

**Performance of a Magnetically Shielded Hall Thruster Operating on Krypton
at High Powers**

by

Leanne Liuyue Su

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

Doctoral Committee:

Professor Benjamin A. Jorns, Chair
Emeritus Professor Alec D. Gallimore
Doctor Richard R. Hofer, Jet Propulsion Laboratory
Professor Ryan D. McBride



“People leading high-entropy lifestyles are always on the run.”
–Cengel & Boles, *Thermodynamics: An Engineering Approach*.

Leanne Liuyue Su
leannesu@umich.edu
ORCID iD: 0000-0002-7859-6807

© Leanne Liuyue Su 2023

DEDICATION

For Grandpa Joe, who taught me to always look up.

And for Pudge, the best kitty, who taught me to always look in her food bowl.

=^ . . ^ =

ACKNOWLEDGEMENTS

In the spirit of PEPL tradition and sticking true to my tendency to be emotional about everything, this is going to be a very long acknowledgments section. First, thank you to the National Science Foundation’s Graduate Student Research Fellowship, the Michigan Institute of Plasma Science and Engineering, and the Zonta International Amelia Earhart Fellowship for putting food on my table and paying my tuition. Thank you to everyone on my committee: Associate Professor Ben Jorns, Provost Alec Gallimore, Dr. Rich Hofer, and Professor Ryan McBride. Rich, thank you for your mentorship, for teaching me so much about Hall thrusters, and for laughing at almost all of my dumb jokes. To my advisor and fearless leader, Ben—thank you so much for your relentless support, constant encouragement (even when it’s just “that’s fine”), and for indulging all my ridiculous acronyms and test campaign names. Even though we may never see eye to eye about the appropriate number of colors to put in a plot (the correct answer is all of them), you have taught me so much about being a good researcher and academic. I sincerely hope that I’ll get to watch you perform Sweet Caroline at karaoke again.

I need to thank my family and especially my parents, who immigrated to this country and watched with vague bewilderment and pride as their daughter decided to do something with pretty colors in space. Mom, thank you for being my biggest cheerleader for my entire life, even when you have no idea what I’m doing. I’m excited to be able to be your cheerleader now that I’m finally done with school and you’ve gone back to it! Hopefully this doctorate makes up for the fact that I keep getting more tattoos. 外婆外公，谢谢你们给我做的饺子。我不在家的时候很想念。 Thank you also to Grandma Gerie and Grandpa Joe. You will always live in my heart alongside my battered copy of *The Stars* by H.A. Rey.

To my friends from back home—Alex, Rita, Jonathan, Matt, Kana, and Tyler—thank you for being an unwavering source of support through all the hardships and emotional traumas of grad school, COVID, the 2021 Seahawks season, the 2021–2022 Kraken season, horrible book club picks (my bad), and everything in between. No matter how far we get from Seattle, I’ll always think of you guys as home.

I’m still dumbstruck that I’ve managed to make the incredible friends in Michigan that I have outside of the lab. Thank you to my fellow keeners, especially Josh and Christina,

for helping me survive (and pass!) my quals. To my Beakers and bubble friends: Jana, Parker, Tom, Julia, Nik, Jenn, Amanda, Julie, Tasha, and Kate—thanks for keeping me sane through It All, or at least descending into madness alongside me. I also would not have been able to finish this degree without brunch from Jefferson Market, ice cream and cheap donuts from Washtenaw Dairy, drinks and free pool at Aut Bar (R.I.P.), coffee/beer/a perfect place to write at Argus, funky sour IPAs from HOMES Campus, and of course, all the happy hours at Bill’s Beer Garden. I hope all these wonderful places continue to keep the graduate students of Ann Arbor alive, fed, and inebriated for decades to come.

I’ve been lucky enough to spend some time with the lovely folks at the NASA Glenn Research Center and Jet Propulsion Laboratory. Scott, thank you for teaching me so much about nested Hall thrusters and for always taking time to thoroughly answer my barrage of questions. Jacob, thanks for making my time at JPL so fun, and for taking the heat for bullying me when really I was the one bullying you. Thank you also to the kind tow truck driver who rescued me and the X3 from Ohio and informed the jerks at Penske that I did indeed have “space shit” in the back.

I want to thank all the people in the Aerospace Engineering department at Michigan that have made my time in FXB more than just being a grad in a cage, especially Ruthie, who helped me every time I lost health insurance (often) or forgot to register for classes (even more often) with the patience of a saint while also bringing some much-needed chaos to the department. I’m also grateful to have been part of a number of student organizations: GEO, GSAC, GoSTEM, GradSWE, and GeMA, thanks for everything you do and for allowing me to meet all the great folks that I have.

PEPL is a magical place that is definitely not a cult despite what everyone else in the department thinks, and its magical powers come entirely from the people in it. Sarah, thank you so much for showing me the ropes when I was still a baby first year. I would know nothing about Hall thrusters without you. I have and will always look up to you, no matter how many cups of coffee you spill on random cars in the JPL parking lot. Ethan, every time I told you I had a problem (which was often), you would leave for five minutes and be back with a little gadget you’d just made that did exactly the thing I needed. Your lab doohickeys and your labelmakers live on eternal. Matt, you are the only one to have made it through being both my fellow TRIPLETeer and KWATRONian, and I’m honored to have suffered through the trenches with you. Wachs, even though your taste in sportsball is garbage, I have to admit that your backflips are impressive and sometimes your jokes are actually pretty funny. Eric, thank you so much for keeping LVTF and the entire PEPL building running despite my best unintentional attempts to destroy it. To everyone else who has come and gone—Marcel, Angela, Josh, Shadrach “Shipwreck” Hepner, Zach, Sophia,

Sal, and Konrad—thank you for helping to make PEPL the wonderful place it is.

To everyone who still remains: first of all, good luck, and don't schedule an 8.2 mile swim three weeks before your defense. (Give yourself at least four.) Madison, you are so kind and so much more capable than you think, in both karaoke and science. Will—how much money would you pay...RIGHT NOW...to be done with your dissertation too?? I will miss your golf swings and your shenanigans. Parker, even though your keyboard is ridiculously loud, thank you for being the best PaL and for your endless optimism. The Hall thruster squad is in great shape. Please continue using absurd muscle acronyms for all the test campaigns. Collin, I'm lucky to call you my friend and birthday twin despite our age gap of approximately 100 years. You are so generous with your time, baked goods, and jams/jellies/conserves. Tate, I'm sorry I broke so many cathodes. Thanks for fixing all of them and for adding a lil' spice to all our lab meetings. Chris, you introduced me to the worlds of D&D and fancy cocktails, and for that I will always be grateful. I will both cherish and be haunted by our memories in the cage forever. Thomas, you've been my classmate, study buddy, labmate, resident computationalist, roommate, Terraforming Mars rival, COVID quarantine pal, coauthor, and most importantly, my friend. I treasure our friendship and your yearly derivation of Ohm's law for me, and I can't wait to continue both for presumably the rest of our lives. To the new kids—Declan, Grace, Ari, Braden, and Miron—grad school is the best of times, the worst of times, and most importantly, the popsiclest of times. Enjoy it while you can, and make sure popsicle time lives on without me.

Finally, this would not have been possible without my partner in crime, euchre, and life. Austin, thank you for supporting me through every emotional breakdown, for throwing me an outdoor Korok birthday party in the middle of COVID, for being the best cat dad, for letting me ask you endless plasma questions about my research even though I still know nothing about yours (SURFACES!!), for putting up with all my bean-related threats. And of course, thank you to Pudge, who likely will never read this because she is a cat and cats cannot read. You are the sweetest and softest kitty, and you'll always be friendly and gregarious in my eyes. I love you both.

If you've gotten this far, I'm going to assume you're a grad student who needed to know how I got the krypton LIF laser working and then got bored and read my acknowledgments. I hope this dissertation proves at least somewhat useful to you. May your thrusters be happy and your cathodes calm.

Leanne Liuyue Su
苏柳越
2023

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	ix
LIST OF TABLES	xv
CHAPTER	
1 Introduction	1
1.1 Problem Statement	1
1.2 Objectives and Contributions	2
1.3 Organization	3
2 Space (Propulsion): The Final Frontier	5
2.1 Introduction	5
2.2 Electric Propulsion	6
2.3 Types of Electric Propulsion	8
2.3.1 Controversy	9
2.4 Plasma Properties	10
2.5 Summary	10
3 The State of the Hall Thruster	11
3.1 Introduction	11
3.2 Physics of Operation	11
3.3 Theoretical Framework for Characterizing Performance	13
3.3.1 Global Performance Metrics	13
3.3.2 Phenomenological Efficiency Model	14
3.3.3 Mass Efficiency Scaling with Plasma Parameters	17
3.4 Historical Designs and Scaling Laws	19
3.4.1 Magnetic Shielding	20
3.4.2 Chamber Material Selection	22
3.5 Krypton Operation	24
3.5.1 Motivation for Alternative Propellants	24
3.5.2 Advantages of Krypton	25
3.5.3 History of Krypton	26

3.5.4	Krypton Performance	26
3.5.5	Strategies for Improving Krypton Performance	28
3.6	High Power Operation	29
3.6.1	Motivation	29
3.6.2	Suitability of Hall Thrusters	30
3.6.3	Strategies for Increasing Thrust Density	31
3.7	Summary	35
4	Experimental Apparatus	36
4.1	Introduction	36
4.2	Thrusters	36
4.2.1	Cathodes	37
4.2.2	Thruster Operation	37
4.2.3	Thruster Monitoring	38
4.3	Facility	38
4.3.1	Pressure Measurements	39
4.3.2	Quartz Crystal Microbalance Sensor	40
4.4	Diagnostics	41
4.4.1	Thrust Stand	41
4.4.2	Probes	43
4.4.3	Probe Integration with Efficiency Model	49
4.4.4	Laser-Induced Fluorescence	54
4.5	Summary	60
5	Comparing Performance of Xenon and Krypton at Baseline Operating Conditions	61
5.1	Introduction	62
5.2	Methodology	63
5.2.1	Thruster and Facility	63
5.2.2	Diagnostics and Data Analysis	64
5.3	Results	67
5.3.1	Global Performance	67
5.3.2	Efficiency Analysis	70
5.3.3	Thruster Health	72
5.4	Discussion	74
5.4.1	Trends in Anode Efficiency with Voltage and Current	74
5.4.2	Comparison of Xenon and Krypton Operation	75
5.4.3	Comparison of Trends in Shielded and Unshielded Thrusters	76
5.4.4	Strategies for Optimizing Performance on Krypton	77
5.5	Summary	78
6	Performance and Operation at Ultrahigh Current Densities	80
6.1	Introduction	81
6.2	Methodology	82
6.2.1	Thruster and Facility	82

6.2.2	Diagnostics and Data Analysis	85
6.3	Results	88
6.3.1	Global Performance	89
6.3.2	Probe Measurements	92
6.3.3	Efficiency Analysis	95
6.3.4	Oscillation Level	99
6.3.5	Thermal	100
6.3.6	Facility	100
6.4	Discussion	102
6.4.1	Trends in Specific Impulse	102
6.4.2	Trends in Efficiency	104
6.4.3	Knudsen Number Analysis	114
6.4.4	Low Oscillation Amplitude	115
6.4.5	Facility Effects at High Power Density	115
6.4.6	Implications for Hall Thrusters on High-Power EP Missions	116
6.5	Summary	118
7	Mass Efficiency Scaling with Increasing Power	121
7.1	Introduction	121
7.2	Methodology	123
7.2.1	Experimental Setup	123
7.2.2	Simulation Setup	125
7.3	Results	129
7.3.1	Experimental Results	129
7.3.2	Simulation Results	133
7.4	Discussion	143
7.4.1	Anomalous Collision Frequency Profiles	143
7.4.2	Overall Efficiency Comparison	144
7.4.3	Efficiency Trends with Increasing Voltage	145
7.4.4	Efficiency Trends with Increasing Current	147
7.5	Summary	148
8	Conclusions and Future Work	151
8.1	Summary of Work	151
8.2	Implications of Work	152
8.3	Future Work	153
8.4	Concluding Remarks	157
	APPENDICES	158
	BIBLIOGRAPHY	174

LIST OF FIGURES

FIGURE		
2.1	SPACE! Comic reposted from <i>xkcd</i> [16].	5
2.2	Comparison of thrust and specific impulse for different types of electric propulsion. Note that the chemical regime is forbidden because it is outside the scope of this dissertation.	9
3.1	Notional side view and cross-sectional diagram of a Hall thruster and principle of operation.	12
3.2	Cross-sectional view of Hall thruster and parameters defining efficiency terms. .	15
3.3	Notional diagram of neutrals flowing through a channel and being ionized. White circles represent neutrals and gray circles represent ions.	18
3.4	Magnetic field topography for unshielded (US) and magnetically shielded (MS) configurations, showing electron temperature and plasma potential a) for US along channel centerline, b) for US along channel wall, and c) for MS along channel wall. Reproduced from Ref. [41].	21
3.5	Diagram showing method by which Hall thruster BN chambers may become coated in graphite. a) An ion from the plume hits the graphite beam dump, knocking off a neutral carbon atom. b) Carbon atom sticks to BN channel. c) Over time, BN channel becomes coated in carbon.	23
3.6	Ionization rate coefficient as a function of electron temperature for xenon and krypton. Data taken from Refs. [32, 33, 34].	27
3.7	Various methods of increasing mass utilization efficiency in a Hall thruster: (1) increasing channel length, (2) increasing ionization cross-section, and (3) increasing plasma density.	28
4.1	a) A photo of LVTF and b) a CAD of LVTF [111] showing the locations of pumps around the chamber. Note that the bottom left corner of the photo corresponds to the bottom right corner of the CAD.	39
4.2	Photos of the a) RPA, b) $E \times B$ probe, and c) Faraday probe used.	43
4.3	RPA data for xenon at a) 300 V and 15 A and b) 300 V and 125 A. The original current trace is shown as a dotted black line, a fitted smoothing spline is shown as a solid black line, the IEDF is shown as a dashed gray line, and the location of the most probable voltage is shown as a vertical dotted line. Note that these plots are corrected for plasma potential.	44

4.4	Langmuir probe data for xenon at 300 V, 125 A. (a) Raw probe I-V trace. (b) Natural log of electron current used to determine electron temperature. (c) Second derivative of current used to determine plasma potential with zero crossing method. In all cases, the electron retarding region is between the black and the gray circles.	45
4.5	Raw E×B trace and two-peak Gaussian fits for all charge species of xenon at a) 300 V and 15 A and b) 300 V and 125 A. Note that the Xe ⁺³ population was too small to detect at 15 A.	47
4.6	Bootstrapping process with RPA trace.	49
4.7	Raw current density trace (blue) and CEX correctional methods (black) (Sec. 4.4.3.3) used for calculating beam current and divergence angle from Faraday probe during operation on xenon at a) 300 V and 15 A and b) 300 V and 125 A.	53
4.8	System overview of laser-induced fluorescence.	56
4.9	Ion velocity distribution functions taken along channel centerline of the H9 operating on krypton at 300 V, 15 A from upstream (pink, low velocities) to downstream (red, high velocities).	58
4.10	Average ion velocities along channel centerline of the H9 operating on krypton at 300 V, 15 A. a) Mean and mode of raw and fitted IVDFs shown in Fig.4.9. b) Mean of fitted ion velocities and electric field. The circular points are the calculated values of mean ion velocities and electric field, while the solid lines represent a smoothing spline fit for the velocity profile and the electric field calculated from that fit. The pink dashed lines indicate the position and width of the acceleration region	59
5.1	ThrusteR Investigation with Probes, LIF, and Erosion Time. Logo design credit shared by T. A. Marks and L. L. Su.	61
5.2	Krypton and Wave Analyses with Time-Resolved Oscillations. Logo design credit shared by T. A. Marks and L. L. Su.	61
5.3	The H9 operating at 300 V and 15 A on a) xenon and b) krypton.	64
5.4	Notional top-down schematic of experimental setup in LVTF, including thruster, thrust stand, ion gauge, probe arm with Langmuir probe, emissive probe, retarding potential analyzer, and Faraday arm, far-field E×B probe, and beam dump. Pumps are located around the walls of the facility but not shown in the diagram.	65
5.5	a) Thrust, b) specific impulse, c) anode efficiency, and d) total efficiency for xenon and krypton as a function of operating condition. Note that the uncertainties on Fig.5.5a are smaller than the point sizes.	68
5.6	Ratio of measured to theoretical limit of specific impulse for xenon and krypton at 15 A as a function of discharge power.	69
5.7	Measurements of the contributions to anode efficiency for xenon inferred from far-field probes.	71
5.8	Measurements of the contributions to anode efficiency for krypton inferred from far-field probes.	72
5.9	Ratio of peak-to-peak oscillation to mean value of discharge current at 300 V for xenon and krypton as a function of operating condition.	73

5.10	Differences in the contributions to anode efficiency between krypton and xenon.	75
6.1	H9 MUSCLE test campaign flag. Logo design credit to T. M. Gill.	80
6.2	H9 MUSCLE operating on a) xenon at 300 V and 125 A and b) krypton at 300 V and 150 A.	83
6.3	Notional top-down schematic of experimental setup in LVTF, including thruster, thrust stand, ion gauge, quartz crystal microbalance sensor, Armored Probe Carrier with $E \times B$ probe, Langmuir probe, and retarding potential analyzer, probe arm with Faraday probe and camera, and beam dump. Note that the QCM was mounted above the ion gauge to prevent thermal interference. Pumps are located around the walls of the facility but not shown in the diagram. . . .	84
6.4	Photos of probe setup inside the vacuum facility showing a a) front view of the APC and probe arm, and b) side view of the APC. The front view shows the graphite shields from the thruster's perspective, as well as the swept Faraday probe and camera. The side view shows the three probes (from nearest to furthest: Langmuir, RPA, $E \times B$) mounted on the motion stage behind the graphite panels.	87
6.5	a) Thrust, b) specific impulse, c) anode efficiency, and b) total efficiency at 300 V for xenon and krypton with increasing discharge current. Note that the uncertainties in Fig. 6.5a are smaller than the marker sizes.	89
6.6	Plasma potential and electron temperature of a) xenon and b) krypton at 300 V with increasing discharge current.	92
6.7	Average ion energy inferred from RPA corrected by the plasma potential for a) xenon and b) krypton at 300 V with increasing discharge current. The discharge voltage is indicated with a dotted black line.	93
6.8	Current fractions of singly, doubly, and triply ionized a) xenon and b) krypton at 300 V with increasing discharge current. Note that the uncertainties are smaller than the marker sizes.	94
6.9	Plume divergence angle of a) xenon and b) krypton at 300 V with increasing discharge current.	94
6.10	Efficiency modes for a) xenon and b) krypton at 300 V with increasing discharge current.	96
6.11	Ratio of peak-to-peak oscillation to mean value of discharge current at 300 V for xenon and krypton with increasing discharge current.	99
6.12	Discharge current and inner front pole temperature throughout a testing day for a) xenon and b) krypton at 300 V. The dashed pink line indicates the temperature limit for the IFP.	100
6.13	a) Carbon deposition rates and b) facility pressures at 300 V for xenon and krypton with increasing discharge current.	101
6.14	Experimentally measured values of specific impulse for a) xenon and b) krypton operating at 300 V compared to theoretical values in which the plasma is assumed to be singly charged (with and without efficiency losses), doubly charged (without efficiency losses), and polydispersive (with efficiency losses).	103

6.15	Experimentally measured values (color) of beam current utilization efficiency for xenon and krypton operating at 300 V compared to theoretical form (black) with a) actual discharge current and b) adjusted current \bar{I}_d with a characteristic current of $I_n = 8$	105
6.16	Experimentally measured values (color) of beam current utilization efficiency for a) xenon and b) krypton operating at 300 V compared to theoretical forms (black) assuming classical scaling with $\alpha_{c,Xe} = 310, \alpha_{c,Kr} = 560$ and Bohm scaling with $\alpha_{B,Xe} = 0.14, \alpha_{B,Kr} = 0.19$	111
6.17	Experimentally measured values of anode efficiency for xenon and krypton operating at 300 V compared to theoretical form with $\eta_{const} = 0.85$ assuming weights of $\beta_c = 0.2, \beta_B = 0.8$ for xenon and $\beta_c = 0.1, \beta_B = 0.9$ for krypton.	113
6.18	Knudsen numbers for xenon and krypton with increasing discharge current. . .	114
6.19	Comparison of H9 MUSCLE a) thrust densities scaled to xenon operation at 4.5 kW, b) thrust densities scaled to krypton operation at 4.5 kW, and c) thrust-to-power ratios to a database of AF-MPD measurements [170].	117
7.1	Parker and Leanne do Science. PaLS features the experiments KLIF-FARD (K rypton L aser- I nduced F luorescence F oray into A cceleration R egion D ynamics) and SCOOBIE (S wirling C athode O scillations from O ptics and B eamdump I mpact E xperiments).	121
7.2	Notional top-down schematic of experimental setup in LVTF, including thruster, motion stages, ion gauge, injection/collection optics, and beam dump. Pumps are located around the walls of the facility but not shown in the diagram. . . .	123
7.3	a) Experimental domain for LIF interrogation (shown as a transparent white rectangle) and b) simulation domain overlaid atop H9 operating on krypton at 300 V, 15 A. Note that in the experimental domain, channel centerline $r = 0$ is represented by the dashed pink line.	125
7.4	Notional anomalous collision frequency profile according to Eq. 7.1, showing locations of nodes (z_i, c_i)	127
7.5	Profiles of mean ion velocity along thruster channel centerline with xenon operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V. 131	131
7.6	Profiles of mean ion velocity along thruster channel centerline with krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	131
7.7	Acceleration region locations and widths as calculated by Eq. 4.19 for xenon and krypton at different operating conditions.	132
7.8	Experimental (black) and simulation (color) ion velocity profiles for each condition. Note that the 600 V, 15 A conditions have a different scale for ion velocity due to the high voltage. For the krypton 600 V, 15 A condition, the dotted line is the condition at which the thrust/efficiency matched and the dashed line is the condition at which the IVR matched.	134

7.9	Calibrated anomalous electron collision frequency profiles along thruster channel centerline with xenon operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	135
7.10	Calibrated anomalous electron collision frequency profiles along thruster channel centerline with krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	136
7.11	Axial profiles of a) electron temperature, b) ionization rate coefficient, c) plasma density, and d) neutral velocity along channel centerline for operation at 300 V, 15 A on xenon and krypton.	138
7.12	Axially-averaged electron temperature along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	140
7.13	Axially-averaged ionization rate coefficient along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	140
7.14	Axially-averaged electron density along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	141
7.15	Mass efficiencies of xenon and krypton operation as calculated from theory with simulated plasma parameters (Eq. 3.15) for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	142
7.16	Axial profiles of a) plasma density, b) electron temperature, and c) ionization rate coefficient along channel centerline for operation at 300 V, 15 A (blue) and 600 V, 15 A (black) on xenon.	145
7.17	Notional mass efficiency scaling for xenon (blue dotted) and krypton (purple dashed), with 15 A and 30 A conditions shown with circular markers for xenon (filled) and krypton (open).	148
B.1	Simulated component efficiencies of xenon and krypton operation as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	171
B.2	Axial profile of electron temperature along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.	172

B.3	Axial profile of ionization rate coefficient along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V. . . .	172
B.4	Axial profile of electron density along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V. . . .	173
B.5	Axial profile of neutral velocity along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V. . . .	173

LIST OF TABLES

TABLE

3.1	Approximate material properties of possible discharge chamber materials. . . .	22
3.2	Gas properties of xenon, krypton, argon, neon, nitrogen, bismuth, and iodine [56, 57]. The conversion from % to ppm is a factor of 1×10^{-4} . *For a monatomic gas. †Price from 2023 Q1 average [5].	24
4.1	Secondary electron emission coefficients for charge states 1 through 4 of xenon and krypton.	52
4.2	Excitation and fluorescence wavelengths used for selected singly-ionized krypton and singly-ionized xenon transitions.	56
5.1	Operating conditions and base pressures. Note that probe data was only collected for the three 4.5 and 6 kW conditions.	63
6.1	Operating conditions and base pressures. The discharge voltage was 300 V and the B-field ratio was 112.5% for all conditions.	85
6.2	Plasma parameter values averaged over all operating conditions and both propellants.	106
7.1	Operating conditions and base pressures.	124
7.2	Numerical parameters employed in this work.	129
7.3	Experimentally measured thrust and anode efficiency for all conditions [71]. *The 300 V, 30 A condition thrusts were linearly extrapolated from 15 to 20 A. . . .	130
7.4	Simulated values of thrust, anode efficiency, and integrated velocity residual for all conditions. *This condition did not meet the IVR requirement. †This condition did not meet the thrust and efficiency requirements.	133
B.1	Xenon performance measurements and discharge current oscillations. *Insufficient data available to determine uncertainties.	164
B.2	Krypton performance measurements and discharge current oscillations. *Insufficient data available to determine uncertainties.	164
B.3	Contributions to xenon efficiency inferred from probe measurements.	165
B.4	Contributions to krypton efficiency inferred from probe measurements.	165
B.5	Xenon performance measurements and discharge current oscillations at 300 V. .	166
B.6	Krypton performance measurements and discharge current oscillations at 300 V.	166
B.7	Xenon efficiencies from probe measurements at 300 V.	167
B.8	Krypton efficiencies from probe measurements at 300 V.	167

CHAPTER 1

Introduction

“It’s a tough galaxy out there, but somebody’s gotta live in it...it might as well be you.”

–Elaine Lee *c/o Galactic Girl Guides c/o Brennan Lee Mulligan, A Starstruck Odyssey.*

1.1 Problem Statement

Electric propulsion devices operating at high powers (>100 kW) with high performance and low specific mass (<5 kg/kW for the system) are a key enabling technology for crewed deep space missions [1, 2, 3, 4]. Hall thrusters are a leading candidate for this purpose as they have decades of in-space heritage in the mid-power (1–10 kW) range. However, scaling to higher powers presents multiple challenges for both in-space operation and ground testing. One area of interest is that of propellant choice. Xenon is typically used for electric propulsion applications due to its high mass, low ionization energy, and non-reactivity. Unfortunately, xenon is rare in the atmosphere and therefore expensive to extract. Additionally, the price of xenon is highly variable, ranging from \$15/L in 2017 to over \$100/L in 2022 [5]. With the high throughput of xenon necessary for both ground validation of high-power thrusters and long-distance space travel, test campaigns and missions can become prohibitively expensive. Krypton is an appealing alternative propellant as another noble gas with a lower mass and about ten times cheaper price compared to xenon.

Historically, krypton has exhibited low performance (5–15% lower efficiencies) compared to xenon at constant operating power [6, 7, 8, 9, 10, 11, 12, 13]. These studies were performed on Hall thrusters with a traditional “unshielded” magnetic topology. The advent of magnetic shielding, a technique by which the field lines in the thruster channel divert energetic ions away from channel walls, has both greatly extended the lifetimes of these devices and changed their internal plasma dynamics [14, 15]. We therefore find it necessary to understand how the performance of krypton and xenon compare on this relatively new magnetic field design. We also expect that as we climb to increasingly high powers, the traditionally-observed dynamics

of xenon and krypton operation may change. We thus see a pressing need to characterize magnetically shielded Hall thruster performance on krypton at high powers.

1.2 Objectives and Contributions

The primary objective of this work was to characterize the performance and scaling of a magnetically shielded Hall thruster operating on krypton at high powers and comparing it to xenon. This characterization was done at both standard operating conditions and current densities an order of magnitude higher than typical of Hall thrusters. We used global performance metrics to evaluate the thruster’s overall performance and far-field probes to measure plasma parameters used to determine an efficiency breakdown at each operating condition. We also used a laser-induced fluorescence (LIF) system to measure ion velocities along channel centerline at a number of operating conditions. These velocity profiles were used to inform multi-fluid simulations that revealed further insight into the plasma dynamics within the thruster channel. With these findings, we constructed and validated general scaling laws for how we expect the mass utilization to change with increasing current for different propellants. Contributions of this work include:

1. **Performance measurements and efficiency breakdown for xenon and krypton operation on a magnetically shielded Hall thruster over a standard throttling table.** We used a thrust stand and a far-field probe suite to perform comprehensive performance measurements on a magnetically shielded Hall thruster operating from 4.5 to 9 kW across a “standard” throttling table. This allowed us to identify the main efficiency modes that drive the gap between xenon and krypton efficiency, as well as point to the difference between unshielded and shielded performance. Notably, we found that the performance gap between xenon and krypton does not close at high voltages as it does with unshielded thrusters.
2. **Krypton LIF measurements.** We set up a new laser-induced fluorescence system for analyzing singly-ionized krypton. With this system, we were able to measure ion velocity distribution functions within the thruster plume along channel centerline. We did so at operating conditions increasing in power by way of voltage and current. These measurements yielded profiles of ion velocity through the channel, lending insight into the position and shape of the acceleration region. These measurements also provided calibration data for simulations.
3. **Operation of a Hall thruster and analysis of efficiency changes at high currents.** We ran a Hall thruster up to ten times its nominal operational current density

with both xenon and krypton as propellant. Throughout this range, we took performance measurements to characterize thrust, global efficiency, and efficiency modes at each condition. We found that although the efficiency does decrease at high voltages, this decrease is minor enough that our Hall thruster is still competitive with state-of-the-art thrusters at ultrahigh currents. Additionally, we found that at sufficiently high currents, the efficiency of krypton is able to overtake that of xenon. This is partially attributed to the ability of krypton to reach 100% mass utilization despite its smaller ionization cross-section.

4. **Calibrated simulations at five operating conditions for both xenon and krypton.** With the velocity profiles procured through LIF, we were able to calibrate ten total operating conditions for the same thruster. These calibrated simulations, found by adjusting the anomalous electron collision frequency profiles, serve as a tool for probing the internal physics of the thruster channel. While the scope of this work was primarily limited to analyzing plasma parameters relevant to efficiency scaling, the calibrated simulations also may have future uses in further analysis of xenon and krypton operation on a shielded Hall thruster.
5. **0D scaling law for efficiency with increasing power on a shielded Hall thruster.** As the culmination of our experimental work, we formulated a 0D scaling law for mass utilization with increasing voltage and current for a given propellant. This scaling law was validated with both experimental and simulation data. Because mass utilization is the primary driver for the efficiency gap between xenon and krypton, characterizing this efficiency mode allows us to identify best methods to improve krypton performance. In this work, we were able to identify the plasma density and not the ionization rate of a given propellant as the primary driver for mass utilization. Additionally, this scaling law is extensible to other propellants, providing a framework with which we may be able to predict the behavior of future thruster operation.

1.3 Organization

In Chapter II, we provide background on the province of electric propulsion and its role in space exploration. We introduce in Chapter III how Hall thrusters work, recent developments in their design, and the theoretical framework we use to characterize their performance. We also discuss the motivation for our interest in both high powers and alternative propellants. Chapter IV outlines the experimental and simulation methods used throughout this dissertation. This includes a discussion of the H9, the primary test article for these experiments,

and the retrofitted H9 MUSCLE. We also describe the diagnostics used and the simulation techniques we leveraged in conjunction with our experimental work.

Chapters V to VII represent the thrust of the dissertation work. In Chapter V, we detail how we used these experimental techniques to characterize the global performance and efficiency breakdown of our thruster operating on both xenon and krypton at its baseline operating conditions. We also present the results from this study and our theories for the different behaviors of xenon and krypton. In Chapter VI, we describe our successful efforts to reach high current density operation on our test device. We compare the performance results of our thruster operating at these conditions on both xenon and krypton, including a discussion of why krypton efficiency overtakes that of xenon at sufficiently high currents. Next, we leverage experimental results—including measurements of ion velocities along channel centerline—with simulation tools in Chapter VII to explore the physics behind the scaling of a shielded Hall thruster operating on these two different propellants. Finally, we discuss the implications of our work in Chapter VIII and potential avenues for future research.

CHAPTER 2

Space (Propulsion): The Final Frontier

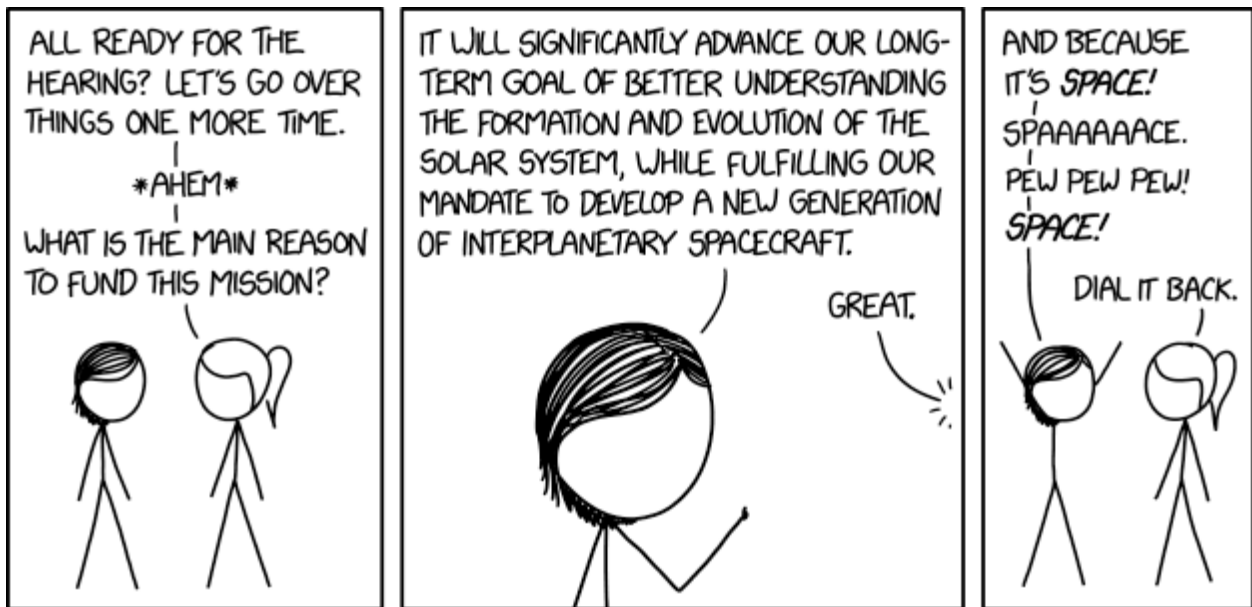


Figure 2.1: SPACE! Comic reposted from *xkcd* [16].

2.1 Introduction

This chapter provides the background information necessary for understanding the general motivation for this work and the larger context of space propulsion. Note that we assume a general knowledge of plasma physics and neglect to go through derivations of the governing laws of plasmas. References [17, 18, 19] are good resources for establishing these fundamentals.

We first outline what electric propulsion is, how it differs from chemical propulsion, and the operational niche it fills. Next, we outline a number of different types of electric propulsion thrusters and their different operational regimes. We conclude with a brief overview

of some of the plasma properties that are useful in understanding and characterizing the physics of electric propulsion thrusters.

2.2 Electric Propulsion

In this section, we introduce the concept of electric propulsion (EP) and how it differs from traditional chemical propulsion. Space propulsion is reactionary, which means we generate force by ejecting mass out the back of a rocket. The classical rocket equation, derived most famously by Russian scientist Konstantin Tsiolkovsky (but also William Moore, Robert Goddard, and Hermann Oberth), describes this relation for a rocket as follows:

$$m_0 = m_f \exp[-\Delta v/v_{ex}], \quad (2.1)$$

where m_0 is the initial mass of the rocket, m_f is the final mass, Δv is the maximum change in velocity, and v_{ex} is the propellant exhaust velocity. The initial and final masses can be related through the propellant mass m_p , $m_0 = m_f + m_p$. This equation indicates how for a given Δv requirement and final mass (typically dictated by the payload), a higher exhaust velocity means less required initial propellant.

One helpful way to characterize exhaust velocity is with the concept of specific impulse, or the amount of impulse (change in momentum) produced per unit of propellant. The specific impulse I_{sp} and exhaust velocity are related through

$$I_{sp} = \frac{v_{ex}}{g_0}, \quad (2.2)$$

where $g_0 = 9.81 \text{ m/s}^2$ is the gravitational acceleration at the surface of Earth. Typical values of specific impulse for electric propulsion devices range from hundreds to thousands of seconds.

Another metric for the performance of a rocket is its thrust T . Typical thrusts for EP devices range from tens to hundreds of millinewtons. Thrust can be related to exhaust velocity and specific impulse through the relation

$$T = \dot{m}_{out} v_{ex} = \dot{m} \frac{I_{sp}}{g_0}, \quad (2.3)$$

where \dot{m}_{out} is the total mass flow out of the device. It becomes apparent from this relation that with a fixed mass flow, the thrust and specific impulse of a thruster are inversely related. We expand upon how these metrics relate to the performance characterization of a thruster

in Sec. 3.3.

Both traditional chemical propulsion and electric propulsion generate thrust for spacecraft. Chemical propulsion does so by breaking chemical bonds to gain energy. This caps the exhaust velocity based on the specific internal energy ε_i of these bonds, i.e.

$$v_{ex}|_{chem} = \sqrt{2\varepsilon_i}. \quad (2.4)$$

This means that chemical propulsion, while capable of generating high amounts of thrust (typically on the order of hundreds of Newtons), is limited to ~ 100 s for specific impulse. This trade-off between thrust and specific impulse is found by considering Eq. 2.3, where it becomes apparent that for a fixed mass flow, the thrust and specific impulse are inversely related.

The specific impulse of electric propulsion, meanwhile, depends on the amount of power able to be imparted through electric and magnetic fields. Its ultimate exhaust velocity is therefore unconstrained by the inherent energy of chemical bonds. Instead, the exhaust velocity of an EP system is

$$v_{ex}|_{EP} = \sqrt{\frac{2P}{\dot{m}}}, \quad (2.5)$$

where P is the power input to the system and \dot{m} is the mass flow input. This relation means that EP fits a niche for high-specific impulse, low-thrust missions.

Electric and chemical propulsion each has its own set of advantages and disadvantages. As the focus of this work is on electric propulsion, we review here some of the major advantages of EP over chemical.

- **High efficiency.** Electric propulsion systems can achieve much higher exhaust velocities compared to chemical rockets. This leads to higher specific impulse, which means they can achieve greater acceleration and velocity with less propellant mass (see Eq. 2.1), making it possible to carry more payload or travel longer distances. Additionally, this can lead to cost savings in launch vehicles and enable more complex mission designs.
- **Higher speeds and extended mission durations.** Due to the efficient fuel economy of electric propulsion, it is possible for these devices to continuously thrust over extended periods and gradually accelerate to high velocities. This makes EP well-suited for long-duration missions such as deep space exploration or station-keeping for satellites.

- **Propellant flexibility.** As previously mentioned, chemical propulsion depends on the internal energy of chemical bonds while electric propulsion depends solely on the mass flow and power input. This means that in theory, anything can be used as propellant for electric propulsion as long as it can be affected by the electric and/or magnetic fields. This increases the potential flexibility of EP, particularly for in-space resource utilization (ISRU) applications.

While electric propulsion has numerous advantages, it is important to note that it also comes with certain challenges. Electric propulsion systems require a power source, such as solar panels or nuclear reactors, which can add complexity and mass to the spacecraft. Additionally, while the low thrusts generated by EP systems are ideal for precise maneuvers, they are ill-suited for rapid changes in trajectory or launches from Earth’s surface. Unlike chemical propulsion, electric propulsion is used for in-space applications only and not with launch. Within the canon of electric propulsion, there are multiple subcategories that fulfill different mission requirements by trading thrust against specific impulse.

2.3 Types of Electric Propulsion

Electric propulsion devices can generally be classified into three groups, dictated by the method through which they produce thrust:

1. **Electrothermal.** These devices gain energy by heating up a propellant using electrical energy before expelling it. This category includes resistojets, arcjets, and electron-cyclotron resonance (ECR) thrusters.
2. **Electrostatic.** These devices gain energy from ions being accelerated through an electric field. This category includes gridded ion thrusters, electrosprays, and Hall thrusters.
3. **Electromagnetic.** These devices gain energy from a plasma being manipulated and accelerated by magnetic fields. This category includes magnetoplasmadynamic thrusters (MPD), pulsed inductive thrusters, ECR thrusters, and Hall thrusters.

Figure 2.2 shows typical operating specific impulses and thrusts for different types of thrusters. The ranges shown here represent a variety of studies, including ones of laboratory devices. We also note that operating power is not shown on this chart and the relationship between thrust and specific impulse is not meant to be representative of constant power. As we can see from this chart, some electric propulsion devices are better-suited to high thrusts

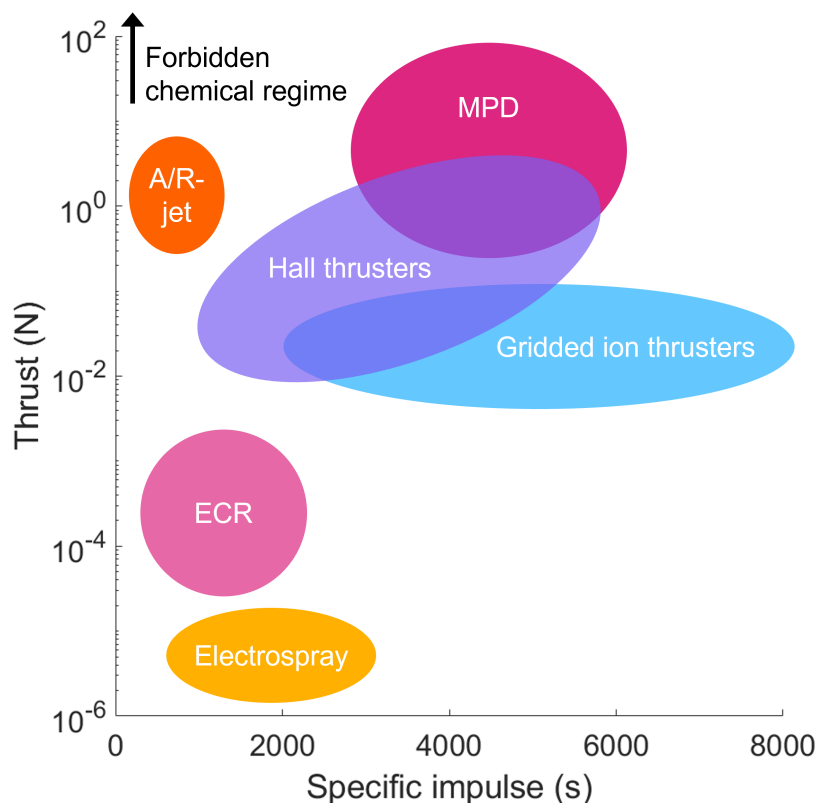


Figure 2.2: Comparison of thrust and specific impulse for different types of electric propulsion. Note that the chemical regime is forbidden because it is outside the scope of this dissertation.

than others. In particular, Hall thrusters and MPD thrusters are capable of reaching higher thrust densities while maintaining the high specific impulses (>1000 s) that make electric propulsion attractive for deep space applications.

2.3.1 Controversy

Some thrusters may fall into multiple classifications. In particular, Hall thrusters can be considered both electrostatic and electromagnetic. The reasons for this will be made clear in the following section when we outline the principles of operation of a Hall thruster. However, in summary, the force accelerating the thrust-producing ions is electrostatic, but the force transfer from the electrons in the plasma to the thruster is electromagnetic. Similarly, thrusters that utilize a magnetic nozzle design—such as ECR thrusters and VASIMR—may be considered both electrothermal and electrostatic.

2.4 Plasma Properties

While providing the reader with a thorough understanding of plasma physics is beyond the scope of this dissertation, we recap here a few key plasma parameters that may prove to be useful in understanding the following sections.

- plasma frequency, $\omega_p = \sqrt{\frac{e^2 n_e}{\varepsilon_0 m_e}}$
- electron cyclotron frequency, $\omega_{c,e} = -\frac{eB}{m_e}$
- ion cyclotron frequency, $\omega_{c,i} = \frac{qB}{m_i}$
- Larmor radius for species s (electron or ion), $r_{L,s} = \frac{v_{i,s}}{\omega_{c,s}}$
- plasma resistivity, $\eta = \frac{m_e \nu_e}{e^2 n_e}$

In these expressions, e is the elementary charge, n_e is the electron density, ε_0 is the permittivity of free space, B is the strength of the magnetic field, m_e is the mass of an electron, q is the charge, m_i is the mass of an ion, and ν_e is the electron collision frequency. We note that in this work, we almost always use η to mean efficiency instead of resistivity. In the cases where η means resistivity, it is explicitly specified and we leave it without a subscript (i.e. η_x for efficiencies).

2.5 Summary

In this chapter, we introduced electric propulsion and the role it plays in the world of space propulsion. We list different types of EP devices and explain the different thrust and specific impulse regimes in which they operate. We overview some of the important plasma parameters that are useful in characterizing the low-temperature plasmas typically seen in EP devices. This introduction gives us the context necessary to understand the motivation for this body of work as a whole. Next, we delve into our device of interest: the Hall thruster.

CHAPTER 3

The State of the Hall Thruster

“Electrons are like cockroaches.”

–Dr. Alec D. Gallimore, founder of the Plasmadynamics and Electric Propulsion Laboratory.

3.1 Introduction

In this chapter, we introduce the Hall thruster, for which we describe its operating principles, historical scaling laws, and recent developments in design and operation. First, we describe the Hall thruster’s principle of operation. Next, we outline a framework we can use to characterize both the global performance and individual efficiency contributions towards Hall thruster performance. This section also includes a derivation for a simplified 0D model of mass utilization in terms of plasma parameters we use to characterize the plume. We then discuss the motivation behind using non-traditional (i.e. non-xenon) propellants for Hall thruster operation, as well as why krypton is usually the first choice. Additionally, we explain strategies for closing the gap between xenon and krypton performance and review previous work done comparing the performance of xenon and krypton. Finally, we review the motivation for high power EP as a whole and discuss why Hall thrusters may be especially well-suited for this operational regime, especially given its recent technological developments.

3.2 Physics of Operation

Hall thrusters are a cylindrical $E \times B$ device with moderate specific impulse and high thrust density within the pantheon of electric propulsion devices (c.f. Fig. 2.2). They are now the most commonly used EP thruster flown, largely owing to the Starlink constellation [20]. More recently, Hall thrusters have also been baselined for deep space missions [21, 22, 23]

The principle of operation of a Hall thruster is shown in Fig. 3.1 with a side view and cross-sectional diagram. In this axisymmetric device, a cathode—typically mounted in the




Figure 3.1: Notional side view and cross-sectional diagram of a Hall thruster and principle of operation.

center along axis of symmetry or outside the thruster—thermionically emits electrons. A discharge voltage V_d is then established between the anode (located at the upstream end of the discharge channel) and cathode, generating a discharge current, I_d , within the channel and an electric field \vec{E} that points axially downstream. Neutrals that are injected through the anode are then ionized through collisions with electrons supplied by the cathode. Normally, electrons would flow back to the anode due to the electric field, but ideally, we would be able to keep them trapped in the channel to continue ionizing the incoming neutrals. To this end, we apply a radial magnetic field orthogonal to the electric field within the channel. Due to the mass differential between electrons and ions, the ions are largely unmagnetized and accelerated out of the channel by the electric field. The electrons, however, are strongly magnetized and confined to an azimuthal $\vec{E} \times \vec{B}$ drift within the channel (shown in Fig. 3.1 as a pink arrow). Specifically, the electrons spiral along the radial magnetic field lines within the channel, bouncing back and forth between channel walls while drifting azimuthally around the channel. The accelerated ions are then neutralized downstream by additional electrons from the cathode.

