole improving cathode coupling. Combined with the improvements to the divergence and current utilization efficiency, there may be gains that can be achieved with other Hall thrusters with centrally mounted cathodes.

The second major research direction was in the application of automation and machine learning. This was split into two main approaches that aimed to expedite thruster development in different ways. The first approach was rapid thruster-in-the-loop optimization, with a primary aim of advancing automation to drastically reduce the required test time for electric propulsion test campaigns. This was motivated by the decreasing costs for access-to-space and the rise in small satellite constellations which have driven the need for a greater number of systems that are cheaper to design, test, and manufacture. The traditional method of characterizing thruster performance across parameter maps is costly in propellant and facility use, and development time. Therefore, a rapid thrust measurement method was developed that significantly reduces the time required to characterize a thruster setpoint. This was combined with full end-to-end computer control and real-time thruster-in-the-loop optimization using Nelder-Mead and Powell optimization schemes.

Two-dimensional Nelder-Mead and Powell optimization schemes were able to converge rapidly to maxima in total efficiency and specific impulse in under 15 test points, or just 11 minutes. Powell's method remained effective in five dimensions with convergence in under 30 test points. Automated optimization in these higher dimensional spaces can expedite many stages of thruster development by rapidly finding key operational areas of interest. The demonstrated plug-and-play nature of the objective functions can further enhance the optimizer's capabilities allowing optimization to any mathematically definable thruster goal.

The second approach to the application of automation and machine learning continued real-time optimization and added more complex modeling capabilities. However, in contrast to the goal of the 'rapid' method, the goal of this approach to machine learning was to characterize more complex systems allowing for greater analysis capabilities. This resulted

in a highly adaptable set of methods, capable of creating surrogate models of any quantifiable aspect of thruster operation using Gaussian Process Regression, and of optimizing that aspect of thruster operation in real time using Bayesian Optimization. These methods were demonstrated by modeling and optimizing the operation of ACME on argon, krypton, and xenon gas mixtures with a series of industry-relevant sample missions. The study of gas mixtures added additional dimensionality to the optimization of thruster performance, resulting in a prohibitively large search space for mapping using conventional approaches, making it an ideal test case for these methods. Due to the added complications of mixed propellant, the control system was improved to allow for startup on any arbitrary mixture of gasses to ensure accurate mixtures. ACME demonstrated over 1500 successful automated restarts during this campaign with over 1000 distinct propellant mixes, another significant achievement outside the primary goal.

Bayesian optimization was demonstrated in real-time on an a sample asteroid-rendezvous mission which converged rapidly on a global minima for mission cost. The surrogate model generated with the optimization data was within 0.2% of the cost found on a larger surrogate model showing excellent localized accuracy. The larger global optimization of both performance and sample mission cost surrogate models accurately captured the operation of ACME in five dimensions. The mission cost surrogate models demonstrated the capability of finding optimal gas mixtures and operational parameters to minimize total mission cost for a given cost model. The final demonstration of the system showed the capability to relate model variables with complex inter-dependencies by optimizing for the output variable across the full range of the input parameters. This was done by investigating the effects of the fluctuating propellant costs due to variable and limited supply of xenon and krypton and of decreasing access-to-space costs. The observed trends were consistent with real world changes to spacecraft propellants, particularly those in large constellations. The combination of surrogate models, Bayesian Optimization and automated test methods allows for more efficient characterization of electric propulsion systems while improving the analysis of complex inter-dependent variables.

Surrogate models can enhance how we research, develop, and use thruster systems. Planned future work includes: the study of mode behavior; creating a digital twin to flight thrusters; and optimizing design or characterizing specific thruster behavior using the slow dimensions of ACME.

While the methods presented in this dissertation are diverse, they are aligned in their goal to advance techniques in Hall thruster research and development.

7.1 *Peer-Reviewed Publications*

The peer-reviewed publications resulting from this PhD are itemized below:

- Influence of Field Topology on Magnetically Shielded Hall Thruster Plume Divergence, 2022, IEEE Aerospace Conference.
- Experimental Investigation of High Velocity Neutral Flow Interaction with a Magnetized Plasma, 2022, Plasma Sources Science and Technology.
- Rapid Thruster-in-the-Loop Optimization for Hall Thrusters, 2025, Journal of Electric Propulsion.
- Hall Thruster Plume Characterization with Spatially Resolved Ion Energy Distribution Measurements, submitted, Journal of Propulsion and Power.
- Surrogate Modeling and Real-Time Optimization of Propellant Mixtures, submitted, Journal of Electric Propulsion.