Despite the confining effect of the magnetic field, we still expect there to be some electron transport back to the anode due to electron-neutral and electron-ion collisions. As these particles collide, the electrons are able to be knocked out of their motion around the magnetic field lines and consequently migrate towards the anode. The small amount of electron current that reaches the anode as a result of this cross-field transport manifests as a loss in the Hall

thruster performance. However, the magnitude of this loss is much greater than what would be predicted with classical collisions [24]. One of the major open questions in Hall thruster physics is why this anomalous electron transport exists and how it may be characterized. For effective Hall thruster operation, it is useful to define a Hall parameter Ω_e , and the requirements for efficient operation:

$$\Omega_e = \frac{\omega_{ce}}{\nu_e} \gg 1. \quad (3.1)$$

This parameter is the ratio between the electron cyclotron frequency ω_{ce} and the electron collision frequency ν_e . The Hall parameter must be sufficiently high in a Hall thruster such that a given electron may complete multiple rotations around a magnetic field line before colliding with another particle or a wall and undergoing cross-field transport.

3.3 Theoretical Framework for Characterizing Performance

In this section, we describe the framework we employed to evaluate the performance of this Hall thruster. We overview key global parameters such as thrust, specific impulse, and anode efficiency, as well as definitions for the constituent efficiency modes of the thruster.

3.3.1 Global Performance Metrics

The primary measurements we used for determining global performance metrics for our thruster are thrust, T , a measure of the total force imparted by the propulsion device, and specific impulse, I_{sp} , a measure of the exhaust velocity of the propellant:

$$I_{sp} = \frac{v_{ex}}{g_0} = \frac{T}{\dot{m}g_0}, \quad (3.2)$$

where v_{ex} is the averaged exhaust velocity, \dot{m} is the total mass flow rate, and g_0 is the gravitational acceleration. We note that given a constant mass flow, there is a direct inverse relationship between thrust and specific impulse. We also define the total efficiency:

$$\eta_{tot} = \frac{T^2}{2\dot{m}P}, \quad (3.3)$$

where the total mass flow includes the anode and cathode flows, $\dot{m} = \dot{m}_a + \dot{m}_c$, and the total power includes the discharge and magnet power, $P = P_d + P_{mag}$. The discharge power is

calculated from $P_d = V_d I_d$. The total efficiency indicates the fraction of total power into the system converted to useful kinetic energy for thrust generation.

In order to characterize how well the discharge power is converted into thrust power without the effect of losses from the magnets and cathode, we also introduce the anode efficiency:

$$\eta_a = \frac{T^2}{2\dot{m}_a P_d}. \quad (3.4)$$

We note here that for high facility background pressures, neutral ingestion can act as an additional source of propellant flow into the channel, \dot{m}_{fac} . This can artificially raise estimates for the thruster efficiency. To correct for this, we follow Refs. [25, 26] in defining a facility correction factor:

$$\eta_{fac} = \frac{\dot{m}_a}{\dot{m}_a + \dot{m}_{fac}}. \quad (3.5)$$

This correction should be applied to calculations of the anode efficiency, total efficiency, and mass efficiency (Sec. 3.3.2). We apply this correction when operating at atypically high pressures; at standard operating pressures, η_{fac} is small enough as to be within the noise of uncertainty for these efficiencies. To determine the facility neutral ingestion for our experimental setup, we employed the model motivated in Ref. [25] updated for the pump configuration we use (Sec. 4.3). We note here that this model does not account for additional “effective” neutral ingestion that may not be captured by the model’s physical description of particle interactions with the facility and pumps. This may become a more prescient source of error at high powers and particularly high currents due to the larger flow rates into the chamber.

3.3.2 Phenomenological Efficiency Model

The anode efficiency can be broken down into various non-ideal processes within the Hall thruster. We quantify five partial efficiencies following the Hall thruster model developed by Hofer et al. [18, 27, 28, 29]:

$$\eta_a = \eta_b \eta_q \eta_m \eta_v \eta_d. \quad (3.6)$$

Here η_b is the beam current utilization efficiency, η_q is the charge utilization efficiency, η_m is the mass utilization efficiency, η_v is the voltage utilization efficiency, and η_d is the plume divergence efficiency. We note here that while the classical nomenclature refers to most

of these as efficiency modes (or components) as “utilization efficiencies,” we use these two terms interchangeably throughout this work (i.e. “mass utilization” meaning the same thing as “mass efficiency”). We also use “beam current efficiency” interchangeably with “beam efficiency” and “current efficiency.” We can use this model to identify how specific processes within the thruster are contributing to performance losses. We describe the physical significance of each of these efficiency modes in this section.




Figure 3.2: Cross-sectional view of Hall thruster and parameters defining efficiency terms.

3.3.2.1 Beam Current Efficiency

The ratio of ion current in the plume to the discharge current is

$$\eta_b = \frac{I_b}{I_d}, \quad (3.7)$$

where I_b is the ion beam current. This metric captures the fact that not all of the discharge current is carried by ions, as some is carried by electrons from the cathode to the anode, i.e. $I_d = I_b + I_e$. Lower values of beam utilization correspond to reduced electron confinement in the thruster channel. This efficiency is almost always the lowest efficiency mode in Hall thrusters and is believed to be a major driver for degraded performance at higher current density. We expand upon this in Sec. 3.6.3.3.

3.3.2.2 Charge Efficiency

The decrease in efficiency from multiply-charged ions in the beam is

$$\eta_q = \frac{\left(\sum_n \frac{\Omega_n}{\sqrt{Z_n}}\right)^2}{\sum_n \frac{\Omega_n}{Z_n}}, \quad (3.8)$$

where Z_n is the charge state of the n^{th} ion species, $\Omega_n = \frac{I_n}{I_d}$ is the current fraction of the n^{th} ion species, and I_n is the total contribution to current from a given ion species. This efficiency encapsulates the fact that the plume is not monoenergetic.

3.3.2.3 Mass Efficiency

The ratio of beam ion mass flow rate to anode neutral mass flow rate is

$$\eta_m = \frac{\dot{m}_b}{\dot{m}_a} = \frac{m_n I_b}{e \sum_n \frac{\Omega_n}{Z_n}} = \xi \eta_b \sum_n \frac{\Omega_n}{Z_n}, \quad (3.9)$$

where \dot{m}_a is the neutral anode mass flow rate, \dot{m}_b is the ion beam mass flow rate, m_i is the ion mass, e is the elementary charge, and $\xi = \frac{I_a m_i}{e \dot{m}_a}$ is a value defined as the exchange ratio. This efficiency loss is due to the fact that not all neutrals are ionized and accelerated out of the channel, therefore not contributing to thrust. We note that like the anode and total efficiencies, the mass utilization is corrected by η_{fac} to account for neutral ingestion.

3.3.2.4 Voltage Efficiency

The conversion of applied voltage into ion velocity is

$$\eta_v = \frac{V_a}{V_d}, \quad (3.10)$$

where V_a is the average acceleration voltage, the potential drop through which ions are accelerated in the channel. This efficiency represents the loss in discharge voltage that is not available for accelerating ions. In a Hall thruster, a small amount of the total discharge voltage (typically a few tens of volts) is needed to extract electrons from the cathode and join the main thruster beam. This cathode coupling voltage means that main beam ions will see a potential drop which is less than the total discharge voltage. The voltage utilization efficiency thus measures the fraction of the discharge voltage which is usable for ion acceleration.

3.3.2.5 Divergence Efficiency

The ratio of radially-directed to axially-directed momentum in the beam is

$$\eta_d = (\cos \theta_d)^2, \quad (3.11)$$

where θ_d is the divergence angle of the beam as measured from channel centerline. This efficiency captures the fact that only axially-directed momentum imparts thrust.

3.3.3 Mass Efficiency Scaling with Plasma Parameters

As we will establish in Sec. 3.5.5, the main driver for the efficiency gap between xenon and krypton is the mass utilization (Eq. 3.9). In this section, we derive a simplified 0D expression for the mass utilization in terms of plasma parameters that allows to better understand what drive trends in this efficiency. First, we rewrite the mass utilization efficiency as

$$\eta_m = 1 - \frac{\dot{m}(L)}{\dot{m}(0)} \approx 1 - \frac{n_n(L)}{n_n(0)}, \quad (3.12)$$

where \dot{m} is the neutral flow and n_n is the radially-averaged neutral density. Here, the values in parentheses denote axial location z , where $z = 0$ is the anode and $z = L$ represents the location where ionization has effectively ended (see Fig. 3.3). We have invoked in this derivation the definition of mass flow rate $\dot{m} = m_n n_n v_n A_{ch}$, where m_n is the mass of the neutral (equivalent to the mass of the ion m_i), v_n is the neutral velocity, and A_{ch} is the channel area. We note here that the approximate equivalence in this expression is because we have assumed both the neutral velocity and channel area to be constant despite both of them experiencing slight variations from $z = 0$ to $z = L$. We expect both of these values to be higher downstream; for velocity, this is due to expansion from the anode to the channel (which has a larger area), and for area, this is due to the chamfer at the exit plane of the thruster that is associated with a larger channel area.

To formulate an expression for the neutral density at the end of the region of interest, we consider the quasi-1D continuity equation for neutrals in the thruster channel at steady-state:

$$v_n \frac{\partial n_n}{\partial z} = -k_{iz} n_e n_n, \quad (3.13)$$

where k_{iz} is the rate coefficient for ionization averaged over a Maxwellian electron energy distribution and n_e is the plasma or electron density. All quantities in this expression are averaged over the channel area. We next make the strong assumption that the plasma




Figure 3.3: Notional diagram of neutrals flowing through a channel and being ionized. White circles represent neutrals and gray circles represent ions.

properties are approximately constant axially in the channel such that we can write

$$n_n(L) \approx n_n(0) \exp \left[-\alpha \frac{\langle k_{iz} \rangle \langle n_e \rangle}{\langle v_n \rangle} L \right], \quad (3.14)$$

where $\langle x \rangle$ denotes a value averaged over the axial range from $z = 0$ to $z = L$. Here we have introduced a scaling factor α to account for the fact that we typically use values of plasma parameters along channel centerline but calculate the mass utilization for the entire channel. The ionization rate coefficients and plasma densities near the wall are lower than on channel centerline, so we expect $\alpha < 1$.

Equation 3.14 physically suggests that the neutral density decays exponentially with distance from the anode. This stems physically from the fact that the neutral population is being depleted by ionization. While the assumptions made in deriving this expression are admittedly strong, the exponential decay of neutral density has been previously observed on Hall thrusters [30, 31]. This lends credibility to our simplifications. Moreover, the simplified scaling of neutral density and its dependence on spatially-averaged properties lends itself to intuitive physical interpretation. To this end, we can substitute this expression into Eq. 3.12 to yield

$$\eta_m \approx 1 - \exp \left[-\frac{L}{\lambda_{iz}} \right], \quad (3.15)$$

where we have defined an ionization mean free path

$$\lambda_{iz} = \frac{\langle v_n \rangle}{\alpha \langle k_{iz} \rangle \langle n_e \rangle}. \quad (3.16)$$

In deriving this expression, we have made the simplifying assumption that we have only singly-charged ions. This result intuitively shows that the mass utilization improves when the neutrals stay in the channel longer (L increases or v_n decreases) or have a higher chance of ionization (k_{iz} or n_e increase). As the ionization mean free path λ_{iz} decreases, the amount of propellant that is ionized increases.

We note that we use ionization rate coefficient instead of ionization cross-section, $k_{iz} = \sigma_{iz} v_e$. While the ionization cross-section is a useful tool in visualizing why ionization improves at higher electron temperatures, the ionization rate coefficient allows us to capture the effect of electron velocity while making fewer assumptions about the thermal speed. The ionization rate, taken by integrating the ionization cross section over a Maxwellian electron energy distribution function, has a direct non-linear dependence on function of electron temperature as later shown in Fig. 3.6. This method follows from Ref. [32] with data taken from Refs. [33, 34].

Equation 3.15 provides a framework for understanding the differences between the scaling of xenon’s and krypton’s mass utilization as well as the various techniques for improving it. In Sec. 3.5, we explore how changes in design and operation may affect this efficiency, as well as how it may scale differently with different propellants.

3.4 Historical Designs and Scaling Laws

There are multiple design variations of the Hall thruster, including the thruster with anode layer (TAL) [35, 36] and the wall-less Hall thruster [37, 38]. In this work, we focus on the “standard” design of a Hall thruster, also known as a stationary plasma thruster (SPT). The typical design of a Hall thruster has electromagnets or permanent magnets supplying the radial magnetic field within the channel, a ceramic discharge chamber, and a conducting anode that doubles as a baffle for the incoming neutral flow.

Traditional scaling laws for Hall thrusters dictate that the ratios of channel length, width, and diameter are approximately constant. The channel must be sized relative to the ions and electrons such that

$$r_{L,e} \ll w_{ch} \ll r_{L,i}, \quad (3.17)$$

where $r_{L,s}$ is the Larmor radius of a given species s and w_{ch} is the width of the discharge channel. This relation ensures that the electrons are free to move in their trapped drift current while the ions are free to be accelerated axially. Higher operating currents require larger thruster sizes and longer discharge channels. Historically, Hall thrusters have been designed such that they operate at an optimal current density of 100–200 mA/cm² [18, 39]. Thrusters with higher current densities have primarily been precluded by concerns with thermal limits, device lifetime, and required mass. Additionally, there is the possibility that beam utilization may significantly decrease at higher current densities. We discuss the validity of this latter concern in Sec. 3.6.3.3.

3.4.1 Magnetic Shielding

Historically, the greatest drawback of Hall thrusters has been their limited lifetime due to channel erosion. During an experiment in the late 2000’s, Aerojet Rocketdyne operated a thruster long enough to reach a zero-erosion state where the channel walls stopped eroding [40, 41]. In turn, researchers at the Jet Propulsion Laboratory designed a new magnetic field topology known as magnetic shielding [14, 15]. With this field design, the magnetic field lines are shaped such that energetic ions are directed away from the walls, therefore reducing channel erosion while maintaining the radial magnetic fields needed in the center of the channel. This technique has been shown to improve the lifetime of Hall thrusters by an order of magnitude over unshielded thrusters with only a minor (<2%) decrease to their efficiency [15].

As detailed in Ref. [14], magnetic shielding relies on properties along the magnetic field lines. In particular, we note that magnetic field lines are isothermal,

$$T_e = T_{e0}, \tag{3.18}$$

and that the potential along field lines is near-constant with low temperatures,

$$\phi = \phi_0 + T_{e0} \ln(n_e/n_{e0}), \tag{3.19}$$

where T_{e0} , ϕ_0 , and n_{e0} are integration constants. Details of these relationships and why they exist may be found in Ref. [14]. When there is a finite electron temperature, the potential along the magnetic field can be variable such that there is a parallel component of the electric field; this in turn leads to ion acceleration to the walls when the magnetic field lines are connected to the surface of these materials. Energetic ions are accelerated by this electric field to the walls, knocking off particles and eventually causing erosion.

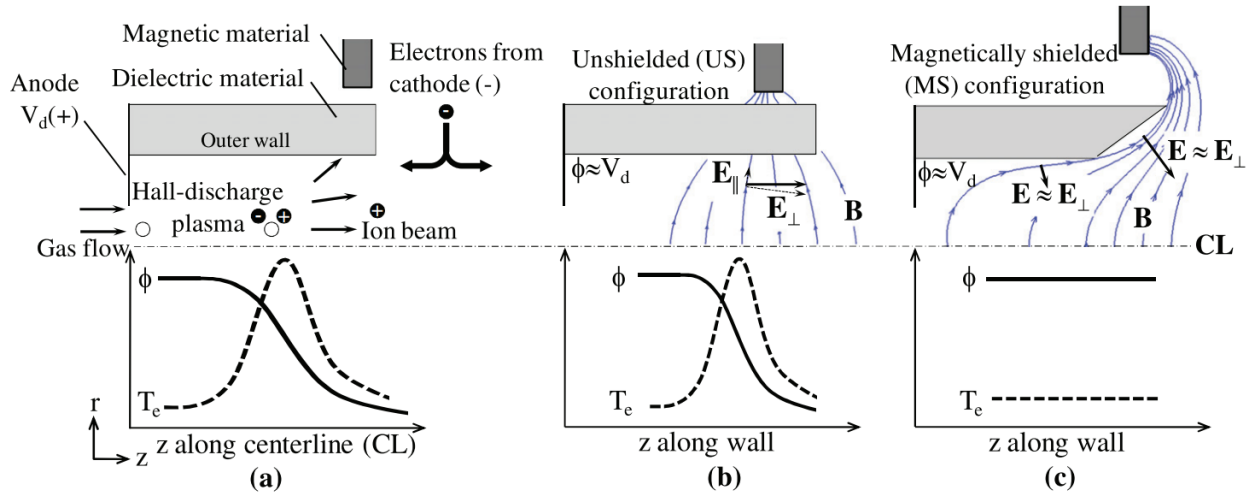


Figure 3.4: Magnetic field topography for unshielded (US) and magnetically shielded (MS) configurations, showing electron temperature and plasma potential a) for US along channel centerline, b) for US along channel wall, and c) for MS along channel wall. Reproduced from Ref. [41].

The magnetic shielding configuration circumvents this issue by curving the magnetic field lines such that the line closest to the channel walls is parallel to the walls and has low electron temperatures (see Fig. 3.4). While no shielded thrusters have yet been flown in space, there are multiple lab devices with magnetic shielding and planned flight thrusters. Laboratory thrusters include the H6MS, the original 6-kW thruster that was modified [14, 15], the H9, its 9-kW class cousin [42, 43], and MaSMi, a low-power Hall thruster [44]. Aerojet Rocketdyne’s Advanced Electric Propulsion System (AEPS) thruster [45], based on its lab counterpart HERMeS [46, 47], is a magnetically shielded flight thruster slated to fly in 2025 on the Artemis mission as part of the Lunar Gateway Station’s power and propulsion element [48].

Between unshielded and shielded thrusters, there are a number of changes in both the plasma properties in the plume as well as overall efficiencies. For example, shielded thrusters have higher electron temperatures along channel centerline compared to unshielded thrusters. Shielded thrusters also have an acceleration region displaced downstream of their unshielded counterparts [15, 49]. In regards to efficiency modes, the voltage and current utilizations increase, while the mass, charge, and divergence efficiencies decrease [49]. The mass utilization is due to the changed plasma formation of shielded thrusters—because the plasma density is concentrated in the middle of the channel, there is a “gap” along the walls of the channel where neutrals are able to leak out. We therefore expect the mass utilization to decrease; experimentally, we indeed see that it does by $\sim 2\%$ between an unshielded and shielded

topology on the same thruster. Concurrently, higher divergence angles are associated with acceleration regions that are further downstream—we therefore see a decrease of $\sim 5\%$ in the divergence efficiency between the unshielded and shielded configurations of the same thruster [49]. However, the overall total efficiency drop between unshielded and shielded thrusters is less than 2% [49], indicating a relatively minor hit to performance compared to a large increase in device lifetime.

3.4.2 Chamber Material Selection

The advent of magnetic shielding enables other changes to the traditional Hall thruster design rules as well. In particular, the significant reduction in wall interactions means that we may be able to explore materials other than boron nitride—one of the more traditional choices—for our discharge chambers. Boron nitride (BN) or borosil (BNSiO₂) are typically used due to their low secondary electron emissivity (SEE), low sputtering yield, and isolating properties. However, their fragility and high price make them non-ideal materials in many cases. The reduced plasma interaction with the walls on shielded thrusters opens up the possibility of using non-insulating materials.

Material	Density (g/cm ³)	SEE at 100 eV	Thermal cond. (W/mK)	Max temp. rating (°C)
Boron nitride	2	1.4–1.6	30	2000
Graphite	1.85	0.5–0.9	90	2000
Stainless steel	8	1.2–1.5	15	1400
Diamond	3.5	0.9–1.1	2000	3000

Table 3.1: Approximate material properties of possible discharge chamber materials.

In Tab. 3.1, we show the relevant properties for some of the candidate materials for Hall thruster discharge channels. In addition to these tabulated values, there are some drawbacks and benefits to each of these materials. For example, while boron nitride has significant heritage in the EP community, it is very fragile and expensive to machine. Graphite may be uniquely well-suited for a number of reasons. Besides its low SEE and high thermal conductivity as shown in Tab. 3.1, graphite is more resistant to breaking and may also present a “fly-like-you-test” situation. As described later in Sec. 4.3, there are chevroned graphite panels at the downstream end of our test facility. Over time, carbon is sputtered off of this beam dump and redeposited all around the vacuum facility, including on the BN discharge chamber of the thruster. This means that even a BN discharge in a Hall thruster will eventually be coated in a thin layer of carbon as shown in Fig. 3.5. Like the operational

choice of tying the cathode electrically to the body of the thruster [50, 51], a point which we expand upon in Sec. 4.2.1, testing with an all-graphite discharge would allow us to test in a way that may also be converted into a flightlike design.




Figure 3.5: Diagram showing method by which Hall thruster BN chambers may become coated in graphite. a) An ion from the plume hits the graphite beam dump, knocking off a neutral carbon atom. b) Carbon atom sticks to BN channel. c) Over time, BN channel becomes coated in carbon.

While the use of graphite has traditionally been limited due to its conductivity and the possibility of arcing to the plasma, magnetic shielding enables its use as a Hall thruster chamber wall material. Indeed, a previous study has already shown that a Hall thruster with a partial graphite chamber is capable of operating stably and without arcing events up to 6 kW [52]. Additionally, studies have indicated that the efficiency of a Hall thruster with a graphite chamber is within 2% that of a Hall thruster with a BN chamber [52, 53, 54]. This modification has previously been demonstrated with multiple thrusters, indicating stable operation with a complete graphite discharge chamber up to 18 kW [54]. Additionally, there was evidence that the wall temperatures were reduced when operating with graphite [52, 54], boding well for expanding Hall thruster operation to a wider power regime.

We note that the last entry in Tab. 3.1 is diamond. While not a common Hall thruster material, it has been used in Hall thruster designs in the past [55]. Diamond's isolating properties and high thermal conductivity make it an alluring material for Hall thrusters, particularly in isolating the anode from the thruster body. Its hardness is also promising from a robust design standpoint. However, diamond is (infamously) expensive, and coating graphite with a diamond layer is not—as we learned the hard way—an easy task. For these reasons, diamonds are not this girl's best friend. With that said, we hope future researchers may find more success in implementing diamond into their Hall thruster designs.

3.5 Krypton Operation

In this section, we make the case for running Hall thrusters on propellants other than xenon and specifically for krypton. We discuss why krypton is an especially appealing alternative and the attempts that have been made to date to improve its performance.

3.5.1 Motivation for Alternative Propellants

The ideal EP propellant for a high thrust density system is a monotonic gas with low ionization energy, high mass, and little reactivity. Having a monotonic gas means that energy is not wasted into dissociation, low ionization energies correspond to less energy required to convert neutrals into ions, and high masses are associated with high thrust. On the contrary, due to the relationship between specific impulse and thrust (Eq. 3.2), propellants with high mass will correspondingly have a lower specific impulse.

Gas	Mass (amu)	First ioniz. energy (eV)	Storage density (g/cm ³)	Atm. abundance (%)	\$/L
Xe	131.3	12.1	1.6	0.0000087	60 [†]
Kr	83.8	14.0	0.5	0.000114	0.80 [†]
Ar	39.9	15.8	–	0.934	<0.01
N ₂	28	14.5*	–	78.1	<0.01
Bi	209.0	7.3	9.8	–	–
I ₂	126.9	10.5*	4.9*	–	–

Table 3.2: Gas properties of xenon, krypton, argon, neon, nitrogen, bismuth, and iodine [56, 57]. The conversion from % to ppm is a factor of 1×10^{-4} . *For a monatomic gas. †Price from 2023 Q1 average [5].

As briefly discussed in Sec. 1.1, xenon has for decades been the main operating propellant for Hall thrusters due to its high mass and low ionization energy (Tab. 3.2). However, there are many factors that may motivate moving away from xenon. First, as previously mentioned, the price of xenon is not only high but also volatile, fluctuating from less than \$20/L in 2020 to a peak of over \$100/L in 2022 [5]. While the severity of the price fluctuation in recent years has been due to the political climate—namely, the first war in Europe since WWII—the limited number of global xenon suppliers still may pose a risk for longer and more ambitious missions [58, 59, 60, 61]. Additionally, its rarity in the atmosphere means that for high-power human missions, obtaining enough xenon for not only the mission itself but also the ground testing becomes a challenge in and of itself. Estimates suggest that a

crewed Mars mission would require more than 10% of the annual global xenon production [60], making xenon an unfeasible propellant choice.

A number of propellants besides xenon have been proposed and in many cases tested on Hall thrusters. This includes argon [20, 62], nitrogen or air [62, 63, 64], bismuth [65, 66, 67, 68], and iodine [56, 69]. We neglect to mention krypton at this point because we are saving the best for last. Some of the more commonly tested propellants and their relevant material properties are tabulated in Tab. 3.1. Note that the storage densities are for near room temperature conditions as per Ref. [56].

Argon is cheap and abundant but historically has very poor performance [62], with recent improvements on a low-power system [20]. Nitrogen is subject to dissociation in addition to ionization, representing an efficiency loss during operation. Finally, while the bismuth and iodine have shown promising performance [56, 65, 66, 67, 69], their high vapor pressures require an initial heating source that represents another loss in the system and an additional test facility challenge.

3.5.2 Advantages of Krypton

Krypton has a number of traits that makes it an attractive alternative to xenon besides the fact that it is purple and very pretty. One factor is its relative similarity to xenon. Both are non-reactive noble gases, simplifying the analysis of their ionization—rather than forming multiple different types of ions (as propellants like nitrogen and water would), xenon and krypton are simply ionized into different charge states of xenon and krypton ions. Located only one row above xenon on the periodic table, krypton is also the noble gas with the closest mass. Although argon is more abundant in the atmosphere than krypton (see Tab. 3.2), its significantly lighter mass translates to significantly lower thrusts and makes it a more distant cousin from xenon than krypton is. Krypton therefore offers the nearest and most direct comparison to xenon.

Besides its relatively low performance (Sec. 3.5.4), there are other issues with krypton as well. Krypton is extracted by the same process as xenon and therefore subject to the same fluctuations in price, albeit to a lesser degree, and has an even worse storage density compared to xenon. This worse storage density (Tab. 3.2) translates to krypton requiring tanks that are three times larger than those required for xenon. This poses a challenge for both crewed missions and robotic deep space missions due to storage volume limitations. While krypton may not be the future of all electric propulsion, understanding the differences between operating on it versus xenon may be extensible to other propellants as well.

3.5.3 History of Krypton

Despite krypton’s name, it has nothing to do with Kal-El, the Kryptonians, or Kryptonite, and in fact is not even green (at least as a plasma). Krypton was in fact discovered before xenon and was named “krypton” for its bashful nature (“kryptos” meaning “hidden” in Greek) [70]. In order to discover krypton, William Ramsay and his student Morris Travers needed to remove the presence of water, oxygen, nitrogen, helium, and argon from air. The remainder produced a spectrum with yellow and green lines, the latter of which likely inspired the color of Kryptonite. Neon (“neo” meaning “new” in Greek) and xenon (“xeno” meaning “other” in Greek) were discovered after, having an even smaller abundance in the atmosphere. The name of xenon would then be meaningless without the discovery of krypton.

3.5.4 Krypton Performance

While krypton is clearly a cooler propellant than xenon and has significant operational benefits, there remain two major limitations for adopting this propellant, particularly for deep space applications. The first is its storage density, which is a third that of xenon at the same pressures and temperatures (Tab. 3.2). This limitation can lead to prohibitively large storage tanks for missions with larger propellant requirements, though this can in part be overcome by storing at higher pressures. The second challenge stems from performance, as previous studies have shown that the anode efficiency of Hall thrusters operating on krypton is 5–15% lower than that of xenon at the same conditions [6, 7, 8, 9, 10, 11, 12]. This difference in efficiency has primarily been attributed to the lower mass utilization of krypton [8, 9, 10, 11, 12, 13]. In practice, this reduction can be a major mission driver, especially for high-power missions where overall system efficiency is critical.

To understand why krypton’s performance has historically been lower than xenon’s, we return to our 0D model for mass utilization. We note from Sec. 3.3.3 that the ionization mean free path λ_{iz} decreases (leading to better mass utilization) with a lower neutral velocity, a higher ionization rate coefficient, and/or a higher plasma density. First, we note that the neutral velocity scales as $v_n \propto 1/\sqrt{m_i}$ assuming that the neutral temperature is the same and approximating the neutral and ion masses as equal. Second, we find an expression for plasma density in terms of ion mass. We first define the ion current density $j_i = qn_i v_i$, where q is charge, n_i is ion density, and v_i is ion velocity. The ion velocity is dependent on ion mass, $v_i \propto 1/\sqrt{m_i}$, so the ion density scales as $n_i \propto \sqrt{m_i}$ when the current density and charge are unchanged. The differences between the current densities and charge states of xenon and krypton are much smaller than the difference in their masses [71], so we approximate these terms as constant between propellants. By invoking quasineutrality, we see that the plasma

(or electron) density in the near-field scales with mass, $n_e \propto \sqrt{m_i}$ [72]. The ratio between ionization mean free paths then becomes

$$\frac{\lambda_{iz(Xe)}}{\lambda_{iz(Kr)}} = \frac{\langle k_{iz(Kr)} \rangle m_{i(Kr)}}{\langle k_{iz(Xe)} \rangle m_{i(Xe)}}. \quad (3.20)$$

The mass of xenon is 131.3 amu in comparison to krypton’s 83.8 amu. In addition, as seen in Fig. 3.6, the ionization rate coefficient for xenon at a given electron temperature is higher than it is for krypton. The ionization rate coefficient, taken by integrating the ionization cross section over a Maxwellian electron energy distribution function [32, 33, 34], has a direct non-linear dependence on function of electron temperature. Given these ratios and subject to the simplifying assumptions we have made, Eq. 3.20 shows that the ionization mean free path of xenon is shorter than that of krypton. This in turn would translate to lower mass utilization and lower performance for krypton.

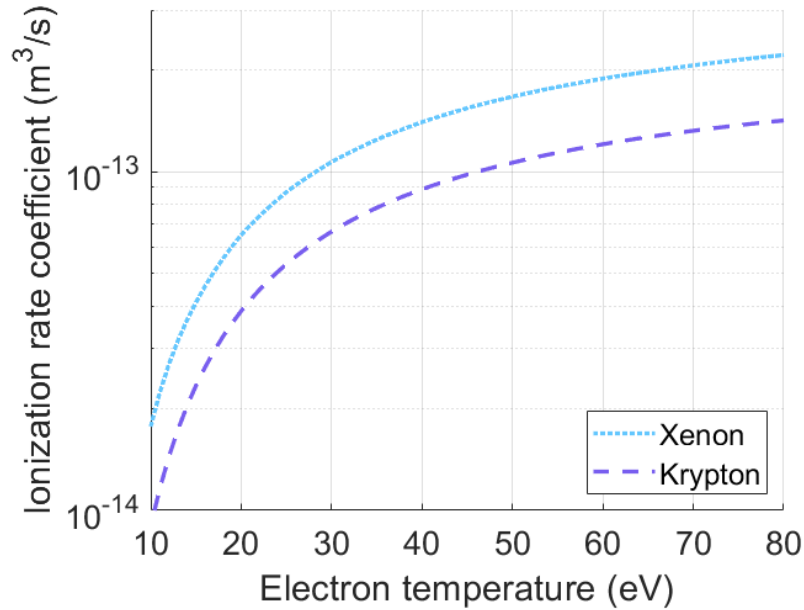


Figure 3.6: Ionization rate coefficient as a function of electron temperature for xenon and krypton. Data taken from Refs. [32, 33, 34].

It is apparent from Eq. 3.20 that at a given operating condition, the mass utilization of krypton will always be worse than that of xenon due to its lower mass. However, because the mass utilization exponentially approaches unity (Eq. 3.15), adopting strategies to improve this parameter for both gases will eventually close the gap. In light of this scaling and given the appeal of using krypton, there are multiple strategies that may be applied to increase the efficiency of krypton.

3.5.5 Strategies for Improving Krypton Performance

In this section, we review different design decisions that may improve the mass utilization of krypton operation and leverage our 0D model to explain why these methods work.

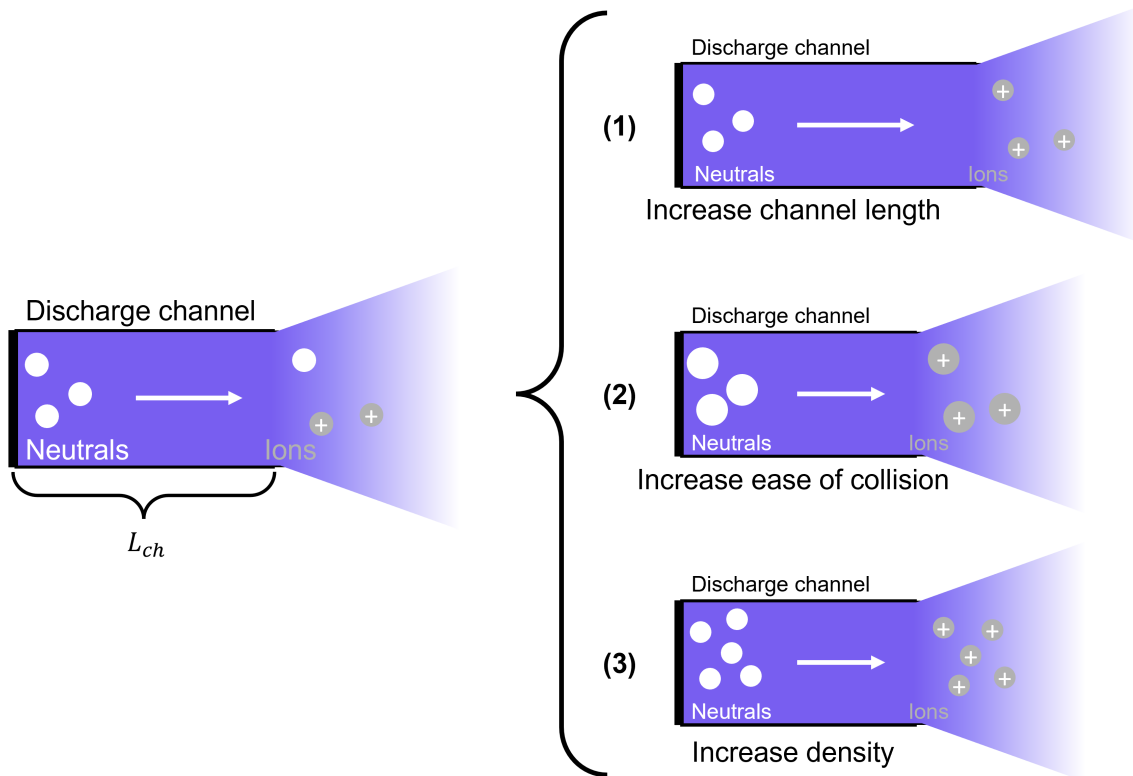


Figure 3.7: Various methods of increasing mass utilization efficiency in a Hall thruster: (1) increasing channel length, (2) increasing ionization cross-section, and (3) increasing plasma density.

Figure 3.7 qualitatively shows a number of different ways to improve the mass utilization. The first method is to simply lengthen the channel. Returning to Eq. 3.15, this has the effect of reducing the ratio of the characteristic ionization length to the channel length. This has the effect of giving the particles within the channel more time to have ionization events, leading to more of the neutrals being ionized by the time they reach the exit plane of the thruster. Lengthening the channel does however introduce more wall losses as the ions transit the channel. Nevertheless, this method has been shown to improve krypton performance at a given operating condition with an otherwise similar Hall thruster design [11, 13].

The second method is to increase the likelihood of a given ionization event taking place. This may be done by increasing the electron temperature of the plasma, which in turns increases the ionization cross-section and ionization rate coefficient. We visualize this in

Fig. 3.7 as the particles becoming larger, i.e. the ionization cross-section improving. Quantitatively, this has the effect of decreasing the ionization mean free path (Eq. 3.16) by making ionization events more frequent. This indeed has been shown to improve mass utilization on both xenon and krypton operation [8, 10, 11] by increasing the discharge voltage, which in turn increases the electron temperature within the channel [18, 73].

Finally, the third method is to increase the density of the incoming neutrals, typically by increasing the amount of current given a constant thruster discharge area. This also effectively increases the likelihood of ionization by increasing the plasma density and therefore reducing the ionization mean free path within the channel as seen in Eq. 3.16. This method has been particularly effective for krypton operation [9, 10, 12].

The exponential nature of mass utilization (Eq. 3.15) explains why we may expect the gap between krypton and xenon to close with all these techniques. While the mass utilization approaches unity in all cases, xenon's starting point is much closer than krypton's. In other words, we would expect xenon's mass utilization to asymptote at 1 earlier than krypton's, meaning that additional increases to channel length, discharge voltage, and/or discharge current would no longer improve xenon's mass utilization even though it would continue to improve krypton's.

3.6 High Power Operation

In this section, we motivate the need for high power electric propulsion at large and how Hall thrusters may be an ideal candidate to fulfill this niche.

3.6.1 Motivation

The development of high-power (>100 kW), lightweight, and high-performance electric propulsion technologies is critical for enabling the next generation of deep-space crewed missions [1, 2, 3, 4]. In particular, we believe they may enable crewed missions to Mars. However, there are a number of challenges associated with achieving these high power levels with electric propulsion. One of the most critical issues is that of in-space power availability. While the amount of in-space power available for spacecraft continues to increase, the megawatt levels necessary for a crewed Mars architecture [1, 74] still pose a challenge in terms of the maximum available power for a given solar array. Another is the issue of facility testing. Due to the long-duration nature (in excess of tens of thousands of hours) of many EP missions, performing a full life cycle on the ground is often unfeasible. Additionally, there are multiple facility effects seen in ground testing that do not extend to space [75],

such as high background pressures [76], facility thermal loading [77], and high backscatter rates from facility materials [47]. A key open question in the EP community is then how best to extrapolate shorter duration ground-based tests to longer periods of operation in space.

While these are pressing issues, the focus of our work is on a third: that of thruster maturity. These include questions of thruster specific mass α (kg/kW), thrust density (N/m^2), and system redundancy. As outlined by the National Academy of Engineering in their 2022 NASA-commissioned study on a crewed nuclear-powered missions to Mars, baseline requirements for future EP systems include an input power greater than 2 MW, specific impulse greater than 2000 s, efficiency greater than 50%, and system specific mass less than 5 kg/kW [1]. Additionally, as pointed out in a study by Dankanich [78], rapid transits to Mars could require thruster specific masses lower than 0.5 kg/kW. In Sec. 2.3, we briefly discussed some of the different types of EP thrusters and their different operating regimes. Here we make the case for Hall thrusters as being particularly well-suited for high powers.

3.6.2 Suitability of Hall Thrusters

While a number of technologies have been proposed to meet these high power requirements [1, 2], the Hall thruster is a particularly promising candidate that has been featured in recent trade studies on NASA’s proposed nuclear electric propulsion plan for Mars [79]. Hall thrusters have extensive development and flight heritage at power levels below 12.5 kW [80, 81, 82, 83, 84, 85], and laboratory models have been demonstrated up to 100 kW [26, 65, 66, 86, 87]. Hall thrusters have also been shown to be generally capable of meeting the specific impulse and efficiency requirements for near-term nuclear architectures in the ~ 1 –10 kW power range [1, 79].

Despite these attractive capabilities, the achievable thrust density and specific mass of Hall thrusters remain a major potential limitations for high-power applications. While achieving high powers in general is critical, doing so while maintaining feasible thruster masses and sizes is equally important. Historically, Hall thruster design has followed general scaling laws [18, 39, 88], including an “optimal” current density in the channel of 100–150 mA/cm² [18]. We discuss the physical basis for this reasoning in Sec. 3.6.3.3. As a direct result, the thrust density of most Hall thrusters has been limited to about 10 N/m² [89]. This is approximately an order of magnitude higher than other mature EP technologies like gridded ion thrusters, which makes Hall thrusters an attractive candidate for most power levels. However, in the 100 kW–1 MW power range envisioned for nuclear architectures, Hall thruster sizes can become prohibitively large. A different type of thruster technology that has an inherently high thrust density, the magnetoplasmadynamic thruster (MPD), has thus historically been

baselined for high-power operation [79].

With that said, there are recent developments in Hall thruster technology that may help pave the way for exploring Hall thruster operation at higher powers and thrust densities than before. As detailed in Sec. 3.4.1, magnetic shielding greatly reduces plasma flux to the interior walls of the discharge channel, a point which we return to in Sec. 3.6.3.3. This means that shielded thrusters may be better able to handle high power densities than their unshielded counterparts. Additionally, the recent pivot towards graphite discharge chambers enabled by shielding (Sec. 3.4.2) may be better for high power as well due to the material's higher temperature tolerance, thermal emissivity, and material robustness compared to the traditional ceramic materials like BN.

3.6.3 Strategies for Increasing Thrust Density

Due to the maturity and flight heritage of Hall thrusters compared to MPDs, it is clear that there is major incentive to explore methods for increasing Hall thruster thrust densities. We also note that in most Hall thruster mission architectures, there are multiple Hall thrusters instead of a singular unit [21, 48] for fault tolerance and redundancy. Additionally, there may be some minor performance benefits beyond just redundancy offered when running clusters of Hall thrusters [90, 91, 92]. We therefore target not the 2-MW power level for a given unit but instead the 100-kW range. In this section, we discuss a few different methods that have been used to increase the thrust density in a Hall thruster.

3.6.3.1 Nested Hall Thrusters

One method to improve thrust density is by utilizing more of the surface area of the thruster. Multiple discharge channels can be nested concentrically to increase thrust density and reduce specific mass [26, 92, 93, 94, 95, 96], resulting in what is known as a nested Hall thruster. By effectively splitting the power across multiple channels of different size, nesting can lower the overall system mass and footprint. For example, at the 200-kW power level, the diameter of a two-channel nested Hall thruster is approximately half that of a single-channel thruster [97]. Additionally, nested Hall thrusters offer the advantage of throttability, as individual channels can be turned on or off to prioritize high specific impulse or high thrust.

This design strategy led to the creation and demonstration of the X3, a 3-channel 100-kW class Hall thruster that achieved a thrust of 5.4 N tested under the NASA NextSTEP program in 2019 [26, 98, 99, 100, 101, 102]. There also may be some performance benefits to operating multiple thruster channels in a nested thruster due to magnetic field interactions and mass ingestion of un-ionized propellant between channels [26, 99, 103]. Unfortunately,

nested Hall thrusters have a number of engineering challenges due to their complexity that have precluded their widespread use [26, 102]. Additionally, even in the ideal limit where the nesting is sufficiently compact such that the entire thruster surface generates thrust, the increase in thrust density over a single channel thruster of the same size is only a factor of two. To this point, for the approximate dimensions of the X3 (Ref. [86]), the thrust density was the same as the nominal $\sim 10 \text{ N/m}^2$ typical of single channel thrusters.

3.6.3.2 Increasing Discharge Voltage

Another way to increase the thrust density of a Hall thruster is to increase the discharge voltage. One of the challenges associated with this is the higher thermal loads at higher powers. The TSNIIMASH DL-160, a laboratory water-cooled thruster tested in the 1970's [65, 66], followed this approach, demonstrating 140 kW at 8 kV discharge voltage on bismuth with a thruster efficiency greater than 70%. With its 160 mm diameter channel, this yielded a thrust density of over 100 N/m^2 , nearly an order of magnitude higher than the typical thrust density of a Hall thruster operating at 300 V [89]. With that said, the high voltages on this system led to a high specific impulse of $\sim 8000 \text{ s}$, which is associated with a thrust-to-power ratio likely too low for a crewed Mars architecture.

In summary, attempts to increase thrust density to date have resulted in prohibitively high specific impulse or only marginal improvements. This is ultimately because the proposed solutions were based on raising thrust density while respecting the historical design rule for “optimal” current density, which we explore in further detail in the next section.

3.6.3.3 Increasing Discharge Current

There are a number of challenges associated with increasing the current density in a Hall thruster. Two of the most critical are high thermal loads and decreased electron confinement. To this first point, operating at high currents (as with high voltages) involves the possibility of overheating due to excess power lost to the thruster body. Indeed, while state-of-the-art thrusters at lower powers are passively cooled and thermally steady, increasing the current density (and by extension power density) may lead to thermal loads that exceed the capability of standard designs to reject heat. One of the dominant loss mechanisms for Hall thrusters is electron thermal flux to the discharge chamber walls. Per Ref. [18], this power can be shown to scale as

$$P_w \propto A_w n_0 T_e^{3/2}, \quad (3.21)$$

where P_w is the power to the walls, A_w is the channel wall area, n_0 is the plasma density in the thruster discharge chamber, and T_e is the electron temperature immediately adjacent to the walls. As Eq. 3.21 shows, as plasma density increases, the power loss to the walls increases. To relate this to current density, we assume the conversion of all potential energy into kinetic energy and again invoke continuity to make the approximation that

$$j \propto n_0 \sqrt{\frac{2qV_d}{m_i}}, \quad (3.22)$$

where V_d is the discharge voltage and m_i is the ion mass. We can substitute this result into Eq. 3.21 to find

$$P_w \propto j A_w T_e^{3/2} \sqrt{\frac{m_i}{2qV_d}}. \quad (3.23)$$

Assuming the electron temperature and discharge voltage remain approximately constant as a function of current density, this result shows that as the current density increases, the power flux to the walls increases. This relationship indicates that at sufficiently high current densities, the thermal load may become prohibitive.

However, as mentioned in Sec. 3.4.1, magnetic shielding help to mitigate this effect. One comparison of temperatures of the unshielded H6 to the shielded H9 indicates that the wall temperatures for the unshielded case is about 17 eV while the shielded case is 3 eV [49, 104]. If we consider this factor of ~ 6 in electron temperature in terms of Eq. 3.23, we could potentially decrease the power loss to the walls by a factor of ~ 15 at a given current density. This in turn would allow us to significantly increase our current densities.

The other concern with high current density operation is loss in electron confinement, leading to a decrease in beam utilization (Eq. 3.7). Dannenmayer and Mazouffre [39] have argued that electron confinement will decrease at higher discharge currents. To motivate this argument, we return to our requirement of Hall parameter in Eq. 3.1 and note that in order for the magnetic field to maintain strong confinement of the electrons, $\frac{\omega_{ce}}{\nu_e} \gg 1$. Physically, this relationship requires that electron collisions, which can allow electrons to cross magnetic field lines, are infrequent on the time scale of electron precession. As we discussed previously, the actual form of electron transport is not known within Hall thrusters; however, in order to arrive at a classical scaling argument, we for now assume that the electron transport scales with the electron-ion collision frequency, i.e. $\nu_e \approx \nu_{ei}$, where

$$\nu_{ei} = 2.91 \times 10^{-6} n_e \ln \Lambda T_e^{-3/2}, \quad (3.24)$$

where all values are in SI units, n_e is electron density, T_e is electron temperature, and $\ln \Lambda$ is the Coulomb logarithm, which is for the most part invariant and about 10 for lab plasmas.

The Hall parameter criteria then becomes

$$\frac{eBT_e^{3/2}}{2.91 \times 10^{-5}m_en_e} \gg 1, \quad (3.25)$$

where e is fundamental charge, B is the magnetic field strength, and m_e is the electron mass. With current density defined as $j = en_e v_e$, we may rewrite this expression as

$$j \ll \frac{e^2 B v_e T_e^{3/2}}{2.91 \times 10^{-5} m_e}. \quad (3.26)$$

If we assume that the electron temperature (and therefore electron velocity) is driven solely by the discharge voltage, and that the discharge voltage and magnetic field remain unchanged, then we can see from Eq. 3.26 that there is an inherent upper bound in current density in order to ensure strong electron confinement. Physically, this relationship illustrates the fact that as current density increases, collisions in the channel increase, reducing electron confinement. It has been suggested that beyond a certain limit, the confinement decreases, and the overall efficiency drops [18, 39]. However, the assumptions underpinning this scaling may need to be re-evaluated with implications for being able to operate at higher currents. We return to this point in Sec. 6; in short, the so-called “optimal” current density may not exist at the value we believe it to due to the non-classical dynamics of electrons in Hall thrusters. Specifically, we assume classical scaling of electron collision frequency in deriving the expression found in Eq. 3.26. Given the existence of anomalous electron transport in Hall thrusters, it is not unreasonable to think that the scaling also does not follow classical trends.

With that said, there has been previous work investigating higher current densities. The SPT-100, nominally designed to operate at 1.35 kW (300 V and 4.5 A), was operated up to 2.4 kW (300 V and 8 A) [105]. This higher-power condition corresponds to a current density of 190 mA/cm². In an even further departure from these “typical” operating current densities, a 30 cm-diameter Hall thruster in the then-USSR was operated up to \sim 1800 mA/cm² [106]. However, the selected propellant for this second case was hydrogen, a very light gas. While the published literature does not include thrust values for this device, we suspect that its thrust-to-power ratios are too low to be usable for applications such as crewed missions to Mars.

3.7 Summary

In this section, we have introduced the motivation for our work. We began by describing what a Hall thruster is and outlining the framework we use for characterizing Hall thruster performance. Next, we review some of the historical choices made in Hall thruster design as well as recent innovations. We go on to discuss the benefits of alternative propellants on Hall thrusters and the work that has been done to improve their performance thus far, particularly for krypton operation. After outlining the benefits of high power and why Hall thrusters may be uniquely advantageous at high powers, we reviewed various methods of scaling Hall thrusters to high powers. This motivation leads us to the structure of this work: first a characterization of how xenon and krypton performance differ at baseline conditions on a magnetically shielded Hall thruster, then how their performances change at atypically high currents, and finally why these trends with increasing voltage and current exist and how they may be extended to other propellants.

CHAPTER 4

Experimental Apparatus

“对牛弹琴。”

–ancient Chinese proverb. [*Lit.*, playing the piano to a cow. *Fig.*, effective communication requires knowing your audience, or in this case, molding your readers into the perfect audience.]

4.1 Introduction

In this section, we describe the test articles, facility, and diagnostics used throughout the experiments detailed in this work. We also detail the methods used to analyze data obtained via these diagnostics and how we use them to infer plasma properties and efficiency modes. All tests were conducted at the Plasmadynamics and Electric Propulsion Laboratory (PEPL) at the University of Michigan (UM).

4.2 Thrusters

The H9 is a 9-kW class magnetically-shielded Hall thruster developed in partnership between the University of Michigan, Jet Propulsion Laboratory (JPL), and the Air Force Research Laboratory (AFRL) [42, 43]. It was built with the intention to serve as a common platform for shared research, a vehicle for investigation into new innovations in Hall thruster technology and physics, and a training instrument in educating the next generation of scientists in the field. The H9 typically employs a centrally-mounted LaB6 cathode with a graphite keeper [107] operating at a fixed 7% cathode flow fraction. As described in Refs. [42, 43, 54], the H9 has a stainless steel anode, graphite pole covers protecting the magnetic circuit, and a magnetically shielded field topology. The H9 has been experimentally characterized on both xenon and krypton as the subject of many performance and plasma property measurements

[42, 43, 54, 71, 108]. The nominal operating envelope of the H9 on xenon is approximately 4.5–12 kW discharge power, 1800–3000 s specific impulse, 290–700 mN thrust, and total efficiencies of 55–70%. The typical operating envelope on krypton is discharge powers of 4.5–9 kW, specific impulses of 1900–2700 s, thrusts of 260–350 mN, and total efficiencies of 45–55%.

We designed a modified version of the H9 known as the H9 MUSCLE to facilitate investigations into high current density operation. First, we replaced the baseline boron nitride channel with an entirely graphite discharge chamber. As discussed in Sec. 3.4.2, this change was enabled by the magnetically shielded topology and helps with the thermal environment of the thruster. For our second modification to the H9, we implemented an active water-cooling loop to maintain safe temperatures for the material of the magnetic circuit. This non-flightlike change alleviated thermal issues related to enhanced plasma flux to the channel walls at high current density and allowed us to focus on the more fundamental question of the role of current density in thruster performance. We note here that the cooling lines were designed to not make thermal contact with any of the propellant lines, therefore not affecting the temperature of the incoming flow to the anode. A similar modification for cooling has been applied in previous laboratory experiments of high-power Hall thruster operation [65, 66].

4.2.1 Cathodes

We use lanthanum hexaboride (LaB_6) cathodes throughout our work, which are favored due to their robustness and high current densities [109]. While in the context of this work, we are guilty of treating cathodes as a necessary evil, other brave individuals have devoted much more time and effort to understanding the physics of these devices [109, 110]. We typically run our thrusters at 7% cathode flow fraction (i.e. the flow through the cathode is 7% of the flow through the anode) and in a “cathode-tied mode,” electrically connecting the cathode to thruster body and isolating them from facility ground [46, 50, 51]. This is a configuration that, unlike tying the cathode to ground, can be replicated in space. Additionally, it has been shown to reduce Hall thruster pole erosion [51].

4.2.2 Thruster Operation

We used Alicat MC-series flow controllers, which have a measurement error of $\pm 0.1\%$ of full scale or $\pm 0.6\%$ of reading, for both anode and cathode flow. These flow controllers were calibrated independently for xenon and krypton before each test campaign with a MesaLabs Bios Drycal Definer 220 flow calibration unit, which have a 1% error in their standardized

measurements. Note that the volumetric flow rates are calculated with STP conditions assuming a temperature of 25°C, a conversion factor from sccm to mg/s of 1/11.12 for xenon and 1/17.48 for krypton. Multiple Magna-Power Electronics power supplies were available for thruster discharge power; we typically used the 60-kW (rated to 1000 V and 60 A) or the 150-kW (rated to 1000 V and 150 A) supplies in our testing. Power to the magnets and cathode heater/keeper was supplied through a series of Lambda TDK DC power supplies. The return from the discharge power was passed through a filter circuit with a 0.3 mH inductor, 47 μ F capacitor, and 10 Ω resistor to protect the power supply from oscillations in the thruster current.

4.2.3 Thruster Monitoring

We used a number of tools to monitor thruster and facility health. One of the most critical of these was thermocouples, which allowed us to monitor temperatures on the thruster, thrust stand or motion stages, cryopump panels, probes, and other important infrastructure. On the thruster, the thermocouple located at the inner front pole (IFP) tended to reach the highest temperatures due to its central nature and proximity to both the plasma and the magnetic circuit. Oscillations in discharge current, cathode current, discharge voltage, and cathode-to-ground voltage were monitored with current guns and a Keysight DSOX 3024A oscilloscope. For thruster health, the primary metric of interest is the relative peak-to-peak oscillation in the discharge current. Thruster and facility health telemetry including real time temperatures, oscillation strengths, facility pressures, and DC supply currents and voltages were monitored via an optically-isolated, automatically-logged data acquisition system (DAQ). Cameras were also mounted around the facility for visual monitoring of the thruster.

4.3 Facility

The Alec D. Gallimore Large Vacuum Test Facility (LVTF), detailed at length in Ref. [111], is a 6 m diameter and 9 m length vacuum chamber. There are eighteen total cryopumps located around the facility. Thirteen of these pumps are off-the-shelf PHPK TM1200i cryopumps capable of pumping \sim 35 kL/s of xenon each, colloquially known as “cryos.” The remaining five pumps are adapted from a custom design at JPL that utilizes Cryomech’s Gifford-McMahon AL600 cryocoolers. These pumps, colloquially known as “thumpers” due to the incessant thumping noises they make during operation (not unlike the bass line of an industrial garage band), were designed with a surface area that allowed them to match the

pumping speed of the PHPK cryos. Details on the sizing of these pumps can be found in Refs. [111, 112]. In total, LVTF is capable of pumping ~ 500 kL/s of xenon and ~ 600 kL/s of krypton.

A beam dump held at facility ground composed of multiple chevroned graphite panels is located at the downstream end of the chamber (see Figs. 5.4, 6.3, 7.2). Due to the energetic ions of the thruster plume, significant sputtering of the chamber can occur at high flow rates and/or long runtimes. By placing a graphite beam dump in the chamber, we are able to mitigate this effect as graphite has a relatively low sputtering yield. We discussed the significance of this in Sec. 3.4.2.




Figure 4.1: a) A photo of LVTF and b) a CAD of LVTF [111] showing the locations of pumps around the chamber. Note that the bottom left corner of the photo corresponds to the bottom right corner of the CAD.

4.3.1 Pressure Measurements

We use MKS Granville-Phillips Stabil ion gauges to measure pressure within our facility as per industry standards, mounted in plane with the thruster 1 m away [113]. These are composed of both a controller and a gauge—throughout this work, we used 370501 controllers and 360116 ion gauges. These controller-gauge systems are traditionally sent to an external company with known pressure measurements to obtain a calibration curve from the reading on the controller to the actual pressure. Typically, we sent them to LACO Technologies and had these systems calibrated on xenon. Pressures measured during krypton operation had a $1.48\times$ factor applied per manufacturer guidelines. There is a 20% uncertainty associated with these gauges [113]. When used in the presence of a plasma, these gauges may be

equipped with a stainless steel neutralizer tube containing grounded metal mesh.

Throughout the tenure of this work, the validity of these gauges and this procedure has been brought into question. According to experts at MKS [M. Kelly, personal communication], the gauges should be replaced every 6–12 months. As this was not expressed when the gauges were initially purchased, this replacement time was not common practice in the lab. Additionally, while the LACO procedure was called a “calibration,” they did not actually calibrate anything internal to the system and instead merely produced a curve of the actual pressure compared to the controller-gauge system’s measured pressure. Oftentimes, we found this fit to be a power law which, once again according to MKS, meant that the unit was no longer calibrated (the behavior for calibrated units is linear). As another concern, we used 370-series controllers with 360-gauges, which while electrically compatible, may have led to minor discrepancies in pressure readings. Finally, while the gauges’ so-called calibration with xenon has been standard practice due to the commonality of using xenon within the EP community, there has been a recent shift towards using alternate propellants. This is particularly true for this body of work.

As an additional measurement tool, we began using Instrutech Hornet ion gauges for some of this work (see Sec. 6.2) when we found that our ion gauges were no longer in calibration; these gauges by default are calibrated for nitrogen. It is our recommendation that in the future, gauges are calibrated/set to nitrogen—typically the default—and converted to the operating pressure for a given gas after testing. For xenon, this conversion is done by dividing the pressure by 2.87, and for krypton, this scalar is 1.94. Additionally, if the Granville-Phillip system continues to be used, the controllers should be regularly sent to MKS for electronic calibration while the gauges should be regularly replaced or properly calibrated by MKS (although it appears that MKS no longer offers this service). The LACO “calibration” process may then be avoided entirely. As a final note, we recommend using copper o-rings (instead of butane) when installing these gauges on vacuum flanges due to the lower temperature limits of butane.

4.3.2 Quartz Crystal Microbalance Sensor

As we described earlier in this section, a significant amount of carbon is backspattered from the graphite beam dump and deposited elsewhere in chamber, including on the thruster itself. To quantify the degree of backspatter in the facility, we may employed a quartz crystal microbalance (QCM) sensor to measure the rate of carbon deposition in the plane of the thruster from the graphite beam dump and probe shield. The QCM estimates the total mass deposition by determining the change in the resonant frequency of a quartz crystal as

material is deposited [114]. We assume in this work that only graphite was deposited onto the thruster and QCM for these measurements. We used a INFICON QCM front load single sensor mounted 1 m from the thruster and aligned with the plane of the thruster and facing downstream (Fig. 6.3). Attached to the sensor were cooling lines connected in parallel to the thrust stand cooling system to regulate the temperature. We also maintained a distance of ~ 1 m between the QCM and ion gauge to prevent the heating of the ion gauge from affecting the QCM. The rate of deposition was calculated by determining the slope of the thickness measurements over time and assuming only carbon deposition.

4.4 Diagnostics

In this section, we describe the various diagnostics used to infer performance metrics and plasma properties about the thruster. We also describe the various analysis methods that may be used to convert raw measurements to meaningful data.

4.4.1 Thrust Stand

Due to the small thrust outputs of electric propulsion systems, designing and implementing a diagnostic capable of measuring these small forces accurately is a non-trivial task. We employed a null-type inverted pendulum thrust stand [77, 92, 115] with spring flexures in the form of double-ended pivot bearings that provided a restoring force against the thruster weight. The thrust stand assembly was surrounded by a radiative shroud which was maintained at 20° C by a chiller pump to reduce thermal drifts. We operated the thrust stand in a null displacement mode by actively controlling the position of the thruster with a null coil solenoid and armature. A fiber-optic sensor monitored the displacement as an analog voltage with a slope of -2700 V/m. Meanwhile, an analog PID controller provided the control signal for the solenoid. We also maintained constant thruster pitch inclination via a stepper motor with software-based control. The thrust stand was calibrated with a series of weights, the range of which depended on the expected thrusts during a given test campaign.

With this approach, we took thrust points by recording the difference in solenoid current before and after shutting down the thruster, $I_{N,on}$ and $I_{N,off}$ respectively. We averaged over a window of time, typically on the order of five seconds, to determine these values of solenoid current. During this shutdown process, we eliminated the impact of cold gas contributions to thrust by cutting off gas flow to both the anode and the cathode with software-actuated solenoid valves. The uncertainty in this measurement is dependent on the variation in $I_{N,on}$ and $I_{N,off}$ as well as the goodness of calibration fit with the known weights. We performed

a least-squares fit to a curve of the form $\Delta I_N = mT + y$, where ΔI_N is the change in null coil current, T is the thrust (or calibration weight) force, and m and y are the slope and intercept of the line of best fit respectively. When taking measurements, we inverted this expression to determine the thrust as a function of the coil current during thruster operation and included an additional uncertainty term, yielding the form

$$T = \frac{[I_{N,on} - I_{N,off}] - y}{m} + U_T, \quad (4.1)$$

where U_T represents uncertainty arising from other sources.

We follow the methodology described in Ref. [116] to characterize these outside sources of uncertainty. In this work, Mackey et al. identified and calculated multiple sources of uncertainty individually and found that the dominant sources of error arose from the displacement drift and inclination drift of the thrust stand. We therefore consider only these two outside sources of uncertainty in the following analysis. We quantified uncertainty in our thrust measurements as

$$\delta T = \sqrt{\frac{\delta I_{N,on}^2 + \delta I_{N,off}^2 + \delta y^2 + T^2 \delta m^2}{m^2} + \delta T_x^2 + \delta T_i^2}. \quad (4.2)$$

Here, $[\delta I_{N,on}, \delta I_{N,off}]$ represent the variance in the null coil currents and $[\delta m, \delta y]$ denote uncertainty in the slope and offset of the fit to the calibration curve. We selected current measurement windows close enough to the shutdown point of the thruster that the timescale of the thermally-induced drifts we observed during thruster operation could be ignored. The δT_x term represents the displacement drift uncertainty [116],

$$\delta T_x = k\delta x, \quad (4.3)$$

where k is the spring stiffness and δx is the variation in displacement location. We obtained δx by converting the measured voltages to distances and taking the standard deviation in displacement distance over the current measurement range (where $I_{N,on}$ and $I_{N,off}$ were calculated). The δT_i term represents the inclination drift uncertainty [116],

$$\delta T_i = mg_0 \sin(\delta\theta_i), \quad (4.4)$$

where mg_0 is the weight of the thruster and $\delta\theta_i$ is the variation in inclination angle. As with δx , we calculated $\delta\theta_i$ by taking the standard deviation in inclination angle over the measurement range. Typically, the displacement drift constituted the largest source of error, followed by the current/slope uncertainty, and finally the inclination drift as the smallest.

4.4.2 Probes

In this section, we describe the basic principles of operation and analysis methods used for different probes. We give dimensions and standard operating conditions for each of these probes within the context of our lab and our experiments. When available, we show two sets of example data for each of the probes: one at typical Hall thruster operating conditions and one at atypically high current density operation. We also briefly discuss how we typically estimate uncertainty in our measurements. Photos of some of the commonly-used probes in these experiments are shown in Fig. 4.2.




Figure 4.2: Photos of the a) RPA, b) $E \times B$ probe, and c) Faraday probe used.

4.4.2.1 Retarding Potential Analyzer

A retarding potential analyzer (RPA) is a gridded electrode plasma diagnostic used to measure the ion energy distribution function (IEDF) of a plasma along the probe axis. These measurements of ion energy inform the calculation of voltage utilization efficiency (Eq. 3.10). During operation, a quasineutral plasma enters the aperture of the RPA and passes through the first grid, which floats to plasma potential and attenuates the plasma density. The second grid is the primary electron suppression grid, which is biased negatively to prevent electrons from passing through it. The third grid is the ion selector grid, which is biased to a user-controlled potential V_{bias} . Ions with energies above V_{bias} are able to pass through the grid while those with lower energies are repelled. The final grid is the secondary electron suppression grid, biased to the same potential as the other suppression grid, which prevents secondary electrons generated from ions impacting the grids from reaching the collector. This process ensures that only ion current is collected. We employed an RPA with an aperture area of 6.45 cm^2 and typically biased both suppression grids to -30 V relative to facility ground.

The ion selection grid voltage V_{bias} was swept from 0 to approximately twice the discharge

voltage. At 0 V, all ions from the plasma were able to make it to the collector, while at the highest voltage, effectively no ions were able to reach the collector. We differentiated the current collected as a function of voltage to yield an estimate of the IEDF expressed as a function of energy per unit charge. The location of the peak in the IEDF corresponds to the most probable ion voltage referenced with respect to ground, V_{MPV} . We in turn use this quantity corrected by measurements of the local plasma potential (see Sec. 4.4.2.2) to determine the acceleration voltage $V_a = V_{MPV} - V_p$, where V_p is the plasma potential. This informs our estimates for the voltage utilization efficiency (Eq. 3.10).

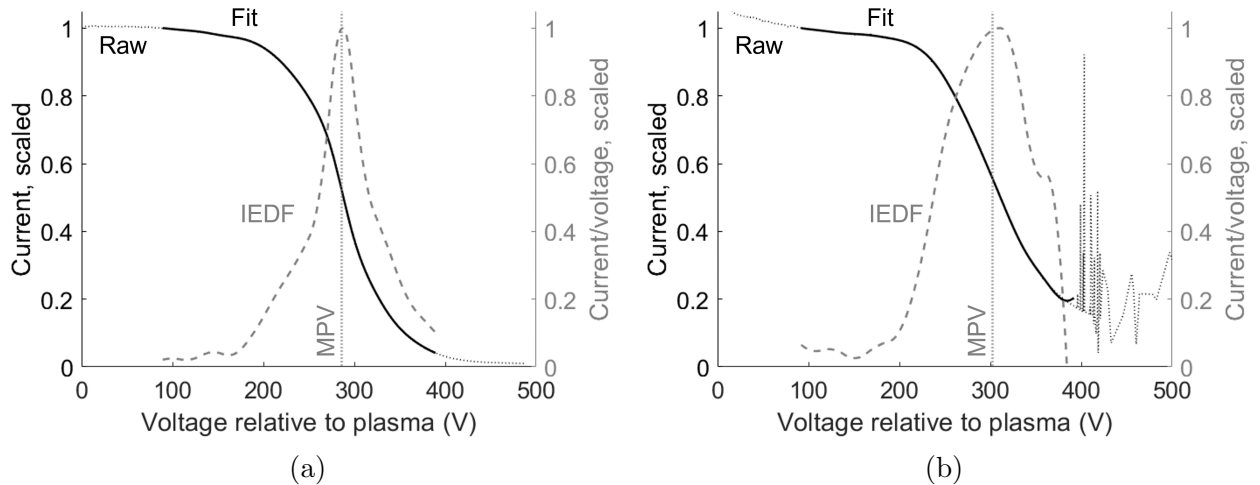


Figure 4.3: RPA data for xenon at a) 300 V and 15 A and b) 300 V and 125 A. The original current trace is shown as a dotted black line, a fitted smoothing spline is shown as a solid black line, the IEDF is shown as a dashed gray line, and the location of the most probable voltage is shown as a vertical dotted line. Note that these plots are corrected for plasma potential.

In Fig. 4.3, we show example smoothing spline fits of the RPA traces, computed IEDFs, and estimated most probable voltages. We employed bootstrapping (described in Sec. 4.4.2.6) to estimate the error in the most probable voltage by sampling a subset of the raw RPA trace.

4.4.2.2 Langmuir Probe

Langmuir probes are an important diagnostic tool for plasma physics due to their simple construction and wide range of uses. Its basic operation is biasing a small electrode in a plasma to some voltage and measuring the resultant current. The current-voltage (or I-V) trace can then be analyzed with a set of limited assumptions to extract plasma parameters

such as electron temperature, plasma potential, plasma densities, and more. While versatile, Langmuir probes also tend to yield measurements with high uncertainties (20–50%) due to the broad assumptions necessitated in their analysis. We recommend Ref. [117] for further details on Langmuir probes and their uses in electric propulsion.

In this section, we summarize the hardware and analysis methods used for the Langmuir probes employed in our experiments. The Langmuir probes we used through all our experiments consisted of a 1 mm diameter by 4 mm length of tungsten wire extending past an alumina tube. We primarily only used Langmuir probes to measure the floating electron temperature and/or plasma potential.




Figure 4.4: Langmuir probe data for xenon at 300 V, 125 A. (a) Raw probe I-V trace. (b) Natural log of electron current used to determine electron temperature. (c) Second derivative of current used to determine plasma potential with zero crossing method. In all cases, the electron retarding region is between the black and the gray circles.

As depicted in Fig. 4.4b, we inferred the electron temperature by subtracting a linear fit to the ion current (the region from approximately -10 V to 0 V) from the raw trace, then fitting a line in log space to the resulting electron current signal I_e as a function of bias voltage in the electron retarding region. The inverse slope of this line yielded the electron temperature T_e and its uncertainty was found using twice the standard deviation in the slope from the linear fit. This methodology for determining electron temperature follows best practices as outlined in Ref. [117]. To infer the electron temperature, we used the inverse slope of the current-voltage Langmuir probe trace as plotted on a log-linear scale.

The plasma potential may be estimated from the Langmuir probe trace alone in two ways. First, we used classical Maxwellian theory for probe analysis to relate the floating potential V_f and electron temperature T_e to the plasma potential,

$$V_p(T_e \text{ correction}) = V_f + T_e \ln \left(\sqrt{\frac{m_i}{2\pi m_e}} \right), \quad (4.5)$$

where m_e is the electron mass. As an alternative method for evaluating the plasma potential, we used the zero crossing in the second derivative (Fig. 4.4c) of electron current, i.e. the “knee” potential $V_p(\text{knee})$. The potentials determined from these methods are shown in Fig. 4.4. Due to noise in the raw Langmuir probe data, a clean estimate of V_p with the “knee” method may not always be obtainable. With our measurements of plasma potential, we corrected the RPA-measured most probable voltage to obtain the acceleration voltage. The plasma potential can be determined directly from the Langmuir probe or by using the electron temperature from the Langmuir probe with the floating potential measured by the emissive probe following the methods outlined later in Sec. 4.4.3.1.

4.4.2.3 Emissive Probe

The purpose of an emissive probe is to measure the local plasma potential with respect to ground. This probe consists of a conducting loop heated to thermionic emission. In our case, we used a 1 mm length thoriated tungsten filament loop. Without any current flow, the probe electrically floats to some potential V_f in the presence of a plasma; to preserve neutrality, a positive sheath forms with a potential drop from the plasma to the probe, measured as a voltage across the probe leads. This attracts ions and repels hotter electrons to maintain zero net current. Once the probe is heated by passing current through it, a smaller potential drop is necessary to maintain this sheath. Eventually, the probe’s floating potential approaches the plasma potential V_p .

Practically, we operate emissive probes by steadily increasing the amount of current until we see a sharp increase in the voltage measured across them. Within a probe trace, the floating voltage is the value that the probe reaches after this jump. Visually, this point can be detected when the emissive probe starts glowing in a slightly alarming manner. Emissive probes, while useful and more accurate in determining plasma potential than Langmuir probes, can burn out very rapidly due to the high currents relative to the thin wire. We therefore often require at least three or four emissive probes where one would technically suffice for redundancy when operating thrusters in vacuum.

There is a small potential offset between the floating potential and the plasma potential due to space charge limitations—we use the electron temperature as measured by the Langmuir probe to correct for this effect, with

$$V_p(\text{emissive}) = V_f + \alpha T_e, \tag{4.6}$$

where α is a scalar factor and T_e is the electron temperature. This correction assumes a Maxwellian electron energy distribution and accounts for space charge effects, with values of

α ranging from 0 to 1.5 [118]. We evaluated both limits of α to provide bounds of $V_p(\text{emissive})$ to average over and take our uncertainty in plasma potential as half the difference in these two bounds.

4.4.2.4 $E \times B$ Probe

The $E \times B$ probe serves as a mass spectrometer, determining the ratios of different charge species within the thruster plume. These fractional charge species inform the calculation of the charge utilization efficiency (Eq. 3.8). Within the $E \times B$ probe are orthogonal electric and magnetic fields. When an ion passes through the probe with some axial velocity, the forces from the electric and magnetic fields acting on the velocity of the ion must be balanced in order for the ion trajectory to be uninterrupted and intercept the collector at the back of the probe. Ions ejected from the thruster have on average a velocity proportional to their square root of charge, so the probe can isolate different charge state populations by varying the strength of the electric field with the magnetic field held constant.

The $E \times B$ probe we employed had an entrance aperture 1.6 mm in diameter, an entrance collimator 7.5 cm long, an exit collimator 15 cm long, and electrical plates spaced 0.97 cm apart. The peak magnetic field was 0.16 T in the center of the probe. We typically swept the applied bias voltage from 0 to an upper limit of 70–100 V depending on how much multiply charged species content we expected to see.

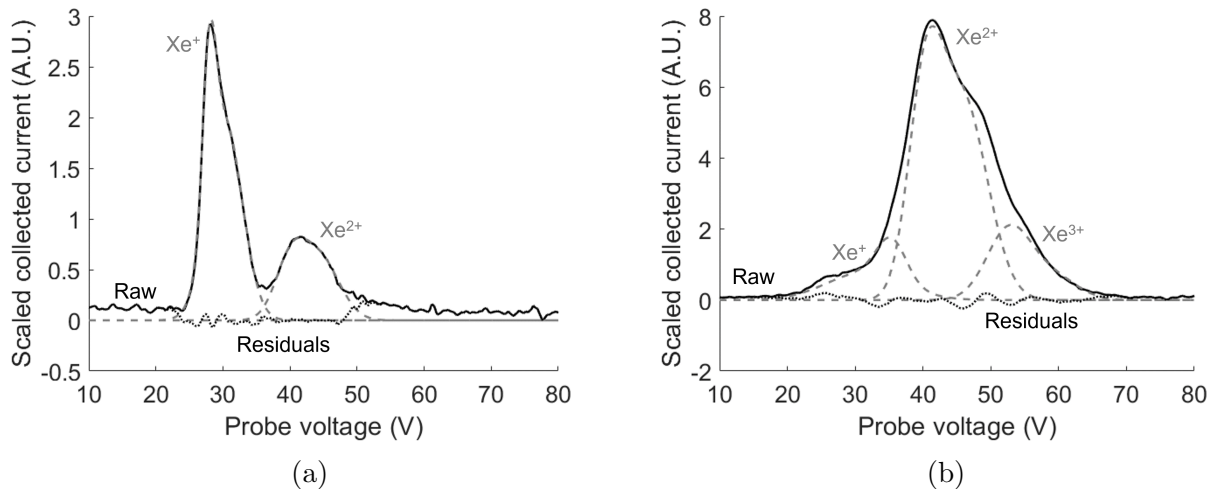


Figure 4.5: Raw $E \times B$ trace and two-peak Gaussian fits for all charge species of xenon at a) 300 V and 15 A and b) 300 V and 125 A. Note that the Xe^{+3} population was too small to detect at 15 A.

We show in Fig. 4.5 examples of raw traces from the probe at two operating conditions.

The peaks in these results correspond qualitatively to the relative magnitude of current carried by each charge species. To determine the relative current fraction Ω_i for each of these species, we follow the approach in Ref. [119] by progressively fitting dual Gaussian distributions to each peak in the raw trace. This previous work also showed that the choice of a dual Gaussian functional form better captures the high-velocity tails frequently exhibited by Hall thruster ion velocity distributions. Indeed, we found that other fit functions that have been proposed previously, such as a single Gaussian, yielded unphysical results for the charge state.

In our procedure, we fit a two-peak Gaussian to the singly-charged peak (the largest peak at standard operating conditions). We then subtracted this fit from the raw trace and repeated this process up to three or four charge states. By integrating the curve fit for each peak and dividing by the total integrated area for all species, we obtained uncorrected current fractions. The uncertainty associated with these values came from uncertainty in the two-peak Gaussian fit. We estimated uncorrected charge fractions from these fits by integrating the area of each peak and comparing it to the total integrated area under the curve. At higher current densities, this procedure was complicated by the larger population of secondly-charged ions; we discuss our specific methodology for this in Sec. 6.2.

4.4.2.5 Faraday Probe

A Faraday probe is a diagnostic that can be used to determine beam plasma characteristics by being biased to a sufficiently negative voltage such that it only collects ion saturation current. These measurements of the ion beam current in turn inform estimates for the beam utilization efficiency (Eq. 3.7), mass utilization efficiency (Eq. 3.9), and divergence efficiency (Eq. 3.11). A guard ring around the central probe collector disc flattens the ion sheath to maintain a constant effective probe collecting area. The Faraday probe we used had a 1.74 cm inner diameter molybdenum collector and 2.38 cm outer diameter molybdenum guard ring, separated by a 0.05 cm gap. By measuring the current collected by the central probe during an azimuthal sweep of the probe through the plasma plume at constant radial distance from the thruster, we can plot a spatial distribution of the current density. During this sweep, the FP is biased to -30 V. Throughout our experiments, we typically swept the FP from 0 to 180 degrees with a resolution of ~ 1 degree. From this, we can determine both the beam current and divergence angle.

The current collected by the probe as a function of angle was converted to current density

j with the relation defined in Ref. [120]:

$$j = \frac{I_{FP}}{A_C + \kappa_G} \kappa_{SEE}, \quad (4.7)$$

where I_{FP} is the current collected by the Faraday probe, A_C is the collector area, and κ_G and κ_{SEE} are correction factors arising from geometric and secondary electron emission (SEE) effects respectively [120]. These correction factors are described in detail in Sec. 4.4.3.3. An example of a Faraday probe trace is shown in Fig. 4.7.

4.4.2.6 Bootstrapping

When possible, we used standard deviations and 95% confidence intervals. Otherwise, we used bootstrapping, which is where we sample a subset of a given trace and perform the same analysis method over multiple iterations to generate a set of values [121]. The standard deviation in the resulting dataset is then taken to be the uncertainty. We show an example of bootstrapping when analyzing RPA data in Fig. 4.6.

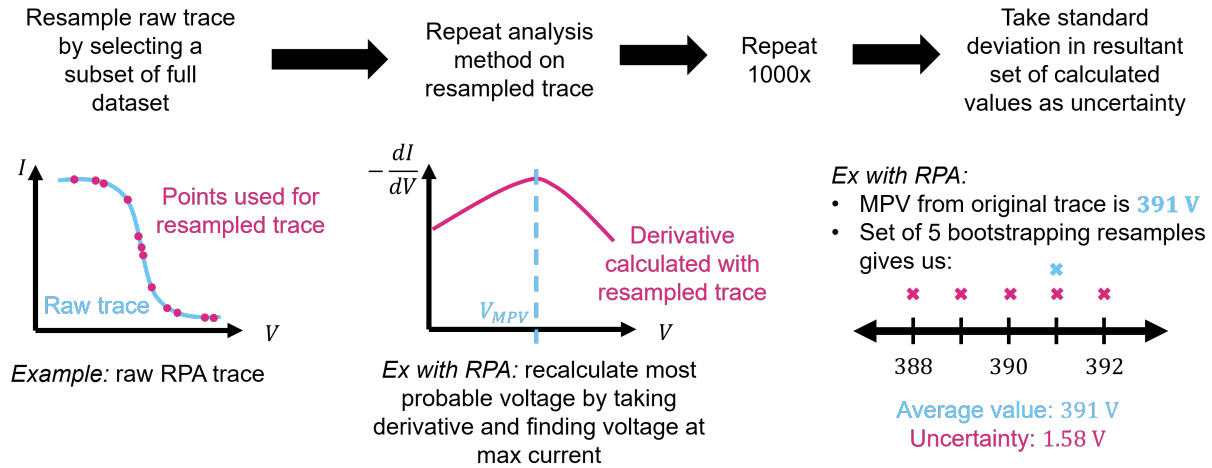


Figure 4.6: Bootstrapping process with RPA trace.

4.4.3 Probe Integration with Efficiency Model

Equipped with the plasma parameters measured by our far-field probe suite, we are now able to calculate the values necessary to evaluate our different efficiency components (Sec. 3.3.2).

4.4.3.1 Acceleration Voltage

We defined the acceleration voltage in Eq. 3.10 as the most probable ion kinetic energy in the far-field plume. To infer this value, we first differentiated the current-voltage trace from the RPA diagnostic. This yielded an estimate for the distribution of ion energy per unit charge (Fig. 4.3), the peak of which corresponds to the most probable ion energy, V_{RPA} . As the RPA was biased with respect to facility ground, we converted this most probable value to an estimation of acceleration voltage by subtracting the local plasma potential,

$$V_a = V_{RPA} - V_p. \quad (4.8)$$

The plasma potential may be determined by using data from both the emissive and Langmuir probes or by the Langmuir probe alone. To determine the plasma potential with the method using both probes, we use Eq. 4.6 with the floating potential as measured by the emissive probe and electron temperature as measured by the Langmuir probe. The plasma potential may also be estimated from the Langmuir probe alone as described in Sec. 4.4.2.2.

4.4.3.2 Current Fractions

Due to the presence of finite background pressure in the facility, we needed to modify the raw current fractions to account for charge exchange (CEX) ions [71, 122, 123]. We show in Fig. 4.5 an example of a typical trace inferred from the E×B probe. The multiple peaks correspond to the different charge states in the plasma. In order to infer the relative fraction, Ω_n , of each of these species, we followed the approach outlined in Ref. [119] by progressively fitting dual Gaussian distributions for each of the species found in the raw trace. This previous work also showed that the choice of a dual Gaussian functional form better captures the high-velocity tails frequently exhibited by Hall thruster ion velocity distributions. Indeed, we found that other fit functions that have been proposed previously, such as a single Gaussian, yielded unphysical results for the charge state. In our procedure, we first fit to the largest peak, assuming it to be representative of the singly-charged species, then subtracted off this fitted curve. We repeated the process for higher charge states up to quadruply-charged ions. Fig. 4.5 shows examples of the fits from this process as well as the residuals, defined as the difference between the raw trace and the fits. We estimated uncorrected charge fractions from these fits by integrating the area of each peak and comparing it to the total integrated area under the curve.

The presence of background neutrals necessitates the consideration of charge-exchange (CEX) collisions in correcting the charge fractions. This was done by modifying the integrated current and density with a term $(j/j_0)_n$ for the n^{th} ion species. This factor corrects

the raw value of each charge species $\Omega_{n,raw}$ as described in Refs. [122] and [123]:

$$\Omega_n = \frac{\Omega_{n,raw}(j/j_0)_n^{-1}}{\sum_k (\Omega_{k,raw}(j/j_0)_k^{-1})}, \quad (4.9)$$

where the summation over k in the denominator is the total area under the fitted curves. Each correction factor is calculated as

$$(j/j_0)_n = \exp(-n_0\sigma_n z), \quad (4.10)$$

where n_0 is the background neutral density, σ_n is the cross-sectional area for the n^{th} ion species, and z is the distance of the E×B probe from the exit plane. We calculated n_0 from measurements of chamber pressure. The expressions for the cross-sectional areas of the n^{th} charge state of xenon are taken from Ref. [124]:

$$\sigma_1|_{Xe} = (87.3 - 13.6 \log(V_1)) \times 10^{-20} \quad (4.11)$$

$$\sigma_2|_{Xe} = (45.7 - 8.9 \log(2V_2)) \times 10^{-20} \quad (4.12)$$

$$\sigma_3|_{Xe} = (16.9 - 3.0 \log(3V_3)) \times 10^{-20}, \quad (4.13)$$

where V_n is the voltage through which the n^{th} ion species is accelerated. For our analysis, we assumed that all values of V_n are equivalent to the calculated acceleration voltage V_a . The ratio of Xe^{4+} was typically small enough (less than $\sim 5\%$) such that we neglected the correctional factor for this state.

For krypton, we used the cross-sectional areas listed in Ref. [125] for the first two charge states:

$$\sigma_1|_{Kr} = (80.7 - 14.7 \log(V_1)) \times 10^{-20} \quad (4.14)$$

$$\sigma_2|_{Kr} = (44.6 - 9.8 \log(2V_2)) \times 10^{-20}, \quad (4.15)$$

where we once again used the acceleration voltage V_a for each species. As there was no tabulated data for the cross-sectional area of Kr^{3+} , we approximated the expression using Eq. 4.13. This was deemed acceptable as the constants for the 1st and 2nd charge states of xenon and krypton are within 10% of each other; additionally, because the ratio of Kr^{3+} is low overall (less than $\sim 4\%$), the impact of this error is minimal. We assumed that $(j/j_0)_4 = 1$ for krypton as we did for xenon. We quantified uncertainty in the final charge species by accounting for two sources of error following the methodology in Ref. [123]. First, we quantified the error from the fit by integrating under the area of the residual curve for each

charge state. Secondly, we quantified the error in the CEX correction by propagating the uncertainty of the neutral density due to the pressure measurement (typically 20% for our gauges [126]). These uncertainties were then propagated through to the charge utilization efficiency (Eq. 3.8) following the methodology in Ref. [123].

4.4.3.3 Divergence Angle, Beam Current, and Beam Ion Flow Rate

For our evaluation of the correction factors in the beam current density shown in Eq. 4.7, we adopted the same formula for geometric correction factor as given in Ref. [120], yielding a value of $\kappa_G = 7.7 \times 10^{-6} \text{ m}^2$. We used the same SEE correction factor from this reference for the first three charge states of xenon. However, for our work, we also included the fourth charge state of xenon and introduced coefficients for correcting krypton SEE up to the fourth charge state. To this end, we note that the SEE correctional factors are defined as

$$\kappa_{SEE} = \frac{1}{1 + \sum_n \frac{\Omega_n \gamma_n}{Z_n}}, \quad (4.16)$$

where the charge fraction Ω_n is inferred from the E×B probe and γ_n is the SEE coefficient of the n^{th} charge state.

Brown et al.’s previous work determined the values for the first and second charge states of xenon by averaging γ_n from 100 to 1000 eV with measurements of xenon impinging on molybdenum from Hagstrum [127]. For the third charge state, they used the ratio of higher to lower γ_n of tungsten from a separate study by Hagstrum [128], as these ratios were not measured for molybdenum. They assumed the same ratio between the third and the second charge state for tungsten as for molybdenum to infer a value for the third charge state of xenon impinging on molybdenum. For our work, we extended this methodology to determine the SEE coefficient for the fourth charge state of xenon. Additionally, we used the same overall method to calculate all four values of γ_n for krypton. The SEE coefficients we found from this technique are tabulated in Table 4.1.

Gas	γ_1	γ_2	γ_3	γ_4
Xe	0.021	0.199	0.697	1.393
Kr	0.069	0.296	0.859	1.889

Table 4.1: Secondary electron emission coefficients for charge states 1 through 4 of xenon and krypton.

We see from Fig. 4.7 that the density is higher closer to centerline, which is where most of the beam is located. The density then decreases with angle. The periphery of the beam

consists of a convolution of ions originating from the thruster as well as those born from charge exchange collisions with background facility neutrals. As with data from the $E \times B$ probe, the current density trace needed to be corrected for effects related to the background pressure in the facility. The dominant factor is the effect of ions from charge exchange, fast neutrals and slow ions resulting from collisions between slow neutrals and fast ions in the beam. The slow ions artificially inflate the current densities in the wings of the plume, leading to an overestimate of beam current [120, 122, 123, 129, 130].

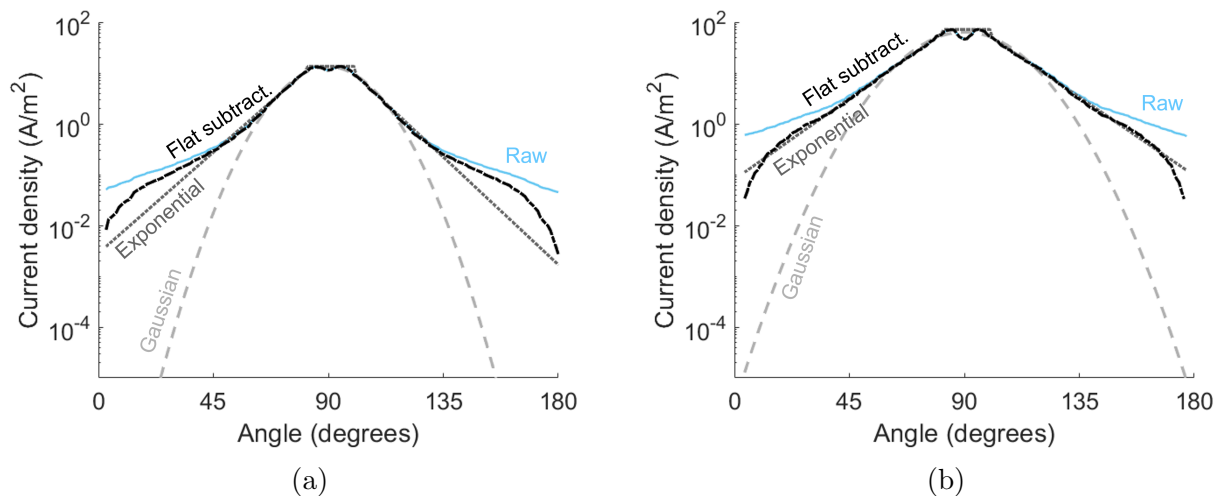


Figure 4.7: Raw current density trace (blue) and CEX correctional methods (black) (Sec. 4.4.3.3) used for calculating beam current and divergence angle from Faraday probe during operation on xenon at a) 300 V and 15 A and b) 300 V and 125 A.

Multiple methods of accounting for the impact of CEX on this beam current calculation have been proposed to date. We review three of them here:

1. **Flat subtraction:** subtracting a value of current density in the wings (i.e. at 0° or 180°) from the entire trace [131],
2. **Gaussian fit:** fitting a Gaussian to the entire trace [126],
3. **Exponential fit:** fitting and piecing together exponential functions to the region 10° to 25° away from centerline (65° – 80° and 100° – 115°) [119].

We show in Fig. 4.7 examples of these fits at two different operating conditions. While all three methods decrease the current density profile as we would expect, the Gaussian method reduces the magnitude significantly more than the exponential and flat subtraction methods. The “best” method to use is often not immediately clear, and the accuracy of a given method

may depend on the thruster and the facility (i.e. background pressure, discharge current, etc.). We determine which methods to use on a case-by-case basis throughout our work. The beam current was calculated from measurements of corrected current density as

$$I_b = 2\pi R^2 \int_0^{\pi/2} j(\theta) \cos \theta d\theta, \quad (4.17)$$

where R is the distance from the thruster exit plane to the probe.

For the divergence angle, we used the methodology provided by Brown et al.[120] and Huang et al.: [132]

$$\cos \theta_d = \frac{\int_0^{\pi/2} j(\theta) \cos \theta \sin \theta d\theta}{\int_0^{\pi/2} j(\theta) \cos \theta d\theta}, \quad (4.18)$$

where θ is the azimuthal location of the probe in radians where 0 is in line with the thruster exit plane (c.f. Fig. 5.4). Physically, this expression represents the ratio of the axially-directed beam current to the total beam current. When possible, the most accurate way to determine divergence angle is to perform this calculation for Eq. 4.18 at multiple distances from the thruster; by taking a linear fit to the exit plane of the thruster, we can obtain an estimate for what the divergence angle at the exit plane of the thruster is [120]. This is done with the “raw” (i.e. uncorrected for CEX) current density traces. However, oftentimes it is unfeasible to perform full Faraday probe sweeps at multiple distances. In these cases, applying Eq. 4.18 to the exponential fit has been shown to produce accurate results [119].

We determined the mass utilization efficiency from Eq. 3.9 by using the measured discharge current, beam current as determined by the FP, and ratios of charge species as determined by the E×B probe. We estimated uncertainty in this efficiency by propagating the individual sources of error from each measurement. The product of the individually-calculated efficiencies was taken to be the “probe-calculated” anode efficiency, $\eta_{a,probe}$, as shown in Eq. 3.6, with the error from individual efficiencies again propagated forward to yield a final estimate of overall uncertainty.

4.4.4 Laser-Induced Fluorescence

Laser-induced fluorescence (LIF) is a non-invasive technique for measuring the ion velocity distribution of a given species. It has been established as a useful tool for evaluating internal plasma properties in a Hall thruster [133]. Ref. [134] provides a thorough overview of LIF and different analysis tools. In this section, we overview the basic theory and detail the hardware and analysis methods used in our setup.

4.4.4.1 Theory

The electronic transition of a metastable state of an ion is non-resonantly excited with a laser at some wavelength; this causes a decay to ground state and the emission of fluorescence of a different wavelength. The intensity of this resulting fluorescence is measured to indicate the density of that population. By detuning the injection laser from the central wavelength, we can take advantage of the Doppler effect to measure the relative density of a species traveling at different velocities based on the intensity of its fluorescence. This results in an ion velocity distribution function (IVDF) at a point in space [134, 135, 136, 137, 138]. The relationship between velocity and wavelength is $v = c(1 - \lambda/\lambda_0)$, where v is velocity, c is the speed of light, λ is the injection wavelength, and λ_0 is the central wavelength of the transition.

Due to the typically low fluorescence signals from the plasma compared to the background light, a method known as homodying is typically used to help isolate the LIF signal. This is done by applying the input signal with a known frequency and searching for data with that same frequency in the measurement. In the LIF setup, this is done with a chopper for injecting the excitation light and a lock-in amplifier for detecting the fluorescence light. As the laser wavelength must be very precise in order to maintain a low uncertainty in velocity, we require both a very precise wavemeter and laser. Finally, operating the laser at different powers affects the strength of the signal as well, adding uncertainty to our density-proxy measurement. We discuss mitigation strategies for this in Sec. 4.4.4.4.

4.4.4.2 Hardware

We used the 728.98 nm in air transition for Kr II that fluoresces at 473.90 nm, a non-resonant transition that has relatively high intensity [139]. For the Xe II setup, we used the 834.72 nm in air transition that fluoresces at 541.91 nm [139]. These transitions have been previously employed for Hall thruster testing [134, 136, 137, 138, 140, 141]. The system starts with a laser head which differs based on injection wavelength. For the krypton transition, we used a TOPTICA TApr0 tunable diode laser and tapered amplifier system with a center wavelength of 729.18 nm, maximum output power of 500 mW, and a mode hop free tuning range of 56 GHz. For the xenon transition, we used a Newport TLB-6700 diode laser and a TA-7600-LN tapered amplifier.

A HighFinesse WS-7 wavemeter precisely measured the wavelength of a probe beam from each laser; this probe beam was emitted directly by the TOPTICA laser and split with a beam splitter from the main beam of the Newport laser. An internal photodiode for the TOPTICA laser and an external photodiode (ThorLabs PDA36A) for the Newport laser monitored output power. The main beam from the laser then passed through a SR540

Species	Excitation λ , air (nm)	Fluorescence λ (nm)
Xe II	834.72	541.91
Kr II	728.98	473.90

Table 4.2: Excitation and fluorescence wavelengths used for selected singly-ionized krypton and singly-ionized xenon transitions.

mechanical chopper operating at 2 kHz before being injected into the chamber through a multi-mode 50 μm diameter fiber with a numerical aperture of 0.22.