BIBLIOGRAPHY

- [1] Ernst Stuhlinger. Ion Propulsion for Space Flight, 1964.
- [2] Dan M. Goebel and Ira Katz. *Fundamental of Electric Propulsion: Ion and Hall Thrusters*. Wiley, 2008.
- [3] Stéphane Mazouffre. Electric propulsion for satellites and spacecraft: Established technologies and novel approaches. *Plasma Sources Science and Technology*, 25(3), 2016.
- [4] K. Holste, P. Dietz, S. Scharmann, K. Keil, T. Henning, D. Zschätzsch, M. Reitemeyer, B. Nauschütt, F. Kiefer, F. Kunze, J. Zorn, C. Heiliger, N. Joshi, U. Probst, R. Thüringer, C. Volkmar, D. Packan, S. Peterschmitt, K. T. Brinkmann, H. G. Zanunick, M. H. Thoma, M. Kretschmer, H. J. Leiter, S. Schippers, K. Hannemann, and P. J. Klar. Ion thrusters for electric propulsion: Scientific issues developing a niche technology into a game changer. *Review of Scientific Instruments*, 91(6), 2020.
- [5] Dan M. Goebel, Richard R. Hofer, Ioannis G. Mikellides, Ira Katz, James E. Polk, and Brandon N. Dotson. Conducting wall hall thrusters. *IEEE Transactions on Plasma Science*, 43(1):118–126, 2015.
- [6] S. Barral, K. Makowski, Z. Peradzyński, N. Gascon, and M. Dudeck. Wall material effects in stationary plasma thrusters. II. Near-wall and in-wall conductivity. *Physics of Plasmas*, 10(10):4137–4152, 2003.
- [7] N. Gascon, M. Dudeck, and S. Barral. Wall material effects in stationary plasma thrusters. I. Parametric studies of an SPT-100. *Physics of Plasmas*, 10(10):4123–4136, 2003.
- [8] Yevgeny Raitses, J Ashkenazy, and G Appelbaum. Experimental investigation of the effect of channel material on Hall thruster characteristics. *25th International Electric Propulsion Conference, Cleveland, Ohio, The United States of America, August 24-28*, (95-06):1–5, 1997.
- [9] Vladimir Kim, Kn Kozubsky, Vjacheslav M Murashko, and Av Semenkin. History of the Hall Thrusters Development in USSR. *30th International Electric Propulsion Conference*, 2007.

- [10] W. Hoskins, R. Cassady, O. Morgan, R. Myers, F. Wilson, D. King, and K. DeGrys. 30 Years of Electric Propulsion Flight Experience at Aerojet Rocketdyne. *International Electric Propulsion Conference*, 2013.
- [11] Kristi De Grys, Alex Mathers, Ben Welander, and Vadim Khayms. Demonstration of 10,400 hours of operation on a 4.5 kW qualification model Hall thruster. *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (July), 2010.
- [12] Ioannis G. Mikellides, Ira Katz, Richard R. Hofer, Dan M. Goebel, Kristi De Grys, and Alex Mathers. Magnetic shielding of the acceleration channel walls in a long-life Hall thruster. *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (July):1–25, 2010.
- [13] Gil Denis, Didier Alary, Xavier Pasco, Nathalie Pisot, Delphine Texier, and Sandrine Toulza. From new space to big space: How commercial space dream is becoming a reality. *Acta Astronautica*, 166(June 2019):431–443, 2020.
- [14] Kenza Bousedra. Downstream Space Activities in the New Space Era: Paradigm Shift and Evaluation Challenges. *Space Policy*, 64, 2023.
- [15] Jeff Foust. SpaceX launches first upgraded Starlink satellites, 2023.
- [16] A I Morozov and V V Savelyev. *Fundamentals of Stationary Plasma Thruster Theory*, pages 203–391. Springer US, Boston, MA, 2000.
- [17] Richard R. Hofer, Dan M. Goebel, Ioannis G. Mikellides, and Ira Katz. Design of a laboratory Hall thruster with magnetically shielded channel walls, Phase II: Experiments. *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit 2012*, 2012.
- [18] Peter Y. Peterson, Hani Kamhawi, Wensheng Huang, John Yim, Daniel Herman, George Williams, James Gilland, and Richard Hofer. NASA HERMeS hall thruster electrical configuration characterization. *52nd AIAA/SAE/ASEE Joint Propulsion Conference, 2016*, pages 1–24, 2016.
- [19] Dan M. Goebel, Benjamin A. Jorns, Richard R. Hofer, Ioannis G. Mikellides, and Ira Katz. Pole-piece interactions with the plasma in a magnetically shielded hall thruster. *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference 2014*, (June 2016), 2014.
- [20] Alejandro Lopez Ortega, Ioannis G. Mikellides, Michael J. Sekerak, and Benjamin A. Jorns. Plasma simulations in 2-D (r-z) geometry for the assessment of pole erosion in a magnetically shielded Hall thruster. *Journal of Applied Physics*, 125(3), 2019.