Inside the chamber, we had an 2-inch diameter injection lens with a focal length of 10 cm mounted approximately 10 thruster diameters downstream pointing at channel centerline in the axial direction and a 3-inch diameter collection lens with a focal length of 20 cm mounted slightly downstream and out of the plume of the thruster. The collection fiber was 1 mm in diameter and had a numerical aperture of 0.39. At atmosphere, we aligned the optics using visible lasers to ensure that the collection and injection optics were focused on the same alignment spot on the thruster. Each optic was equipped with a small motion stage that allowed us to adjust their positions when the chamber was pumped down, improving our alignment to a spot we marked on the thruster. The thruster was mounted on a motion stage that allowed us to vary the interrogation point, as shown in Fig. 7.2; instead of either the injection or collection optic moving, the thruster itself was able to move in the r and z directions (defined in Fig. 7.3a). Our interrogation point was approximately 1 cubic mm in size. An example of an LIF setup implemented into a vacuum chamber can be seen in Fig. 7.2.




Figure 4.8: System overview of laser-induced fluorescence.

Once collected by the collection optic, the fluoresced light was passed through a bandpass filter centered at 473 nm for krypton and two bandpass filters at 540 and 546 nm for xenon to reduce the noise in our signal. This light was converted to current using a Hamamatsu E717-500 photomultiplier tube and converted to voltage using a Oriel 70710 transimpedance amplifier. Finally, an SRS 810 lock-in amplifier tied to the frequency of the chopper was used to distinguish the fluorescence from the background light. We used an integration time

of 300 ms on the lock-in amplifiers and a frequency of 2 kHz for the chopper. This entire system is summarized in Fig. 4.8.

4.4.4.3 Laser Troubleshooting

While providing full details on how to operate and troubleshoot a diode laser is outside the scope of this dissertation, we offer here some advice and lessons learned from setting up the krypton LIF system. First, safety should be your top priority. The diode lasers we use in this work are Class IV, which mean they are hazardous to eyes and other organic materials (like skin). Especially when first setting up a laser, proper precautions such as laser curtains, eye protection (with the correct optical density), and lab warnings need to be set up so people (including you, presumed graduate student) do not accidentally blind themselves. Make sure you have the proper safety training before attempting to operate any laser.

That said, there are many nuances to operating a laser, and they frequently are only applicable to a given model. For example, the Toptica laser used for krypton LIF has both fine tuning and coarse tuning to access different wavelengths. The range of wavelengths used in this work required two ranges of coarse tuning; this required a physical knob to be turned with an Allen key by just a few degrees. However, we can still provide a few generalized points of advice regarding laser operation. First, if the laser system has a tapered amplifier and a laser head, it is much easier to troubleshoot them one at a time. Second, calling the laser manufacturer is often the best way to resolve an issue, even if you get put on hold for a few hours. Third, aligning a beam is a tedious process and an unavoidable rite of passage in working with optics—make sure you move downstream starting from the beam moving slowly and methodically. Also, using a phone camera allows you to see in near-infrared; this can be surprisingly helpful in aligning the beam, although proper care needs to be taken to ensure that you yourself do not look at the beam with your bare eyes. Finally, a number of issues with the laser can simply be resolved with better temperature regulation. Make sure you have the proper cooling (e.g. a water-cooled metal plate) or heating (e.g. a lizard tank heater) to maintain your laser’s preferred operating temperature.

4.4.4.4 Analysis Methods

Throughout our work, we only took LIF data along channel centerline in the axial direction. At a given location, we measured an ion velocity distribution function (IVDF) by varying the input wavelength, changing the range of velocities based on where the peak of the fluorescence intensity was located. We power-corrected this signal by dividing it by the power input from the laser; this isolated the intensity of our fluorescence measurements from the intensity of




Figure 4.9: Ion velocity distribution functions taken along channel centerline of the H9 operating on krypton at 300 V, 15 A from upstream (pink, low velocities) to downstream (red, high velocities).

the injection signal. Hall thruster ion populations are typically cold [142] and often have two distinct populations [134, 143]. This latter fact is due to the oscillatory nature of Hall thrusters and the time-averaged nature of our LIF measurements; it is possible to detect both a slow and a fast ion population at a given point in space. For these physical reasons, we selected a two-peak Gaussian fit for each of the IVDFs. For each fit, we calculated the mean and mode for the raw data as well as for the fit. These fits are shown in an example trace for the H9 operating on krypton at 300 V, 15 A in Fig. 4.9. In our work, we use the mean value of the two-peak Gaussian fit to the IVDF as our average velocity measurement at a given location.

We neglect to account for broadening of the IVDF Zeeman splitting and hyperfine splitting in our analyses. Zeeman splitting is when spectral lines split due to the presence of magnetic fields. Previous LIF work performed on singly-charged xenon has found this effect to contribute less than 1% uncertainty to the velocity moment at the exit plane and downstream of the thruster channel [134, 144, 145], so we neglect it in our analysis. Hyperfine splitting is due to the interaction between the nucleus and electrons within an atom which cause the energy levels to shift. Because this splitting is symmetric for the species we

investigate, we neglect this effect as well [133, 134].




Figure 4.10: Average ion velocities along channel centerline of the H9 operating on krypton at 300 V, 15 A. a) Mean and mode of raw and fitted IVDFs shown in Fig.4.9. b) Mean of fitted ion velocities and electric field. The circular points are the calculated values of mean ion velocities and electric field, while the solid lines represent a smoothing spline fit for the velocity profile and the electric field calculated from that fit. The pink dashed lines indicate the position and width of the acceleration region

Characterizing the electric field profile is useful for studying the acceleration region, the region of the thruster channel in which the majority of the ions are accelerated through the potential drop. To calculate the electric field based on our velocity profile, we use

$$E(z) = \frac{m_i}{q} u(z) \frac{du}{dz}, \quad (4.19)$$

which assumes no ionization in the acceleration region. While this is not entirely true, as there has been significant overlap between the ionization (where the majority of ion are ionized) and acceleration region observed in the past [146], we use this assumption to easily characterize the electric field and acceleration region. We fit a smoothing spline to the fitted mean velocity profile and apply Eq. 4.19 to obtain the electric field profile. The “location” of the acceleration region is defined as the peak of the electric field, while the “width” is defined as the region over which the electric field is over half of its maximum value.

We performed a bootstrapping analysis to obtain fit uncertainty estimations for each IVDF, with the error bars representing two standard deviations (95%) of the distribution. With this bootstrapping method, described in Sec. 4.4.2.6, we randomly sampled each trace and performed a fit on the resultant profile, repeating this process a thousand times and

treating the distribution in the mean and mode velocities as the uncertainty for each [121].

4.5 Summary

In this chapter, we cover all the test apparatus and diagnostic tools used throughout this work. We begin by discussing the thrusters used as test articles and the facilities we test them in. We then cover all the different diagnostics we used and how we analyze the measurements, including for thrust stands, various probes, and laser-induced fluorescence. Armed with these tools, we may now begin discussing the experiments performed for this dissertation.

CHAPTER 5

Comparing Performance of Xenon and Krypton at Baseline Operating Conditions

“The magnetic field lines are often called ‘lines of force.’ They are not lines of force. The misnomer is perpetuated here to prepare the student for the treacheries of [their] profession.”

–Francis F. Chen, *Introduction to Plasma Physics and Controlled Fusion* [17].

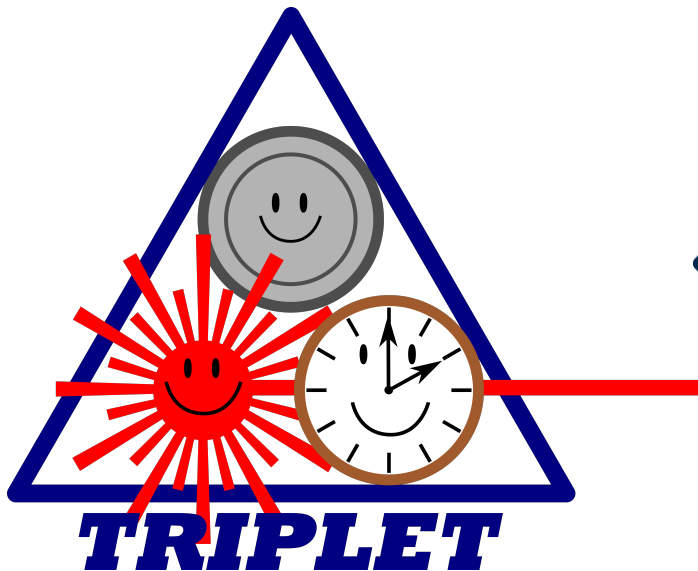


Figure 5.1: ThrusteR Investigation with Probes, LIF, and Erosion Time. Logo design credit shared by T. A. Marks and L. L. Su.




Figure 5.2: Krypton and Wave Analyses with Time-Resolved Oscillations. Logo design credit shared by T. A. Marks and L. L. Su.

5.1 Introduction

This chapter covers work initially published in Ref. [71, 104].

As motivated in Sec. 3.5, there is major incentive to investigate and improve the performance of Hall thrusters operating on krypton. To this end, multiple studies have already been performed on comparing xenon and krypton performance, generally indicating that the efficiency gap is the result of krypton’s low mass utilization [8, 9, 10, 11, 12, 13].

One major caveat regarding these previous comparisons of xenon and krypton performance is that they were conducted on Hall thrusters with a more traditional “unshielded” magnetic field configuration (see Sec. 3.4.1). Magnetic shielding is achieved with a specially-designed field shape that ensures low electron temperatures and high electric potentials at the walls of the thruster discharge channel [14, 15]. While this geometry has the benefit of reducing erosion-causing energetic ion flux to the walls, it also can lead to changes in the distribution of the plasma state and therefore thruster performance. For example, magnetically-shielded (MS) thrusters exhibit higher electron temperatures along centerline and a downstream shift of the ion acceleration region compared to unshielded (US) thrusters [15, 49]. Similarly, a previous comparison of US and MS operation on the H6, a 6-kW class Hall thruster, indicated that efficiency modes such as voltage and current utilization increased when changing from US to MS thrusters, while other efficiency terms like mass utilization and divergence decreased, leading to a 1.5% decrease in anode efficiency [49].

The change in plasma configuration that results from the differences in the US and MS designs may in turn lead to differences in the performance gap between krypton and xenon. For example, given the higher peak electron temperatures along centerline in MS thrusters, we may expect the mass utilization for krypton to be higher, thus reducing the efficiency loss. Prior to this work, however, it has yet to be shown experimentally how shielding may impact the transition from xenon to krypton. In light of the increasing use of magnetically-shielded Hall thrusters and the potential advantages afforded by krypton operation, there is a pressing need for this type of detailed investigation.

The goal of this study is to compare the operation of a magnetically-shielded Hall thruster with xenon and krypton propellants and to identify the key drivers of performance. To this end, this paper is organized in the following way. In Sec. 5.2, we describe the experimental setup, including the thruster and facility used for these tests. This section also includes a summary of the diagnostics and methodology used in context of our performance metrics and efficiency model (Sec. 3.3.2). In Sec. 5.3, we present the results of our study, including a summary of measured efficiencies, thrusts, and specific impulses for both xenon and krypton at all operating conditions. Finally, in Sec. 5.4, we discuss the potential reasons for the

gap between xenon and krypton performance on a shielded thruster and draw comparisons to unshielded thrusters. We also explore possible mitigation strategies for future krypton-operated magnetically-shielded Hall thrusters.

5.2 Methodology

In this section, we describe the test article and facility used for our experiment, our chosen operating conditions, the diagnostics used to collect data of the plasma parameters, and the data analysis methods we employed to calculate our performance metrics.

5.2.1 Thruster and Facility

For this study we used the H9 thruster (Sec. 4.2) with its cathode electrically connected to the thruster body and isolated from facility ground [46]. All tests were conducted at the University of Michigan in LVTF (Sec. 4.3) with between 14 and 16 pumps (11–13 cryos and 2–3 thumpers) on at all times. The thruster was operated on xenon and krypton at background pressures of 4.8-6.2 Xe- μ torr and 4.5-5.8 Kr- μ torr as measured with a Stabil ion gauge Fig. 5.4. We note that due to our misguided faith in Stabil ion gauges at the time (Sec. 4.3.1), we assumed a 10% uncertainty rather than 20% in our calculations of background density for our current fraction analysis (Sec. 4.4.3.2).

Species	V_d (V)	I_d (A)	P_d (kW)	B-field ratio (%)	\dot{m}_a (mg/s, sccm)	Pressure (μ torr)
Xe	300	15	4.5	100	14.8, 165	4.8
Xe	300	20	6	100	18.5, 206	5.8
Xe	400	15	6	100	15.4, 172	5.0
Xe	500	15	7.5	100	16.1, 179	5.3
Xe	600	15	9	100	16.3, 181	6.2
Kr	300	15	4.5	100	11.8, 207	4.6
Kr	300	20	6	87.5	15.3, 268	5.8
Kr	400	15	6	112.5	11.5, 201	4.5
Kr	500	15	7.5	112.5	12.0, 209	4.7
Kr	600	15	9	112.5	12.5, 218	4.9

Table 5.1: Operating conditions and base pressures. Note that probe data was only collected for the three 4.5 and 6 kW conditions.

Table 5.1 shows the operating conditions we tested for this study. The magnetic field shape remained the same in each case, but the magnitudes, which are referenced to the nominal magnitude at the 4.5-kW condition, were adjusted to minimize oscillations and discharge

current. We note that while we made global performance measurements (thrust, specific impulse, and efficiency) at all operating conditions, we only collected probe data at the 300 V and 15 A, 300 V and 20 A, and 400 V and 15 A conditions. This stemmed from concerns about thruster stability at higher powers where we observed discharge current oscillations over 160–210% when the thruster was operated with krypton (App. B.1). The time to generate a complete set of probe data under these conditions was judged to be prohibitively long and unsafe for thruster health. We also note that the flow rate and pressure for the xenon 600 V, 15 A condition are different from what we previously published in Ref. [71]; the value in this previous work was incorrect and an errata has been submitted.

The H9 can be seen in operation on both propellants in Fig. 5.3. We employed a 400 sccm mass flow controller for anode flow, a 50 sccm mass flow controller for cathode flow, and a 60-kW power supply rated to 1000 V and 60 A for discharge power. More details may be found in Sec. 4.2.2.




Figure 5.3: The H9 operating at 300 V and 15 A on a) xenon and b) krypton.

5.2.2 Diagnostics and Data Analysis

We used a null-type inverted pendulum thrust stand to collect thrust data and evaluate anode efficiency at all conditions. We calibrated the thrust stand with known masses corresponding to a range up to 500 mN. The thrust stand had an uncertainty of $\pm 1\text{--}3$ mN and

a resolution of ~ 1 mN. This uncertainty arose from standard deviations in solenoid current over the averaging window used to obtain thrust before and after the thruster was shut down (Sec. 4.4.1). We note here that at the time of data collection, we had not accounted for the uncertainties introduced by the drift in thrust stand displacement and inclination. The uncertainties calculated in this chapter are therefore likely an underestimate of the real values. Based on our other analyses with this thrust stand (Sec. 6.3), we expect the actual uncertainties to be about twice as high as the values we use in this analysis.




Figure 5.4: Notional top-down schematic of experimental setup in LVTF, including thruster, thrust stand, ion gauge, probe arm with Langmuir probe, emissive probe, retarding potential analyzer, and Faraday arm, far-field $E \times B$ probe, and beam dump. Pumps are located around the walls of the facility but not shown in the diagram.

To evaluate the plume properties that contribute to the terms in the efficiency model (Sec. 3.3.2), we used a probe suite consisting of a retarding potential analyzer (RPA), emissive probe (EP), Langmuir probe (LP), Faraday probe (FP), and $E \times B$ probe. The RPA, EP, LP, and FP were mounted on an azimuthal probe arm with an axis of rotation above the exit plane of the thruster (Fig. 5.4). These probes yielded data at a location 10.25 thruster diameters (D_T) downstream, with the FP taking additional sweeps at 9.25 and 8.25 D_T . The RPA, EP, and LP each provided a single trace at 90 degrees, directly facing the centerline of the thruster, while the FP was swept from 0 to 180 degrees with a resolution of ~ 1 degree.

The E×B probe was mounted $\sim 10.5 D_T$ downstream of the thruster and aligned to the channel centerline.

Both the primary and secondary electron suppression grids of the RPA were set to -30 V, and the ion selection grid was swept from 0 to $\sim 2\times$ the discharge voltage. The emissive probe we used consisted of a 1 mm length thoriated tungsten filament loop, heated to thermionic emission such that the probe potential approached the plasma potential. We employed a Langmuir probe that was comprised of a 4 mm length tungsten wire routed through a ceramic tube. This probe was biased from -5 to 15 V. We swept the applied bias voltage to the E×B plates from 0 to 70 V to fully map the charge states in the plume. Like the electron suppression grids of the RPA, the Faraday probe was biased to -30 V.

We note here that in this work, we treat our uncertainties as one standard deviation or a 68% confidence interval. This differs from our future work (Sec. 6), where we treat the uncertainty as a 95% confidence intervals, corresponding to two standard deviations. With this in mind, we outline the methods for estimating the acceleration voltage, current fractions, beam current, and divergence angle.

The acceleration voltage, discussed in Sec. 4.4.3.1, was calculated using the electron temperature from the Langmuir probe, floating potential from the emissive probe, and most probable voltage between plasma and ground along channel centerline from the RPA (Eq. 4.6, 4.8). This uncertainty in the most probable voltage measured by the RPA from bootstrapping was compounded with the uncertainty stemming from the variation in the correctional factor α (Eq. 4.6) to yield a total estimate for error in the reported values of V_a and η_v . Details of this process may be found in Sec. 4.4.2. We obtained current fractions following the process outlined in Sec. 4.4.2.4 and 4.4.3.2 using four two-Gaussian fits for each propellant (assuming our highest charge ion is quadruply charged).

We follow the methods outlined in Sec. 4.4.2.5 and 4.4.3.3 to obtain estimates of beam current and divergence angle. We took Faraday probe traces at three locations, 8.25, 9.25, and 10.25 D_T axially downstream from the thruster. To calculate the beam current, we adopted a method based on placing upper and lower bound estimates for the beam current calculations. The integrated value of I_b from the raw trace was treated as the upper bound for the beam current, while for the lower bound, we performed the flat subtraction method by subtracting off the value of j at the point furthest from centerline (i.e. 0 or 180 degrees) following Ref. [131]. This was motivated by the assumption that any charge collected at this peripheral location is purely due to ambient ions without any beam ions. Figure 4.7 shows an example of these “raw” and “subtracted” traces. We averaged the six estimates of beam current (an upper and lower bound at each of the three distances) to determine our average value for I_b . The error is reported as the standard deviation in these six values.

We inferred the divergence angle at each of the three downstream locations with Eq. 4.18 as outlined in Eq. 4.4.3.3 with the raw (i.e. uncorrected for CEX) traces of current density. We then applied a linear fit to the exit plane to determine the divergence angle at the thruster. We determined error in this angle as the standard deviation of the fit parameters. The beam current and divergence angle give us estimates for beam efficiency, divergence efficiency, and mass efficiency as outlined in Sec. 3.3.2.

5.3 Results

In this section, we first compare the global performance metrics calculated for xenon and krypton at each operating condition, including thrust, specific impulse, and anode efficiency. We then present the phenomenological efficiency contributions as described in Sec. 3.3.2 for each gas individually.

5.3.1 Global Performance

Figure 5.5 shows the thrust, specific impulse, and anode efficiency as a function of operating condition for both xenon and krypton. These results are broadly in keeping with previous investigations into the performance of unshielded Hall thrusters on these propellants [6, 7, 8, 9, 10, 11, 49].

5.3.1.1 Thrust

We see that the thrust (Fig. 5.5a) increases linearly with power for both gases, with xenon producing ~ 60 -100 mN more than krypton at each power level. The general trend of increasing thrust with power results from the increasing acceleration voltage (and therefore ion exhaust velocity) while maintaining an approximately constant mass flow rate. The higher thrust levels of xenon relative to krypton, which has been observed in previous comparison studies of unshielded thrusters [6, 7, 10, 12], can be understood in light of the difference in their masses. For both gases, the 20 A condition exhibits higher thrust than the 15 A condition at 6 kW. This result stems from the fact that the acceleration voltage remains approximately constant while we increased the flow rate to achieve the higher discharge current.

5.3.1.2 Specific Impulse

Figure 5.5b shows the specific impulse as a function of propellant and power. We see that in both cases, this value monotonically increases with power, consistent with previous para-

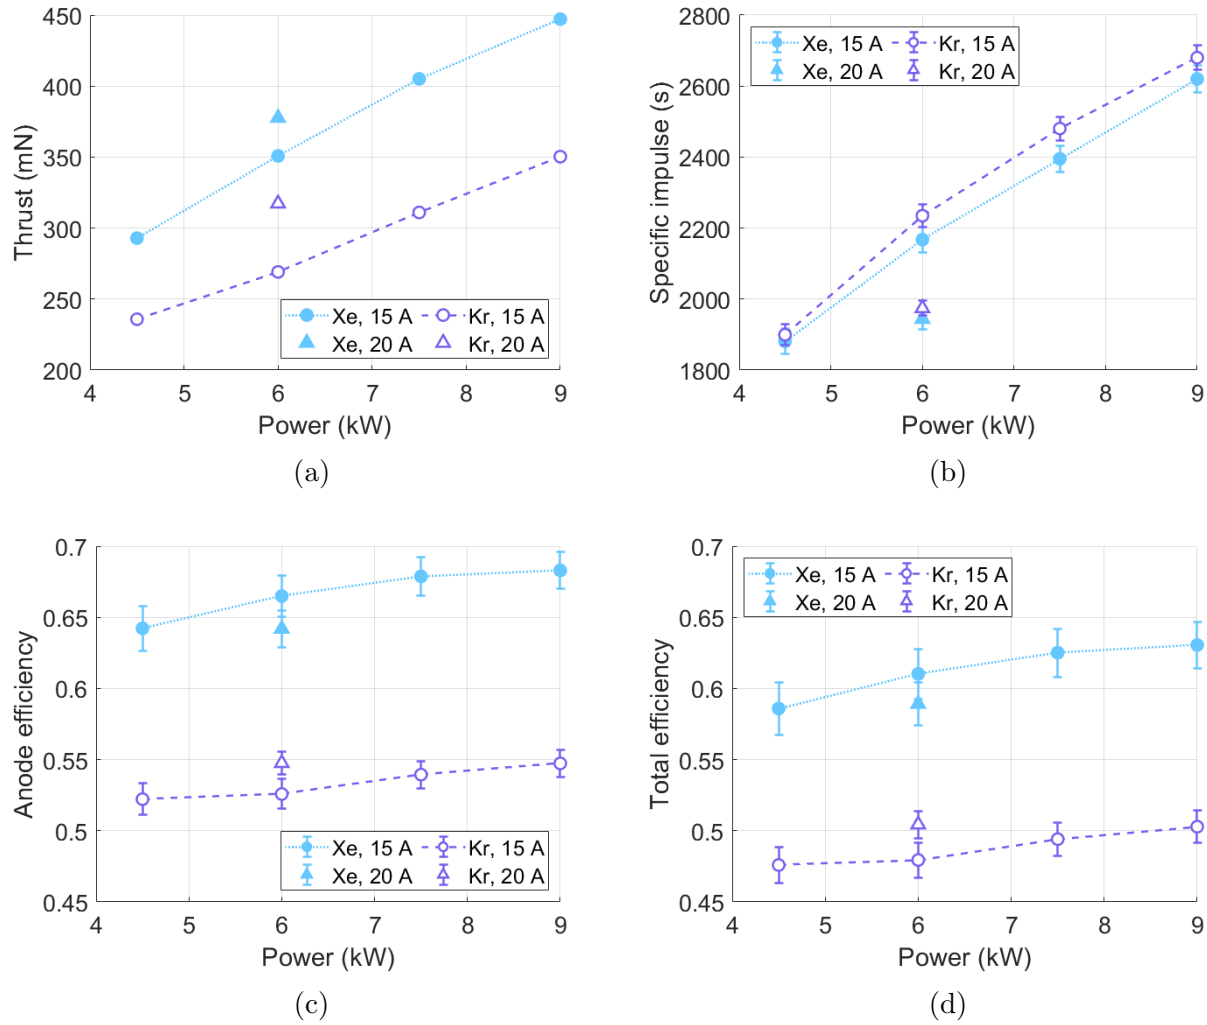


Figure 5.5: a) Thrust, b) specific impulse, c) anode efficiency, and d) total efficiency for xenon and krypton as a function of operating condition. Note that the uncertainties on Fig.5.5a are smaller than the point sizes.

metric studies performed on unshielded thrusters [6, 7, 10, 11, 147]. This trend is due to the fact that as the acceleration voltage increases with discharge voltage, the ions are accelerated to a higher effective exhaust velocity (Ref. [18]). Notably, the specific impulse at the 20 A condition is higher for both gases in comparison to the 15 A condition at the same discharge voltage, implying that thruster efficiency increases with current at a fixed voltage. However, the specific impulses at 300 V, 20 A are lower than those at 400 V, 15 A due to the primary scaling of specific impulse with acceleration voltage.

At all conditions, the specific impulse of krypton is marginally higher (~ 20 – 200 s) than that of xenon. The difference in specific impulse between gases can be understood in part

from the lower mass of krypton, leading to a higher velocity for a given acceleration voltage. With that said, ideally, the specific impulse ratio should scale inversely with the square root of mass. Assuming singly-charged ions, an acceleration voltage equivalent to the discharge voltage, and a fully collimated beam (i.e. $\eta_q = 1, \eta_d = 1, \eta_a = 1$), the theoretical limit of specific impulse is

$$I_{sp,th} = \frac{\sqrt{2qV_d}}{\sqrt{m_i g_0}}. \quad (5.1)$$

This yields values of ~ 2100 – 3000 s for xenon and ~ 2700 – 3800 s for krypton.

The ratio of the measured specific impulse to the theoretical limit of specific impulse is shown in Fig. 5.6. We see that the specific impulse of krypton is $\sim 70\%$ of its theoretical limit while xenon is $\sim 90\%$ of its theoretical limit, suggesting that the efficiency of krypton is generally lower than that of xenon. This matches the trends shown in Fig. 5.5c. We note here that these trends showing the disparity in higher specific impulse and lower efficiency of krypton operation are not unique to the H9 and are in family with the results of previous studies on unshielded thrusters [6, 7, 8, 9, 10, 11].




Figure 5.6: Ratio of measured to theoretical limit of specific impulse for xenon and krypton at 15 A as a function of discharge power.

5.3.1.3 Efficiency

Figure 5.5c shows the anode efficiencies of both xenon and krypton as functions of operating condition. We see that this parameter increases with power in both cases, with anode values improving from 64.2–68.3% for xenon and from 52.2–54.8% for krypton. Since current remains fixed with the exception of the 20 A case, this trend indicates that increasing discharge voltage facilitates more efficient thrust generation. The general trend of increasing efficiency with voltage has been observed in a previous study of an unshielded thruster where the magnetic field was tailored to regulate electron current at higher discharge voltage. Ultimately, the improvements in performance with higher voltage were attributed to the increasing beam, mass, and voltage utilization efficiencies [147].

Quantitatively, the anode efficiency on xenon is higher than krypton at all conditions with a difference ranging from $9\text{--}14 \pm 2\text{--}3\%$. Qualitatively, the efficiencies of xenon and krypton increase at approximately the same rate from 4.5 to 9 kW. Notably, while the efficiency gap on unshielded thrusters was similar over a comparable voltage range (5–15%) to the gap we report here [7, 8, 9, 10, 11], the trends in voltage differed. Specifically, the efficiency of krypton increased more rapidly with voltage than it did for xenon in previous studies [8, 10, 11], while we saw the efficiency of krypton increase at the same rate as xenon.

As a final observation, we see that the efficiency gap between the two gases is the smallest at the high-current 300 V and 20 A condition. A similar decrease in the efficiency gap with current density has been previously observed on unshielded thrusters [9, 10, 11, 12, 13]. This trend suggests that increasing the discharge current rather than the voltage is one potential method of closing the gap between xenon and krypton operation at high powers. We discuss this possibility in further detail in Sec. 5.4.

The trends in total efficiency closely mirror those of anode efficiency at $\sim 5\%$ lower values. This difference is due to the accounting of flow to cathode and power to the magnets in the total efficiency. Despite the slightly different magnetic fields for the higher-power krypton conditions (Tab. 5.1), this change is not drastically evident when comparing Fig. 5.5c and 5.5d due to the small percentage of total power that comes from the magnets.

5.3.2 Efficiency Analysis

To complement the global analysis from the previous section, we present here probe-based measurements of the individual contributions to anode efficiency. This data provides insight into the factors that are driving trends in performance with power and gas choice. To this end, Fig. 5.7 shows the thrust-calculated and probe-calculated anode efficiencies as well as the efficiency breakdown for xenon performance at three operating conditions. The numerical

values and uncertainties are reported in tabular form in the Appendix. Across all points, we see that the anode efficiencies as measured by the thrust and from the probe suite match to within 2.4%, with an uncertainty of $\pm 5\text{--}6\%$ on the probe-calculated efficiencies. The magnitudes of these results are consistent with previous measurements for a shielded Hall thruster operating in this range [49]. In particular, the largest efficiency losses stem from the current utilization and divergence.

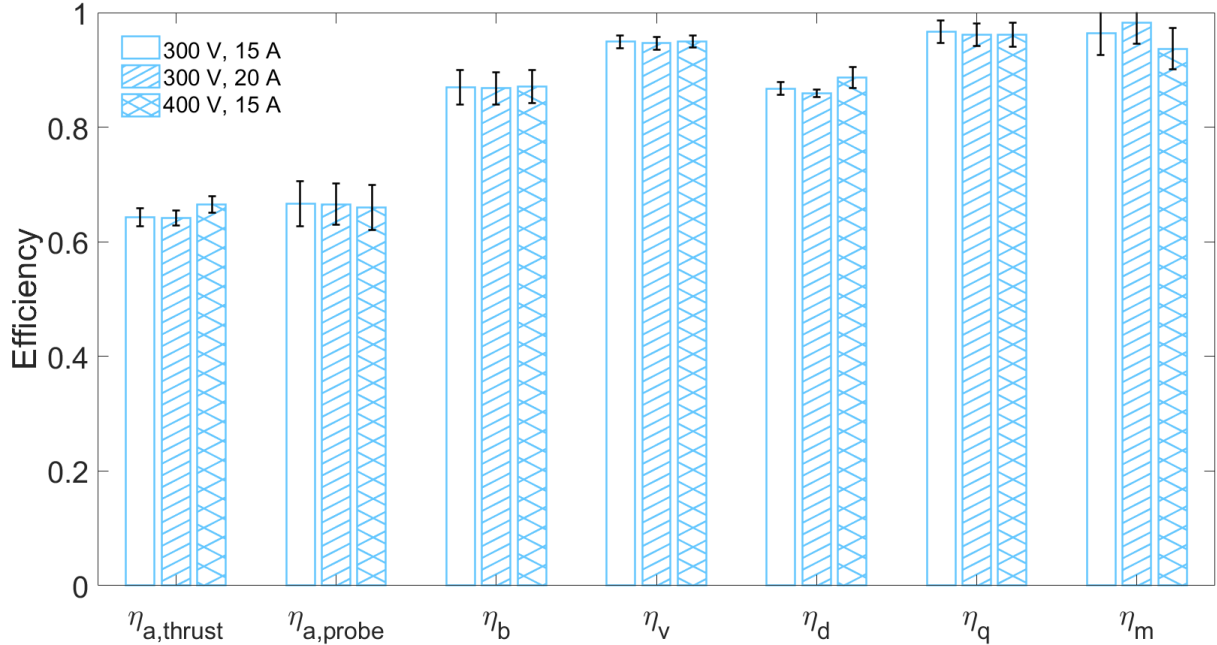


Figure 5.7: Measurements of the contributions to anode efficiency for xenon inferred from far-field probes.

Figure 5.8 shows a breakdown of efficiency for krypton at the same operating conditions that we analyzed for xenon. The values and uncertainties are also reported in App. B.1. The difference between thrust-calculated and probe-calculated anode efficiencies match to within 2.1% with an uncertainty of $\pm 4\text{--}7\%$ on the probe-calculated values. As can be seen from Fig. 5.8 and as is consistent with previous studies of unshielded thrusters operating on krypton [8], the largest losses are in the current utilization, divergence, and mass utilization efficiencies. Of these, the largest difference from xenon is in the mass utilization. This loss may stem from the lower mass and ionization rate coefficient at a given electron temperature. We discuss this trend and its implication in more detail in the following section.

In summary, in this section we have reported what is to our knowledge the first detailed comparison of the performance of a magnetically-shielded Hall thruster operating on xenon versus krypton. We ultimately have found, consistent with earlier work on unshielded

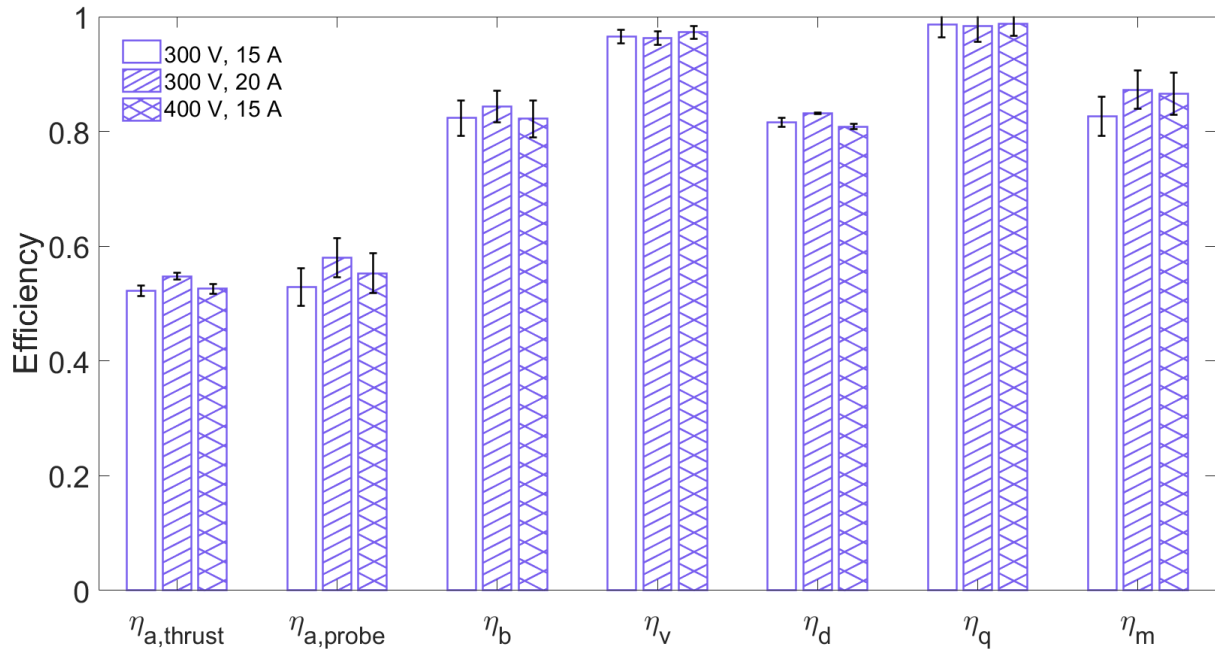


Figure 5.8: Measurements of the contributions to anode efficiency for krypton inferred from far-field probes.

thrusters, that the efficiency when operating on krypton is lower. We similarly have shown from an analysis of the phenomenological efficiencies that the main contributor to the efficiency difference stems from the mass utilization. This is also consistent with previous work reported on unshielded thrusters. In comparing the performance gap between these propellants on shielded versus unshielded thrusters, we have found that the overall difference in efficiency remains approximately the same (5–15% for US versus 9–14% for MS). However, the trend in this gap differs, as we did not see the discrepancy between xenon and krypton efficiency decrease at higher voltages as it does on unshielded thrusters [8, 10, 11]. In the following section, we elaborate on the differences between xenon and krypton performance and discuss potential reasons for the behavior at high voltages on unshielded and shielded thrusters.

5.3.3 Thruster Health

In this section, we briefly discuss the oscillation levels and temperatures of the thruster operating on each propellant. Figure 5.9 shows the variation in peak-to-peak current oscillation at different operating conditions. For xenon, we see that the oscillations remain under 100% (smaller than the mean value of discharge current, which is 15 or 20 A) at all conditions and

exhibit no clear trend with increasing power. For krypton, the oscillation strength increases with voltage, with a particularly sharp jump from $\sim 40\%$ to $\sim 160\%$ between the 400 V, 15 A and 500 V, 15 A conditions. Krypton operation at 600 V, 15 A reaches a maximum value of $\sim 210\%$ current oscillations. We also note here that at 400 V and higher, krypton has a stronger magnetic field strength than at the baseline 300 V, 15 A condition. This choice was primarily made to mitigate these strong oscillations—when the thruster was operated at 400 V, 15 A at a lower field strength (100% of nominal B-field instead of the 112.5% used for the experiment), the oscillations were $\sim 160\%$ instead of $\sim 40\%$. We occasionally observed a “mode-hopping” behavior in the thruster between this low-oscillation and high-oscillation mode without changing any operational parameters. These very oscillatory conditions may be dampened by operating at higher currents (Sec. 6.3.4) or heating the propellant in the anode [148].




Figure 5.9: Ratio of peak-to-peak oscillation to mean value of discharge current at 300 V for xenon and krypton as a function of operating condition.

We monitored the temperature at the inner front pole of the thruster, the location at which we usually see the hottest temperatures during thruster operation. For each propellant, we operated to thermal steady-state (i.e. less than 1 °C change per minute) at 300 V, 15 A. Xenon’s inner front pole reached a steady-state temperature of 400 °C, while krypton reached 475 °C. The higher temperature of krypton may be attributed to its lower mass and therefore higher mobility. From an operational standpoint, this may pose challenges for long-duration testing with krypton at higher powers.

5.4 Discussion

In this section, we first discuss potential drivers for the observed trends in anode efficiency with varying voltage and current. We then examine possible physical mechanisms underlying the difference in various efficiency modes between xenon and krypton, as well as how and why the trends may differ from unshielded thrusters. Finally, we leverage our findings to comment on strategies for improving the performance of a magnetically-shielded Hall thruster operating on krypton.

5.4.1 Trends in Anode Efficiency with Voltage and Current

As noted in Sec. 5.3, the anode efficiency increases with voltage for both gases (Fig. 5.5c). Since our efficiency analyses of individual contributions (Fig. 5.7 and 5.8) do not extend to voltages higher than 400 V, we cannot explicitly point to the dominant drivers for this increase. Moreover, we see only minor changes in each efficiency mode (i.e. within uncertainty) between 300 V and 400 V. Qualitatively, however, we can draw upon previous studies of unshielded thrusters to explain physically why our efficiencies increase with voltage [147, 149]. These studies have suggested that both current and mass utilization increase with voltage. The reason underlying the former trend is not well-understood; this stems from the fact that beam utilization is linked to electron mobility within the channel, which is non-classical and poorly understood in these devices [24]. The latter trend in mass utilization can be explained in part by the higher electron temperatures that occur at higher voltages, leading to improved ionization rates and therefore mass utilization [150].

A notable behavior we see on the H9 that matches what has been previously observed on unshielded Hall thrusters is the improvement in krypton efficiency with current density [9, 10, 11, 12]. One study has attributed krypton's higher efficiencies at high current densities to increases in the mass utilization, current utilization, and beam divergence [9]. Although we only have two conditions to compare, these are the same three efficiencies for which we see the gap between xenon and krypton close when comparing the 300 V, 15 A condition to the 300 V, 20 A condition (Fig. 5.10). Additionally, as seen in Fig. 5.5c, the anode efficiency of krypton at the 20 A condition is 54.8%, the highest observed for krypton; at 600 V, 15 A, we measured an anode efficiency of 54.7% for krypton. This behavior may point us towards potential optimization strategies for krypton operation on shielded thrusters.

5.4.2 Comparison of Xenon and Krypton Operation

We look in this section to the three conditions for which we have probe data to help elucidate major drivers for the gap in efficiency between xenon and krypton. To this end, we show in Fig. 5.10 the differences between the two gases at each condition, $\Delta\eta = \eta_{Xe} - \eta_{Kr}$. Positive values indicate that the xenon efficiency is higher than that of krypton, while the reverse is true for negative values. As we discussed qualitatively in Sec. 5.3, the efficiency where krypton shows the largest disparity from xenon is the mass utilization, where the gap ranges from 7–14% with an uncertainty of about $\pm 5\%$ for all three conditions. The possible reasons for this gap are discussed in detail in Sec. 5.4.3. The current and divergence efficiencies also contribute to the gap between xenon and krypton, with $\Delta\eta_b$ ranging from 2–5% with an uncertainty of about $\pm 4\%$ and $\Delta\eta_d$ ranging from 3–8% with an uncertainty of ± 1 –2%.




Figure 5.10: Differences in the contributions to anode efficiency between krypton and xenon.

The second-highest driver of the lower krypton efficiency is the divergence. This is likely due to differences in the location of the acceleration region, defined as the narrow axial length over which the majority of the ions are accelerated. An acceleration region located further downstream of the exit plane is more curved, therefore contributing to divergence losses in the beam as ions can be accelerated radially as well axially [49, 151]. It has previously been shown that on an unshielded Hall thruster, the acceleration region moves downstream for krypton relative to xenon [152]. On an MS thruster, the peak of the acceleration zone is

usually displaced slightly downstream compared to US thrusters as a result of the magnetic field geometry [49]. Further downstream movement of this region, assuming the trend for krypton on US thrusters remains the same for MS thrusters, would only exacerbate this efficiency loss. We do note here that later LIF measurements indicate that this is not the case and that the acceleration region actually shifts upstream with increasing voltage on krypton (Fig. 7.7). Resolving this discrepancy is left as an exercise to the reader.

Finally, the third highest difference between the two propellants is in beam utilization efficiency. Low beam current utilization is an indication of reduced electron confinement [18]. Our result therefore suggests that the electron mobility may be higher for krypton operation than it is for xenon at these conditions, increasing the electron current in the channel (see Eq. 3.7). As electron transport in Hall thrusters is still a poorly-understood phenomenon [24], we are limited in our speculation of why this difference in efficiency exists.

Krypton holds a slight advantage in charge utilization of 2–3% with an uncertainty of about $\pm 3\%$ over xenon. This trend is unsurprising given the higher ionization energy for both singly and multiply charged states of krypton relative to xenon [153], resulting in a lower population of multiply-charged ions for krypton operation. This higher charge utilization efficiency of krypton also matches results reported in previous studies of unshielded thrusters [9, 147].

5.4.3 Comparison of Trends in Shielded and Unshielded Thrusters

We mention in our presentation of the results in Sec. 5.3 that we observed a gap between xenon and krypton efficiency that does not close at high voltages, contrary to what has been reported in previous studies of unshielded thrusters. We expand upon potential drivers of this difference here.

One important consideration is that there are factors beyond the change in magnetic field geometry that could lead to this difference. For example, a number of previous studies on unshielded thrusters were conducted in constant-density mode instead of constant-power mode; i.e., instead of matching current between xenon and krypton, they matched neutral density at the anode [11, 147, 152]. Some of the data observed at various operating conditions on these thrusters may therefore be difficult to directly compare with our results. Additionally, there are a number of design differences beyond the shape of the magnetic field between the various thrusters from these previous studies and the one we present here. These include the dimensions of the channel, the position of the cathode used, and the sizing of the thruster. These features can influence the overall efficiency. For example, the dimensions of the channel impact electron confinement, affecting the current utilization, and the position of the

cathode is tied to the coupling voltage, affecting the voltage utilization. The implication is thus that shielding may not be the only reason for the larger efficiency gap observed at high voltages between krypton and xenon.

With that said, changes between the unshielded and shielded configuration ultimately may still be a driving contributing factor. Notably, there are changes in the shape and magnitude of the electron temperature profiles of unshielded and shielded thrusters [49], in turn leading to differences in the profiles of ionization rate coefficient. As the electron temperature in Hall thrusters typically scales with discharge voltage ($V_d \propto T_e$ [18]) and this temperature directly impacts the ionization, we anticipate that the mass utilization in shielded and unshielded thrusters may respond differently to variations in voltage. There may also be changes in the nature of the plasma density profile within the channel of the thruster—in particular, the fact that shielded thrusters have lower plasma densities along the walls may also change the density profile along centerline. Both the changes in ionization rate coefficient and plasma density may have a significant impact on the scaling of the mass utilization with increasing voltage for xenon and krypton. Absent measurements of these internal plasma properties, however, we reserve our discussion for the specifics of this trend to our later work in Chap. 7.

5.4.4 Strategies for Optimizing Performance on Krypton

Our results from Fig. 5.10 indicate that the largest gaps between xenon and krypton efficiency are in mass utilization, beam utilization, and divergence. Improving these efficiencies is therefore the critical driver for krypton to become competitive with xenon as a high-performance Hall thruster propellant. We discuss in the following section potential strategies for increasing the efficiency of each of these contributors.

In order to increase the beam utilization efficiency, it is ultimately necessary to reduce the electron current in these devices. However, as we have remarked previously, the processes driving electron current remain poorly understood [24]. It is therefore difficult to determine why this discrepancy exists and in turn strategies to mitigate it. With that being said, electron current typically is reduced in these devices by adjusting the magnetic field strength. In our study, we adjusted the magnetic field to minimize discharge oscillation strength at each set point. This minimum is not always coincident with the minimum in electron current [29]. There is then potentially some margin for optimizing the magnitude of the magnetic field to reduce beam utilization losses.

The challenges with improving divergence efficiency mirror those with beam utilization. Indeed, since the beam divergence is linked with the location of the acceleration zone, which in turn is tied to the electron dynamics, it is difficult to identify physics-based strategies for

closing this gap. We do note however that recent work has shown that increasing the cathode flow fraction has the effect of pushing the acceleration zone upstream in the thruster, thereby lowering divergence losses [154, 155]. This behavior potentially recommends a strategy for closing this efficiency gap on a krypton thruster, although operationally, the increase in cathode flux will represent a new overall efficiency loss to the thruster.

One established method for improving mass utilization is to extend channel length, therefore increasing the residence time for ionization [11, 13]. The theoretical justification for this approach is illustrated in Eq. 3.15 where we see a linear dependence on the ionization length. We note, however, that this technique may not have the same degree of efficacy on shielded thrusters compared to unshielded thrusters. This stems from the fact that the magnetic shielding topology results in a downward shift of the plasma discharge with respect to the thruster exit plane. The ionization zone can therefore partially extend downstream of the channel. Increasing the channel length on a shielded thruster ultimately may not have as pronounced of an impact on improving ionization.

In addition to changing geometry, based on our results in Fig. 5.5c and 5.10, we see that the efficiency gap for krypton and xenon appears to close at the high-current condition. Equation 6.2 offers a justification for this behavior: as discharge current density I_d/A increases, so too does the plasma density and the mass utilization efficiency is also expected to increase. This suggests that at sufficiently high currents, the entire mass flux will be ionized, ultimately closing the dominant driver of the efficiency gap between krypton and xenon. We explore this possibility in the next chapter.

One final comment on the operation of a magnetically-shielded Hall thruster on krypton is regarding the high oscillations seen at high voltages. At 500 V, we observed oscillations of 160% in the discharge current, and this increased to 210% at 600 V. While we do not yet have conclusive strategies to dampen these oscillations, it is a critical factor to consider for thruster stability when operating these devices on krypton at high voltages.

5.5 Summary

In this work, we have measured global performance metrics and the various contributions to anode efficiency of xenon and krypton at discharge voltages ranging from 300–600 V and discharge currents of 15 and 20 A on a 9-kW class magnetically-shielded Hall thruster. We have found that xenon operation reached a maximum anode efficiency of $68.3 \pm 1.4\%$ at 600 V and 15 A, while krypton reached a maximum anode efficiency of $54.8 \pm 0.8\%$ at 300 V and 20 A. We have shown that generally, the anode efficiency of krypton was worse than that of xenon at the same condition by 9–14% with an uncertainty of 2–3%. This gap is in

family with the 5–15% that has been observed from previous studies of krypton operation on unshielded thrusters.

We have discussed potential physical causes underlying the trends in our global performance parameters. As the discharge voltage increased, so did the acceleration voltage and the electron temperatures along centerline, improving the voltage and mass utilization. This resulted in the improved efficiencies seen at higher voltages. The specific impulse was improved by the higher acceleration voltages as well. For increasing current at a fixed discharge voltage, we saw higher thrust due to the higher anode flow rate and an improvement in anode efficiency due to the improved mass utilization at high current densities for both propellants. We contextualized these trends in anode efficiency by using far-field probe data to evaluate the major efficiency terms for the thruster operating at 300 and 400 V. From these results, we identified mass utilization as the primary detractor to krypton performance in comparison with xenon. This behavior also matches what has been seen on previous studies of unshielded thrusters.

Finally, we have discussed potential strategies for closing the efficiency gap between krypton and xenon operation on magnetically-shielded Hall thrusters. These include changes to design and operation that address the specific efficiencies where krypton is worse than xenon, namely in mass utilization, beam utilization, and divergence. Some of the proposed strategies for improving krypton performance include increasing the magnetic field strength, increasing cathode flow fraction, lengthening the thruster channel, operating at higher current densities, and operating at higher voltages.

In summary, we have shown in this work that an MS thruster optimized for xenon can operate on krypton at the expense of a reduction in performance. This behavior ultimately does not represent a significant departure from the conclusions reached on previous studies of unshielded thrusters. One difference in trend is in the unfavorable scaling of mass utilization for krypton compared to xenon, a point we return to in Chap. 6. With that said, our comprehensive analysis in this work of the factors driving these efficiency trends can be leveraged to improve future designs and iterations of magnetically-shielded, krypton-operated Hall thrusters. This may enable a new paradigm of higher-power, longer-lifetime thrusters operating on this attractive, alternative propellant.

CHAPTER 6

Performance and Operation at Ultrahigh Current Densities

“...these go to eleven.”

–Rob Reiner c/o Nigel Tufnel, *This Is Spinal Tap*.




Figure 6.1: H9 MUSCLE test campaign flag. Logo design credit to T. M. Gill.

6.1 Introduction

This chapter covers work initially published in Ref. [72, 108, 156].

As discussed in Sec. 3.6, there is strong incentive to investigate the perceived upper limit on current density in Hall thrusters. One common justification for a physics-based limit is that further increasing the current density decreases electron confinement and therefore leads to a degradation in performance. However, this argument is based on the assumption that the electron dynamics in Hall thrusters scale classically [39], whereas in practice, the electron dynamics in $E \times B$ devices are driven by non-classical effects [24]. To this point, Simmonds et al. recently performed a theoretical study of the thrust density in Hall thrusters allowing for non-classical transport [89]. These authors ultimately concluded that the theoretical limit on thrust density is on the order of 1000 N/m^2 , two orders of magnitude higher than the currently accepted 10 N/m^2 limit.

To this end, we conducted a preliminary study on the standard H9 Hall thruster, a magnetically shielded 9-kW class Hall thruster, operating on only xenon [108]. We ran the H9 up to 40 A (a factor of $2.7 \times$ higher than its nominal current of 15 A) with the magnetic field strength unchanged from normal operation and held a constant voltage of 300 V. Due to the limited nature of this preliminary study, particularly the fact that we only operated on xenon, we only summarize here key findings from this work. Further details may be found in Ref. [108]. First, we saw a linear increase in thrust with increasing current to a maximum of 700 mN at 12 kW. Second, despite concerns regarding decreasing performance at high current densities, we actually saw the anode efficiency generally increase with increasing current (from a minimum of $\sim 62\%$ at 20 A to a maximum of 66% at 40 A). Finally, from probe measurements taken at 15, 20, and 40 A, we saw an increase in the mass utilization but a decrease in the beam and voltage utilizations. We caveat this last point with the fact that although we did see trends in these efficiency components, the differences fell within uncertainty of each other. Based on these takeaways, we found the path forward for high current density testing to be promising. In particular, it appeared that we had not yet reached a maximum efficiency for xenon operation as it was still increasing at 40 A.

As additional motivation to test at high currents, we saw in our previous experiment that the efficiency gap between xenon and krypton may close with increasing current due to mass utilization approaching unity. In light of these theoretical arguments, previous exploratory studies, and the major advantages of being able to operate Hall thrusters at higher current density, the need is apparent to systematically explore Hall thruster operation in this high current density regime. In other words, we need to take Hall thruster current density and turn it up to eleven.

As a secondary consideration, as we mention in Sec. 3.6, there remain major technical challenges in the testing of any high-power Hall thruster (or other EP devices) in ground-based facilities. These include the influence of background pressure on thruster performance and stability [76], thermal loading on the facility [77], and backsputter rates that can obscure the results of wear tests [47]. With these challenges in mind, a NASA Space Technology Research Institute, the Joint Advanced Propulsion Institute (JANUS), was recently formed to study the problem of ground testing for higher power concepts [75]. One major obstacle for investigating facility effects at high power is that our experience at these power levels is limited. Cumulative testing time for Hall thrusters on the order of 100 kW across the entire EP community in published work is on the order of just a few days. An additional goal of this effort was therefore to gain experience with the issues that may be encountered while testing at higher powers. These insights may inform follow-on work at JANUS.

With these two objectives in mind—to investigate fundamental limits of current density in Hall thrusters and to characterize facility effects at high-power operation—this work is organized in the following way. In Sec. 6.2, we give an overview of the experimental setup including the test article, facility, and operating conditions, as well as the diagnostics and analysis methods used. In Sec. 6.3, we present the key results of our study. Finally, in Sec. 6.4, we discuss explanations for our findings, their impacts on our understanding of high current density Hall thruster operation, and their implications in the broader context of high-power electric propulsion.

6.2 Methodology

In this section, we describe the experimental setup for this test campaign. We first overview the test article and the modifications we employed to adapt it for testing at increased power density. We then describe the test facility and thruster operating conditions.

6.2.1 Thruster and Facility

To modify the H9 for high powers, we replaced the boron nitride discharge chamber with a graphite chamber and implemented water cooling in the thruster. This modified version of the H9 is the “H9 MUSCLE” (Fig. 6.2). Details on the design and history of both these thrusters can be found in Sec. 4.2. The pathfinding campaign with the H9 is described in Ref. [108] but used the same setup shown in Fig. 5.4 and analysis methods outlined in Sec. 5.2 from the previous chapter. All methodologies described in this chapter are exclusively for the H9 MUSCLE campaign.




Figure 6.2: H9 MUSCLE operating on a) xenon at 300 V and 125 A and b) krypton at 300 V and 150 A.

We ran this test in LVTF at the University of Michigan with 16 pumps on at all times (13 cryos and 3 thumpers as described in Sec. 4.3). Due to the high flow rates we were anticipated for this effort, we also expected to see high facility pressures. This correlates to a risk of the pumps failing from overheating (i.e. attempting to pump too much gas). As a precaution, we closely monitored the temperatures of the pumps on our data acquisition system at all times during thruster operation to ensure that they did not rapidly rise in temperature.

A top-down view of the chamber and experimental setup of the H9 MUSCLE campaign is shown in Fig. 6.3. We mounted the thruster on a thrust stand and oriented it along the centerline of the chamber directed toward the downstream probe suite and beam dump. We aligned a Faraday probe and camera vertically to thruster centerline and swept them azimuthally at a constant radius with respect to the thruster. We note here that although our ion gauge's pressure readings were slightly higher than pressures previously reported for similar operating conditions on the H9 (Tab. 5.1), we found comparable values to the pressures measured by the Stabil with a Hornet ion gauge from InstruTech during this test campaign. We used a 3000 sccm mass flow controller for anode flow, a 400 sccm mass flow controller for cathode flow, and a 150-kW supply for discharge power. More details regarding this infrastructure can be found in Sec. 4.2.2.




Figure 6.3: Notional top-down schematic of experimental setup in LVTF, including thruster, thrust stand, ion gauge, quartz crystal microbalance sensor, Armored Probe Carrier with E×B probe, Langmuir probe, and retarding potential analyzer, probe arm with Faraday probe and camera, and beam dump. Note that the QCM was mounted above the ion gauge to prevent thermal interference. Pumps are located around the walls of the facility but not shown in the diagram.

Table 6.1 shows the operating conditions we experimentally characterized in this campaign. We operated the thruster at a constant discharge voltage of 300 V for all points and varied the anode flow rate to achieve the target discharge current. The cathode flow fraction was held at 7% for all operating conditions and we employed a magnetic field strength 112.5% of its nominal value for this thruster. The thruster body was electrically tied to facility ground, a necessary configuration resulting from the cooling system we employed. As a general approach, we progressively increased thruster power during the test campaign by increasing the flow to control current. At each new operating condition, we waited until the mean discharge current achieved a constant value before performing a far-field probe sweep. We then shut down the thruster to take a thrust measurement (Sec. 4.4.1). Finally, we restarted the thruster and rapidly increased power by way of flow rate and current to the next operating condition. This process typically resulted in a 5–15 minute dwell time at each operating point higher than 15 A (Fig. 6.12). While we attempted to visit the same

discharge current conditions for both gases, we were unable to reach the 150 A condition for xenon due to a later-resolved issue in the electrical harnessing.

Species	I_d (A)	P_d (kW)	\dot{m}_a (mg/s, sccm)	Pressure (μ torr)
Xe	15	4.5	15.2, 169	5.7
Xe	50	15	35.9, 400	12.3
Xe	75	22.5	46.7, 519	16.4
Xe	100	30	57.9, 644	20.3
Xe	125	37.5	68.1, 757	23.6
Kr	15	4.5	11.4, 199	5.9
Kr	50	15	30.7, 537	14.5
Kr	75	22.5	41.5, 726	19.2
Kr	100	30	50.6, 885	23.7
Kr	125	37.5	59.5, 1040	28.3
Kr	150	45	68.3, 1193	31.8

Table 6.1: Operating conditions and base pressures. The discharge voltage was 300 V and the B-field ratio was 112.5% for all conditions.

6.2.2 Diagnostics and Data Analysis

In this section, we describe the diagnostics we used to characterize the thruster performance. We begin with a description of the thrust stand we employed to measure the global performance. We then describe the far-field probes we used to measure plasma properties and how they relate to calculations of the various efficiency modes. We conclude with an overview of the diagnostics used to monitor the health of the thruster and facility.

We used our standard null-type inverted pendulum thrust stand (Sec. 4.4.1), calibrated with a series of known masses corresponding to thrusts ranging from 10 mN to 4.8 N. The measured thrusts and mass flow rates from the thrust stand and the DAQ were used to obtain values of specific impulse, anode efficiency, and total efficiency. Uncertainties in mass flow rate and thrust were calculated following the methodology outlined in Sec. 4.4.1. We note here that although we did take into account η_{fac} , the artificial efficiency boost from neutrals in the facility being ingested into the Hall thruster channel, this represented less than a 0.5% change at all conditions. This is much smaller than the impact that neutral ingestion had on X3 testing at similar currents [26], which we attribute to the fact that the larger surface area of the X3 allows for more ingestion. Additionally, as mentioned in Sec. 3.3.2, Eq. 3.5 may not account for some sources of “effective” neutral ingestion that may also be artificially inflating the efficiencies. A more thorough characterization of these

facility effects may be performed by varying pressure and taking thrust measurements; this would make an ideal candidate for the subject of future work.

We employed a standard set of far-field probes for this experiment including a swept Faraday probe, RPA, Langmuir probe, and $E \times B$ probe. With the exception of the Faraday probe, which was mounted on an azimuthally swept arm, these probes were located ~ 14 thruster diameters downstream of the thruster exit plane. As our goal was to investigate thruster performance at abnormally high current densities, we had anticipated that thermal loading and erosion may have posed a challenge for the protection of these probes. To attempt to mitigate these issues, we implemented a probe shield, the Armored Probe Carrier (APC), which consisted of a linear motion stage shielded by three large graphite plates (Fig. 6.4). The RPA, Langmuir probe, and $E \times B$ probe were mounted at different positions along this translation stage (Fig. 6.4b). The shields were thermally isolated from the main structure of the APC with stainless steel washers to prevent the motion stage from overheating at high thruster powers. A small circular cutout was located in the central graphite shield and aligned to thruster channel centerline at the 3 o'clock position. When not actively using one of the probes located on the APC, the suite was in a rest position where all probes were fully shielded by the graphite panels. During the data collection phase, each probe was translated to the appropriate location such that it had a line of sight to the thruster channel through the cutout.

We used an RPA as described in Sec. 4.4.2.1 to measure average ion energies between the plume and ground, sweeping the collector from 0 to 500 V. We note here that at high discharge currents, the RPA trace exhibited marked noise in the signal at ion selector voltages above 400 V (Fig. 4.3b). We believe this was due to the increasingly high current density in the diagnostic at these operating conditions, leading to internal arcing of the grids. As this probe was originally designed for standard Hall thruster operating conditions, this type of breakdown at high current densities was not anticipated. It does suggest, however, that additional consideration should be given to the design of diagnostics for this operational regime. We return to this point in Sec. 6.4.5.

Figure 4.4 represents an example of raw and processed data for xenon operation at 300 V and 125 A. We typically swept the potential of the LP from a bias voltage of -10 to +30 V and measured the resultant current from the plasma via a picoammeter. During some of the high current density conditions, the electron current collected near $V_{bias} = +30$ exceeded 2 mA, the maximum input of the picoammeter. In these cases (such as in Fig. 4.4), the over-range signals in the electron saturation region were removed. At low currents, we were unable to obtain a clean estimate of V_p with the “knee” method at every condition. We therefore used only the Maxwellian method to determine the plasma potential at all

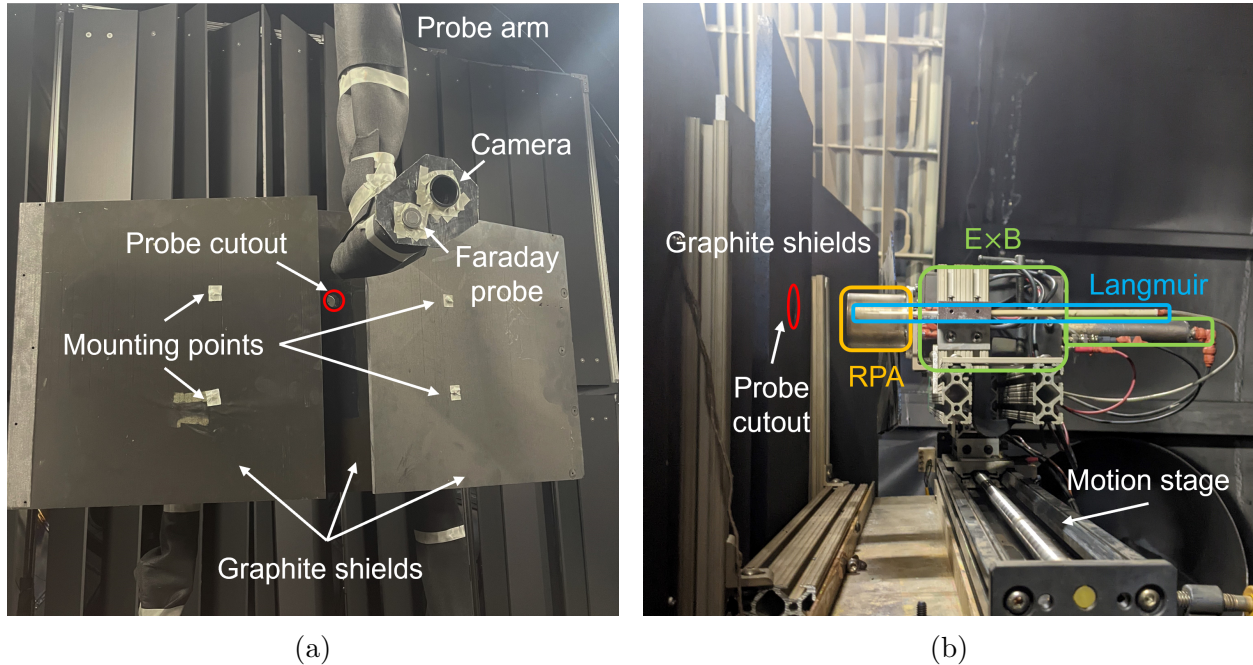


Figure 6.4: Photos of probe setup inside the vacuum facility showing a a) front view of the APC and probe arm, and b) side view of the APC. The front view shows the graphite shields from the thruster’s perspective, as well as the swept Faraday probe and camera. The side view shows the three probes (from nearest to furthest: Langmuir, RPA, $E \times B$) mounted on the motion stage behind the graphite panels.

conditions and quantified a characteristic uncertainty for a given propellant as the average difference between the two methods for calculating V_p at conditions where $V_p(\text{knee})$ was able to be determined. We note here that while using an emissive probe to estimate plasma potential is preferable, emissive probes are also very prone to burning out during operation. Due to the already-perilous environment of the high-current testing regime, we opted to go with the simpler and more robust Langmuir probe.

The $E \times B$ probe was a critical diagnostic for this campaign specifically due to the high populations of multiply-charged we expected to see given our results from 40 A [108]. We swept the bias voltage to the plates of our $E \times B$ probe from 0 to 80 V at lower current densities and up to 100 V at higher current densities to capture all the multiply charged species content. We processed this data generally following the methodology outlined in Sec. 4.4.2.4 and 4.4.3.2 for up to three charge states. For the high current conditions, we fitted the two-peak Gaussian to the location we anticipated the singly charged species to be at for a given propellant based on the 15 A trace where singly charged ions dominate (Fig. 4.5a).

The Faraday probe we used in our experiment was mounted on a rotational probe arm

located ~ 10.5 thruster diameters downstream of the thruster exit plane (Fig. 6.4). During operation, the Faraday probe was swept from 0 to 180 degrees and back again with 90 referenced to thruster centerline (see Fig. 6.3). The range of angles we used for our calculations of beam current and divergence angle were taken from sub-samples of the total sweep: 0–90, 90–180, 180–90, and 90–0 degrees. We used two different methods to correct for the CEX effect in the current measurement (Sec. 4.4.3.3): a Gaussian fit for calculating beam current [120] and an exponential fit for calculating divergence angle [119]. For each of the four sweep segments, we fit a Gaussian curve over the entire beam and used Eq. 4.17 to obtain a value of beam current. The uncertainty was then determined to be twice the standard deviation in this set of four values. We used these resulting values of beam current in the estimates for beam utilization and mass utilization efficiencies.

We found the divergence angle using the methodology outlined in Sec. 4.4.3.3 with Eq. 4.18. In calculating this parameter, we used only the corrected current density trace from the exponential fit method. This is because the other CEX correction methods have not been shown to yield accurate estimations of divergence angle without extrapolating back to the exit plane of the thruster [71, 120]. As we did not take Faraday probe traces at multiple radial distances, fitting the downstream divergence back to the thruster exit plane was not possible. The exponential fit method has been used previously to calculate divergence angle without any additional extrapolation [119], so this method alone was selected for divergence angle calculations. Once we found the beam divergence, we used it to estimate the divergence efficiency (Eq. 3.11).

We implemented a series of diagnostics to monitor thruster health in real time. Due to the atypically high-current operating regime, we paid particular attention to the thermocouples located around the thruster and the chamber (Sec. 4.2.3), especially the one located at the inner front pole (IFP) because we anticipated it to reach the highest temperatures. One thermocouple also monitored the temperature on the motion stage of the Armored Probe Carrier (Fig. 6.4a) to ensure that it did not overheat at the high current conditions. For visual inspection of the thruster, we mounted an Ethernet camera (Sanyo model VCC-HD2500P) on the azimuthal probe arm adjacent to the Faraday probe (Fig. 6.4). The camera was placed behind a sacrificial glass optic and had no difficulty operating continuously in vacuum and briefly in the thruster plume as the arm was swept.

6.3 Results

We present in this section the results of our experimental characterization of the thruster operating on both xenon and krypton at 300 V and discharge currents ranging from 15 to

150 A. We first report on the global performance metrics including thrust, specific impulse, and efficiency. We then present the processed results of individual probe measurements and thruster efficiency modes. We conclude with an overview of key parameters related to thruster and facility health including oscillation levels, temperature data, carbon deposition rates, and facility pressures.

6.3.1 Global Performance

Figure 6.5 presents the thrust, specific impulse, anode efficiencies, and total efficiencies for both propellants at all conditions. Exact values are tabulated in App. B.2.

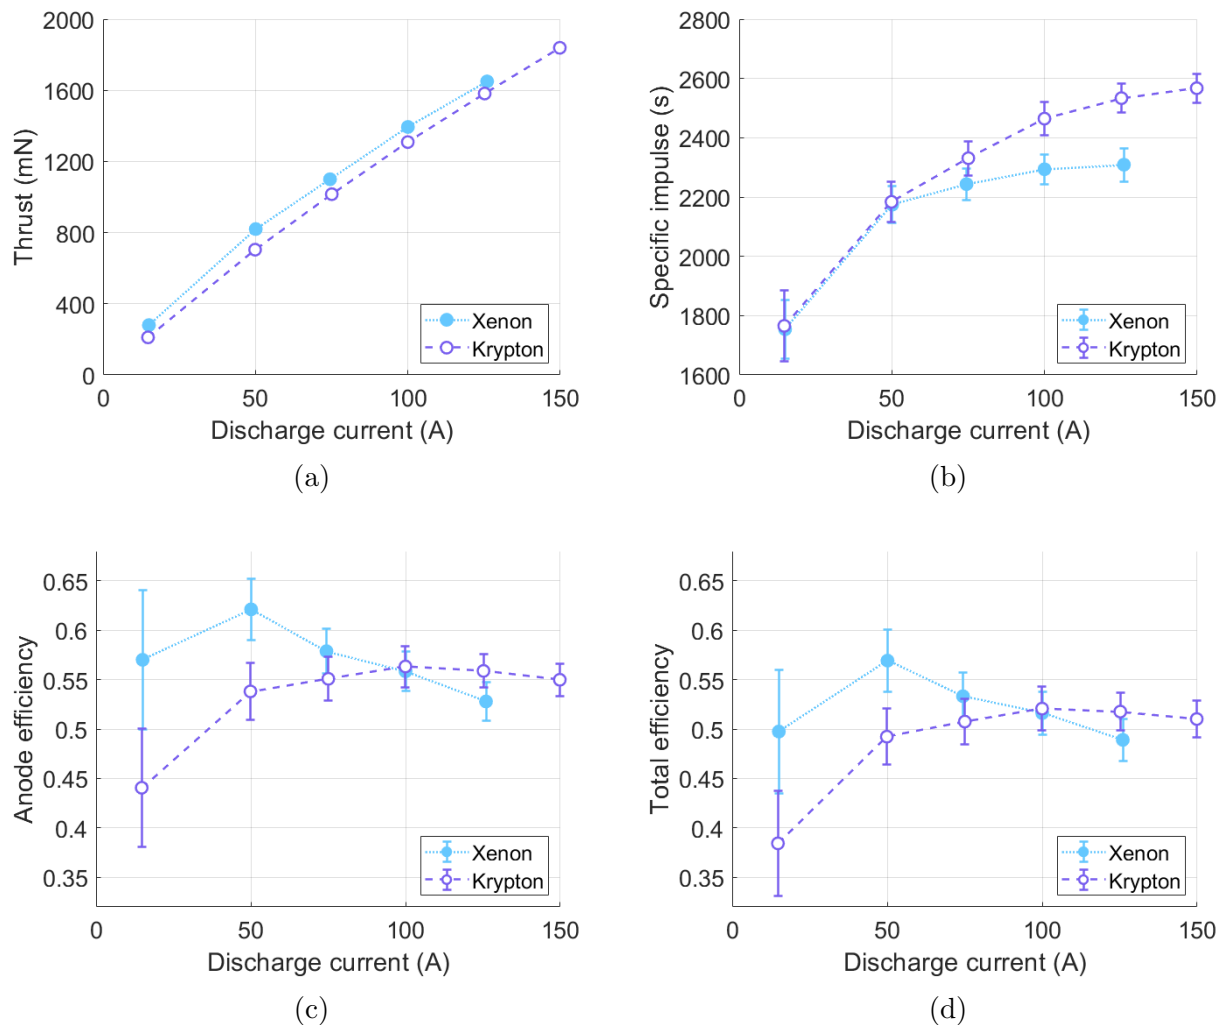


Figure 6.5: a) Thrust, b) specific impulse, c) anode efficiency, and b) total efficiency at 300 V for xenon and krypton with increasing discharge current. Note that the uncertainties in Fig. 6.5a are smaller than the marker sizes.

6.3.1.1 Thrust

As can be seen from Fig. 6.5a, the thrust trends approximately linearly across the range of currents for each propellant. The xenon thrust exceeds krypton at all discharge currents; at the maximum current where we tested both gases, 125 A, xenon reaches 1650 ± 30 mN and krypton 1582 ± 16 mN. Krypton achieves the maximum overall thrust recorded from this effort with a value of 1839 ± 18 mN at the 150 A condition. We note that the thrust for both propellants measured at the 4.5 kW condition are about 10–20 mN (or 4–9%) lower than those previously measured on the H9 [71]. We attribute this discrepancy to differences between the baseline H9 and the modified H9 MUSCLE, such as changes in material and electrical configuration (e.g. body tied to ground instead of body tied to cathode). Most saliently, this may have been in part caused by the switch from a BN to graphite chamber, which has been shown to slightly decrease performance on other shielded thrusters [52, 54]. Additionally, the effective channel length of the H9 MUSCLE is slightly shortened compared to the H9 due to design changes required to accommodate electrical isolation between the anode and thruster body. This could significantly reduce the mass utilization efficiency, which we do in fact see decrease between the H9 to the H9 MUSCLE at the same operating condition. Krypton’s mass utilization at 300 V and 15 A decreases from 83% to 68%, while xenon’s decreases from 96% to 84% (see App. B.1 and B.2).

6.3.1.2 Specific Impulse

As shown in Fig. 6.5b, the specific impulses measured during krypton performance are all higher than that of xenon at the same discharge current, although the difference is within uncertainty at the 15 A and 50 A conditions. At 150 A, krypton operation yields a specific impulse of 2567 ± 48 s. At 125 A, krypton has a specific impulse of 2534 ± 48 s and xenon has 2309 ± 56 s. This disparity can be attributed to the fact that krypton is a lighter atom—since the discharge voltage remained the same at each test point, the approximately constant amount of electrostatic acceleration resulted in higher exit velocities and specific impulse for krypton, the lighter gas. This follows from the relation that the theoretical limit for specific impulse with only singly charged ions is $I_{sp,th} = \frac{\sqrt{2eV_a/m_i}}{g_0}$. Similar comparisons between krypton and xenon specific impulse have been reported in previous studies [6, 7, 8, 9, 10, 11, 12, 13, 71]. The disparity in specific impulse is related to the differences in thrust exhibited in Fig. 6.5a. Thrust-to-power scales inversely with specific impulse, so xenon—which has a heavier mass per atom—has a lower specific impulse and higher thrust at fixed power.

We note here that for a beam of only singly charged ions at 300 V, the theoretical limit where all efficiencies are unity is 2000 s for xenon and 2500 s for krypton. However, as seen

in Fig.6.5b, xenon surpasses this theoretical limit at the 50 A condition, while krypton does so at 125 A. We attribute this to the high population of multiply charged species in the beam, particularly for xenon at higher currents (see Sec. 6.3.2). We discuss this trend in further detail in Sec. 6.4.1.

6.3.1.3 Efficiency

Figure 6.5c shows the anode efficiency for both propellants. For each gas, the efficiency exhibits non-monotonic behavior, peaking at a given power before decreasing. This peak is more pronounced for xenon than for krypton, which reaches more of a plateau (i.e. efficiencies within uncertainty of each other) between 50–150 A. The maximum for xenon occurred at 50 A with an efficiency of $62.1 \pm 3.1\%$, while the maximum for krypton occurred at 100 A with $56.3 \pm 2.1\%$. We note that the peak in xenon efficiency at 50 A is consistent with the behavior we saw in our preliminary study on the H9, where the efficiency monotonically increased from 15 to 40 A [108]. Practically, we remark that although it is evident that the efficiency does decrease at higher currents, the efficiency across the current range is still competitive with state-of-the-art Hall thrusters that operate at lower currents. Indeed, the lowest anode efficiency for krypton operation occurs at the nominal operating current of 15 A. The increase in efficiency by increasing current above 15 A for krypton operation is in line with previous observations [9, 71]. The lack of a marked decrease in efficiency with current (and current density) is an indication that the scaling laws historically governing Hall thruster sizing may not be absolute. This has direct implications for the achievable thrust density of Hall thrusters at higher powers, a point we expand upon in Sec. 6.4.6.

We conjecture that the non-monotonic trend in efficiency is the result of a trade between improved mass utilization and reduced beam utilization efficiencies. Krypton’s efficiency may peak at a higher current due to its relatively lower mass utilization at nominal current density conditions. We discuss this in further detail in Sec. 6.4.2. Also of note is that at the 100 and 125 A conditions, the efficiency of krypton becomes higher than that of xenon. Although we were not able to take data at the 150 A condition for xenon, we postulate that the downwards trend continues and the gap between xenon and krypton may widen with krypton now exhibiting higher performance.

The trends in total efficiency, shown in Fig. 6.5d, closely mirror the trends in anode efficiency. The magnitudes of η_{tot} are $\sim 4\text{--}7\%$ lower, however, because we now account for the flow through the cathode and the power supplied to the magnets. We note here again that neutral ingestion, η_{fac} , is taken into account for calculations of both η_a and η_{tot} . Xenon’s total efficiency ranges from $48.9 \pm 2.1\%$ to $56.9 \pm 3.1\%$, while krypton’s ranges from $38.4 \pm 5.3\%$ to $52.1 \pm 2.2\%$. Compared to the trends in anode efficiency, the maxima in total

efficiency are located at the same currents (50 A for xenon and 100 A for krypton) and the crossover point between propellants remains the same at 100 A.

6.3.2 Probe Measurements

We present in this section the processed results of our far-field probe measurements. We first show in Fig. 6.6 the plasma potential and electron temperature as calculated from Langmuir probe traces using the methodology outlined in Sec. 4.4.2.2. The electron temperature and the plasma potential generally trend downwards after achieving a maximum value with increasing discharge current. This maximum value occurs at 75 A (with a decline from 15 to 50 A) for xenon and at 50 A for krypton. This trend suggests that for the most part, the electrical coupling between the cathode and main discharge is improving with higher discharge currents [157]. This trend in xenon plasma potential decreasing with increasing current is in line with behavior seen on this thruster in previous work up to 50 A [108].




Figure 6.6: Plasma potential and electron temperature of a) xenon and b) krypton at 300 V with increasing discharge current.

We show in Fig. 6.7 the average ion energy inferred from the RPA for both xenon and krypton. In this case, we have made a correction to the RPA measurement with respect to ground by subtracting the local plasma potential as determined by the LP. The ion energy per unit charge, within error bars, is on the order of the applied voltage of 300 V. This is consistent with the electrostatic acceleration imparted on the ions. Notably, the average ion energy increases monotonically with xenon while generally trending down with krypton. This may suggest that the acceleration process for each propellant is impacted in opposing ways by the increasing current density in the thruster. We note here that the large uncertainties

associated with ion energy at higher currents may be due to arcing between the RPA grids, a point discussed in Sec. 4.4.2.1.






Figure 6.7: Average ion energy inferred from RPA corrected by the plasma potential for a) xenon and b) krypton at 300 V with increasing discharge current. The discharge voltage is indicated with a dotted black line.

In Fig. 6.8, we present the current fractions as calculated from the $E \times B$ data at all conditions. For both propellants, we see that the fraction of the singly charged state generally decreases while the populations of multiply charged states increase. Indeed, xenon exhibits a population inversion (i.e. where the fraction of doubly charged ions surpasses that of singly) between 50 and 75 A. While krypton never reaches this point, we see in Fig. 6.8b that it approaches a population inversion at the highest current condition of 150 A. These trends likely result from increasing current densities which in turn translate to mass utilization efficiencies approaching unity (see Sec. 6.3.3). With a mass utilization of one, all the inlet propellant is effectively ionized. Consequently, as current density increases beyond this point, there are additional collisions between the singly charged ions and electrons. This may explain the higher densities of multiply charged species. We also note that the triply charged population appears to be steadily increasing for xenon, although for krypton it remains approximately constant. This may be attributed to the higher required ionization energies to generate triply charged krypton compared to xenon.

Figure 6.9 shows the plume divergence angles as calculated by the exponential fit method (Sec. 4.4.3.3). For xenon, this angle monotonically increases with current from $22.5^\circ \pm 1.6^\circ$ at 15 A to $28.2^\circ \pm 1.5^\circ$ at 125 A. This range is consistent with previously reported values on the H9 [71, 151]. Physically, our result indicates that the beam became increasingly




(a)




(b)

Figure 6.8: Current fractions of singly, doubly, and triply ionized a) xenon and b) krypton at 300 V with increasing discharge current. Note that the uncertainties are smaller than the marker sizes.

divergent with increased current in the xenon case. In contrast, for krypton, there is a less well-defined trend as discharge current increases. Indeed, the divergence angle stays nearly constant for all currents, only ranging from a minimum of $26.5^\circ \pm 1.2^\circ$ at 50 A to a maximum of $27.9^\circ \pm 1.7^\circ$ at 150 A. We discuss this further in Sec. 6.4.2.



(a)



(b)

Figure 6.9: Plume divergence angle of a) xenon and b) krypton at 300 V with increasing discharge current.

6.3.3 Efficiency Analysis

We show in Fig. 6.10 the trends in calculated efficiency modes as a function of discharge current for both propellants. These estimates are based on the far-field plasma measurements reported in the previous section. The uncertainty in these values is from the propagation of experimental error through the governing equations of the efficiencies (Sec. 3.3.2). Additionally, we compare the anode efficiency as measured directly from thrust, $\eta_{a,thrust}$ (Eq. 3.4), to the anode efficiency as calculated by the product of partial efficiencies, $\eta_{a,probe}$ (Eq. 3.6). The values of each efficiency mode are shown in Appendix B in Tables B.7 and B.8. We discuss in the following section key trends in each component of the efficiency.

6.3.3.1 Anode Efficiency

The gray line in Fig. 6.10 indicates the anode efficiency as determined from direct thrust measurements (Eq. 3.4), while the black line indicates the anode efficiency as calculated as the product of partial efficiencies (Eq. 3.6). Generally, $\eta_{a,probe}$ is 4–9% lower than $\eta_{a,thrust}$, but these values are typically within uncertainty and follow the same trends with increasing current. This indicates that the efficiency model outlined in Fig. 3.3.2 is an accurate representation for the performance of the thruster. The discrepancy between these methods of calculation implies that we are underpredicting one or more of the partial efficiency modes that contribute to $\eta_{a,probe}$. As we elaborate on later in this section, the mass utilizations (and therefore beam utilizations) are unlikely to be an underestimation due to how close they already are to 100%. This points to the remaining efficiency modes—voltage, divergence, and charge—as the potential culprits. Indeed, absent Faraday probe traces are multiple radii, we were only able to calculate the divergence angle using the exponential fitting method (Sec. 4.4.3.3), potentially leading to an overestimation of divergence angle. Additionally, the difficulty of deconvolving the various current species (Fig. 4.5), particularly at high currents, may lead to slightly erroneous species fractions. Practically, however, our two methods of calculating anode efficiency follow the same trends with increasing current, pointing to the qualitative behavior of the efficiency modes (if not the exact quantitative values) as being an accurate reflection of the physics underlying thruster operation. For xenon, both anode efficiencies peak at 50 A, and for krypton, both anode efficiencies peak between 100–125 A but generally have a plateau from 50–150 A.

6.3.3.2 Beam Current Efficiency

The beam current utilization steadily decreases for xenon with increasing current. With krypton, the current utilization efficiency increases from 15 to 50 A before slightly decreasing.




Figure 6.10: Efficiency modes for a) xenon and b) krypton at 300 V with increasing discharge current.

ing throughout the higher current conditions. This downwards trend in beam utilization efficiency has been observed in previous studies of Hall thruster efficiency at high powers, and the values ($\sim 70\%$) are in line with what is typically observed on Hall thrusters [18, 27, 108, 130, 132, 149, 158, 159, 160] albeit lower than our previously reported measurements for the baseline condition [108]. This discrepancy may in part be due to the different methods used to calculate beam current in this work compared to our previous work. The trends in beam utilization suggest that electron confinement does decrease at higher currents, leading to a higher relative contribution of electron current to the total current. This behavior is broadly consistent with past phenomenological explanations that have been proposed (c.f. Refs. [18, 39]) for the impact of plasma density within the discharge channel on Hall thruster operation. Indeed, it has been suggested in these works that the higher plasma density associated with higher current density will lead to enhanced electron collisionality and therefore higher degrees of cross-field electron transport. Notably, however, we see that the decrease in the beam utilization efficiency is at most $\sim 7\%$ over an order of magnitude increase in current. This suggests that the loss of electron confinement is not driving the overall performance to non-competitive values. The relatively gradual downwards trend in beam utilization largely explains why the thruster performance has not suffered markedly at these high current conditions. This behavior is further explored in Sec. 6.4.2.

6.3.3.3 Charge Efficiency

As can be seen from Fig. 6.10, the charge utilization is nearly unity for the range of investigated currents. We see that for xenon, there is a minor decrease in charge utilization in the range from 15–75 A before increasing again. This behavior can be explained by the relative disparity of charge states—charge utilization reaches 100% when the beam is monoenergetic and deviates from this when the beam is polydispersive, i.e. containing differently charged species. Ultimately, the change in charge utilization is minimal (<2%) throughout the entire range and therefore not a major contributor to the trends in anode efficiency.

6.3.3.4 Mass Efficiency

The mass utilization efficiencies for xenon and krypton are 84% and 68% respectively at the 15 A discharge current conditions. For krypton in particular, the mass utilization at this current is its lowest partial efficiency. These values are in line with previous results indicating low mass utilization for krypton at typical Hall thruster current densities [8, 9, 10, 11, 13, 71]. With increasing current for both gases, the mass utilization reaches 100% within uncertainty. This behavior is physically consistent with the interpretation that the plasma density at these conditions is sufficiently high such that all inlet neutral gas is ionized. Xenon reaches full ionization at 50 A, while krypton approaches this level at 100 A. This difference is likely due to the heightened difficulty of ionizing krypton, which is mainly attributed to its smaller ionization cross section (see Sec. 6.4.2). However, as the current density and therefore plasma density in the channel increases, the ionization mean free path becomes small enough such that all the neutrals are ionized even for krypton.

We note here that some of our values for mass utilization exceed 100%. These non-physical values occur at 50–75 A for xenon and 100–150 A for krypton. The average values that exceed 100% may be attributed to the limitations of our attempts to deconvolve the impact of CEX ions or to the validity of key assumptions underlying our calculations. Indeed, the flow rates into the facility are exceedingly high (by up to an order of magnitude) compared to standard testing, and it thus may not be unexpected that standard convolution methods may be compromised. As an alternative explanation, we note that it is common to assume that cathode ions have a negligible contribution to the downstream current density. This stems from the fact that the majority of current from the cathode is carried by electrons and therefore do not contribute to thrust. With that said, if ions from the cathode are accelerated through a sufficient potential drop to contribute to the downstream current density, this could manifest as a higher-than-unity efficiency in the mass utilization. This process may be a possibility given the atypical operating conditions for the thruster, and

indeed, it would explain why we only see the mass utilization exceed 100% when the discharge current exceeds the nominal maximum 20 A of the H9. With that said, we do note that all measurements are within experimental uncertainty of 100%. In our subsequent discussion, we thus work under the assumption that practically, the total mass utilization efficiency is effectively 100% whenever the measured value exceeds this value, particularly due to the large associated uncertainties.

6.3.3.5 Voltage Efficiency

As discussed in Sec. 6.3.2, we determine the voltage utilization by using the RPA ion energy per unit charge that is corrected for the plasma potential measured by a Langmuir probe. Figure 6.10 shows that for xenon, the voltage utilization monotonically increases as current increases. This behavior is consistent with the trends exhibited in Fig. 6.7 where we see the average ion energy increase with discharge current. For krypton, all points are within uncertainty of each other (a minimum of $91.1 \pm 4.0\%$ and a maximum of $96.7 \pm 1.7\%$), indicating no evident trend in this efficiency mode with increasing current.

6.3.3.6 Divergence Efficiency

As seen in Fig. 6.9a, the plume of the thruster operating on xenon continuously expands as the current increases, manifesting in the plume divergence efficiency monotonically decreasing (Fig. 6.10a). The behavior of krypton's divergence angle, which remains approximately constant throughout the range of currents, is also reflected in the nearly-constant divergence efficiency of krypton (Fig. 6.10b). The plume divergence represents the second-largest detractor to the overall efficiency behind beam utilization for both propellants at all conditions excepting 15 A, where the mass utilization is sufficiently low as to be in the same range as the beam and divergence efficiencies. We discuss the possible physical causes for this low divergence efficiency in Sec. 6.4.2.

6.3.3.7 Summary of Efficiency Trends

In this section, we have presented the partial efficiencies for operation on both xenon and krypton and discussed how they trend with increasing current. At higher-than-nominal currents (>15 A), the lowest efficiency mode for all conditions is the beam utilization efficiency, which generally decreases with increasing current and ranges between 66–71% excluding the 15 A condition for both propellants. The next lowest efficiency mode at high currents is the divergence efficiency, which ranges from 78–82% in 50–150 A range. The charge and voltage utilization were both above 90% at all conditions for both propellants. Finally, the

mass utilization efficiency varied the most, generally with a dramatic increase from low to high currents. In Sec. 6.4.2, we focus on the physical drivers behind the mass and beam utilization efficiencies and discuss how they shape the overall trends in thruster performance with increasing current.

6.3.4 Oscillation Level

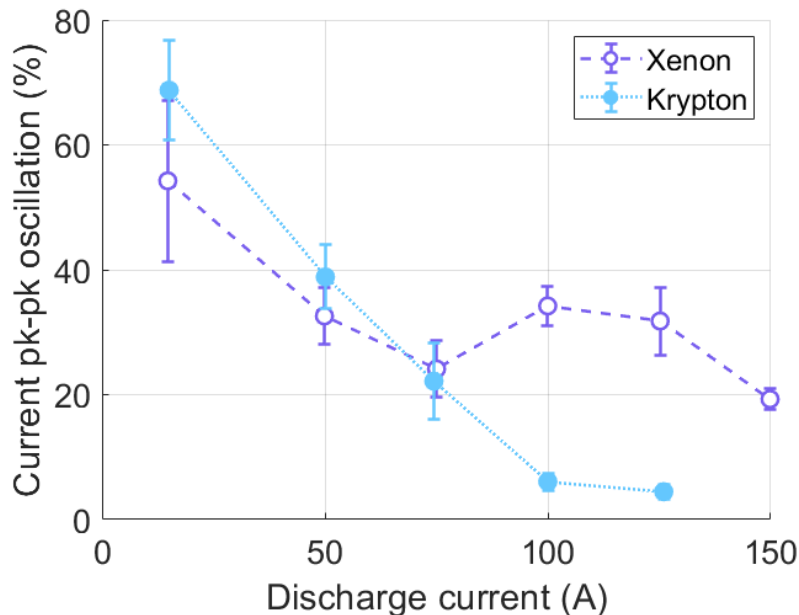


Figure 6.11: Ratio of peak-to-peak oscillation to mean value of discharge current at 300 V for xenon and krypton with increasing discharge current.

We show in Fig. 6.11 the magnitude of the peak-to-peak oscillations in the discharge current relative to the mean value. The uncertainty of these oscillations was determined by taking twice the standard deviation in the peak-to-peak oscillation strength over ten seconds. For xenon, we see that the current oscillations monotonically decreased with discharge current, starting at 69% at 15 A and ending at 4% at 125 A. This latter value is exceptionally low for Hall thrusters [148]. For krypton, the current oscillations generally decreased with the exception of an increase from 75 to 100 A. The oscillation levels for this propellant ranged from a high of 54% at 15 A to low of 19% at 150 A. From a practical perspective, these measured oscillation levels show that the thruster did not exhibit unusually high oscillation modes with increasing current density. This suggests that, at least in this configuration, a high current density Hall thruster not only maintains comparatively high performance but is also stable. We return to this discussion in Sec. 6.4.4.

6.3.5 Thermal

In Fig. 6.12, we show the temperature of the inner front pole (IFP) and corresponding discharge current throughout a testing day for each propellant. The times at which the current is zero are thrust points or shutdowns. The rate of temperature increase steepens at higher currents as is expected due to the increased power deposition to the walls and other surfaces of the thruster. During xenon operation, we were able to maintain IFP temperatures below an internally-defined limit. This limit is based on concerns for the health of the thruster magnetic circuit. For krypton at 150 A, the IFP temperature did briefly cross this limit, but this was deemed an acceptable risk to collect all necessary data at this condition and did not result in any permanent damage to the thruster. The typical ramp time to achieve each power level was 5–10 minutes, and the total time spent at each power condition varied from 5–15 minutes with a total of 10 minutes at the maximum power of 150 A and 300 V (45 kW) on krypton.




Figure 6.12: Discharge current and inner front pole temperature throughout a testing day for a) xenon and b) krypton at 300 V. The dashed pink line indicates the temperature limit for the IFP.

6.3.6 Facility

We present in this section two metrics for evaluating the response of the facility to the thruster at high current density operation. First, we show in Fig. 6.13a the carbon deposition rate as a function of discharge current as measured by the QCM. For comparison, we show as a horizontal line on this plot a value of $60 \mu\text{m}/\text{kh}$. This is a typical erosion rate reported for

the inner front pole cover on the HERMeS 12.5-kW class magnetically shielded Hall thruster operating at 300 V in a chamber with comparable background pressures [47]. As expected, the key insight from our measurements is that the deposition rate increases with total current from the thruster. Physically, this behavior is because the sputter rate is proportional to total incident flux on the facility surfaces. Xenon shows a marginally higher rate of deposition at each power level, which could be a result of the difference in the mass flux of the ions. The point at 150 A for krypton has been excluded due to possible sensor malfunction—the frequency of the crystal changed between testing the 125 and 150 A conditions such that it was no longer in the nominal operating window.




Figure 6.13: a) Carbon deposition rates and b) facility pressures at 300 V for xenon and krypton with increasing discharge current.

Practically, we see that at high currents, the deposition rates become on par with typical thruster erosion rates previously reported for magnetically shielded Hall thrusters. This highlights a potential problem with thruster testing at these unprecedented high power levels where the rate of deposition from the facility may mask erosion on the surfaces. We qualify this observation with the fact that erosion rates may scale with current density such that the relative ratio may remain unchanged, which in turn would also decrease lifetime. For example, the QCM used in the experiment with the HERMeS thruster measured only 2.2 $\mu\text{m}/\text{kh}$ of deposition while it was operating at ~ 21 A [47]. This relationship between erosion and deposition should be the subject of future study [75].

As a second notable trend, we show in Fig. 6.13b the facility pressure as a function of discharge current. We have included uncertainty bars of 20%, a typical value for the Stabil ion gauges we used [126]. We also show for comparison on this plot multiple standards of

maximum operating pressure. The 10 $\mu\text{Torr-Xe}$ value is the threshold of pressure below which linear extrapolation to zero-pressure conditions are valid as tested on the HERMeS thruster [132]. The Randolph standard for probe measurements is shown at 13 $\mu\text{Torr-Xe}$ and the Randolph standard for thrust measurements is shown at 50 $\mu\text{Torr-Xe}$ [161]. It is immediately evident from this result that testing powers above 15 kW (50 A) are associated with pressures in the facility that exceed the zero-extrapolation and Randolph probe standards. These high pressures may be partially responsible for the discrepancy between our anode efficiency as measured from thrust and as inferred from probe measurements (Sec. 6.3.3). We may also expect that as we reach even higher powers, we will eventually surpass the Randolph standard for thrust measurements as well. This again underscores the challenges of extrapolating ground test results of high power systems to be representative of flight [75].