- [21] James E. Polk, Robert B. Lobbia, Arthur Barriault, Pablo Guerrero, Ioannis G. Mikellides, and Alejandro Lopez Ortega. Inner Front Pole Cover Erosion in the 12.5 kW HERMeS Hall Thruster Over a Range of Operating Conditions. *35th International Electric Propulsion Conference*, pages IEPC–2017–409, 2017.
- [22] I. D. Boyd. Review of Hall Thruster Plume Modeling. *Journal of Spacecraft and Rockets*, 38(3):381–387, 2001.
- [23] Ioannis G. Mikellides, Ira Katz, Richard R. Hofer, and Dan M. Goebel. Magnetic shielding of a laboratory Hall thruster. I. Theory and validation. *Journal of Applied Physics*, 115(4), 2014.
- [24] Richard R. Hofer, Peter Y. Peterson, and Alec D. Gallimore. Optimization of Hall thruster magnetic field topography. *IEEE International Conference on Plasma Science*, (6):240, 2000.
- [25] Richard Robert Hofer and Robert S Jankovsky. The Influence of Current Density and Magnetic Field Topography in Optimizing the Performance , Divergence , and Plasma Oscillations of High Specific Impulse Hall Thrusters. *International Electric Propulsion Conference*, (December):142, 2003.
- [26] Yushi Hamada, Junhwi Bak, Rei Kawashima, Hiroyuki Koizumi, Kimiya Komurasaki, Naoji Yamamoto, Yusuke Egawa, Ikkoh Funaki, Shigeyasu Iihara, Shinatora Cho, Kenichi Kubota, Hiroki Watanabe, Kenji Fuchigami, Yosuke Tashiro, Yuya Takahata, Tetsuo Kakuma, Yusuke Furukubo, and Hirokazu Tahara. Hall thruster development for Japanese space propulsion programs. *Transactions of the Japan Society for Aeronautical and Space Sciences*, 60(5):320–326, 2017.
- [27] Jesse A. Linnell and Alec D. Gallimore. Internal plasma potential measurements of a Hall thruster using plasma lens focusing. *Physics of Plasmas*, 13(10), 2006.
- [28] Hirokazu Tahara, Takashi Fujioka, Takahisa Kitano, Atsushi Shirasaki, Takao Yoshikawa, Kenji Fuchigami, Fujio Iinoya, and Fumihiko Ueno. Optimization on magnetic field and acceleration channel for low power Hall thrusters. *Proc. of the 28th International Electronic Propulsion Conference*, 2003.
- [29] Lou Grimaud, Julien Vaudolon, Stéphane Mazouffre, and Claude Boniface. Design and characterization of a 200W Hall thruster in “magnetic shielding” configuration. *52nd AIAA/SAE/ASEE Joint Propulsion Conference, 2016*, 2016.
- [30] Daniel Pérez-Grande, Pablo Fajardo, and Eduardo Ahedo. Evaluation of erosion reduction mechanisms in Hall effect thrusters. *34th International Electric Propulsion Conference (Kobe, Japan)*, (September):IEPC–2015–280/ISTS–2015–b–280, 2015.

- [31] Ryan W. Conversano, Dan M. Goebel, Richard R. Hofer, and Nitin Arora. Performance enhancement of a long-life, low-power hall thruster for deep-space smallsats. *IEEE Aerospace Conference Proceedings*, 2017.
- [32] Ryan W. Conversano, Dan M. Goebel, Richard R. Hofer, Ioannis G. Mikellides, and Richard E. Wirz. Performance analysis of a low-power magnetically shielded hall thruster: Experiments. *Journal of Propulsion and Power*, 33(4):975–983, 2017.
- [33] Ryan W. Conversano, Dan M. Goebel, Ioannis G. Mikellides, Richard R. Hofer, and Richard E. Wirz. Performance analysis of a low-power magnetically shielded hall thruster: Computational modeling. *Journal of Propulsion and Power*, 33(4):992–1001, 2017.
- [34] Ryan W Conversano, Robert B Lobbia, Kimberlyn C Tilley, Dan M Goebel, Sean W Reilly, Ioannis G Mikellides, and Richard Robert Hofer. Development and Initial Performance Testing of a Low-Power Magnetically Shielded Hall Thruster with an Internally-Mounted Hollow Cathode. *35th International Electric Propulsion Conference*, pages IEPC–2017–64, 2017.
- [35] Zhaoyu WANG, Hong LI, Chao ZHONG, Yanlin HU, Yongjie DING, Liqiu WEI, and Daren YU. Matching characteristics of magnetic field configuration and chamfered channel wall in a magnetically shielded Hall thruster. *Plasma Science and Technology*, 23(10):104008, 2021.
- [36] Jinwen Liu, Chao Zhong, Hong Li, Yongjie Ding, Liqiu Wei, and Daren Yu. The influence of the inclination of strong magnetic field lines on the performance and plume divergence of a magnetically shielded Hall thruster. *Journal of Physics D: Applied Physics*, 53(24), 2020.
- [37] Daren Yu, Jie Li, Hong Li, Yong Li, and Binhao Jiang. Experimental study on the effects of magnetic field configuration near the channel exit on the plume divergence of hall thrusters. *Plasma Science and Technology*, 11(6):714–720, 2009.
- [38] Yegor V. Plyashkov, Andrey A. Shagayda, Dmitrii A. Kravchenko, Alexander S. Lovtsov, and Fedor D. Ratnikov. On Scaling of Hall-Effect Thrusters Using Neural Nets. *Journal of Propulsion and Power*, 38(6):935–944, 2022.
- [39] Maximilian Mörtl, Aaron Knoll, Vinay Williams, Peter Shaw, Vasileios Argyriou, Jack Zamattio, and Luca Pugliese. Enhancing Hall Effect Thruster Simulations with Deep Recurrent Networks. *International Electric Propulsion Conference*, 2019.