6.4 Discussion

We discuss in this section key implications of our results. We first provide explanation for the monotonically increasing specific impulse and its relation to multiply charged species. Next, we explore the physical drivers behind trends in efficiency, particularly for the mass and beam utilizations and briefly for divergence. We then discuss the low oscillations in discharge current and the challenges of facility effects that have been highlighted by this study. We conclude with a discussion regarding how our results impact our fundamental understanding of the limits of Hall thrusters operating at high powers and their comparative advantages to other technologies.

6.4.1 Trends in Specific Impulse

As shown in Fig. 6.5b, one of the notable findings from of our study is the 30–40% increase in specific impulse with current exhibited by each propellant. To explain this trend, we consider a theoretically informed expression for the specific impulse:

$$I_{sp,th} = \frac{1}{g_0} \eta_c \eta_m \sqrt{\eta_d \eta_v} \frac{\sum_n \frac{\Omega_n}{\sqrt{Z_n}}}{\sum_n \frac{\Omega_n}{Z_n}} \sqrt{\frac{2eV_d}{m_i}}, \quad (6.1)$$

where we have assumed that the thrust is $T = \dot{m}_b v_i \cos \theta_d$, where v_i is ion velocity, to account for the divergence angle. We note that for the case of an entirely singly-charged beam with no efficiency losses, this equation simplifies to Eq. 5.1. As this result shows, the specific impulse depends on some of the different efficiency modes in the plasma. In Fig. 6.14, we show Eq. 6.1 plotted as a function of current where we have employed our measured

efficiencies from Sec. 6.3.3. We consider the cases of a singly charged beam (dotted colored line), a multiply charged beam (dashed colored line), and the limiting cases where the beam is entirely singly and doubly charged with 100% efficiency (horizontal black lines).

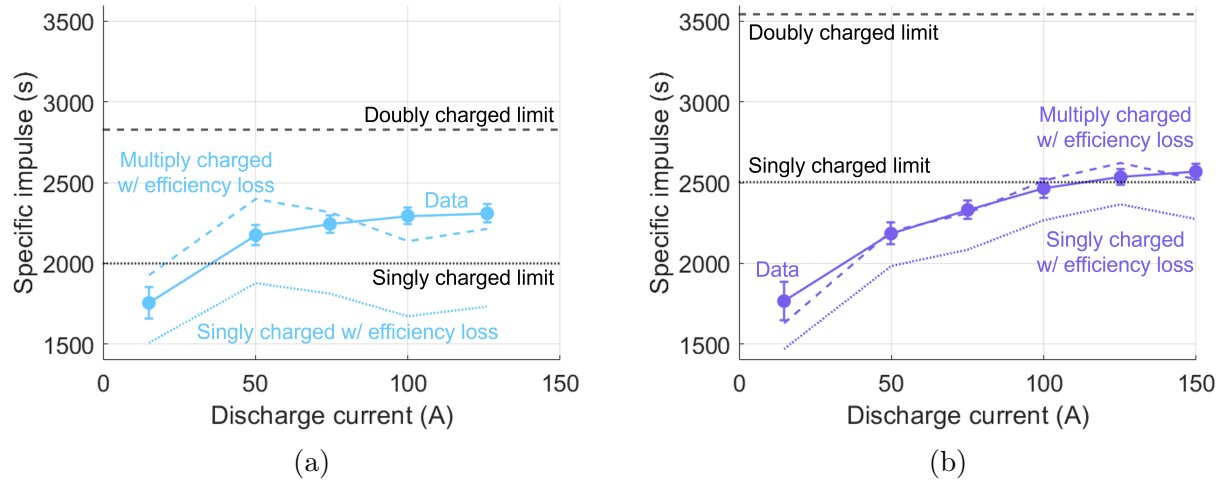


Figure 6.14: Experimentally measured values of specific impulse for a) xenon and b) krypton operating at 300 V compared to theoretical values in which the plasma is assumed to be singly charged (with and without efficiency losses), doubly charged (without efficiency losses), and polydisperse (with efficiency losses).

As this result shows, the trend of increasing specific impulse with discharge current is largely explained—particularly in the case of krypton—by the increase in mass utilization. This is caused by enhanced ionization at higher currents leading to more of the inlet gas being converted to plasma, which in turn can be accelerated to generate more thrust. We note, however, that the magnitude of the specific impulse cannot be explained purely by the acceleration of singly charged ions. In fact, the maximum theoretical specific impulse for singly charged species at 300 V is exceeded at the higher current cases for both gases. This result suggests that the polydisperse nature of the plasma is responsible for the higher specific impulse values. Indeed, accounting for the large presence of multiple charge states brings the magnitude of the theoretical value (Eq. 6.1) in line with measurement.

The physical reason behind this polydisperse state comes from the effective ionization of all neutrals (i.e. $\eta_m = 1$) at high currents. As the current and therefore plasma density in the channel continue to increase, the resulting singly charged ions can have additional ionization events, thereby gaining additional charge. A notable implication of this result is that in these high power density operating regimes where the plasma is dominated by higher charge states, the effective specific impulse can receive a boost that brings it to values greater than the typical level exhibited by lower current density thrusters [76, 122, 162] without needing

to increase the discharge voltage. This has potential benefits particularly for deep space missions where high specific impulse operation is desirable [1, 18]. However, multiply charged ions also erode thruster surfaces more quickly, posing an additional potential challenge to thruster lifetimes at high current densities.

6.4.2 Trends in Efficiency

We comment here on physical explanations for the key notable trends in our efficiency results. We specifically focus on the trends in mass utilization with current and gas species, the trends in beam utilization with current, the variations in divergence efficiency, and the existence of a maximum in overall efficiency.

6.4.2.1 Trends in Mass Efficiency

In this section, we refer back to our 0D model for mass utilization derived in Sec. 3.3.3 and recast it in terms that we can evaluate for our performance measurements. First, we take Eq. 3.15 and make the additional assumption that the neutral velocity is equal to its initial value into the channel,

$$\eta_m = 1 - \frac{\dot{m}(L)}{\dot{m}(0)} \approx 1 - \frac{n_n(L)}{n_n(0)} \approx 1 - \exp \left[-\frac{\langle k_{iz} \rangle \langle n_e \rangle L}{v_n(0)} \right]. \quad (6.2)$$

As previously discussed, this result shows that as the dwell time of neutrals in the channel (L/v_n) or the probability of ionization (dictated by $\langle k_{iz} \rangle \langle n_e \rangle$) increases, the mass utilization will improve.

To relate Eq. 6.2 to discharge current, we invoke quasineutrality between species to find the relationship

$$n_e = n_i = \frac{I_b}{q \langle v_i \rangle A} = \frac{\eta_b I_d}{A} \sqrt{\frac{m_i}{e^3 V_a}} \sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}, \quad (6.3)$$

where we have included the superposition of differently charged beams. We also have introduced a factor of $\sqrt{1/2}$ in relating the ion velocity to discharge voltage for each charge species to account for the fact that the density is averaged over the channel, i.e. $v_i = \sqrt{eV_a/m_i}$ instead of $v_i = \sqrt{2eV_a/m_i}$. Leveraging this expression for density, we can simplify Eq. 6.2 to

$$\eta_m \approx 1 - \exp \left[-\frac{\bar{I}_d}{I_c} \right], \quad (6.4)$$

where we have introduced the variables

$$\bar{I}_d = \frac{m_i}{m_{i(Xe)}} \frac{\langle k_{iz} \rangle}{\langle k_{iz(Xe)} \rangle} I_d \quad (6.5)$$

$$I_c = \frac{1}{m_{i(Xe)} \langle k_{iz(Xe)} \rangle} \frac{A}{\eta_b L} \sqrt{\frac{2e^3 V_a k_B T_g}{\pi}} \frac{1}{\sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}}. \quad (6.6)$$

Here, \bar{I}_d is an adjusted current that depends on the mass and ionization rate coefficient for each species. This parameter, according to this simplified theory, should be universal in dictating the mass utilization for different propellants. The parameter I_c is a ‘‘characteristic current’’ where T_g is the temperature of the neutral gas at the inlet, $m_{i(Xe)}$ is the xenon mass, and $\langle k_{iz(Xe)} \rangle$ is a constant given by the ionization rate coefficient for xenon at a reference temperature of $T_e = 30$ eV. This characteristic current depends on channel geometry, charge composition, discharge voltage, and temperature of the neutral gas. Some of these parameters vary with discharge current. However, our experimental results have shown that the relative change in discharge current over the current range is ten times greater than the variations in beam utilization efficiency, summation over charge states, and the acceleration voltage. We therefore approximate in Eq. 6.4 the characteristic current I_c as weakly varying and thus approximately constant with current.

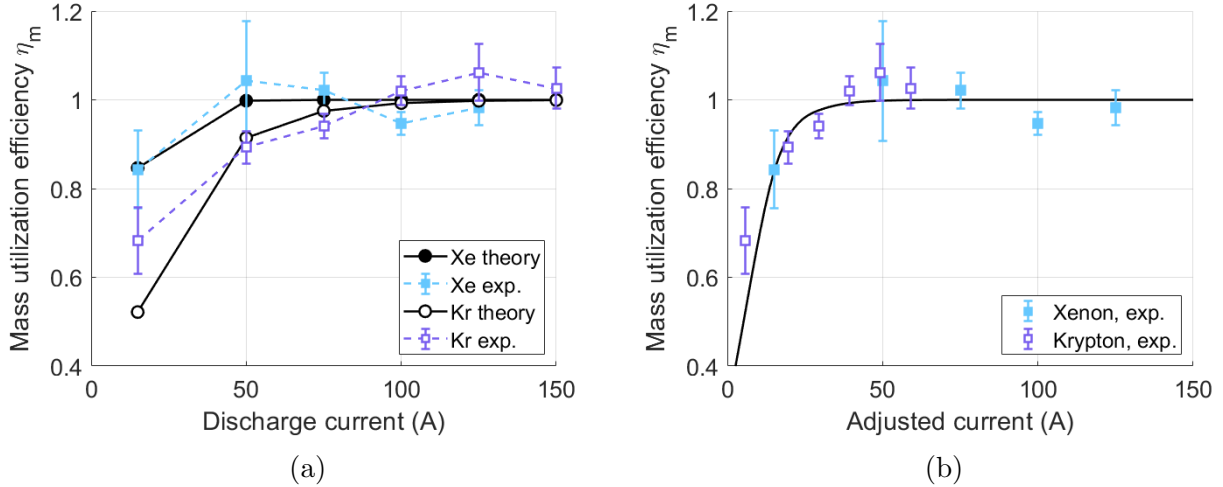


Figure 6.15: Experimentally measured values (color) of beam current utilization efficiency for xenon and krypton operating at 300 V compared to theoretical form (black) with a) actual discharge current and b) adjusted current \bar{I}_d with a characteristic current of $I_n = 8$.

We plot the experimental measurements of mass utilization as a function of the adjusted current \bar{I}_d in Fig. 6.15b, with mass utilization as a function of unadjusted current I_d for

comparison in Fig. 6.15a. In order to evaluate this current for each gas, we have made the assumption that the electron temperature in eV scales with the discharge voltage by a factor of 0.1 (Refs. [18, 163, 164, 165, 166]) such that $T_e = 30$ eV for both gases. For comparison, we also plot Eq. 6.4 where we have empirically calibrated $I_c = 8$ A to achieve a best fit of the data. This result illustrates that the xenon and krypton data collapse onto the same curve when we use \bar{I}_d , and the agreement between the shape of the curve and data lends support to this simplified but physics-motivated description of mass utilization efficiency.

As additional validation, to evaluate if the fit value for our characteristic current is plausible, we use our experimental measurements of the terms in Eq. 6.6 to calculate a theoretical value of I_c . We obtained these measured values, shown in Table 6.2, by averaging the data shown in Figs. 6.6, 6.8, and 6.10 over all discharge currents for both propellants. We also use a neutral temperature of $T_g = 400$ K. With these averaged values, we find a value of $I_c = 2.1$ A. This is within a factor of 4 of the best fit value. The discrepancy between the calculated and best fit parameters can likely be attributed to the several simplifications made in our analysis. For example, the $T_e \approx 0.1V_d$ relation is typically used to approximate the peak electron temperature in the channel. Since we are averaging across the channel, we may therefore expect the effective electron temperature (and therefore ionization rate coefficient) to be lower than 30 eV. This would increase our calculated value of characteristic current I_c to be more in line with our best fit value of 8 A. Our later work does indeed show the axially averaged electron temperature to be about half its peak value (Fig. B.4, 7.14)).

Parameter	Value
$\eta_{b,const}$	0.7
Ω_1	0.6
Ω_2	0.35
Ω_3	0.05
V_a	290 V
$\langle T_e \rangle$	30 eV
$\langle k_{iz(Xe)} \rangle$	1.079×10^{-13} m ³ /s
$\langle k_{iz(Kr)} \rangle$	6.651×10^{-14} m ³ /s

Table 6.2: Plasma parameter values averaged over all operating conditions and both propellants.

There are two key physical insights that emerge from our simplified but validated description for the mass utilization. The first is that the increase in mass utilization with current can be attributed to the fact that the number density of particles in the channel also increases with current. This higher density of particles facilitates more ionization of the inlet neutral gas. With sufficiently high discharge current for a given gas, all particles are ionized,

and the mass utilization approaches unity. This is reflected by the exponential dependence of Eq. 6.4 on current.

The second insight is that the necessary current to achieve a fixed level of mass utilization is species dependent. For example, we saw experimentally that the mass utilization of xenon reaches 95% at ~ 30 A (Fig. 6.10). For krypton, this occurs at ~ 75 A, yielding a ratio of $\frac{I_{d(Kr)}}{I_{d(Xe)}} \approx 0.4$. This is consistent with the theoretical calculation of current ratios for the 300 V operating voltage:

$$\frac{I_{d(Kr)}}{I_{d(Xe)}} = \frac{m_{i(Kr)} \langle k_{iz(Kr)} \rangle}{m_{i(Xe)} \langle k_{iz(Xe)} \rangle} \approx 0.39. \quad (6.7)$$

The need for higher current to achieve full ionization with krypton is firstly because it is more difficult to ionize with its smaller cross section and secondly because of its shorter residence time in the channel due to its smaller mass. As an extension from this result, we see that the increase in current density appears to be able to compensate for the traditional mass utilization losses associated with gases with lower ionization cross sections. This invites the possibility that operating Hall thrusters at higher current densities may enable more efficient performance on alternative gases. This approach has been suggested in previous work [64, 87].

6.4.2.2 Trends in Beam Current Efficiency

As noted in the preceding Sec. 6.3.3, the beam utilization decreases with increasing current, a likely indication that the confinement of electrons in the channel is being curtailed. We discuss in this section possible physical explanations for this trend. To this end, we begin by rewriting the definition of beam utilization as

$$\eta_b = \frac{1}{1 + \frac{I_e}{I_b}}, \quad (6.8)$$

where I_e denotes the electron current in the channel across the applied magnetic field,

$$I_e = en_e v_e A, \quad (6.9)$$

where n_e is the electron density, v_e is the electron velocity, and A is channel area. For the electron current in this expression, we invoke a generalized Ohm's law averaged over the

channel area (App. A.3) to find

$$I_e = e \langle n_e \rangle \frac{\langle \nu_e \rangle}{\langle \omega_{ce} \rangle} \frac{V_a}{\langle B \rangle L} A, \quad (6.10)$$

where ν_e is the electron collision frequency, ω_{ce} is the electron cyclotron frequency, B is the magnetic field strength, and $\langle x \rangle$ again denotes a plasma property averaged over the channel length. In arriving at this expression, we have neglected electron pressure and approximated the electric field in the axial direction in the channel as $E \approx \frac{V_a}{L}$. We also have assumed a large Hall parameter $\Omega_e = \frac{\omega_{ce}}{\nu_e} \gg 1$, which is typical for Hall thrusters.

For the ion current, we can adapt Eq. 6.3 to find

$$I_b = \langle n_i \rangle A \sqrt{\frac{e^3 V_a}{m_i}} \frac{1}{\sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}}. \quad (6.11)$$

With these expressions for ion and electron current, we can write the ratio of currents as

$$\begin{aligned} \frac{I_e}{I_b} &= \frac{\sqrt{m_i V_a}}{\sqrt{e \langle B \rangle}} \frac{1}{L} \frac{\langle \nu_e \rangle}{\langle \omega_{ce} \rangle} \sum_n \frac{\Omega_n}{\sqrt{Z_n^3}} \\ &= \frac{r_{Li}}{L} \frac{\langle \nu_e \rangle}{\langle \omega_{ce} \rangle} \sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}, \end{aligned} \quad (6.12)$$

where we have invoked the definition for average ion Larmor radius in the channel as the ratio of ion velocity to ion cyclotron frequency, $r_{Li} = \frac{\langle v_i \rangle}{\langle \omega_{ci} \rangle}$, again incorporating a factor of $\sqrt{1/2}$ in the ion velocity, and assumed quasineutrality. Physically, this expression shows the intuitive result that as the electrons become progressively demagnetized, i.e. $\frac{\langle \nu_e \rangle}{\langle \omega_{ce} \rangle}$ increases, the cross-field confinement is reduced. In turn, per Eq. 6.8, the beam utilization efficiency will decline.

As a next step, we express the ratio of currents explicitly as a function of variables that are dependent on propellant choice and discharge current, and substitute into the beam utilization expression (Eq. 6.8):

$$\eta_b = \frac{1}{1 + a \sqrt{m_i} \langle \nu_e \rangle}, \quad (6.13)$$

where a can be explicitly written as

$$a = \sqrt{\frac{V_a}{e^3}} \frac{m_e}{\langle B \rangle^2 L} \sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}. \quad (6.14)$$

As with our discussion in the preceding section, we again assume V_a and Ω_i to be weakly varying with current such that this parameter, a , is constant. We find a to be $5.8 \times 10^4 \sqrt{\frac{\text{m}}{\text{N}}}$ for our operating conditions (see Table 6.2). The key question for relating beam utilization to discharge current is therefore how the average electron collision frequency $\langle \nu_e \rangle$ depends on plasma properties. This is an active area of research within the Hall thruster community, as the measured electron mobility in Hall thrusters is much higher than can be predicted via classical collisions [24]. Since the actual mechanisms driving this transport and the resulting form for collision frequency remain open questions, we consider in this discussion two models which effectively bound the limits of collision frequency: classical electron-ion collisions and Bohm scaling.

The classical electron-ion collision frequency can be written as

$$\nu_{ei,c} = 2.91 \times 10^{-12} n_e \ln \Lambda T_e^{-3/2}, \quad (6.15)$$

where all units are SI except for the electron temperature, which is expressed in eV. $\ln \Lambda$ is the Coulomb logarithm, which is for the most part invariant and on the order of 10 for our plasma conditions. We combine this expression with Eq. 6.3, approximate $V_a \approx V_d$, and again use the $T_e \propto 0.1V_d$ relation to find that this expression scales as

$$\nu_{ei,c} = c_c \sqrt{m_i} \eta_b I_d, \quad (6.16)$$

where we have introduced the parameter

$$c_c = \alpha_c \frac{9.2 \times 10^{-10}}{AV_a^2 \sqrt{e^3}} \sum_n \frac{\Omega_n}{\sqrt{Z_n^3}}. \quad (6.17)$$

For the values assumed in our work, we find that $c_c = \alpha_c 8.3 \times 10^{15} \frac{1}{\text{C}\sqrt{\text{kg}}}$. In this expression, we have used the same approximations to relate current density to discharge current and electron temperature to voltage as in the previous discussion, and we have introduced a parameter α_c . This latter constant accounts for the fact that the actual collision frequency in the channel is orders of magnitude higher than the classical [167]. This parameter in turn allows us to capture this effect while still retaining the classical scaling of collision frequency with current. As with our discussion of mass utilization, we assume that the plasma parameters in this constant are weakly varying with discharge current. We therefore approximate c_c as a constant.

We can substitute this expression for the classical collision frequency into Eq. 6.13 to

express the beam utilization efficiency with classical scaling,

$$\eta_{b,c} = \frac{\sqrt{4ac_c m_i I_d + 1} - 1}{2ac_c m_i I_d}. \quad (6.18)$$

Physically, this result suggests that as the discharge current increases, the beam efficiency decreases. We plot this theoretical result against experimental data in Fig. 6.16 with a best fit value of $\alpha_c = 310$ for xenon and $\alpha_c = 560$ for krypton. As can be seen, this scaling qualitatively captures the decrease in confinement with discharge current—indicated by a decreasing beam utilization efficiency—exhibited by xenon. This is intuitive because the number of particles in the channel increases with current, raising the collision frequency and reducing confinement. However, we note that the scaling is too strong, predicting more of a reduction in confinement than exhibited by the plasma. Notably, the beam utilization for xenon is only gradually decreasing with current and thus not captured by solely considering classical scaling. This disagreement is even more apparent for krypton, where the beam utilization efficiency changes only by 2% throughout the range as opposed to the >30% predicted by classical scaling.

Another common method of modeling electron transport in Hall thruster simulations is using a Bohm-like scaling [24]. For this case, the collision frequency scales with the electron cyclotron frequency, $\nu_{e,B} = \alpha_B \frac{\omega_{ce}}{16}$, where α_B is a constant smaller than 0.1. In this case, we find that the beam utilization has the form

$$\eta_{b,B} = \frac{1}{1 + ac_B \sqrt{m_i}}, \quad (6.19)$$

where the parameter $c_B = \alpha_B \frac{\omega_{ce}}{16}$ is simply the Bohm frequency. Unlike in the case of classical scaling, the beam utilization is invariant to discharge current with a Bohm-like scaling. In Fig. 6.16, we show this form compared to the experimental measurement where we can see that we best match the magnitude of data with $\alpha_B = 0.14$ for xenon and $\alpha_B = 0.19$ for krypton, values that are in line with previous studies [24, 89]. This yields a value for c_B on the order of 10^7 . Comparing the theory to the experimental trends, however, we see that the Bohm scaling for xenon does not capture the behavior exhibited by xenon as the experimental measurements steadily decline. The beam utilization of krypton, on the other hand, has a more Bohm-like trend with constant values within uncertainty throughout the range of currents. This indicates that the transport for krypton may in fact be more Bohm-like in that it is not strongly dependent on changes in discharge current.

In general, we see that the transport for both gases is non-classical and correspondingly exhibits a weaker dependence on current. This diverges from what has previously been

proposed for conventional scaling laws in Hall thrusters [39] and in part explains why the performance did not drop as precipitously as conventional scaling laws may suggest. Indeed, based on our results, we may expect that operating at higher currents beyond what was explored during this effort would only yield moderate decreases in performance.




Figure 6.16: Experimentally measured values (color) of beam current utilization efficiency for a) xenon and b) krypton operating at 300 V compared to theoretical forms (black) assuming classical scaling with $\alpha_{c,Xe} = 310$, $\alpha_{c,Kr} = 560$ and Bohm scaling with $\alpha_{B,Xe} = 0.14$, $\alpha_{B,Kr} = 0.19$.

With that said, the transport is also not entirely Bohm-like—particularly for the case of xenon—as the beam utilization does exhibit some dependence on current density. This suggests that the upper bound in performance for these devices may be lower than the limits suggested in recent scaling laws studies that explored a Bohm-like dependence [89]. Regardless, from the perspective of informing future efforts, the arguments presented in this section suggest that future efforts should be guided by the assumption that the performance decrease with current density can be bounded by assuming the transport falls between these two limits.

6.4.2.3 Trends in Divergence Efficiency

As seen in Fig. 6.10, the divergence efficiency is generally the second-lowest efficiency mode besides beam utilization, though the values are consistent with other state-of-the-art magnetically shielded Hall thrusters [15, 54, 71, 132]. We note here that magnetically shielded thrusters are known to have slightly higher plume divergence angles and therefore lower divergence efficiencies compared to unshielded thrusters. This is because the location of

peak magnetic field has been pushed downstream in the shielded topology [15], therefore increasing the divergence.

Xenon and krypton exhibit different trends with increasing current—the xenon divergence angle increases monotonically, while the krypton divergence angle stays within of 2° throughout the current range. Correspondingly, we see a decreasing divergence efficiency for xenon and an approximately constant divergence efficiency for krypton. This trend for xenon might be attributed to the acceleration region shifting downstream, thereby increasing plume divergence. This type of correlation between the location of the acceleration zone and divergence angle has previously been noted in pressure studies of Hall thrusters [76, 151, 168, 169]. However, absent measurements of the velocity profile within the channel, we are unable to conclusively attribute the changes in divergence angle to this effect. Additionally, the physical reason for the different trends between the two propellants is unclear. We do note however that the trends in divergence efficiency closely mirror those in beam efficiency for each propellant (i.e. decreasing for xenon, approximately constant for krypton), suggesting that the two efficiency modes may be linked.

6.4.2.4 Trends in Overall Anode Efficiency

Having discussed the independent trends in mass utilization and beam utilization for different anomalous scaling methods in the previous sections, we now focus on the behavior of the anode efficiency which scales with the mass and beam utilizations ($\eta_a \propto \eta_b \eta_m$). In particular, we discuss the existence of a pronounced optimum in the xenon efficiency and a plateau with a slight optimum in the krypton efficiency. We can explain these as the result of a trade between mass utilization and beam utilization, noting that we neglect the effect of the divergence efficiency in this discussion. As noted in Sec. 6.3.3, the krypton efficiency peaks at a higher current than xenon does. This is primarily explained by the fact that the mass utilization asymptotes at a higher current for krypton than for xenon. When convolved with the decreasing beam utilization efficiency, the higher discharge current required for krypton to reach 100% mass utilization (as discussed in Sec. 6.4.2.1) has the effect of pushing the maximum value for overall efficiency of krypton to occur at a higher discharge current than xenon.

One other notable effect of this trade is that once the mass utilization has reached 100% efficiency, the beam utilization dominates the overall efficiency curve, leading to an eventual decrease in performance. Krypton’s η_b declines at a more gradual rate than xenon’s, which is responsible for krypton slightly outperforming xenon at sufficiently high currents. Additionally, while there is a pronounced maximum followed by a decrease in xenon’s efficiency, krypton’s anode efficiency reaches more of a plateau with the efficiencies between 50–150 A

varying by no more than 3%. This may be due to the weaker dependence of krypton’s beam utilization on changing discharge current as apparent from our measurements.




Figure 6.17: Experimentally measured values of anode efficiency for xenon and krypton operating at 300 V compared to theoretical form with $\eta_{const} = 0.85$ assuming weights of $\beta_c = 0.2, \beta_B = 0.8$ for xenon and $\beta_c = 0.1, \beta_B = 0.9$ for krypton.

To make this trend more clear, we assume that the beam utilization for each propellant scales as a weighted average of classical and Bohm scaling, i.e. $\eta_{b,avg} = \beta_c \eta_{b,c} + \beta_B \eta_{b,B}$, where $\beta_c + \beta_B = 1$. Next we assume that the anode efficiency scales primarily with the current and mass utilizations, i.e. with all other efficiencies constant with propellant and current, $\eta_a \propto \eta_m \eta_b \eta_{const}$. This allows us to compare the experimentally-measured anode efficiency to the theoretical anode efficiency as the product of the average beam utilization, theoretical mass utilization, and efficiency scalar. We empirically find a value of $\eta_{const} = 0.85$ with weights of $\beta_c = 0.2, \beta_B = 0.8$ for xenon and $\beta_c = 0.1, \beta_B = 0.9$ for krypton. In Fig. 6.17 we show this comparison. We see that with even with this simple expression, our theory is able to capture key components of the experimental trends; namely, that xenon efficiency peaks at 50 A while krypton peaks at 100 A, and that krypton’s efficiency overtakes that of xenon’s at sufficiently high currents.

For a final remark, we note that while previous work has suggested that losses in beam utilization may prohibit operation at higher current densities [39, 18], our results indicate this is not the case. Indeed, although beam utilization does decrease with current density, the performance remained competitive with typical metrics of efficiency for Hall thrusters

reported at lower power. This finding is backed both by our experimental measurements and a theoretical understanding of the anomalous transport in Hall thrusters as scaling somewhere between how it would classically and how it would with Bohm diffusion. An additional implication is that efficiencies for both propellants would continue to decrease at higher discharge currents as the beam utilization now dominates the trends. However, it is possible that with even lighter gases such as Ar or N₂, the peak efficiency would be at an even higher current and may have values that overtake krypton and xenon’s at sufficiently high powers. Additionally, increasing the magnitude of the magnetic field would improve the beam utilization by maintaining electron confinement at higher current densities.

6.4.3 Knudsen Number Analysis




Figure 6.18: Knudsen numbers for xenon and krypton with increasing discharge current.

The Knudsen number is the ratio between the molecular mean free path and characteristic length scale of a system, $Kn = \lambda/L$. It serves a useful way to characterize flow regimes, with high values (>10) corresponding to free molecular flow (i.e. entirely collisionless) and low values (<0.01) corresponding to collisional flow. The plasma within the discharge of a Hall thruster is usually considered to be non-collisional. However, this characterization is based on the typical current densities of Hall thrusters and may no longer hold at the high current density regime we operate at in this work. To elucidate this relationship, we calculate the Knudsen number using the ionization mean free path as defined in Eq. 3.16 and treat L as

the channel length. We use the values in Tab. 6.2 for the ionization rate coefficients and for calculating the plasma density n_e from Eq. 6.3. Finally, because the scaling between the centerline values we use and the radially-averaged values is unknown, we selected a range of values for α in Eq. 3.16 that spanned from 0.1 to 0.4. (For comparison, we use $\alpha = 0.25$ in Sec. 7.4.)

In Fig. 6.18, we show how the Knudsen number trends with increasing current. As we would expect, this number decreases with increasing current due to the higher plasma densities. We note that although the Knudsen number does change from the transitional to the more collisional slip regime, it never reaches below 0.01, the threshold for fully collisional flow. This means that the mass utilization reaching 100% is not a phenomenon that requires a fully collisional flow, as xenon reaches this point at 50 A and krypton does so at 100 A; both of these points fall within the slip/transitional regime.

6.4.4 Low Oscillation Amplitude

In Fig. 6.11, we see that the peak-to-peak oscillation strength of the discharge current generally decreases as power increases. This is especially true for xenon, where at currents over 100 A, we see oscillations that are $< 6\%$ of the total current. While krypton does not follow an identical trend to xenon, the oscillations still do generally decrease with increasing current. Compared to the $> 150\%$ oscillations seen on krypton operating at high voltage conditions [71], all conditions had fairly mild ($< 50\%$) oscillation strengths. As a possible explanation for this behavior, we note that recent work has suggested that for fixed discharge voltage, thruster stability can improve with higher discharge current. This stems from the fact that the higher local current densities and corresponding high ionization rates help damp the dominant breathing mode in the thruster [148]. Furthermore, this same work showed that decreasing the residence time of neutrals in the anode can also improve stability margin. Given the high thermal fluxes to the thruster walls and anode and the resulting high neutral temperatures (Fig. 6.12), this may also be a contributing factor to the relatively low oscillation levels.

6.4.5 Facility Effects at High Power Density

In addition to examining the operational limits of a Hall thruster at atypically high current densities, our results also help provide exploratory insight into the challenges of ground testing high-power Hall thrusters and other EP devices. While the impact of carbon backsputter and facility pressure were already anticipated before this study, we have been able to provide a new dataset of direct measurements (Fig. 6.13) for LVTF, a highly capable test facility.

There ultimately remain outstanding questions as to how these elevated facility pressures and rates of sputtered material impact performance, and as a direct result, we recognize that the validity of the measurements we report here may ultimately need to be qualified.

Ongoing efforts under JANUS must use datasets like the one we have collected throughout this test campaign to inform strategies for deconvolving or mitigating facility effects in order to improve confidence in the transition of ground measurements like ours to flight [75]. We also remark here that our results have highlighted some of the more nuanced challenges with high power operation. For example, we have found that standard designs for probes such as RPAs begin to break down under these new high current operating limits as we saw the grids shorting due to the high current densities within the probe. This invites the question of how to adapt these tools to be better suited for this new environment.

As a final comment, we remark that while alternative concepts such as liquid metal applied-field MPD thrusters do not in principle have the same challenges with facility pressure as gas-fed Hall thrusters, the problem of diagnostic development and backscatter largely remain the same. Additionally, as we show in the next section, our thrust densities are comparable to those expected from applied-field MPDs. We therefore anticipate that the facility effects observed during this test campaign may also pose a testing challenge for other high-power electric propulsion devices.

6.4.6 Implications for Hall Thrusters on High-Power EP Missions

We conclude our discussion with a comment on the implications of our results for Hall thrusters as a candidate technology for high-power applications. As discussed in the introduction, the traditionally perceived limit to thrust density of Hall thrusters has been a major limitation for scaling to higher power when contrasted with fully electromagnetic concepts like MPD thrusters. Our results have demonstrated that these traditional scaling laws may ultimately not be absolute. Indeed, we have shown that the thruster efficiency does not precipitously drop with an approximately 10 times increase in current density and, correspondingly, 6–9 times increase in thrust density compared to the nominal operating condition of our Hall thruster.

To put this result in the context of alternative technologies for high-power electric propulsion, we show in Fig. 6.19 the normalized thrust density of our device for xenon and krypton compared to thrust densities compiled from Princeton University’s Electric Propulsion and Plasmadynamics Laboratory database of applied-field MPD thrusters [170]. In this case, we determined thrust density using the solenoid radius for AF-MPDs and the outer discharge channel radius for the H9 MUSCLE. All results are normalized to the thrust density of the

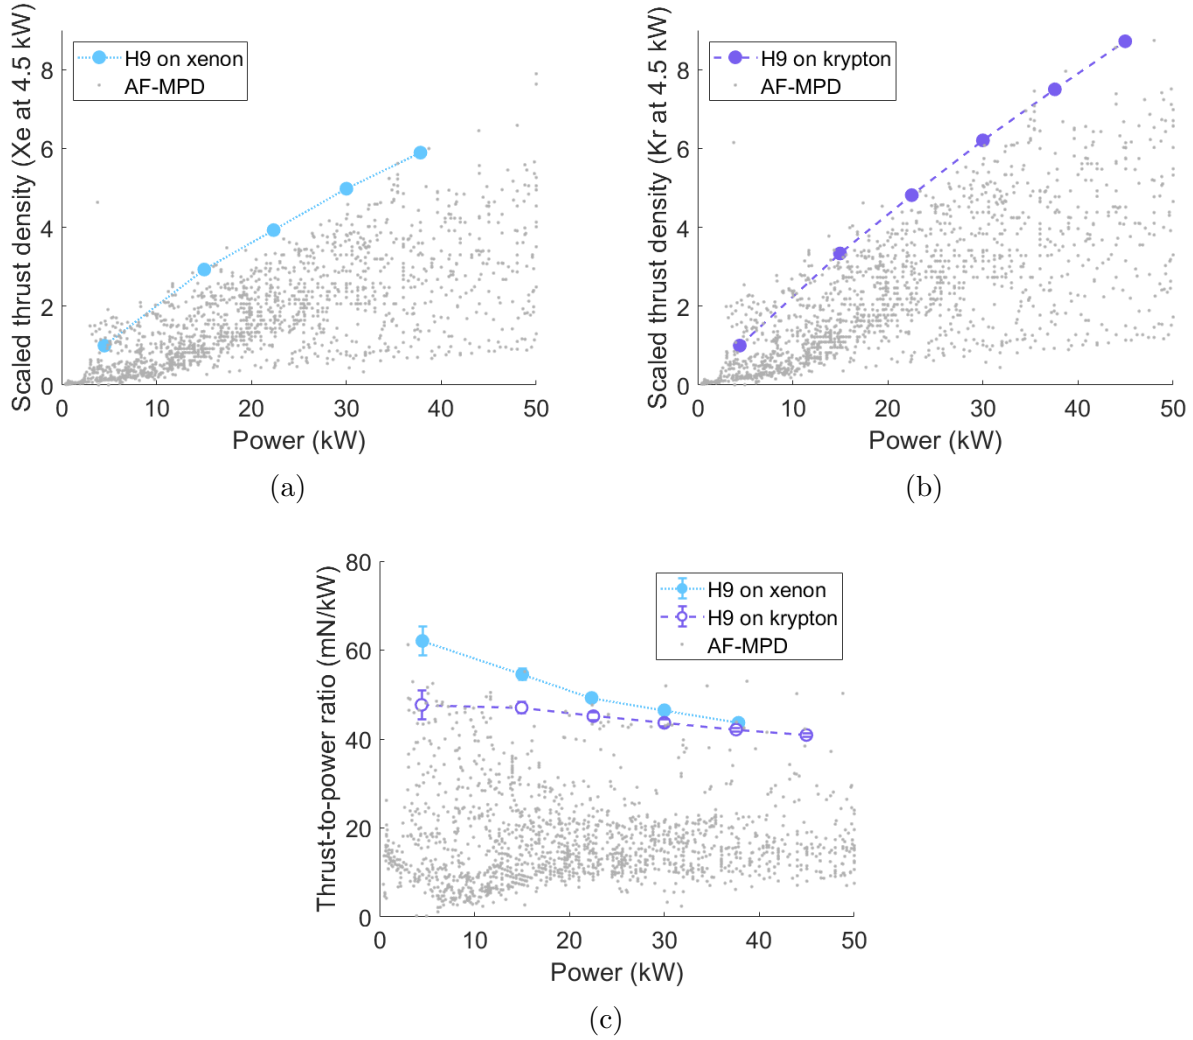


Figure 6.19: Comparison of H9 MUSCLE a) thrust densities scaled to xenon operation at 4.5 kW, b) thrust densities scaled to krypton operation at 4.5 kW, and c) thrust-to-power ratios to a database of AF-MPD measurements [170].

H9 MUSCLE operating on xenon and krypton at 300 V and 15 A. This result shows that the H9 MUSCLE operating at these high current densities is not only competitive with but can actually exceed AF-MPD thrust density in this power range. Moreover, as shown in Fig. 6.19c, the thrust-to-power ratios of the H9 MUSCLE exceed most of the values extracted from the the AF-MPD database. This difference is primarily attributed to the fact that MPDs typically operate on lower-mass propellants than xenon and krypton, therefore leading to higher specific impulses and lower thrust-to-power ratios. We do note here that we neglect self-field MPDs in this analysis, which typically are able to reach higher thrust densities.

The major implication of this comparison is that over the 15–45 kW power regime, Hall thrusters can in principle offer competitive thrust density scaling compared to other high-power concepts. We caveat this result, however, with the fact that it is unclear how well this scaling holds as the discharge power continues to increase. We may ultimately find, as the trends in our efficiency results suggest (especially the decrease in beam utilization), that performance eventually decreases to undesirable limits. However, there may be some margin before this limit is reached, and design decisions such as increased magnetic field strength and alternative propellants may help offset these potential losses.

In the context of recent mission studies, the demonstrated performance (Fig. 6.5) of the H9 MUSCLE on xenon from 50 to 125 A and krypton from 75 to 150 A both satisfy the >50% total efficiency, >2000 s specific impulse requirements outlined for near-term Mars missions [1]. While the demonstrated thruster power (45 kW) is under the 2 MW required for these missions, it may be possible to achieve these levels with a single larger unit or by arraying multiple thrusters of the same size. Indeed, a recent study of nuclear architectures suggested that the ideal power level for a 2 MW Hall thruster system may be \sim 100-200 kW per thruster [79]. In terms of more far term applications, we note that the specific mass of the thruster ranges from 0.56 to 1.67 kg/kW (compared to \sim 2.4 kg/kW typically exhibited by Hall thrusters [95, 171]). This metric is approaching the 0.5 kg/kW threshold identified for rapid Mars transit [78], though we note that the mass of the power supply itself remains the dominant challenge in these types of rapid mission architectures.

As a final comment, we emphasize here that this work was fundamental in nature, performed on a laboratory thruster with a controlled, actively cooled system. Our goal was to test the validity of the traditional arguments about electron confinement at higher current. In practice, transitioning this physics-based experiment to flight has many major technical engineering challenges related to thermal limits as well as magnetic field design.

6.5 Summary

The goal of this effort has been to investigate the fundamental limits of current density scaling in a Hall thruster. This work has been motivated by the increasing interest in developing high-power technologies that may enable nuclear architectures for missions like crewed explorations of Mars. In this context, while Hall thrusters are considered a leading candidate technology, one of their primary historic limitations for MW-class applications has been their achievable thrust density.

Motivated by recent experimental findings and theoretical studies, we pushed the limits of the traditional scaling laws for maximum current density in a Hall thruster. We made

modifications to a 9-kW class magnetically shielded laboratory thruster to support this test and systematically characterized this device’s performance on xenon and krypton from its nominal discharge current of 15 A to a maximum of 150 A. We employed global performance measurements to characterize thrust, specific impulse, and efficiency, as well as far-field probes to measure the phenomenological efficiency modes of the thruster. Key findings include that the specific impulse increases monotonically for both gases with discharge current; the efficiency decreases at the highest current densities but remains competitive with the performance metrics for conventional Hall thrusters operating at lower powers; and the overall efficiency exhibits an optimum with xenon at 50 A and krypton at 100 A.

We have discussed and explained key aspects of these trends with discharge current qualitatively in the context of simplified scaling laws for thruster operation. For example, we have shown that propellants with lower mass and ionization cross section, such as krypton in comparison to xenon, are still capable of reaching 100% mass utilization at sufficiently high currents. We have also examined the dependence of beam utilization on current, which historically was believed to be the major limitation for operation at high current density. Our findings indicate that the dependence of beam utilization does decrease with discharge current but not so precipitously as to preclude competitive efficiencies. Indeed, our results suggest that the effective electron collision frequency in the channel falls within the limits of a classical and Bohm-like scaling. We in turn have shown that the increasing mass utilization with discharge current combined with the gradual but persistent decrease in beam utilization is what results in a local peak in efficiency for both propellants. The ratio of masses and ionization rate coefficients for these two propellants is responsible for krypton’s peak efficiency occurring at a higher current than xenon. In parallel, we have used this test to highlight key challenges with testing higher power electric propulsion devices, quantifying the degree of backscatter and background pressure as well as showing the limitations of conventional probing methods at high powers.

From a practical perspective, our findings may have implications for future developmental efforts on high-power electric propulsion systems. We have discussed, for example, how operating at a higher current density may enable efficient operation on propellants that have lower ionization cross sections than xenon. We have also shown that over the power range we investigated, the thrust densities of a Hall thruster can exceed those of a competing technology, the applied-field magnetoplasmadynamic thruster, while still maintaining attractive specific impulse and efficiency. We caveat both of these findings, however, with the remark that it is not certain if these advantages remain at even higher current densities. It is not clear if and how much performance will decrease as the thruster operating envelope continues to expand. Indeed, this study was a laboratory demonstration on a laboratory device, and

there remain several potentially major challenges with a more practical flightlike implementation. With that said, ultimately, this work represents a critical step in demonstrating how Hall thrusters may be a key enabler for new and exciting capabilities of high-power electric propulsion.

CHAPTER 7

Mass Efficiency Scaling with Increasing Power

“Rocket science[, like finding Hall thruster scaling laws,] is tough, and rockets[, like Hall thruster scaling laws,] have a way of failing.”

–Sally Ride. Information in brackets added to make quote germane to this work.



Figure 7.1: **P**arker and **L**eanne do **S**cience. PaLS features the experiments KLIF-FARD (**K**rypton **L**aser-**I**nduced **F**luorescence **F**oray into **A**cceleration **R**egion **D**ynamics) and SCOOBIE (**S**wirling **C**athode **O**scillations from **O**ptics and **B**eamdump **I**mpact **E**xperiments).

7.1 Introduction

This chapter covers work initially published in Refs. [140, 172].

In Chap. 5, we showed that the performance gap between xenon and krypton operation at baseline conditions is worse on shielded thrusters than it is on unshielded [71]. In particular, we found that the efficiency gap widens at high voltages on a shielded thruster whereas it narrows at high voltages for an unshielded thruster [8, 10, 11, 13]. We proposed that this is due to the non-linear relationship between ionization cross-section and electron temperature combined with the higher electron temperatures seen on shielded thrusters (Sec. 5.4). However, in the pursuit of improving krypton performance, we also observed in Chap. 6 (Ref. [72]) that at ultrahigh current densities (i.e. an order of magnitude higher than the traditionally-accepted limit of ~ 100 mA/cm²), the efficiency gap between xenon and krypton closes. Indeed, at sufficiently high currents with a constant discharge voltage of 300 V, the efficiency of krypton overtakes that of xenon.

This argument is supported by our findings in Chap. 6 in which we saw the mass utilization of a shielded Hall thruster operating at constant voltage on xenon reach 100% mass utilization at 50 A. When operating on krypton, this point was not until 100 A. That said, we still lack direct measurements of the plasma density, electron temperature/ionization rate coefficient, and neutral velocity necessary to verify these scalings on shielded thrusters. In this chapter, we outline how we obtain these measurements and how we use them to evaluate our 0D model for mass utilization.

Given the potential benefits of krypton as a Hall thruster propellant and the physical changes between unshielded and shielded thrusters, there is an apparent need to investigate the underlying physics of krypton and xenon operation on shielded Hall thrusters. The goal of this chapter is to generate these types of internal measurements and relate them to our understanding of the thruster performance trends. To this end, we adopt an approach using multi-fluid simulations that are calibrated against limited experimental data to infer the internal plasma properties [173]. Our decision to use this method stems from the fact that internal probes are perturbative and may obscure physical trends [174]. This chapter is organized as follows. In Sec. 7.2 we describe the experimental setup, including the thruster and facility, the non-invasive diagnostics, and the numerical code which we employed. Next, in Sec. 7.3, we present the results of our experimental measurements and calibrated simulations. We also evaluate our 0D model for mass utilization, described in Sec. 3.3.3, from our simulated measurements of plasma parameters and compare it to experimental trends in mass utilization. In Sec. 7.4, we discuss our findings by using the simulation outputs to interrogate the physical origins of the difference in performance between krypton and xenon.

7.2 Methodology

In this section, we describe the experimental apparatus used in our work, including descriptions of the thruster and operating conditions as well as details regarding our laser-induced fluorescence system. Then, we discuss how we integrated our work with a multi-fluid Hall thruster code to aid in our understanding in the underlying physics of the thruster operation.

7.2.1 Experimental Setup

We ran this test in LVTF (Sec. 4.3) with 17 total pumps (13 cryos and 4 thumpers). An overview of the thruster and diagnostic setup is shown in Fig. 7.2. We employed the same thruster infrastructure as described in Sec. 5.2; in brief, we used a mass flow controllers rated to 400 sccm and 50 sccm for the anode and cathode respectively and a 60 kW power supply for the discharge. We operated the thruster with a constant cathode flow fraction of 7% and cathode electrically tied to body [51].