- [40] Benjamin Jorns. Predictive, Data-Driven Model for the Anomalous Electron Collision Frequency in a Hall Effect Thruster. *Plasma Sources Science and Technology*, 27(10), 2018.
- [41] Hirotaka Fuchigami, Yusuke Egawa, Taichi Morita, and Naoji Yamamoto. Prediction of the Thruster Performance in Hall Thrusters Using Neural Network. In *International Electric Propulsion Conference*, 2017.
- [42] Masato Kawazu, Naoji Yamamoto, Masatoshi Chono, Hirotaka Fuchigami, and Taichi Morita. Prediction of Thruster Performance in Hall Thrusters Using Neural Network with Auto Encoder. *Transactions of the Japan Society for Aeronautical and Space Sciences, Aerospace Technology Japan*, 19(5):760–765, 2021.
- [43] Sridevi Bhat. Prediction of Liner Erosion and Life Estimation of Stationary Plasma Thrusters using Machine Learning. *International Electric Propulsion Conference*, 2019.
- [44] K V Korolev. Hall Effect Thruster Design via Deep Neural Network for Additive Manufacturing. 2023.
- [45] William Hurley, Thomas Marks, Alex Gorodetsky, and Benjamin A. Jorns. Application of Bayesian Inference to Develop an Air-Core Magnetic Circuit for a Magnetically Shielded Hall Thruster. *International Electric Propulsion Conference*, 2022.
- [46] Benjamin N. Wachs and Benjamin A. Jorns. Optimization of an ECR Thruster using Single, Two Frequency, and Pulsed Waveforms. *AIAA Propulsion and Energy Forum*, 2021.
- [47] Hani Kamhawi, David Manzella, Luis Pinero, Thomas Haag, Alex Mathers, and Howard Liles. In-Space Propulsion High Voltage Hall Accelerator Development Project Overview. In *Joint Propulsion Conference*, number August, 2009.
- [48] Dan Lev, Raanan Eytan, Gal Alon, Abraham Warshavsky, Leonid Appel, Alexander Kapulkin, and Maxim Rubanovych. CAM200 Hall thruster - Development overview. *Proceedings of the International Astronautical Congress*, pages 7516–7520, 2015.
- [49] Peter Y. Peterson, Hani Kamhawi, Wensheng Huang, John Yim, Daniel Herman, George Williams, James Gilland, and Richard Hofer. NASA HERMeS Hall Thruster Electrical Configuration Characterization. *Joint Propulsion Conference*, 2016.
- [50] Hani Kamhawi, Wensheng Huang, Thomas Haag, John Yim, and Daniel Herman. Characterization Tests of NASA 's Hall Effect Rocket with Magnetic Shielding Thruster. *Joint Propulsion Conference*, 2019.

- [51] I. D. Boyd. Review of Hall Thruster Plume Modeling. *Journal of Spacecraft and Rockets*, 38(3):381–387, 2001.
- [52] Peter Peterson, David Jacobson, David Manzella, and Jeremy John. The Performance and Wear Characterization of a High-Power High-Isp NASA Hall Thruster. (July), 2005.
- [53] Michael R Nakles, William A Hargus, Jorge J Delgado, and Ronald L Corey. A Performance Comparison of Xenon and Krypton Propellant on an SPT-100 Hall Thruster The 32nd International Electric Propulsion Conference. *Iepc-2011*, page 003, 2011.
- [54] A Neumann, J Sinske, and HP Harmann. The 250mN Thrust Balance for the DLR Goettingen EP Test Facility. *International Electric Propulsion Conference*, pages 1–10, 2013.
- [55] James E. Polk, Anthony Pancotti, Thomas Haag, Scott King, Mitchell Walker, Joseph Blakely, and John Ziemer. Recommended Practice for Thrust Measurement in Electric Propulsion Testing. *Journal of Propulsion and Power*, 33(3):539–555, 2017.
- [56] T Randolph, V Kim, H Kaufman, K Kopzubsky, V Zhurin, and M Day. Facility effects on stationary plasma thruster testing. *International Electric Propulsion Conference*, 1993.
- [57] Fernando Mier Hicks. Characterization on a magnetically levitated testbed for electro-spray propulsion systems. *MIT Masters Thesis*, pages 1–82, 2014.
- [58] T. W. Haag. Thrust stand for high-power electric propulsion devices. *Review of Scientific Instruments*, 62(5):1186–1191, 1991.
- [59] Trevor Moeller and Kurt A. Polzin. Thrust stand for vertically oriented electric propulsion performance evaluation. *Review of Scientific Instruments*, 81(11), 2010.
- [60] Ugur Kokal and Murat Celik. Development of a mili-Newton level thrust stand for thrust measurements of electric propulsion systems. *Proceedings of 8th International Conference on Recent Advances in Space Technologies, RAST 2017*, pages 31–37, 2017.
- [61] L. Boccaletto and L. D’Agostino. Design and testing of a micro-newton thrust stand for FEED. *35th Intersociety Energy Conversion Engineering Conference and Exhibit*, (July), 2000.
- [62] Aj Jamison and Ad Ketsdever. Accurate Measurement of Nano-Newton Thrust for Micropropulsion System Characterization. *27th International Electric Propulsion Conference*, pages 1–14, 2001.

- [63] Manuel Gamero-Castano. A torsional balance for the characterization of microNewton thrusters. *Review of Scientific Instruments*, 74(10):4509–4514, 2003.
- [64] Brett Tartler. Construction and Performance of an Inverted Pendulum Thrust Balance. (May 2010), 2010.
- [65] Leonard Cassady, Andrea Kodys, and Edgar Choueiri. A Thrust Stand for High-power Steady-state Plasma Thrusters. *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference Exhibit*, 2002.
- [66] N. Nagao, S. Yokota, K. Komurasaki, and Y. Arakawa. Development of a two-dimensional dual pendulum thrust stand for Hall thrusters. *Review of Scientific Instruments*, 78(11):8–11, 2007.
- [67] John K Ziemer. Performance Measurements Using a Sub-Micronewton Resolution Thrust Stand. *27th International Electric Propulsion Conference*, 2001.
- [68] Peter Thoreau and Justin Little. Development of the SPACE Lab Thrust Stand for Millinewton Thrust Measurement. *International Electric Propulsion Conference*, 2019.
- [69] Rohit Shastry, Richard R. Hofer, Bryan M. Reid, and Alec D. Gallimore. Method for analyzing $E \times B$ probe spectra from Hall thruster plumes. *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, pages 1–22, 2008.
- [70] L. L. Su and B. A. Jorns. Performance comparison of a 9-kW magnetically shielded Hall thruster operating on xenon and krypton. *Journal of Applied Physics*, 130(16), 2021.
- [71] James E. Pollard, Kevin D. Diamant, Vadim Khayms, Lance Werthman, David Q. King, and Kristi H. de Grys. Ion flux, energy, and charge-state measurements for the BPT-4000 hall thruster. *37th Joint Propulsion Conference and Exhibit*, (July), 2001.
- [72] Mitchell L.R. Walker, Richard R. Hofer, and Alec D. Gallimore. The effects of nude faraday probe design and vacuum facility backpressure on the measured ion current density profile of hall thruster plumes. *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (July), 2002.
- [73] Daniel L. Brown, Mitchell L.R. Walker, James Szabo, Wensheng Huang, and John E. Foster. Recommended practice for use of faraday probes in electric propulsion testing. *Journal of Propulsion and Power*, 33(3):582–613, 2017.

- [74] Wensheng Huang, Rohit Shastry, George C. Soulas, and Hani Kamhawi. Farfield Plume Measurement and Analysis on the NASA-300M. *60th JANNAF Propulsion meeting*, (December), 2013.
- [75] Leanne L. Su, Alexander R. Vazsonyi, and Benjamin A. Jorns. Performance of a 9-kw magnetically-shielded hall thruster with krypton aiaa-2020-3617. *AIAA Propulsion and Energy 2020 Forum*, pages 1–18, 2020.
- [76] I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 2002.
- [77] Käthe Dannenmayer and Stéphane Mazouffre. Elementary Scaling Relations for Hall Effect Thrusters. 27(1), 2011.
- [78] Eunkwang Lee, Younho Kim, and Hodong Lee. Scaling Approach for Sub-Kilowatt Hall-Effect Thrusters. 35(6), 2019.
- [79] John W. Dankanich, Mitchell Walker, Michael W. Swiatek, and John T. Yim. Recommended practice for pressure measurement and calculation of effective pumping speed in electric propulsion testing. *Journal of Propulsion and Power*, 33(3):668–680, 2017.
- [80] Mitchell Louis Ronald Walker and Alec Gallimore. Effects of Facility Backpressure on the Performance and Plume of a Hall Thruster. *27th International Electric Propulsion Conference*, (October):IEPC–01–045, 2001.
- [81] Ryan W Conversano, Dan M Goebel, and Richard R Hofer. Magnetically Shielded Miniature Hall Thruster : Development and Initial Testing. *IEPC-2013-201 Presented*, pages 1–21, 2013.
- [82] Vladimir Kim, Vyacheslav Kozlov, and Garry Popov. Plasma Parameter Distribution Determination in SPT-70 Plume. *28th International Electric Propulsion Conference*, pages 1–10, 2003.
- [83] Ryan W Conversano, Dan M Goebel, Richard R Hofer, Ioannis G Mikellides, Ira Katz, and Richard E Wirz. Magnetically Shielded Miniature Hall Thruster: Design Improvement and Performance Analysis. *34th International Electric Propulsion Conference*, pages IEPC–2015–100, 2015.
- [84] Ryan K. Ham, John D. Williams, Casey C. Farnell, Seth J. Thompson, Scott J Hall, Timothy R. Verhey, and Gabriel F. Benavides. Characterization of a Fixed-Volume Release System for Initiating an Arc Discharge in a Heaterless Hollow Cathode. *International Electric Propulsion Conference*, pages 1–10, 2019.