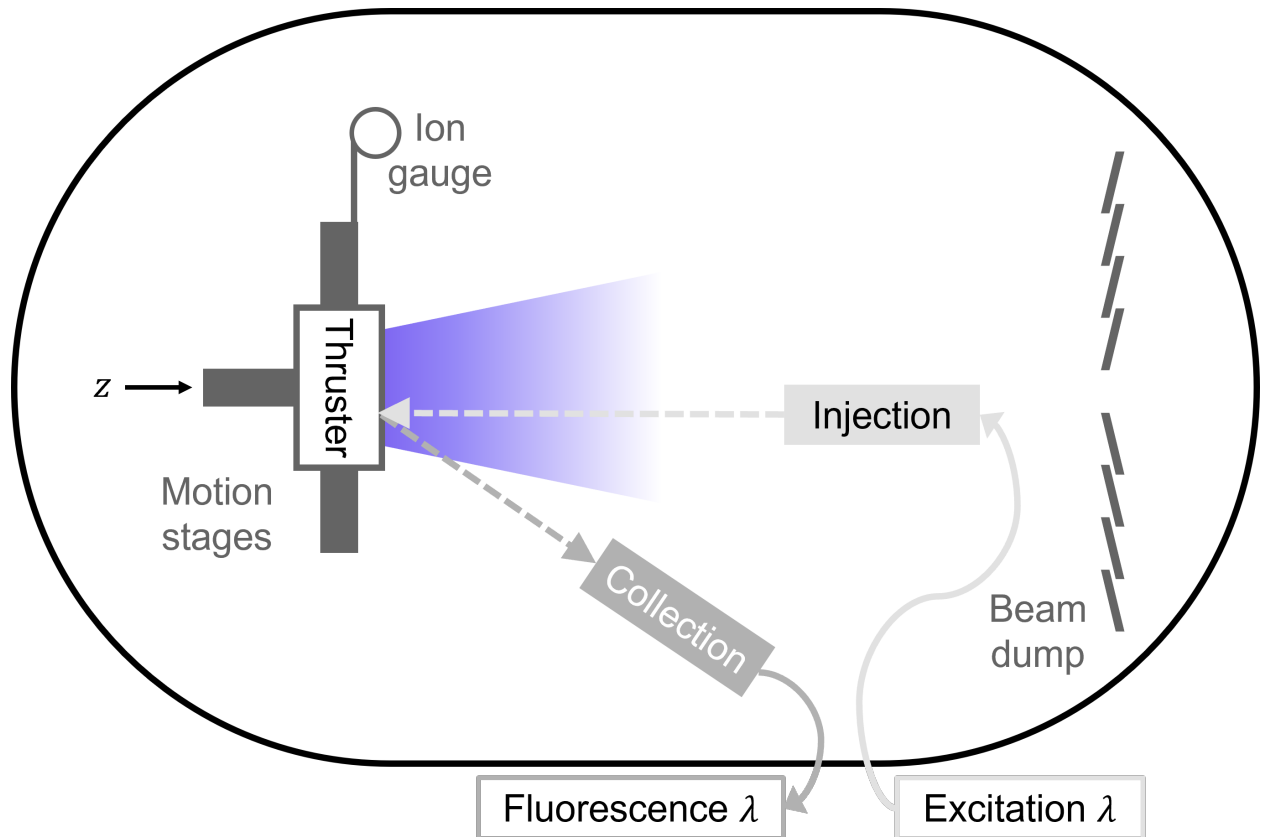


Figure 7.2: Notional top-down schematic of experimental setup in LVTF, including thruster, motion stages, ion gauge, injection/collection optics, and beam dump. Pumps are located around the walls of the facility but not shown in the diagram.

In Tab. 7.1, we show the operating conditions used throughout the LIF test campaign. Beginning with our baseline condition at 300 V and 15 A (4.5 kW), we scaled both voltage and current while holding the other constant to 6 kW and 9 kW. As was the case in our previous work (Chap. 5), we adjusted the strength of the magnetic field strength to minimize oscillations and discharge current. We note that we see slightly different flow rates and pressures in some of these values compared to Tab. 5.1. While we tabulate these values here for completion, the actual data used to calibrate Hall2De performance data was from our previous campaign (Tab. 5.1).

Species	V_d (V)	I_d (A)	P_d (kW)	B-field ratio (%)	\dot{m}_a (mg/s, sccm)	Pressure (μ torr)
Xe	300	15	4.5	100	14.7, 164	5.4
Xe	400	15	6	100	15.5, 172	5.8
Xe	600	15	9	100	16.3, 181	6.4
Xe	300	20	6	100	18.5, 206	6.6
Xe	300	30	9	100	25.3, 281	8.8
Xe	300	15	4.5	100	11.2, 195	5.4
Kr	400	15	6	112.5	11.2, 196	5.6
Kr	600	15	9	112.5	12.4, 217	6.4
Kr	300	20	6	87.5	14.5, 253	6.8
Xe	300	30	9	87.5	20.9, 366	7.8

Table 7.1: Operating conditions and base pressures.

In our LIF setup, we use the xenon and krypton transitions from Tab. 4.2 and follow the methodology outlined in Sec. 4.4.4. For each condition, we performed axial LIF along channel centerline from approximately 0.13 thruster channel lengths L_{ch} upstream to 0.53 channel lengths downstream of the exit plane as shown in Fig. 7.3a. The spacing of these points varies for each condition based on the difficulty of collecting data—for instance, the 600 V krypton condition had rapidly-increasing thruster temperatures, so we took fewer data points to limit operation at this condition. We took the average velocity to be the mean of the two-peak Gaussian fit to each profile (Sec. 4.4.4.4) and performed bootstrapping on each IVDF to obtain uncertainties in this value. It should also be noted that due to the stark two-peak distribution seen at locations near the exit plane for the 600 V, 15 A xenon condition, slight changes were made to initial parameters in the two-peak Gaussian fit function to force a fit with two distinct peaks, leading to a larger uncertainty in the fit.

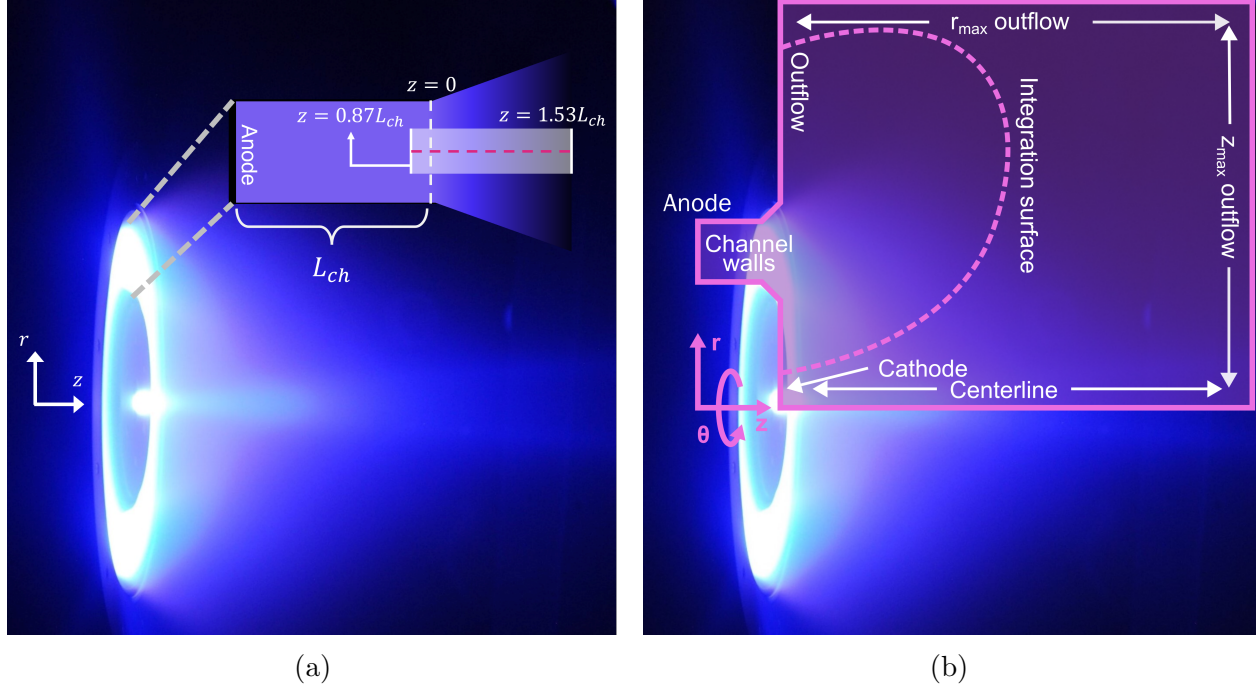


Figure 7.3: a) Experimental domain for LIF interrogation (shown as a transparent white rectangle) and b) simulation domain overlaid atop H9 operating on krypton at 300 V, 15 A. Note that in the experimental domain, channel centerline $r = 0$ is represented by the dashed pink line.

7.2.2 Simulation Setup

We used Hall2De, a 2D axisymmetric multi-fluid/particle-in-cell (PIC) Hall thruster code developed at the Jet Propulsion Laboratory [173], as our virtual diagnostic to infer internal plasma properties. In this section, we first provide an overview of the physics incorporated into the code. Next, we discuss the need for ad hoc transport coefficients to model electron dynamics within the thruster. Finally, we outline the metrics we used to determine whether the transport coefficients yield a “calibrated” simulation.

7.2.2.1 Model Physics

We briefly describe the code here. Additional information can be found in Refs. [173] and [19]. The code has different solvers for each species. For the electrons, Hall2De employs a generalized Ohm’s law formulation for the momentum balance and an equation for the internal electron energy. These electron equations of motion are solved on a mesh aligned with the applied magnetic field. This magnetic-field-aligned mesh (MFAM) helps reduce numerical diffusion which results from the anisotropy of electron dynamics in a magnetized

plasma.

For the ions, Hall2De has the capability to model them kinetically or as a series of fluids. In this present work, we used the fluid version of the code, which solves fluid mass, momentum, and energy equations for several ion species, differentiated by the charge state and origin. We include three ion charge states (singly-, doubly-, and triply-charged), for each of two ion fluids, yielding a total of six ion populations. The first ion fluid tracks ions born in the main thruster beam, while the second comprises ions emitted from the cathode or born in the cathode plume. As ions in a Hall thruster are unmagnetized by design, their equations of motion are solved on a more standard rectilinear/Cartesian grid. For the neutrals, the code uses a collisionless line-of-sight view-factor algorithm on the same mesh as the ions.

7.2.2.2 Anomalous Electron Transport

As with all fluid- or hybrid-based approaches to the thruster modeling [24], Hall2De under-predicts the electron current in these devices with classical forms for the electron dynamics. This is accounted for by introducing an effective or “anomalous” electron collision frequency into the generalized Ohm’s law: $\nu_e = \nu_{e,class} + \nu_{e,anom}$, where ν_e is the total electron collision frequency, $\nu_{e,class}$ is the electron collision frequency due to classical electron-neutral collisions, and $\nu_{e,anom}$ is the anomalous collision frequency. By adjusting this collision frequency, it is possible to increase the effective cross-magnetic field electron transport and therefore bring the electron current more in line with what is measured experimentally. The version of Hall2De we use for this work employs a static model in which the anomalous collision frequency is prescribed as a piecewise function of axial distance from the anode along channel centerline:

$$\log \frac{\nu_{e,anom}}{\omega_{ce}} = \begin{cases} c_1 & z \leq z_1 \\ c_i + (c_{i+1} - c_i) \frac{z - z_i}{z_{i+1} - z_i} & z_i < z \leq z_{i+1} \\ c_4 & z > z_4 \end{cases} \quad (7.1)$$

In the above, i is an index that ranges from 1 to 4, z denotes axial distance from anode, (z_i, c_i) denote adjustable free parameters, and ω_{ce} is the electron cyclotron frequency. The spatial dependence of the anomalous collision frequency is determined by these four “nodes”, (z_i, c_i) , with the magnitude of the anomalous collision frequency between the nodes determined by a logarithmic interpolation. This yields a five-zone model with eight free parameters. In Fig. 7.4, we show a notional example of such a collision frequency profile defined by Eq. 7.1. With the values on channel centerline determined, we then assume the anomalous collision frequency is constant in the radial direction. In practice, the user can adjust the eight free

parameters in this model to match the experimentally measured quantities of interest such as current and thrust. This piecewise approach has been applied with Hall2De over a wide range of Hall thrusters [167, 175].

Because the form for the anomalous collision frequency is not known a priori, Hall2De and other Hall thrusters codes are not fully predictive. They cannot, given only the geometry and operating conditions of a thruster, accurately capture its plasma properties and performance. However, once a collision frequency profile is prescribed, the physics in the model as represented by the governing equations for the electrons, ions, and neutrals are closed and self-consistent. Our overarching assumption is then that if the collision frequency can be calibrated by matching measurable quantities of interest, the other plasma properties produced by the code also will represent those of the actual thruster. In this sense, the model can be used as a virtual diagnostic for plasma parameters within the thruster and is thus an invaluable tool for accessing difficult-to-measure quantities.




Figure 7.4: Notional anomalous collision frequency profile according to Eq. 7.1, showing locations of nodes (z_i, c_i) .

7.2.2.3 Metrics for Calibrating Anomalous Collision Frequency

As discussed in the preceding section, it is necessary to have experimental measurements to match to simulation outputs. For our effort, we use the discharge current, thrust, anode efficiency, and ion velocity profile. We characterize a “calibrated” simulation as one that matches experimental data within some margin. The current and thrust are direct outputs of

the code. The simulated thrust from Hall2De, however, is known to be lower than experimental data [167, 176]. We adjusted the thrust for all simulations by an empirically-determined constant factor of 1.09 before comparing with experimental data. We then calculate the anode efficiency with this adjusted thrust. Hall2De is also capable of generating ion velocity as a function of axial distance from the anode. This is a property that we are able to measure non-invasively with laser-induced fluorescence (Sec. 4.4.4). These spatially-resolved data has become one of the standard references for calibrating the electron collision frequency as it varies axially [175]. To quantify the agreement between the simulation and the experimentally-measured ion velocities, we define the integrated velocity residual (IVR) [167]:

$$\text{IVR} = \sqrt{\frac{\int_{z_0}^{z_N} (u_{i,exp}(z) - u_{i,sim}(z))^2 dz}{\int_{z_0}^{z_N} u_{i,exp}^2(z) dz}}. \quad (7.2)$$

When we calibrate the collision frequencies in the code, we manually adjust the coefficients in Eq. 7.1 and systematically compare the code outputs to these four metrics. The requirements we adopt for a “calibrated” simulation are

1. discharge current, I_d , within 1%,
2. thrust, T , within uncertainty of experimental data,
3. anode efficiency, η_a , within uncertainty of experimental data,
4. IVR below 0.12.

7.2.2.4 Simulation Domain

In Fig. 7.3b, we show the simulation domain for Hall2De superimposed on an image of the H9 thruster, including the boundaries and coordinate system. The total simulation domain for the thruster extended approximately eight channel lengths downstream of the anode and eight channel lengths radially outward from the center of the thruster.

We summarize in Tab. 7.2 simulation parameters for Hall2De used in this investigation. We used a field-aligned mesh with 1949 cells and a rectilinear grid with 1684 cells. This grid resolution was selected to balance numerical accuracy with speed, owing to the need to run large numbers of simulations sequentially in order to calibrate the anomalous collision frequency profiles. As with experimental testing, the cathode flow fraction was set to 7%. The ionization fraction and electron temperature at the cathode orifice were set to 5% and 3 eV, respectively. These values are in line with experimental measurements of the

cathode employed in the H9 [110]. We used two ion fluids in the simulation differentiated by their origin, one from the beam and from the cathode. Lastly, we set the temperature of the thruster walls and anode to 400 K. We assumed the neutrals were emitted from the anode with an average speed given by a half-Maxwellian distribution at this temperature, $v_n = \sqrt{2k_B T / \pi m_i}$, resulting in a velocity of 129.1 m/s for xenon and 161.6 m/s for krypton.

Parameter	Value
Number of cells (MFAM)	1949
Number of cells (rectilinear grid)	1684
Maximum charge state	3+
Number of fluids	2
Cathode flow fraction	7%
Cathode ionization fraction	5%
Cathode electron temperature	3 eV
Wall temperature	400 K
Timestep	20–50 ns
Max simulated time	3 ms

Table 7.2: Numerical parameters employed in this work.

We ran each simulation for 3 ms of simulated time with a timestep between 20 and 50 ns. This was a sufficient amount of time for the ion velocity, discharge current, and thrust to converge in a time-averaged sense. We then evaluated time-averaged properties using the last millisecond of simulation time. Three milliseconds of simulation time corresponded to between 5 and 20 hours of wall time per simulation using 8 CPU cores. While Hall2De is designed to run on ordinary workstation computers, we conducted the majority of the simulations using the Great Lakes supercomputing cluster at the University of Michigan. This enabled us to run many simulations in parallel, reducing the amount of user time required to calibrate each condition.

7.3 Results

In this section, we present experimental results as well as calibrated simulations and their outputs for each condition.

7.3.1 Experimental Results

In this section, we summarize the global performance metrics for our thruster and present the ion velocity profiles obtained through LIF. The performance measurements are taken

Species	Condition	Thrust (mN)	η_a (%)
Xe	300 V, 15 A	292.9 ± 3.5	64.2 ± 1.6
Xe	400 V, 15 A	350.8 ± 3.4	66.5 ± 1.5
Xe	600 V, 15 A	447.2 ± 3.0	68.3 ± 1.3
Xe	300 V, 20 A	377.6 ± 3.8	64.2 ± 1.3
Xe	300 V, 30 A	547.0*	–
Kr	300 V, 15 A	235.8 ± 2.5	52.2 ± 1.1
Kr	400 V, 15 A	269.1 ± 2.5	52.6 ± 1.1
Kr	600 V, 15 A	350.3 ± 2.5	54.7 ± 0.9
Kr	300 V, 20 A	317.3 ± 2.0	54.8 ± 0.8
Kr	300 V, 30 A	480.3*	–

Table 7.3: Experimentally measured thrust and anode efficiency for all conditions [71]. *The 300 V, 30 A condition thrusts were linearly extrapolated from 15 to 20 A.

from Chap. 5 and are summarized in Tab. 7.3. All the data in this table is taken directly from the work in App. B.1 and reprinted here for easy access. We note here that we did not have experimental thrust and anode efficiency measurements of the 30 A conditions. Instead, we estimated the values for those conditions by linearly extrapolating thrust from 15 to 20 A.

The thrust and efficiency results in Tab. 7.3 show both metrics increasing monotonically with increasing power. This occurs for both voltage and current. As we noted previously (Chap. 5), the efficiency gap between xenon and krypton remains approximately constant between the 300, 15 A and 600 V, 15 A conditions. This is in contrast to trends with increasing voltage observed on unshielded thrusters where the efficiency gap closes [8, 10, 11]. On the other hand, we note that the gap decreases as current increases from 15 to 20 A with voltage held constant at 300 V. This is more in line with previous measurements on unshielded thrusters [9, 10, 12].

As described in Sec. 7.2, we measure the velocity at a given point axially along channel centerline as the mean of a two-peak Gaussian fit to the IVDF. In Fig. 7.5, we show the ion velocity profiles for different voltages and currents. At higher voltages, the ions reach a faster final velocity due to the greater acceleration voltage. We also note that with increasing voltage, the profile moves upstream, consistent with previous measurements [136, 138, 177]. However, there is minimal change from 400 to 600 V. With increasing current, the final velocity remains approximately constant, and the acceleration region shifts downstream. This downstream shift in acceleration region is consistent with previous measurements on both unshielded [136] and shielded [138, 151, 177] thrusters.

In Fig. 7.6 we present the velocity profiles for krypton as a function of voltage and




Figure 7.5: Profiles of mean ion velocity along thruster channel centerline with xenon operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

current. We first note that krypton reaches a higher final velocity than xenon at a given condition. This stems from the fact that krypton is a lighter gas and therefore achieves a faster final velocity for a given acceleration voltage. In terms of trends with voltage, the velocity profiles exhibit similar behavior to xenon in that they generally move upstream and steepen. However, while the 400 and 600 V profiles nearly lie on top of each other for xenon,

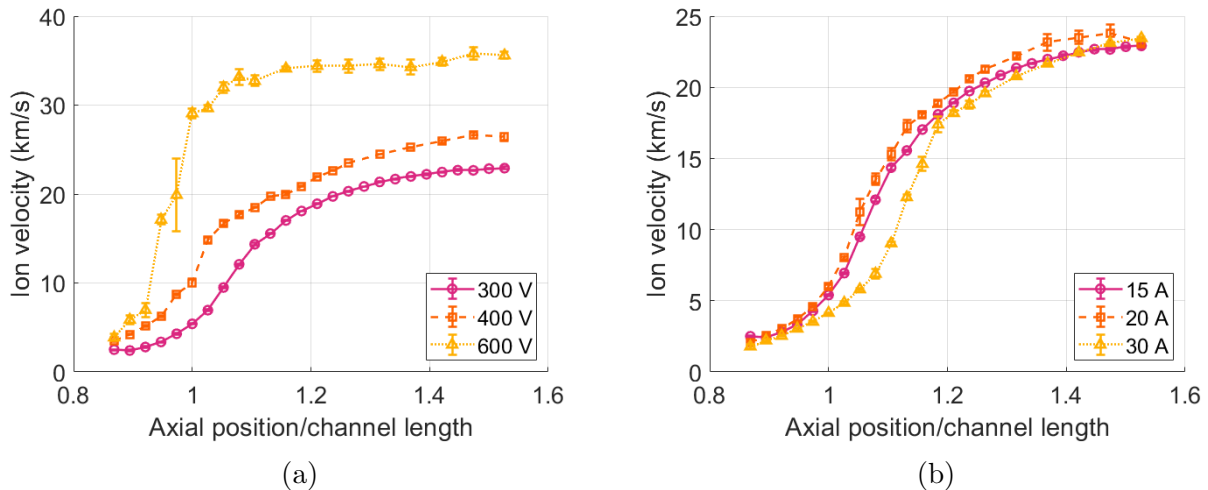


Figure 7.6: Profiles of mean ion velocity along thruster channel centerline with krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

the 400 V condition for krypton lies between the 300 and 600 V conditions. With increasing current, there is a similar downstream shift in the acceleration region for krypton operation as exhibited by the thruster operating on xenon from 15 to 30 A. However, the 20 A profile in the case of xenon lies between the 15 and 30 A profiles, while for krypton it is slightly upstream of the 15 A profile.

Using the definitions outlined in Sec. 4.4.4, we show in Fig. 7.7 the locations and widths for krypton and xenon at each operating condition. Note that “error bars” represent the width of the acceleration region, not uncertainty in the data. It becomes more apparent from Fig. 7.7 that the width of the acceleration region for krypton is distinctly decreasing as we increase current. This trend is not the case for xenon, where the acceleration region width stays relatively constant with increasing current. Figure 7.7 also emphasizes the trend in the krypton profile of narrowing and moving upstream continuously with higher voltage, while the xenon profile between 400 and 600 V is nearly identical. The upstream shift in krypton’s acceleration region is opposite to trends previously observed on unshielded thrusters [152] and non-intuitive given the worsening divergence we see for krypton at high voltages (Fig. 5.8). We encourage the reader to find an answer to this riddle. The furthest upstream acceleration zone location relative to all other conditions with the same propellant is at 600 V, 15 A for krypton and 400 V, 15 A for xenon, while the furthest downstream location is at 300 V, 30 A for both propellants.




Figure 7.7: Acceleration region locations and widths as calculated by Eq. 4.19 for xenon and krypton at different operating conditions.

7.3.2 Simulation Results

We present in this section the results from our calibrated simulations, highlighting the metrics that illustrate acceptable agreement with experimental data. We then show 1D profiles of internal plasma properties relevant to the mass utilization and convert them to 0D spatial averages. Finally, we use these average values to compare the trends in mass utilization as informed by theory to trends observed experimentally.

7.3.2.1 Demonstrating Model Calibration

In Fig. 7.8, we show comparisons for the velocity profile as measured experimentally through LIF and as produced through calibrated Hall2De simulations. In general, the simulated profiles are slightly broader than the experimental profiles. This is particularly true for the higher-voltage conditions. However, we are again within the requirements for IVR as defined by Eq. 7.2 for all but one condition.

Species	Condition	Thrust (mN)	η_a (%)	IVR
Xe	300 V, 15 A	291.7	63.9	0.07
Xe	400 V, 15 A	349.3	66.0	0.11
Xe	600 V, 15 A	450.7	69.2	0.08
Xe	300 V, 20 A	377.4	64.2	0.09
Xe	300 V, 30 A	543.2	64.7	0.08
Kr	300 V, 15 A	236.8	52.8	0.07
Kr	400 V, 15 A	270.7	53.1	0.05
Kr	600 V, 15 A	352.9	55.3	0.41*
Kr	600 V, 15 A	382.8 [†]	65.0 [†]	0.10
Kr	300 V, 20 A	318.6	55.3	0.05
Kr	300 V, 30 A	477.6	60.6	0.07

Table 7.4: Simulated values of thrust, anode efficiency, and integrated velocity residual for all conditions. *This condition did not meet the IVR requirement. [†]This condition did not meet the thrust and efficiency requirements.

We summarize key global metrics (Sec. 7.2.2.3) from these simulations in Tab. 7.4, where we compare the simulation thrust, anode efficiency, and ion velocity residual (Eq. 7.2) for each condition. In comparing to Tab. 7.3, we note that all conditions satisfy our criteria outlined in Sec. 7.2.2.3 with two exceptions. First, we again note that there was no thrust data experimentally measured at 30 A. For these two simulation cases, we neglected condition 3. and modified condition 2. from Sec. 7.2.2.3 such that we matched the extrapolated thrust of the 30 A condition to within 1%, approximately equal to the standard experimental error. Second, we were not able to satisfy all calibration requirements for krypton at 600 V, 15

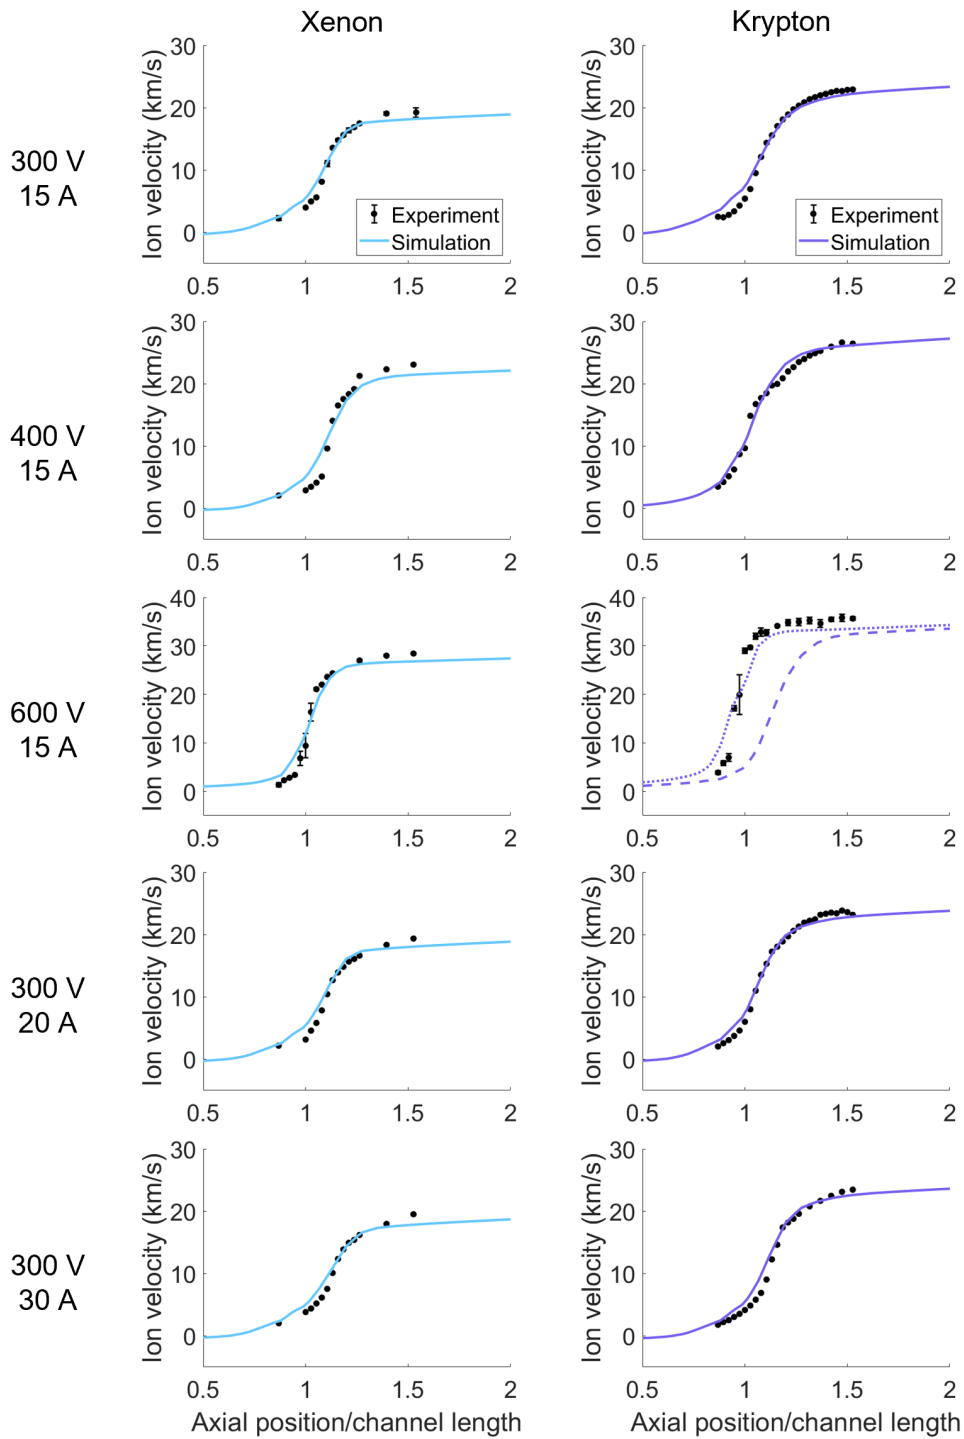


Figure 7.8: Experimental (black) and simulation (color) ion velocity profiles for each condition. Note that the 600 V, 15 A conditions have a different scale for ion velocity due to the high voltage. For the krypton 600 V, 15 A condition, the dotted line is the condition at which the thrust/efficiency matched and the dashed line is the condition at which the IVR matched.

A. We could reliably produce a profile that satisfied the thrust and efficiency requirement but had a high IVR ($\text{IVR} > 0.4$) and another profile that matched the velocity profile ($\text{IVR} = 0.1$) but had a thrust 30 mN too high. We show both values for performance metrics in Tab. 7.4 and both ion velocity profiles in Fig. 7.8. We suspect our inability to satisfy the criteria may be due to the simplified form of anomalous collision frequency profile that we have adopted, which makes it difficult for the simulated ion velocity profile to capture the steepness of the experimentally measured profile. To account for our lack of convergence to an exact match, we consider both profiles in the following analysis.

7.3.2.2 Anomalous Transport Profiles

In Fig. 7.9 we show the calibrated anomalous transport profiles for xenon with increasing voltage and increasing current. The classical collision frequency over the channel of a Hall thruster ranges from 10^5 to 10^7 s^{-1} . In general, the anomalous collision frequency begins at values near the classical frequency before steadily increasing until slightly downstream of the exit plane. At this point, there is a sharp decrease in the anomalous frequency roughly coincident with the location of the peak magnetic field strength. Finally, the frequency rises to a higher value—on the order of 10^9 s^{-1} —and declines as we move further downstream. This shape, particularly the minimum in anomalous collision frequency, is typical for Hall thrusters [28].




Figure 7.9: Calibrated anomalous electron collision frequency profiles along thruster channel centerline with xenon operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

In terms of trends with increasing power, the anomalous transport profiles remain largely unchanged with increasing current, excepting a slight downstream shift from 20 to 30 A. This behavior intuitively follows from the observed slight downstream shift from 20 A to 30 A in the velocity profile as seen in Fig. 7.5b. With increasing voltage, there is a much more pronounced change in the transport profile. Notably, the anomalous frequency near the anode before the sharp decrease increases with higher voltages.

In Fig. 7.10 we present the calibrated anomalous collision frequency profiles for krypton with increasing current and voltage. The general shape of the anomalous collision frequency profiles for krypton is the same as that of xenon, with an initial increase, a sharp trough leading into another increase, and finally a steady decline. As was the case with xenon, there is very little change in the anomalous profile with increasing current. Although the 20 A condition has a sharper decrease than the other conditions, the profile largely remains the same. With increasing voltage, we see that the initial anomalous collision frequency increases for krypton in a manner similar to that of xenon. One minor difference is that with increasing voltage on krypton operation, the downstream collision frequency decreases while for xenon it remains approximately constant between conditions.

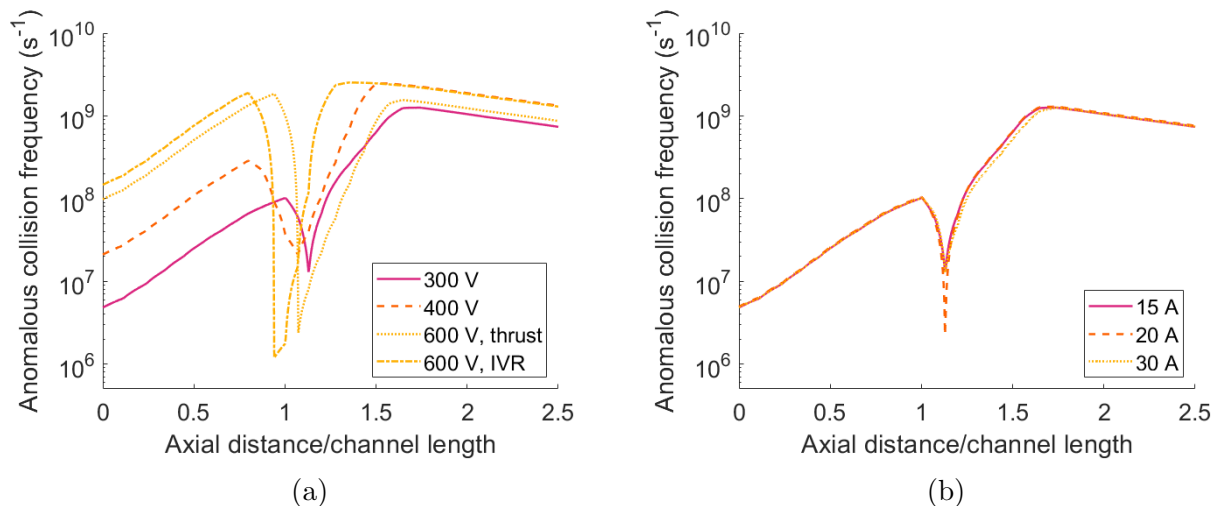


Figure 7.10: Calibrated anomalous electron collision frequency profiles along thruster channel centerline with krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

While these observations are for now purely empirical, the constancy of trends between xenon and krypton show promise for how extensible they are to other propellants. We may find it easier to calibrate high-current or high-voltage conditions for different propellants in the future, given only a “baseline” condition at a lower voltage and current, by employing

the same general trends. In particular, the transport profile for current may remain entirely unchanged and still yield close agreement between experiment and simulation. While the trends for voltage are less clear, increasing the anomalous collision frequency near the anode as voltage increases may be a useful guideline for hand-tuning simulations in the future.

7.3.2.3 1D Profiles of Plasma Properties

We show in Fig. 7.11 examples of the 1D simulated profiles of key plasma properties along channel centerline. These parameters are electron or plasma density n_e , electron temperature T_e , ionization rate coefficient k_{iz} , and neutral velocity v_n . While we only show profiles for the 300 V, 15 A conditions, the trends are broadly representative for all cases. These profiles are shown from the anode ($z = 0$) to the end of the integration region a half-channel length downstream of the exit plane ($z = 1.5$).

We can see from Fig. 7.11a that the electron temperature steadily increases through the region before peaking slightly downstream of the exit plane and decreasing. We can attribute the increase in temperature from 1.5 to $1.2L_{ch}$ to Ohmic heating as the electrons transit upstream. We expect this to peak at the location of peak magnetic field where the electrons are most strongly impeded [49]. The subsequent decrease in temperature from the exit plane to the anode can be attributed to ionization and wall losses. The electron temperature peaks at about 40 eV for both propellants, a scaling of approximately $T_e \propto 0.13V_d$. This is a departure from the assumption of $T_e \propto 0.2V_d$ we made in our previous work on shielded thrusters (Chap. 5) and is closer to the scaling of $T_e \propto 0.1V_d$ exhibited by unshielded thrusters [18, 163, 164, 165, 166]. Physically, this scaling arises from the fact that the discharge voltage is the source for electron heating.

We also note that the electron temperature of krypton is slightly higher than that of xenon throughout the range, in keeping with trends previously seen on unshielded thrusters [6, 152]. This may be attributed to the need for hotter temperatures with krypton to maintain ion production in the thruster, which in turn may be the result of krypton's lower ionization rate coefficient at a given temperature (Fig. 3.6). Indeed, we see from Fig. 7.11b that xenon and krypton have very similar ionization rate profiles. We return to this discussion point in Sec. 7.4.3. We do comment here, however, that this is a somewhat surprising result given that in our previous work, we attributed krypton's lower mass utilization to its lower ionization rate coefficient at a given temperature [71].

Figure 7.11c shows that the plasma density for both propellants increases immediately after the anode before decaying throughout the rest of the channel. This behavior can be understood by again considering the movement of electrons towards the upstream anode. The density increases from the exit plane (at L_{ch}) to $0.1L_{ch}$ due to ionization and subse-

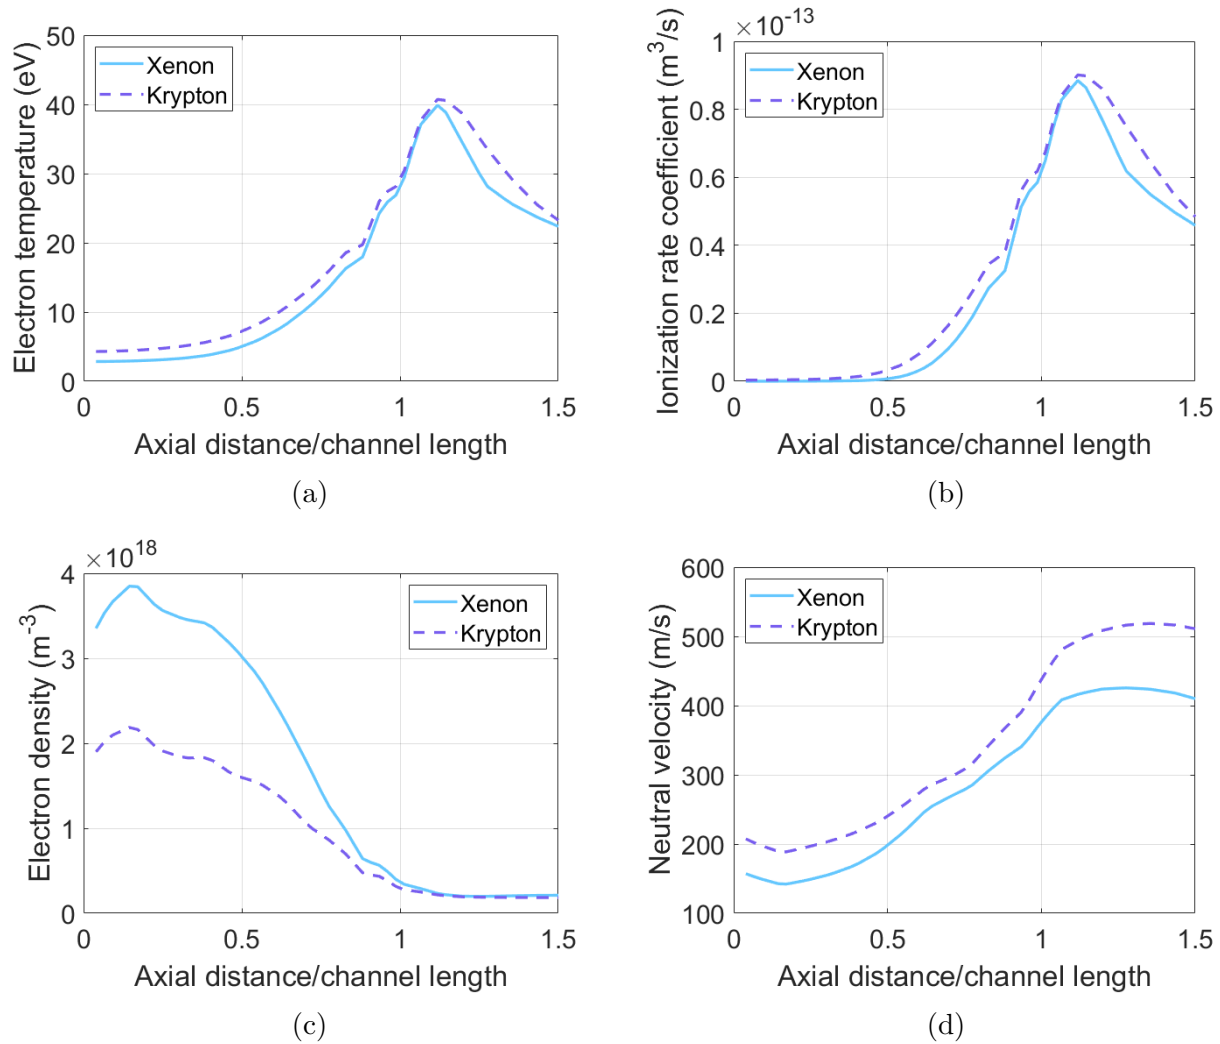


Figure 7.11: Axial profiles of a) electron temperature, b) ionization rate coefficient, c) plasma density, and d) neutral velocity along channel centerline for operation at 300 V, 15 A on xenon and krypton.

quently decreases from $0.1L_{ch}$ to the anode due to the existence of pre-sheath effects directly downstream of the anode. We also note that the peak plasma density of xenon is approximately a factor of two higher than that of krypton. This is likely due to krypton’s higher mobility—because of its lower mass, krypton reaches higher velocities at a given discharge voltage than xenon (c.f. Fig. 7.5b and 7.6b). Since the krypton population is moving faster while a constant current is being maintained, particles spend less time in a given region, leading to a lower particle density.

Finally, we see in Fig. 7.11d that the neutral velocity gradually increases through the channel before plateauing near the exit plane at L_{ch} . This may be attributed to the expansion of the gas as it exits from the anode, which has a smaller surface area than the channel itself. Additionally, ionization may preferentially deplete slower neutrals, leaving a larger population of fast neutrals further downstream. This accelerating behavior is reflected by experimental measurements of neutral xenon velocity in the channel of a Hall thruster [?]. We also see that krypton’s neutral velocity is higher than xenon’s, again owing to its lower mass. The ratio between these velocities is approximately $\frac{v_{n(Kr)}}{v_{n(Xe)}} = 1.22$, which is nearly the same as the inverse root ratio of masses, $\sqrt{\frac{m_n(Xe)}{m_n(Kr)}} = 1.25$.

7.3.2.4 Spatially Averaged Plasma Properties

While the 1D profiles of plasma properties are useful for understanding changes through the channel, averaged (0D) values allow us to use the framework outlined in Sec. 3.3 to understand how the mass utilization scales. To this end, we define an average value of a given property x over the axial direction z as

$$\langle x \rangle = \frac{\int_0^L x(z) dz}{L}, \quad (7.3)$$

with a trapezoidal numerical integration, where $z = 0$ is the location of the anode and $z = L$ is the length of integration. For this work, we used $L = 1.5L_{ch}$ as we found this to be location at which for all conditions the neutral density has decayed to near-zero, implying that the bulk of the ionization is complete. For the krypton 600 V, 15 A condition (see Sec. 7.3.2.1), we evaluated Eq. 7.3 for both cases and averaged them, evaluating uncertainty as the difference between cases. Armed with this framework for analyzing the outputs of our simulation, we present here our 0D results as a function of increasing discharge voltage and current for each propellant.

In Fig. 7.12a, we see that the spatially averaged electron temperature steadily increases with discharge voltage. The electron temperature of krypton is 2–3 eV higher than that of




Figure 7.12: Axially-averaged electron temperature along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

xenon at all voltages. The scaling of electron temperature with voltage is expected from a consideration of energy balance [18]. As shown in Fig. 7.12b, the electron temperature remains nearly constant with increasing current. With the exception of a slightly-higher electron temperature of 18 eV for krypton at 15 A, all other average electron temperatures for both propellants are ~ 16 eV.




Figure 7.13: Axially-averaged ionization rate coefficient along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

In Fig. 7.13 we plot the trends in channel-averaged ionization rate coefficient for increasing voltage and current. These closely mirror the results for electron temperature (Fig. 7.12). With increasing voltage, the ionization rate increases. With increasing current, we once again see very little variation in the ionization rate coefficient for both propellants.




Figure 7.14: Axially-averaged electron density along channel centerline of xenon and krypton operation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

From Fig. 7.14, we see that plasma density is higher with both increasing voltage and current, although the increase with voltage is less dramatic than with current. The increase in plasma density with current is intuitive—as more neutrals flow into the channel, more charge carriers are made available, and the plasma density increases. The dependence on voltage, however, is not expected based on conventional scaling laws. For a constant current and the higher velocities at increased voltages, we would anticipate that the plasma density decreases. This behavior is not reflected here. We comment further on this in Sec. 7.4.3. We also note that at all conditions, xenon’s plasma density is $\sim 50\text{--}70\%$ higher than krypton’s.

We note that we do not show plots for the channel-averaged neutral velocity at all conditions because we found that this value varied by less than $\sim 2\%$ across all conditions (voltage and current) for a given propellant. The channel-averaged neutral velocity was approximately 282 m/s for xenon and 341 m/s for krypton.

7.3.2.5 Mass Efficiency

We present in Fig. 7.15 our calculated values of mass utilization as inferred from evaluating Eq. 3.15 with the 0D plasma properties presented in the previous section. In all cases, we

used a value of $\alpha = 0.25$ to achieve the best agreement with experimental measurements of mass utilization. While an internal study showed that averaging across radial position at each axial location yielded the same trends as using centerline values, we use this scaling factor to account for the difference in magnitude between the centerline value and the radially-averaged value. This result qualitatively shows that mass utilization increases with both voltage and current. This is consistent with the theory we outlined in Sec. 3.3 and with previous work on Hall thrusters. Using this theoretical framework in turn allows us to identify which plasma parameters are driving the increases in mass utilization. With voltage, the increases in both electron temperature (Fig. 7.12a) and plasma density (Fig. 7.14a) serve to decrease the ionization mean free path. With current, the higher plasma density (Fig. 7.14b) drives a shorter ionization mean free path.




Figure 7.15: Mass efficiencies of xenon and krypton operation as calculated from theory with simulated plasma parameters (Eq. 3.15) for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

Comparing the values at the lower-power conditions to experiment, we see that our efficiencies for both xenon and krypton are broadly consistent with previous measurements (Sec. 5.3). We see a theoretical mass utilization range of 91–98% for xenon compared to an experimental range of 93–98% (Fig. 5.7), and a theoretical range of 74–84% for krypton compared to an experimental range of 82–87% (Fig. 5.8). For both gases, we have excluded the 600 V, 15 A and 300 V, 30 A conditions as we do not have experimental measurements of mass utilization here. The slightly lower values of experimental krypton mass utilization at low currents compared to experiment may be attributed to the simplifications made in our model (Sec. 3.3) or uncertainty in the measurements ($\sim 4\%$ as reported in Chap. 5.3).

At all conditions, we see that the mass utilization of xenon is higher than that of krypton. This result is consistent with our previous experimental measurements of mass utilization (c.f. Fig. 5.7 and 5.8). As we discuss in the following Sec. 7.4, the key driver of this efficiency gap is xenon’s higher average plasma density in all cases. We note that the efficiency gap closes at higher powers for both voltage and current, i.e. krypton’s mass utilization improves more than xenon’s. This finding is in contrast to the experimentally-observed anode efficiency gap between xenon and krypton, which was shown to not close at high voltages (Fig. 5.5c). This discrepancy is likely attributed to the fact that there are more contributors to the anode efficiency than just mass utilization, a point we return to in Sec. 7.4.3.

Ultimately, the efficiency gap closes at high powers because of the exponential dependence of mass utilization (Eq. 3.15). Xenon, which has a higher plasma density at lower powers than krypton, exhibits a mass utilization closer to the asymptotic 100% even at its lowest-power condition. Therefore, as the ionization mean free path decreases, xenon’s mass utilization shows a smaller increase compared to krypton. We further explore this behavior in Sec. 7.4.4.