- [85] Ryan W. Conversano, Dan M. Goebel, Ioannis G. Mikellides, Richard R. Hofer, Taylor S. Matlock, and Richard E. Wirz. Magnetically Shielded Miniature Hall Thruster: Performance Assessment and Status Update. *Joint Propulsion Conference*, 2014.
- [86] Richard R. Hofer, Dan M. Goebel, Ioannis G. Mikellides, and Ira Katz. Design of a laboratory Hall thruster with magnetically shielded channel walls, Phase II: Experiments. *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit 2012*, (July), 2012.
- [87] Peter Y. Peterson, Hani Kamhawi, Wensheng Huang, John Yim, Daniel Herman, George Williams, James Gilland, and Richard Hofer. NASA HERMeS hall thruster electrical configuration characterization. *52nd AIAA/SAE/ASEE Joint Propulsion Conference, 2016*, 2016.
- [88] Daniel Roberts. Enhancing Neutral Propellant Flow Uniformity in Hall Thrusters via Anode Design. 2024.
- [89] John Steven Snyder, Jeff Baldwin, Jason D. Frieman, Mitchell L. R. Walker, Nathan S. Hicks, Kurt A. Polzin, and James T. Singleton. Recommended Practice for Flow Control and Measurement in Electric Propulsion Testing. *Journal of Propulsion and Power*, 33(3):556–565, 2017.
- [90] Daniel L. Brown, C. William Larson, Brian E. Beal, and Alec D. Gallimore. Methodology and historical perspective of a hall thruster efficiency analysis. *Journal of Propulsion and Power*, 25(6):1163–1177, 2009.
- [91] Lain D. Boyd and Michael L. Falk. A review of spacecraft material sputtering by hall thruster plumes. *37th Joint Propulsion Conference and Exhibit*, (c), 2001.
- [92] Robert B. Lobbia and Brian E. Beal. Recommended practice for use of langmuir probes in electric propulsion testing. *Journal of Propulsion and Power*, 33(3):566–581, 2017.
- [93] Bryan M. Reid, Rohit Shastry, Alec D. Gallimore, and Richard R. Hofer. Angularly-resolved $E \times B$ probe spectra in the plume of a 6-kW Hall thruster. *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (July), 2008.
- [94] Peter Thoreau and Justin Little. Influence of Field Topology on Magnetically Shielded Hall Thruster Plume Divergence. *IEEE Aerospace Conference Proceedings*, 2022.
- [95] Peter Thoreau and Justin Little. Influence of Field Topology on Magnetically Shielded Hall Thruster Performance. *International Electric Propulsion Conference*, 2022.

- [96] K K Jameson, D M Goebel, R R Hofer, and R M Watkins. Cathode Coupling in Hall Thrusters. *International Electric Propulsion Conference*, pages 1–20, 2007.
- [97] Jason D. Sommerville and Lyon B. King. Hall-effect thruster-cathode coupling, Part II: Ion beam and near-field plume. *Journal of Propulsion and Power*, 27(4):754–767, 2011.
- [98] B. Van Reijen, S. Weis, A. Lazurenko, J. Haderspeck, A. Genovese., P. Holtmann, K. Ruf, and N. Püttmann. High Precision Thrust Vector Determination through Full Hemispherical RPA Measurements assisted by Angular Mapping of Ion Energy Charge State Distribution. *33rd Internatinal Electric propalsion conference, The George Was- ington University, USA, October 6-10*, pages 1–6, 2013.
- [99] Thibault Hallouin and Stéphane Mazouffre. Far-field plume characterization of a 100-W class Hall Thruster. *Aerospace*, 7(5):1–21, 2020.
- [100] Daniel L. Brown and Alec D. Gallimore. Evaluation of facility effects on ion migration in a hall thruster plume. *Journal of Propulsion and Power*, 27(3):573–585, 2011.
- [101] Zhe Zhang, Zun Zhang, Haibin Tang, William Yeong Liang Ling, Zhiyuan Chen, Junxue Ren, and Jinbin Cao. Measurement of the distribution of charge exchange ions in a Hall-effect thruster plume. *Plasma Sources Science and Technology*, 29(8):0–23, 2020.
- [102] Jason D. Sommerville and Lyon B. King. Hall-effect thruster-cathode coupling, Part I: Efficiency improvements from an extended outer pole. *Journal of Propulsion and Power*, 27(4):744–753, 2011.
- [103] Dan M. Goebel, Kristina K. Jameson, and Richard R. Hofer. Hall thruster cathode flow impact on coupling voltage and cathode life. *Journal of Propulsion and Power*, 28(2):355–363, 2012.
- [104] Guangchuan ZHANG, Junxue REN, Wei LIANG, Ning OUYANG, Chao LU, and Haibin TANG. Coupling plasma plume of a low-power magnetically shielded Hall thruster with a hollow cathode. *Chinese Journal of Aeronautics*, 33(12):3018–3026, 2020.
- [105] J. Yamasaki, Shigeru Yokota, and K. Shimamura. Performance enhancement of an argon-based propellant in a Hall thruster. *Vacuum*, 167:520–523, 2019.
- [106] Dongho Lee, William P. Brabston, Dan Lev, and Mitchell L.R. Walker. Argon admixture-driven enhanced ionization and performance of a 5 kW Hall thruster on krypton. *Journal of Physics D: Applied Physics*, 57(32), 2024.

- [107] Guangchuan Zhang, Junxue Ren, Haibin Tang, Yibai Wang, Zhongkai Zhang, Jiubin Liu, Ruojian Pan, Zun Zhang, and Jinbin Cao. Plasma diagnosis inside the discharge channel of a low-power Hall thruster working on Xe/Kr mixtures. *Acta Astronautica*, 204(February):389–401, 2023.
- [108] Peng Zheng, Jianjun Wu, Yu Zhang, and Biqi Wu. A comprehensive review of atmosphere-breathing electric propulsion systems. *International Journal of Aerospace Engineering*, 2020(4), 2020.
- [109] Francesco Marchioni and Mark A. Cappelli. Extended channel Hall thruster for air-breathing electric propulsion. *Journal of Applied Physics*, 130(5):1–12, 2021.
- [110] Peter Thoreau, Aria R Johansen, Michael R. Holmes, and Justin M. Little. Modeling and Optimization of Propellant Mixtures for Hall Thrusters. In *International Electric Propulsion Conference*, number September, London, UK, 2025.
- [111] Thomas F. Munro-O’Brien and Charles N. Ryan. Performance of a low power Hall effect thruster with several gaseous propellants. *Acta Astronautica*, 206(January):257–273, 2023.
- [112] Davic C. Meeker. Finite Element Method Magnetics,, 2019.
- [113] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 1965.
- [114] M. J. D. Powell. An Efficient Method for Finding the Minimum of a Function of Several Variables Without Calculating Derivatives. *The Computer Journal*, 7(2):155–162, 1964.
- [115] Richard P. Brent. An algorithm with guaranteed convergence for finding a zero of a function. *The Computer Journal*, 14(4):422–425, 1971.
- [116] Rafael A. Martinez, Hoang Dao, and Mitchell L.R. Walker. Power Deposition into the Discharge Channel of a Hall Effect Thruster. *Journal of Propulsion and Power*, 30(1):209–220, 2014.
- [117] Carl F. Book and Mitchell L.R. Walker. Effect of Anode Temperature on Hall Thruster Performance. *Journal of Propulsion and Power*, 26(5):1036–1044, 2010.
- [118] Takeshi Furukawa, Takeshi Miyasaka, and Toshi Fujiwara. Methods of Controlling Low-Frequency Oscillation in a Hall Thruster. *International Electric Propulsion Conference*, 2001.

- [119] Jacob B Simmonds, Richard R Hofer, Samad Firdosy, Takuro Daimaru, Eric F Smith, Scott Roberts, Tomas Wexler, Peter Dillon, and Dan M Goebel. Thermal Management and Krypton Performance of the H10 High Power Density Hall Thruster. *International Electric Propulsion Conference*, 2024.
- [120] Wensheng Huang, Hani Kamhawi, Robert B. Lobbia, and Daniel L. Brown. Effect of Background Pressure on the Plasma Oscillation Characteristics of the HiVHAc Hall thruster. *Joint Propulsion Conference*, 2014.
- [121] J Mockus and L Mockus. *Bayesian Approach to Global Optimization and Applications*. 1987.
- [122] Joshua D Eckels, Thomas A Marks, Madison G Allen, Benjamin A Jorns, and Alex A Gorodetsky. Hall thruster model improvement by multidisciplinary uncertainty quantification. *Journal of Electric Propulsion*, 2024.
- [123] F Taccogna and Laurent Garrigues. Latest progress in Hall thrusters plasma modelling To cite this version : HAL Id : hal-02326278 Latest progress in Hall thrusters plasma modelling HAL Id : hal-02326278. *Reviews of Modern Plasma Physics*, 3(1), 2019.
- [124] Richard Hofer. *Development and Characterization of High-Efficiency, High-Specific Impulse Xenon Hall Thrusters*. PhD thesis, The University of Michigan, 2004.
- [125] Vlad George Tirila, Alain Demairé, and Charles N. Ryan. Review of alternative propellants in Hall thrusters. *Acta Astronautica*, 212(August):284–306, 2023.
- [126] Amazon. Inside Project Kuiper’s Florida hub: Preparing satellites for Amazon’s space network, 2025.
- [127] Vladimir Kim, Vladimir Kim, Garry Popov, Garry Popov, Vyacheslav Kozlov, Vyacheslav Kozlov, Alexander Skrylnikov, Alexander Skrylnikov, Dmitry Grdlichko, and Dmitry Grdlichko. Investigation of SPT Performance and Particularities of it’s Operation with Kr and Kr/Xe Mixtures *+. *Society*, pages 15–19, 2001.
- [128] B.A. Arkhipov, A.I. Koryakin, V.M. Murashko, A.N. Nesterenko, I.A. Khoromsky, V. Kim, V.I. Kozlov, G.A. Popov, and A.I. Skrylnikov. The Results of Testing and Effectiveness of the Kr-Xe Mixture Application in SPT. *October*, pages 15–19, 2001.
- [129] Cosimo Ducci, Tommaso Andreussi, Alexey Arkhipov, Andrea Passaro, Mariano Andrenucci, Alexandra Bulit, and Clive Edwards. Investigation of a 5 kW class Hall-effect thruster operating with different xenon-krypton mixtures. *34th International Electric Propulsion Conference (Hyogo-Kobe, Japan)*, pages IEPC–2015–126, 2015.

- [130] Xingyu Liu, Hong Li, Zezhong Peng, Yongjie Ding, Liqiu Wei, Daren Yu, and Wei Mao. Experimental study on the discharge of a xenon-assisted krypton Hall thruster. *Acta Astronautica*, 216(November 2023):295–303, 2024.
- [131] Leanne L. Su, Thomas A. Marks, and Benjamin A. Jorns. Trends in mass utilization of a magnetically shielded Hall thruster operating on xenon and krypton. *Plasma Sources Science and Technology*, 33(6), 2024.
- [132] J. B. Simmonds, M. P. Byrne, and V. H. Chaplin. Xenon tracers for cost effective laser induced fluorescence of alternative propellant Hall thrusters. *Review of Scientific Instruments*, 95(2), 2024.
- [133] J. P. Chiles and P Chauvet. Kriging : A method for cartography of the sea floor. *The International Hydrographic Review*, 1974.
- [134] Thomas Bayes. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions*, 1763.
- [135] A Stuart and K Ord. *Kendall's Advanced Theory of Statistics: Volume I – Distribution Theory*. Edward Arnold, 6th editio edition, 1994.
- [136] Xilu Wang, Yaochu Jin, Sebastian Schmitt, and Markus Olhofer. Recent Advances in Bayesian Optimization. *ACM Computing Surveys*, 55(13s), 2023.
- [137] Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected Improvements to Expected Improvement for Bayesian Optimization. *Advances in Neural Information Processing Systems*, 36(NeurIPS), 2023.
- [138] R. M. Jones. Comparison of potential electric propulsion systems for orbit transfer. *Journal of Spacecraft and Rockets*, 21(1):88–95, 1984.
- [139] SpaceX. SpaceX Rideshare Program.
- [140] Richard P. Welle. Propellant Storage Considerations for Electric Propulsion. *22nd International Electric Propulsion Conference*, (14-17):IEPC–1991–107, 1991.
- [141] Olivier Duchemin, Dominique Valentian, and Nicolas Cornu. Cryostorage of propellants for electric propulsion. *45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (August):1–11, 2009.

- [142] Kenneth G Unfried, R Joseph Cassady, W Andrew Hoskins, James D Marschke, and Aerojet Rocketdyne. Electric Propulsion Propellant Considerations: Can We Still Use Xenon and Krypton? *38th International Electric Propulsion Conference*, pages IEPC-2024-195, 2024.
- [143] Yongjie Ding, Haotian Fan, Dong Ma, Lei Wang, Boyang Jia, Hezhi Sun, Weifeng Xu, Hongbo Su, Liqiu Wei, Hong Li, and Daren Yu. *Extending service life of hall thrusters: recent progress and future challenges*, volume 3. Springer Singapore, 2019.
- [144] John R Brophy. NASA’s deep space 1 ion engine (plenary). In *Review of Scientific Instruments*, volume 73, page 1071, 2002.
- [145] Chapman Snowden. Odin’t: A Complete Debrief of Our Deep Space Mission, 2025.
- [146] Michael R. Nakles, William A. Hargus, J J Delgado, and Ronald L Corey. A Performance Comparison of Xenon and Krypton. *32nd International Electric Propulsion Conference*, 13(IEPC-2011-003), 2011.
- [147] Mariano Andrenucci, Francesco Battista[^], and Pietro Piliero[^]. Hall Thruster Scaling Methodology. pages 1–17, 2005.
- [148] Jesse A. Linnell and Alec D. Gallimore. Efficiency analysis of a hall thruster operating with krypton and xenon. *Journal of Propulsion and Power*, 22(6):1402–1412, 2006.
- [149] Harry W Jones. The Recent Large Reduction in Space Launch Cost. *48th International Conference on Environmental Systems*, (July 2018):81, 2018.
- [150] Pierre Lionnet. SpaceX and the categorical imperative to achieve low launch cost, 2024.
- [151] Archit Bapat, Pramod B. Salunkhe, and Aakash V. Patil. Hall-Effect Thrusters for Deep-Space Missions: A Review. *IEEE Transactions on Plasma Science*, 50(2):189–202, 2022.

Appendix A
HALL THRUSTER DIAGRAMS




Figure A.1: The three sub-assemblies that make up ACME. Version 3 is shown.




Figure A.2: The parts of the hollow cathode and keeper assembly.

Table A.1: Hollow cathode parts list.

Part Number	Name	Material
C1	Swagelock Nut	316 Stainless Steel
C2	Plasma Controls 1/8" Hollow Cathode	Tantalum, Tungsten, BaO
C3	Rear Isolator	Alumina
C4	Cathode Electrical Connection	316 Stainless Steel
C5	Bolt Isolators	Alumina
C6	Middle Isolator	Alumina
C7	Tube Isolator	Alumina
C8	Keeper Body	316 Stainless Steel
C9	Keeper Electrode/Orifice	Graphite
C10	Retaining Ring	Stainless Steel
C11	Forward Isolator	Alumina




Figure A.3: Numbered parts of the ACME Hall thruster showing the separate inner and outer poles.

Table A.2: ACME parts list.

Part Number	Name	Material
HT1	Inner Pole Cap	Alumina
HT2	Inner Channel Wall	Boron Nitride
HT3	Inner Hiperco Pole	Hiperco 50A
HT4	Inner Hiperco Screen	Hiperco 50A
HT5	Inner Hiperco Circuit	Hiperco 50A
HT6	Inner Hiperco Base	Hiperco 50A
HT7	Lower Stainless Structure	316 Stainless Steel
HT8	Upper Stainless Structure	316 Stainless Steel
HT9	Outer Pole Cap	Alumina
HT10	Outer Hiperco Pole	Hiperco 50A
HT11	Outer Channel Wall	Boron Nitride
HT12	Outer Hiperco Screen	Hiperco 50A
HT13	Outer Hiperco Circuit	Hiperco 50A
HT14	Outer Hiperco Base	Hiperco 50A

Table A.3: Cooling block parts list.

Part Number	Name	Material
W1	Thruster Radiator / Plasma Shield	101 Copper
W2	Keeper Heat-shield	101 Copper
W3	Plug for Manufacturing Access	Brass
W4	Cooling Block	101 Copper
W5	Cooling Loop Inlet Swagelock	Brass
W6	Cooling Loop Outlet Swagelock	Brass
W7	Thruster Mounting Plate	6061 Aluminium




Figure A.4: Numbered parts of the cooling block showing the separate inner and outer poles.