7.4 Discussion

In this section, we elaborate on the behaviors of plasma parameters and mass utilization for xenon and krypton with increasing power. First, we discuss the implications of trends in the anomalous collision frequency profiles of our calibrated simulations. Then, we investigate the physical cause of xenon’s higher mass utilization compared to krypton. Finally, we discuss plasma parameter trends in the context of how they relate to the behavior of the efficiency gap between these propellants at high voltages and high currents.

7.4.1 Anomalous Collision Frequency Profiles

While discussing the potential physical causes of the scaling of the anomalous collision frequency is outside the scope of this work, we note that the trends observed may be conducive to easier future calibration of similar conditions. For instance, we observed that the anomalous collision frequency profiles needed to match data changed only slightly with increasing current density. This may indicate that the anomalous collision frequency scales weakly or not at all with plasma density. One consequence of this is that the electron confinement may not be severely impacted at high power densities. This is supported by our recent work in Chap. 6 (Ref. [72]) in which we operated a modified version of the H9 at up to 125 A on xenon and 150 A on krypton with only minor decreases in beam utilization efficiency. Obtaining LIF measurements of the ion velocity profiles at these higher current conditions

will help to validate or refute this hypothesis.

With respect to the scaling of the anomalous transport with discharge voltage, we found that the anomalous collision frequency near the anode needed to be increased as the discharge voltage increased. The physical reasons for this behavior are unclear, but if this trends holds, it could be used to infer anomalous collision frequency profiles for discharge voltages at which no LIF data is available. Generalizing these scaling trends across more propellants, operating conditions, and geometries could assist efforts to develop predictive models of the anomalous collision frequency in Hall thrusters.

7.4.2 Overall Efficiency Comparison

A key finding from our analysis is that as expected, the mass utilization is higher for xenon than it is for krypton at all conditions. This is broadly consistent with multiple previous studies [6, 7, 9, 10, 11, 12, 13, 71]. With that said, in these previous works, the cause for the difference in efficiency was largely attributed to the role of ionization rate [8, 9, 13, 71, 178, 179]. Assuming that xenon and krypton have the same electron temperature at a given condition, xenon would have a higher ionization rate coefficient (Fig. 3.6). In contrast, we have found that the electron temperature of krypton is higher than xenon by a large enough margin that its ionization rate coefficient is comparable to xenon’s (c.f. Fig. 7.12a and Fig. 7.13a). The implication of this result is that xenon’s higher mass utilization is primarily attributed to its higher plasma density and lower neutral velocity, both of which are dependent on xenon’s mass and not its ionization rate.

In retrospect, the difference in temperature and similarity in ionization rate coefficient between gases may not be surprising given how charge production scales with plasma properties. At a given condition, we have prescribed that the total current remains the same. By invoking quasineutrality and following the same rationale outlined in Sec. 3.5.4, we find that this discharge current is proportional to the neutral density, plasma density, and ionization rate coefficient, i.e. $I_d \propto n_n n_e k_{iz}$. We have shown in our previous work that the neutral densities for these two propellants is approximately equal at the exit plane [104], and we know the plasma density to be higher for xenon due to its higher mass (Sec. 3.3). Since the ionization rates of krypton are lower at a given temperature, we thus may expect higher temperatures to compensate. The fact that the ionization rate coefficients for each propellant are within 15% of each other (with krypton consistently higher) at all conditions lends credence to this interpretation. With that said, this qualitative explanation overlooks the potential roles of other factors, such as neutral density and plasma density, that also influence charge production.

7.4.3 Efficiency Trends with Increasing Voltage

Figure 7.15a shows that the mass utilization increases with voltage for both propellants. This can be attributed to both higher electron temperatures (Fig. 7.12a) and therefore ionization rate coefficients (Fig. 7.13a) as well as higher plasma densities (Fig. 7.14a). This trend in ionization rate is not surprising given the known scaling between voltage and electron temperature [18]. However, as we introduced in Sec. 7.3.2.4, the increase in plasma density with voltage is less intuitive. We expect that as the voltage increases, the speed of the ions does as well, leading to shorter residence times in the channel. Per established scaling laws for ion current density, we expect that for a fixed discharge current, the plasma density will scale inversely with the ion velocity and therefore voltage. In Fig. 7.16a we show how the plasma density profile changes between 300 and 600 V for xenon, and we see that this behavior is indeed true downstream of the exit plane. However, from the anode until $\sim 0.2L_{ch}$ upstream of the exit plane, the plasma density is higher with increasing voltage. Indeed, it is apparent from Fig. 7.14 that when averaging over our axial range of interest, we see the net plasma density increase with voltage. This in turn improves the mass utilization.

To help elucidate the underlying cause of the increasing plasma densities at high voltage, we also show in Fig. 7.16 comparisons of the electron temperature and ionization rate coefficient profiles at different voltages for xenon. It is apparent from these profiles that increasing the voltage also increases both of these plasma parameters in the near-anode region. Small increases in electron temperature (Fig. 7.16b) correspond to larger increases in ionization rate coefficient (Fig. 7.16c) due to the non-linear relationship between electron temperature and ionization rate (Fig. 3.6). Consequently, there is a marked increase in the ionization rate near the anode at high voltages, in turn producing a denser plasma. This may explain the higher plasma densities we see in the near-anode region and overall.

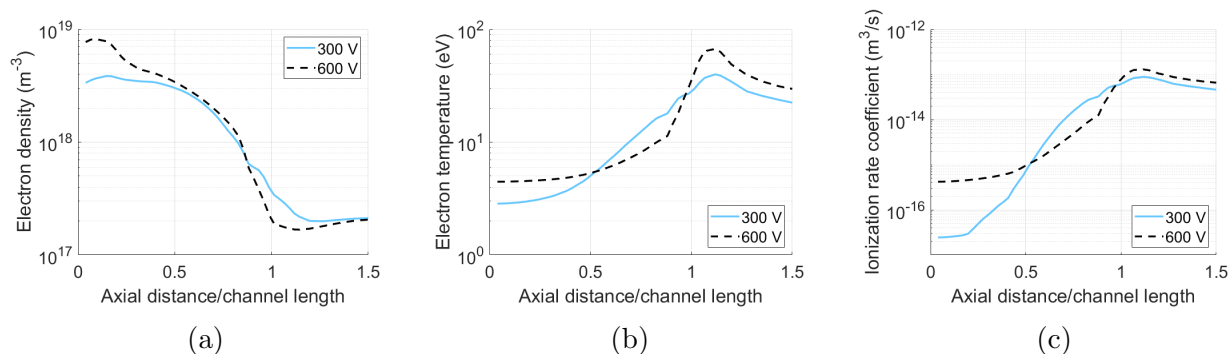


Figure 7.16: Axial profiles of a) plasma density, b) electron temperature, and c) ionization rate coefficient along channel centerline for operation at 300 V, 15 A (blue) and 600 V, 15 A (black) on xenon.

The increase in plasma density has a significant impact on the increase in mass utilization. While the ionization rate coefficient only increases by 20–23% for both gases from 300 to 600 V, the plasma density increases by 40–47%. Given the equal dependence of mass utilization on these two parameters (Eq. 3.15), this points to higher plasma densities, not ionization rates, as the primary cause for the increase in mass utilization with voltage.

Finally, we may address why we see the efficiency gap between xenon and krypton closes significantly on unshielded thrusters but not on this shielded thruster. In our previous work, we attributed this to the range of electron temperatures being higher for shielded thrusters than for unshielded, leading to a larger increase in the ionization rate coefficient of krypton compared to xenon over a given change in electron temperature. However, given that krypton and xenon are at different electron temperatures for a given condition and that we consider the axially-averaged electron temperature along channel centerline (instead of just the peak value), this argument is unfounded. We in fact see that krypton has a higher ionization rate than xenon at a given condition, and as the voltage increases, the ionization rate coefficient of both gases increases by about the same factor (Fig. 7.13a). The plasma density also increases by approximately the same factor for both propellants (Fig. 7.14a)—by about 27% from 300 to 400 V and 11% from 400 V to 600 V. As we expand on in the next section, a larger increase in plasma density is more beneficial in terms of mass utilization for a propellant that begins at lower density due to the exponential nature of η_m . This is why we see the gap between xenon and krypton’s efficiency narrow at high voltage.

While the efficiency gap in mass utilization as predicted by our theory does close, the efficiency gap we observed experimentally at 600 V (Ref. [71]) is the same as it is at lower voltages. Given the discrepancy between the behavior of the mass efficiency and the anode efficiency, we need to reconsider the effect of other efficiency losses. Indeed, as we have previously noted [71], the second largest driver of the performance gap between xenon and krypton is the divergence efficiency. Considering how changing from an unshielded to a shielded topology pushes the acceleration region downstream and increases divergence [49], it is possible that the acceleration region and therefore divergence may respond differently at higher voltages and for each propellant. These effects are outside the scope of our mass utilization analysis.

In summary, the behavior of mass efficiency with increasing discharge voltage is not driven only by the increase in electron temperature and ionization rate coefficient as we initially thought, but instead primarily by the increase in plasma density. This phenomenon explains the higher mass utilization of xenon despite its lower ionization rate coefficient as well as the decreasing mass utilization gap between xenon and krypton efficiency with increasing voltage. The approximately constant difference in anode efficiency between the two gases

may instead be attributed to other efficiency losses, such as the divergence of the plume. Given the historical dominance of mass utilization as the primary driver for the performance gap between xenon and krypton, we may eventually expect the anode efficiency gap to close as the mass utilization of both propellants approaches unity with increasing voltage. That said, if the other efficiency modes indeed respond differently with increasing voltage on these shielded thrusters than on unshielded, we may see unanticipated behaviors that lead to an efficiency gap that remains unchanged even at higher voltages.

7.4.4 Efficiency Trends with Increasing Current

Having established that increasing discharge voltage is an ineffective method of closing the efficiency gap between xenon and krypton, we turn now to increasing current. Figure ?? shows that increasing discharge current does indeed close the gap between xenon and krypton efficiency. While the simulated anode efficiency of xenon increases by less than 1% from 15 to 30 A, the anode efficiency of krypton increases by $\sim 8\%$ (Tab. 7.4). This behavior may once again be understood in context of mass utilization. As seen in Fig. 7.15b, the mass utilization increases for both propellants, but the relative increase for krypton ($\sim 20\%$ from 15 to 30 A) is far greater than that of xenon (less than 10% from 15 to 30 A). This trend closely mirrors what we observed in our previous work [72] when we increased current density by an order of magnitude over the traditional nominal value for Hall thrusters of 100–150 mA/cm² [18]. In this previous work, we saw the mass utilization of krypton increase more over a given change in current than xenon.

To understand the more rapid gains in mass utilization for krypton at these currents, we turn once again to the scaling given by Eq. 3.15 and the channel-averaged plasma parameter outputs from simulation presented in Sec. 7.3.2.4. First, in contrast with the voltage trends seen in Fig. 7.13a, the ionization rate coefficient remains nearly constant over the entire range of currents for both propellants. Taking the average over both propellants and all three current conditions, we find the mean ionization rate coefficient to be 3×10^{-14} m³/s with a standard deviation of less than 1%. Given this term and the nearly-constant neutral velocity for a given propellant, the main scaling for mass utilization with increasing current comes from the increasing plasma density.

Figure 7.17 shows a mass utilization curve for each propellant given its average neutral velocity (Sec. 7.3.2.4) and assuming the same ionization rate coefficient. We note once again that we included a factor of 0.25 in the mass utilization efficiency (Eq. 3.15) to account for the values used being along centerline instead of averaged radially across the channel. The plasma densities for each propellant at 15 A and 30 A are indicated with circular markers




Figure 7.17: Notional mass efficiency scaling for xenon (blue dotted) and krypton (purple dashed), with 15 A and 30 A conditions shown with circular markers for xenon (filled) and krypton (open).

on the plot with shaded regions between them indicating the range of plasma densities for a given propellant. The xenon curve approaches unity more rapidly, which means that at a given plasma density, the mass utilization of xenon is higher than krypton's. Both xenon and krypton's plasma density increase by a factor between 2.2 and 2.5 from 15 to 30 A (Fig. 7.14b), a much larger increase than what is seen from 300 to 600 V at a constant 15 A. However, the range of plasma densities for xenon is sufficiently higher than krypton's such that there are diminishing returns within that range. The amount of increase in η_m decreases for a given Δn_e at higher values of n_e due to the exponential nature of the mass utilization curve. This is why krypton, which begins at a lower plasma density and mass utilization, sees more of a benefit from increasing plasma density than xenon.

7.5 Summary

We operated and obtained performance measurements for a magnetically shielded Hall thruster from 4.5 to 9 kW by increasing both discharge current and discharge voltage for xenon and krypton, resulting in ten total conditions (five for each propellant). By empirically hand-tuning the anomalous electron transport profile in a multifluid Hall thruster code to velocity profiles obtained from laser induced fluorescence, we were able to match

the experimental thrust and anode efficiency (with a constant scalar) and the shape of the velocity profile. These calibrated simulations yielded simulated measurements of component efficiencies and plasma parameters through the channel as a function of axial distance.

To pinpoint the cause of the efficiency trends for each gas with increasing power, we invoke an expression for mass utilization that is dependent on the ionization rate coefficient (in turn dependent on electron temperature), plasma density, and neutral velocity. We find that in contrast to our previous understanding, the mass utilization of xenon is higher than that of krypton primarily due to its higher plasma density and lower neutral velocity, not its higher ionization rate. The ionization rate of krypton is actually higher than that of xenon due to krypton's hotter electron temperature. In turn, unlike what we had postulated in our previous work [71], the increase in mass utilization with voltage is driven not only by higher electron temperatures but also plasma densities. That said, based on our theoretical framework for mass utilization, we would still anticipate a decrease in the performance gap between gases at high voltage that we did not observe in the anode efficiency experimentally. This behavior may be due to the influence of other efficiency loss mechanisms that behave differently on shielded thrusters compared to unshielded thrusters.

In contrast, by increasing the discharge current by a factor of two, we were able to close the theoretical gap between xenon and krypton anode efficiency from 11% to 4%. We attribute this to a $\sim 20\%$ increase in krypton's mass utilization compared to $<10\%$ in xenon's. As the ionization rate coefficient and neutral velocity for a given propellant are virtually unchanged, this increase in mass efficiency is in turn caused by the lower range of plasma densities exhibited by krypton and the exponential behavior of mass utilization. Additionally, because krypton's beam utilization has a weaker dependence on increasing current than xenon's, the overall anode efficiency of krypton increases from 15 to 30 A while that of xenon remains approximately constant. This behavior of mass and beam utilization closely mirrors trends seen in our previous work (Chap. 6). In addition to the insight gained into improving krypton operation, this work has also yielded calibrated simulations of multiple operating conditions on a magnetically shielded Hall thrusters, in particular the anomalous electron collision frequency profiles needed to match experimental results. The trends we see in the anomalous electron transport profile may help guide future modeling efforts in anomalous electron transport as well as provide guidance in the hand-tuned calibration process.

Our results show that with both increasing voltage and increasing current, the behavior mass utilization is driven by plasma density. A disparity in plasma densities is also responsible for the efficiency gap between xenon and krypton in general, with xenon exhibiting average plasma densities 50–70% higher than those of krypton. We find that plasma density increases with voltage, leading to an efficiency gap between xenon and krypton that should

close at higher voltages. Given the lack of change in the anode efficiency gap that has been observed experimentally for this shielded thruster, we suggest that other efficiency modes may be in effect. Additionally, we believe the larger increase in plasma density at higher currents to be what allows the gap between xenon and krypton efficiency to close at these conditions. The findings in this work represent a marked step in our understanding of the physics behind xenon and krypton operation on a shielded Hall thruster and a path forward for continuing to improve the performance of krypton operation.

CHAPTER 8

Conclusions and Future Work

“I cannae’ change the laws of physics!”

–Montgomery “Scotty” Scott, *Star Trek*.

8.1 Summary of Work

We first introduce and outline this work in Chapter 1. In Chapter 2, we provide context for our work by describing space propulsion at large before focusing in on electric propulsion. Next, in Chapter 3, we introduce the principle of operation for Hall thrusters and review some of its major design elements. We also outline a framework used to characterize its performance and describe the current state of affairs for the Hall thruster. We motivate and review work that has been done to date on both krypton and high power operation. In Chapter 4, we introduce the experimental apparatus we used throughout our work, including the thrusters, probes, and integration with simulation.

We begin reviewing the major experiments performed during work in Chapter 5, starting with our comparison of xenon and krypton performance at baseline conditions on a magnetically shielded Hall thrusters. We found that unlike with unshielded Hall thrusters, the efficiency gap between xenon and krypton does not close with increasing voltage. We postulated that this is due to the behavior of the mass utilization efficiency and how it differs between unshielded and magnetically shielded thrusters. Having found that increasing voltage did not close the gap between xenon and krypton efficiency as we had hoped, we turn instead to the knob of increasing current in Chapter 6. Through this work, we were indeed able to improve the efficiency of krypton operation; at sufficiently high currents, we even found that the efficiency of krypton exceeds that of xenon. This behavior was attributed to the higher currents required for krypton to reach 100% mass utilization and to the weaker dependence of its electron confinement on discharge current compared to xenon. Finally, in

Chapter 7, we used simulations as a tool to investigate the underlying physics of the performance scaling for each propellant. This work confirmed the experimentally-observed scalings of krypton and xenon with voltage and current and confirmed that mass utilization was the main driver for these trends. Additionally, our findings pointed to the electron density as being the main culprit for the efficiency gap between propellants as well as the main driver for why the gap does not close with increasing voltage but does with increasing current.

8.2 Implications of Work

There are a few major takeaways from this work that we would like to highlight. First, we achieved what we set out to do: close the gap between xenon and krypton performance. This was accomplished by operating a Hall thruster at unprecedentedly high current densities (order of magnitude higher than traditional). The implication of this is that krypton is able to reach 100% mass utilization despite its lower mass and higher required ionization energies. This may be applied to other propellants as well, paving the way for efficient operation on even more exotic propellants.

The ability to operate at high current densities is in and of itself noteworthy. We were able to show that there was a minor (or for krypton, nonexistent) decrease in efficiency when increasing current from 15 to 125–150 A while operating on a magnetically shielded Hall thruster. This finding could enable the use of high-thrust density Hall thrusters in crewed deep space architectures.

Another key finding is that the plasma density drives the efficiency gap between xenon and krypton to a much greater degree than the ionization rate coefficient. This implies that the mass differential and not the difference in ionization rate is what causes krypton’s mass utilization to be so much worse than xenon’s. Extending this to other propellants, we may expect lower masses and not higher required ionization energies to be the main driver of the efficiency gap.

In terms of methods, we extended a number of probe correction factors for $E \times B$ and Faraday probes that had already been established for xenon to be applicable to krypton as well. On the hardware side, we sourced and set up a new krypton LIF system, giving PEPL the capability to further study krypton operation. As the industry shifts towards propellants other than xenon, having the diagnostic tools and capabilities to keep up becomes increasingly important.

8.3 Future Work

It is the nature of research and science as a whole that the more we learn, the less we feel that we know. This is why senior graduate students are so sad. While we hope that the body of this work has made a notable impact in our understanding of Hall thruster operation on krypton at high powers, there are multiple questions that remain unaddressed or have emerged as a result of our work. In this section, we suggest some potential avenues for future research.

8.3.0.1 Thruster Temperature, Stability, and Lifetime

The bulk of this work focused on the efficiencies and global performance metrics of Hall thrusters operating on krypton. However, practically, there are a number of other operational concerns equally (if not more) important as performance. Two critical metrics of thruster health are temperatures and stability. As discussed in Sec. 5.3.3, the steady-state temperature of krypton operation is nearly 100 °C higher than that of xenon. This may pose challenges in terms of thermal margin when designing thrusters for flight applications. In terms of oscillations, we saw peak-to-peak oscillations of well over 100% when operating on krypton at high voltages (Fig. 5.9). This may pose a risk to power processing units and other subsystem components. These challenges are points that need to be addressed when considering operation on different propellants.

Additionally, quantifying the lifetimes of these shielded devices on alternative propellants is outside the scope of this work but a vital aspect of characterizing thruster operation. Indeed, there has been some work suggesting that operating a shielded thruster on krypton has less effective shielding than with xenon due to higher electron temperatures [179]. We also did not address the challenge of pole erosion that has become increasingly prescient with shielded thrusters [144, 180, 181]; given the differences between xenon and krypton operation, there may be significant disparities in the rate of pole erosion as well. We challenge the reader to further delve into these fascinating problems.

8.3.0.2 Divergence Efficiency

As discussed in Sec. 5.4.2 and 7.4.3, the divergence efficiency plays a significant role in the efficiency gap between xenon and krypton. While the focus of our work has been on the primary driver of this gap, the mass utilization, it is apparent that the less-collimated krypton plume also contributes to its lower efficiency. This could be due to the lighter and therefore more mobile krypton atom; however, absent measurements such as ion velocity vectors (i.e. from LIF performed along multiple axes instead of just axially), we cannot conclusively

point to a physical cause for this behavior. Additionally, the different behavior of krypton's acceleration region with increasing voltage (Fig. 7.6a) and the potentially detrimental trends in divergence efficiency at high voltage (Sec. 7.4.3) both warrant further investigation.

8.3.0.3 More Alternative Propellants

We believe that much of the work done here may be extended to other propellants as well. In particular, the behavior we saw at high currents—where krypton's efficiency overtook that of xenon's—would likely extend to propellants with lighter mass, such as argon and nitrogen. It is apparent that with sufficiently high current density, the mass utilization of even low-ionization cross-section propellants can approach unity. Given that krypton's efficiency peaks at a current density twice that of xenon, we may also expect other propellants to peak at even higher currents. This is especially true given the weaker dependence of krypton's beam utilization on current. Detailed studies on the current densities of other propellants and where their optima exist would help elucidate these trends. Preliminary work from Hurley et al. [182] has indeed shown that argon's efficiency on the H9 MUSCLE continues trending upwards past 100 A, where the optimum for krypton lies.

Another avenue of exploration is that of propellant mixtures. We typically use high-purity gases in our work, with 99.9995% xenon and krypton. From a scientific perspective, this makes it easier to analyze differences in performance between propellants. Imagine, for example, how much of a pain it would be to align both a xenon and krypton LIF laser to the same point in a plasma plume. However, there may be advantages to operating with a mixture of gases. From a cost perspective, it is possible that a mix of xenon and krypton is cheaper than high-purity concentrations of either due to the difficulty of separating them. There could also be some performance benefits given xenon's higher performance at low current densities and krypton's higher performance at high current densities.

8.3.0.4 Increasing Discharge Voltage

All of the high current density work presented in this dissertation was done at 300 V. In the spirit of transparency, this was not entirely voluntary; while we attempted to operate at higher voltages, our harnessing was insufficient to support these voltages and we ended up with significant arcing issues. However, there is good reason to believe that operating at high voltages has benefits for krypton performance beyond just increases in overall power.

One phenomenon that we observed but were unable to thoroughly investigate is the very high oscillations seen in krypton operation at high voltages. At 600 V and 15 A, the peak-to-peak oscillations in discharge current reached over 200%, making operation not only unstable

but also difficult to probe. Increasing the magnetic field helped to damp these oscillations; at 400 V, we saw very oscillatory (>100%) operation at a standard magnetic field strength, but by increasing it by 12.5%, we were able to significantly reduce the oscillation strength (<50%). However, we were unable to reduce the oscillations at 500 and 600 V even with higher magnetic field strengths.

From our high current work at constant voltage, we saw that the magnitudes of current oscillations generally decreased with increasing current (Fig. 6.11). This is true for both xenon and krypton. We then postulate that a high-voltage, high-current operating condition on krypton may have less violent oscillations than those seen at high-voltage, low-current. By studying this operational regime, we may be able to isolate whether the relatively low performance of krypton we saw at high voltages is due to its oscillatory nature.

8.3.0.5 Increasing Discharge Current

While we did reach current densities a factor of ten higher than typical, we believe that we can reach higher currents still. Practically, the two main limitations on our ability to reach higher currents were our power supply (capped at 150 A) and thermal concerns with the thruster. This first issue is easily addressed by switching power supplies or running multiple in parallel. This second concern may be addressed by changing the materials used in the thruster. Indeed, while the water cooling system we implemented helped maintain safe temperatures during testing, this design is not flightlike and we would ideally use passive rather than active cooling. Ongoing efforts at PEPL include using a graphite anode instead of stainless steel, allowing this component to reach higher temperatures, and adding multi-layer insulation between the discharge and the magnetic material of the thruster, preventing the hot environment of the discharge plasma from heating up the sensitive ferromagnetic material of the thruster.

For krypton operation, it would be interesting to see if the overall decrease in anode efficiency continues to drop off slowly or if it reaches a point where it begins to sharply decline. Understanding this behavior, which is primarily dictated by the beam utilization, could help us understand the scaling of anomalous electron transport at higher currents (i.e. if it continues in a Bohm-like fashion or it begins to scale more classically). We may also expect that the peak efficiency of other lower-mass propellants to occur at current densities in excess of what we observed at 150 A on the H9 MUSCLE.

8.3.0.6 Varying Magnetic Field Strength

Throughout this work, we used the same shielded magnetic field topology and selected field strengths that minimized oscillations and mean values of discharge current. However, this sometimes had the effect of obfuscating possible trends in efficiencies or differences between propellants. For example, while we expect the beam utilization to decrease at higher currents, increasing the magnetic field strength may lessen the impact of this. To fully characterize the effects of magnetic field strength of various efficiency modes, a more thorough parametric study is necessary. In particular, we recommend performing the full suite of efficiency measurements varying only the field strength while maintaining constant voltage and current (or flow rate, dealer’s choice).

8.3.0.7 Nested Channel Hall Thrusters

While we briefly discussed nested Hall thrusters in Sec. 3.6 and addressed some of their shortcomings, we believe that they may still provide a path forward to higher-power thruster operation in combination with the higher current densities we tested in our work. While NHTs do involve more complex design than standard Hall thrusters, they also offer some amount of built-in redundancy and throttleability.

The N30, a magnetically shielded Hall thruster designed by Cusson et al. [92, 96], is particularly well-suited for studying the potential benefits of high current density operation in nested Hall thrusters. In addition to its shielded design, its inner channel is sized to the same dimensions as the H9. This means we can directly compare the performance of the N30 with the H9 and, to a lesser extent, the H9 MUSCLE. The N30 is also a good candidate to switch to fully graphite channels like the H9 MUSCLE given the difficulty and risk of machining large BN pieces like those required for its outer channel. There are a number of challenges associated with this work, like managing the thermal environment with two regions of plasma contributing heat to the thruster than than just one, but we still believe it to be a worthwhile endeavor.

8.3.0.8 Predictive Simulations

As shown in Sec. 7.3.2.2, we were able to observe some trends in anomalous transport profile with increasing voltage and current that were consistent between xenon and krypton. This may make it easier to scale simulations with less data in the future. For example, we do not have velocity profiles for the high-current conditions (>30 A) we tested; this would be exceedingly difficult to obtain, given both the sensitivity of optics to high temperatures in the range of what we observed in the chamber at high current densities (Fig. 6.12) as well

as the decreasing population of singly charged ions and the different wavelengths required to target different charge states with LIF (Fig. 6.8, Sec. 4.4.4). However, given that we saw very little change in the anomalous transport profile from 15 to 30 A, we may be able to obtain well-calibrated (in terms of matching performance) simulations of higher currents without significant adjustments to the transport profile.

Additionally, while we limited our analysis of Hall2De’s output to the global performance metrics and plasma parameters along channel centerline, there is plenty of additional information associated with each calibrated simulation. Now that there are hand-tuned anomalous transport profiles for ten total conditions across two propellants and multiple powers (both high voltage and high current), future researchers may perform further investigation into the internal plasma physics may be done. Other available outputs of Hall2De, such as erosion rate predictions and thermal models, may also serve to be useful with planning future test campaigns and even designing new thrusters.

8.4 Concluding Remarks

This work represents an important step forward in proliferating the use of krypton on magnetically shielded Hall thrusters operating at high powers. We have successfully identified a strategy by which we can close the efficiency gap between krypton and xenon: increasing the current density until both propellants reach 100% mass utilization. Additionally, we have found that the plasma density and not the ionization rate coefficient plays an outsized role in the scaling of this efficiency with increasing power. Ultimately, this work expands our ability to operate Hall thrusters on alternative propellants and at atypically high powers.

APPENDIX A

Derivations

Here we derive the phenomenological efficiency model, singly- and multiply-charged limits for specific impulse, and electron velocity from Ohm's law.

A.1 Phenomenological Efficiency Model

Anode efficiency is defined as the following:

$$\eta_a = \frac{T^2}{2\dot{m}_a P_d} = \frac{\eta_{tot}}{\eta_c \eta_p}, \quad (\text{A.1})$$

where T is thrust, \dot{m}_a is neutral flow into the anode, and P_d is discharge power. The cathode flow efficiency is $\eta_c = \frac{\dot{m}_a}{\dot{m}}$ and the magnet power efficiency is $\eta_p = \frac{P_d}{P}$.

Thrust is defined as:

$$T = \dot{m}_b v_{ex} \cos \theta_d, \quad (\text{A.2})$$

where \dot{m}_b is ion mass flow rate in the beam, v_{ex} is exhaust velocity, and θ_d is the plume divergence from channel centerline; only momentum in the axial direction is considered thrust.

Discharge power is defined as

$$P_d = V_d I_d, \quad (\text{A.3})$$

where I_d is discharge current and V_d is discharge voltage.

The beam current is defined as the sum of current contributions from all ion species within

the plume (singly-ionized, doubly-ionized, etc.):

$$I_b = \sum_n I_n, \quad (\text{A.4})$$

where n is a given species. We define a current fraction for each species as

$$\Omega_n = \frac{I_n}{I_b}. \quad (\text{A.5})$$

The beam flow rate is then

$$\dot{m}_b = \frac{m_i}{e} \sum_n \frac{I_n}{Z_n} = \frac{m_i I_b}{e} \sum_n \frac{\Omega_n}{Z_n}. \quad (\text{A.6})$$

where I_n is the current per species, Z_n is the charge per species (1, 2, etc.), e is the fundamental charge, and m_i is propellant/ion mass.

Assuming that all of the electric potential energy in the thruster for a given ion (in terms of charge and acceleration voltage V_a) is converted into kinetic energy (in terms of propellant mass and exhaust velocity), we can use energy conservation to determine an expression for the exhaust velocity:

$$\frac{1}{2} m_i v_{ex}^2 = \sum_n e Z_n V_a \implies v_{ex} = \sqrt{\frac{2eV_a}{m_i}} \sum_n \sqrt{Z_n}. \quad (\text{A.7})$$

Equations A.6 and A.7 can now be plugged back into Eq. A.2, yielding

$$\begin{aligned} T &= \frac{m_i I_b}{e} \sum_n \frac{\Omega_n}{Z_n} \sqrt{\frac{2eV_a}{m_i}} \sum_n \sqrt{Z_n} \cos \theta_d \\ &= \sqrt{\frac{2m_i V_a}{e}} I_b \cos \theta_d \sum_n \frac{\Omega_n}{\sqrt{Z_n}}. \end{aligned} \quad (\text{A.8})$$

We can now define a few efficiencies that will help us separate out individual processes. First, we define current utilization efficiency as the ratio between the beam current and the discharge current:

$$\eta_b = \frac{I_b}{I_d}. \quad (\text{A.9})$$

We can also define mass utilization efficiency as the ratio between the ion mass flow rate in

the beam and the neutral mass flow rate into the anode:

$$\eta_m = \frac{\dot{m}_b}{\dot{m}_a}. \quad (\text{A.10})$$

The discharge power can then be defined as:

$$P_d = \frac{V_d I_b}{\eta_b}, \quad (\text{A.11})$$

and the anode mass flow can be defined as (with Eq. A.6):

$$\dot{m}_a = \frac{\dot{m}_b}{\eta_m} = \frac{m_i I_b}{e} \frac{\sum_n \frac{\Omega_n}{Z_n}}{\eta_m}. \quad (\text{A.12})$$

Plugging Eqs. A.8, A.11, and A.12 back into Eq. A.1 yields the following monster:

$$\begin{aligned} \eta_a &= \frac{\frac{2m_i V_a}{e} I_b^2 \left(\sum_n \frac{\Omega_n}{\sqrt{Z_n}} \right)^2 (\cos \theta_d)^2}{2 \frac{I_b m_i}{e} \frac{1}{\eta_m} \sum_n \frac{\Omega_n}{Z_n} \frac{V_d I_b}{\eta_b}} \\ &= \eta_m \eta_b \frac{V_a}{V_d} \frac{\left(\sum_n \frac{\Omega_n}{\sqrt{Z_n}} \right)^2}{\sum_n \frac{\Omega_n}{Z_n}} (\cos \theta_d)^2. \end{aligned} \quad (\text{A.13})$$

We define a few more efficiency modes; we start with the voltage utilization efficiency, the ratio between the acceleration voltage and the discharge voltage:

$$\eta_v = \frac{V_a}{V_d}. \quad (\text{A.14})$$

Next we define the charge utilization efficiency, which takes into account multiply-charged ions:

$$\eta_q = \frac{\left(\sum_n \frac{\Omega_n}{\sqrt{Z_n}} \right)^2}{\sum_n \frac{\Omega_n}{Z_n}}. \quad (\text{A.15})$$

Lastly we define the plume divergence efficiency, which takes into account the diffusion of the plume off-centerline:

$$\eta_d = (\cos \theta_d)^2. \quad (\text{A.16})$$

The anode efficiency can now be defined purely in terms of phenomenological processes:

$$\eta_a = \eta_m \eta_b \eta_v \eta_q \eta_d. \quad (\text{A.17})$$

A.2 Specific Impulse Limits

Specific impulse is defined

$$I_{sp} = \frac{T}{\dot{m} g_0}, \quad (\text{A.18})$$

where we consider the total mass flow into the thruster (i.e. including cathode flow). First we note that

$$\frac{1}{\dot{m}} = \frac{\eta_c}{\dot{m}_a} = \frac{\eta_c \eta_m}{\dot{m}_b}. \quad (\text{A.19})$$

We then return to our expression for thrust in Eq. A.8 and plug it into our I_{sp} equation with the modification for \dot{m} and with some efficiencies pulled out:

$$I_{sp} = \frac{1}{g_0} \eta_c \eta_m \sqrt{\eta_d \eta_v} \frac{1}{\dot{m}_b} \sqrt{\frac{2m_i V_d}{e}} I_b \sum_n \frac{\Omega_n}{\sqrt{Z_n}}. \quad (\text{A.20})$$

Next, we substitute in our expression for \dot{m}_b from Eq. A.6:

$$I_{sp} = \frac{1}{g_0} \eta_c \eta_m \sqrt{\eta_d \eta_v} \frac{m_i I_b}{e} \frac{1}{\sum_n \frac{\Omega_n}{Z_n}} \sqrt{\frac{2m_i V_d}{e}} I_b \sum_n \frac{\Omega_n}{\sqrt{Z_n}}. \quad (\text{A.21})$$

Simplifying, we arrive at

$$I_{sp,q>1} = \frac{1}{g_0} \eta_c \eta_m \sqrt{\eta_d \eta_v} \frac{\sum_n \frac{\Omega_n}{\sqrt{Z_n}}}{\sum_n \frac{\Omega_n}{Z_n}} \sqrt{\frac{2eV_d}{m_i}}. \quad (\text{A.22})$$

This is our expression for specific impulse based on measurements attainable by probes (versus via the thrust stand). If we assume an entirely singly-charged plume, this expression may be simplified to

$$I_{sp,q=1} = \frac{1}{g_0} \eta_c \eta_m \sqrt{\eta_d \eta_v} \sqrt{\frac{2eV_d}{m_i}}. \quad (\text{A.23})$$

Finally, without any efficiency losses, we can observe the maximum specific impulse for a

monoenergetic beam:

$$I_{sp,max} = \frac{1}{g_0} \eta_c \sqrt{\frac{2eZV_d}{m_i}}, \quad (\text{A.24})$$

where Z is the given charge state.

A.3 Ohm's Law

We first solve for the electron velocity by beginning with Ohm's law:

$$0 = -en_e(\vec{E} + \vec{v}_e \times \vec{B}) - m_e n_e \nu_e \vec{v}_e, \quad (\text{A.25})$$

assuming steady-state and no electron pressure gradient, but including a collision term. We can then break this into its three coordinate directions: azimuthal ($\hat{\theta}$), axial (\hat{z}), and radial (\hat{r}). The electron velocity we care about is in the axial direction, and as will become evident, the radial direction does not actually need to be solved for.

We first solve Ohm's law in the azimuthal direction and note that there is no electric field in this direction. We can solve by invoking our definitions of electron cyclotron frequency and Hall parameter:

$$\begin{aligned} 0 &= -en_e v_{e\theta} B - m_e n_e \nu_e v_{e\theta} \\ v_{e\theta} &= -\frac{eB}{m_e \nu_e} v_{ez} \\ v_{e\theta} &= \frac{\omega_{ce}}{\nu_e} v_{ez} = \Omega_e v_{ez}. \end{aligned} \quad (\text{A.26})$$

We next solve Ohm's law in the axial direction and substitute in Eq. A.26:

$$\begin{aligned} 0 &= -en_e E_z - en_e v_{e\theta} B - m_e n_e \nu_e v_{ez} \\ 0 &= -eE - e \frac{\omega_{ce}}{\nu_e v_{ez} B - m_e \nu_e v_{ez}} \\ v_{ez} &= -\left(\frac{e}{e \frac{\omega_{ce}}{\nu_e} B + m_e \nu_e} \right) E \\ &= -\frac{e}{\Omega_e^2 + 1} E. \end{aligned} \quad (\text{A.27})$$

Noting that $\Omega_e \gg 1$ typically in Hall thrusters, this simplifies to

$$\begin{aligned} v_{ez} &= -\frac{e}{m_e \nu_e \Omega_e^2} E \\ &= -\frac{e \nu_e}{m_e \omega_{ce}^2} \\ &= \frac{\nu_e}{\omega_{ce} B} E = \frac{1}{\Omega_e} \frac{E}{B}. \end{aligned} \tag{A.28}$$

At long last, we have our expression for electron velocity.

APPENDIX B

Tabulated and Auxiliary Data

Here we tabulate data for your viewing pleasure.

B.1 H9 Baseline Data

Baseline xenon and krypton performance comparison (Chapter 5).

V_d (V)	I_d (A)	P_d (kW)	Thrust (mN)	Isp (s)	η_a (%)	η_{tot} (%)	I_d osc. (%)
300	15	4.5	292.9 ± 3.5	1881 ± 35	64.2 ± 1.6	58.6 ± 1.8	45 ± 1
300	20	6	377.6 ± 3.8	1944 ± 29	64.2 ± 1.3	58.9 ± 1.5	34 ± 1
400	15	6	350.8 ± 3.4	2168 ± 36	66.5 ± 1.5	61.0 ± 1.8	98 ± 5
500	15	7.5	405.1 ± 3.4	2395 ± 37	67.9 ± 1.4	62.5 ± 1.7	73*
600	15	9	447.2 ± 3.0	2620 ± 38	68.3 ± 1.3	63.1 ± 1.6	72*

Table B.1: Xenon performance measurements and discharge current oscillations. *Insufficient data available to determine uncertainties.

V_d (V)	I_d (A)	P_d (kW)	Thrust (mN)	Isp (s)	η_a (%)	η_{tot} (%)	I_d osc. (%)
300	15	4.5	235.8 ± 2.5	1900 ± 29	52.2 ± 1.1	47.6 ± 1.2	31 ± 3
300	20	6	317.3 ± 2.0	1974 ± 21	54.8 ± 0.8	50.4 ± 1.0	41 ± 4
400	15	6	269.1 ± 2.5	2235 ± 33	52.6 ± 1.0	47.9 ± 1.2	42 ± 3
500	15	7.5	311.0 ± 2.2	2480 ± 33	53.9 ± 0.9	49.4 ± 1.2	163*
600	15	9	350.3 ± 2.5	2680 ± 35	54.7 ± 0.9	50.3 ± 1.1	207*

Table B.2: Krypton performance measurements and discharge current oscillations. *Insufficient data available to determine uncertainties.

V_d (V)	I_d (A)	P_d (kW)	$\eta_{a,thrust}$ (%)	$\eta_{a,probe}$ (%)	η_b (%)	η_v (%)	η_d (%)	η_q (%)	η_m (%)
300	15	4.5	64.2 ± 1.6	66.6 ± 3.9	86.9 ± 3.0	94.9 ± 1.1	86.7 ± 1.1	96.6 ± 1.9	96.4 ± 3.8
300	20	6	64.2 ± 1.3	66.6 ± 3.6	86.8 ± 2.8	94.6 ± 1.1	85.9 ± 0.7	96.1 ± 2.0	98.2 ± 3.6
400	15	6	66.5 ± 1.5	66.0 ± 4.0	87.1 ± 2.9	94.9 ± 1.0	88.6 ± 1.8	96.1 ± 2.1	93.7 ± 3.6

Table B.3: Contributions to xenon efficiency inferred from probe measurements.

V_d (V)	I_d (A)	P_d (kW)	$\eta_{a,thrust}$ (%)	$\eta_{a,probe}$ (%)	η_b (%)	η_v (%)	η_d (%)	η_q (%)	η_m (%)
300	15	4.5	52.2 ± 0.9	52.9 ± 3.3	82.3 ± 3.1	96.6 ± 1.2	81.6 ± 0.8	98.6 ± 2.3	82.7 ± 3.4
300	20	6	54.8 ± 0.6	58.0 ± 3.4	84.4 ± 2.8	96.3 ± 1.2	83.2 ± 0.1	98.3 ± 2.7	87.3 ± 3.3
400	15	6	52.6 ± 0.8	55.3 ± 3.5	82.2 ± 3.2	97.3 ± 1.1	80.9 ± 0.4	98.8 ± 2.2	86.6 ± 3.7

Table B.4: Contributions to krypton efficiency inferred from probe measurements.

B.2 H9 MUSCLE High Current Data

High current density xenon and krypton performance comparison (Chapter 6).

I_d (A)	P_d (kW)	Thrust (mN)	Isp (s)	η_a (%)	η_{tot} (%)	I_d osc. (%)
15	4.5	280 \pm 15	1754 \pm 98	57.0 \pm 7.0	49.7 \pm 6.2	69 \pm 4
50	15	820 \pm 19	2174 \pm 61	62.1 \pm 3.1	56.9 \pm 3.1	39 \pm 3
75	22.5	1100 \pm 19	2244 \pm 53	57.9 \pm 2.3	53.3 \pm 2.4	22 \pm 3
100	30	1394 \pm 21	2294 \pm 50	55.8 \pm 2.0	51.6 \pm 2.1	6 \pm 1
125	37.5	1650 \pm 30	2309 \pm 56	52.8 \pm 2.0	48.9 \pm 2.1	4 \pm 1

Table B.5: Xenon performance measurements and discharge current oscillations at 300 V.

I_d (A)	P_d (kW)	Thrust (mN)	Isp (s)	η_a (%)	η_{tot} (%)	I_d osc. (%)
15	4.5	211 \pm 14	1765 \pm 120	44.1 \pm 6.0	38.4 \pm 5.3	43 \pm 6
50	15	704 \pm 19	2184 \pm 69	53.8 \pm 2.9	49.2 \pm 2.9	33 \pm 2
75	22.5	1016 \pm 19	2331 \pm 57	55.1 \pm 2.2	50.8 \pm 2.3	24 \pm 2
100	30	1309 \pm 22	2465 \pm 57	56.3 \pm 2.1	52.1 \pm 2.2	34 \pm 2
125	37.5	1582 \pm 16	2534 \pm 48	55.9 \pm 1.7	51.8 \pm 1.9	32 \pm 3
150	45	1839 \pm 18	2567 \pm 48	55.0 \pm 1.6	51.0 \pm 1.9	19 \pm 1

Table B.6: Krypton performance measurements and discharge current oscillations at 300 V.

I_d (A)	P_d (kW)	$\eta_{a,probe}$ (%)	η_b (%)	η_v (%)	η_d (%)	η_q (%)	η_m (%)
15	4.5	49.3 \pm 8.1	73.2 \pm 5.7	95.4 \pm 1.2	85.4 \pm 2.0	98.0 \pm 0.1	84.3 \pm 8.8
50	15	56.6 \pm 10.6	71.0 \pm 9.0	95.6 \pm 2.6	82.3 \pm 1.6	97.0 \pm 0.4	104.4 \pm 13.5
75	22.5	52.0 \pm 3.4	69.1 \pm 2.2	96.1 \pm 1.9	79.4 \pm 1.9	96.3 \pm 0.9	102.4 \pm 4.0
100	30	48.0 \pm 2.3	68.1 \pm 1.0	98.1 \pm 1.6	78.2 \pm 1.6	96.8 \pm 1.5	95.0 \pm 2.5
125	37.5	49.5 \pm 3.8	66.3 \pm 2.2	100.7 \pm 3.8	77.7 \pm 2.2	96.8 \pm 2.0	98.6 \pm 4.0

Table B.7: Xenon efficiencies from probe measurements at 300 V.

I_d (A)	P_d (kW)	$\eta_{a,probe}$ (%)	η_b (%)	η_v (%)	η_d (%)	η_q (%)	η_m (%)
15	4.5	34.9 \pm 6.0	69.2 \pm 5.9	95.4 \pm 1.2	79.0 \pm 2.0	97.8 \pm 0.1	68.3 \pm 7.5
50	15	48.1 \pm 3.2	70.8 \pm 2.2	96.7 \pm 1.7	80.1 \pm 1.6	98.1 \pm 0.2	89.4 \pm 3.6
75	22.5	48.8 \pm 2.6	70.7 \pm 1.4	94.9 \pm 2.0	79.0 \pm 1.6	97.6 \pm 0.4	94.3 \pm 2.8
100	30	51.8 \pm 5.1	69.3 \pm 1.8	94.4 \pm 7.7	79.2 \pm 2.0	97.7 \pm 0.6	102.3 \pm 3.3
125	37.5	51.9 \pm 5.2	69.0 \pm 3.9	91.1 \pm 4.0	79.3 \pm 1.9	97.7 \pm 1.0	106.5 \pm 6.5
150	45	51.2 \pm 3.9	68.6 \pm 2.8	95.4 \pm 1.8	78.1 \pm 2.4	97.1 \pm 1.7	103.1 \pm 4.7

Table B.8: Krypton efficiencies from probe measurements at 300 V.

B.3 Hall2De Output Data

In this section we present Hall2De output data that did not make the cut.

B.3.1 Simulated Component Efficiencies

In this study, we employ the phenomenological efficiency model described in Sec. 3.3.2. The version of Hall2De employed for this study calculates each of these component efficiencies as the simulation runs by integrating plasma quantities of interest over a control volume. In this study, we choose a control volume defined by the last closed magnetic field line in the domain. This gives the best agreement with experiment and represents a natural point at which the ability of the near plume to couple magnetically back to the thruster body is diminished. In Fig. 7.3b, we indicate the integrating surface in the context of the Hall2De domain.

The mass utilization (Eq. 3.9) is computed by integrating the net flux of ions of all species and charge states across the bounding magnetic field line. In Hall2De, we compute the ion beam current (Eq. 3.7) by integrating the ion current density over the previously-defined integration surface. We obtain the acceleration voltage for voltage efficiency (Eq. 3.10) in Hall2De by recording the velocity v_i at which singly-charged ions along the discharge channel centerline cross the integration surface. The acceleration voltage can then be computed as

$$V_a = \frac{m_i v_i^2}{2e}, \quad (\text{B.1})$$

where e is the fundamental charge and m_i is the mass of the ion species in question. We note here in calculating the divergence efficiency (Eq. 3.11) that the version of Hall2De used throughout this work did not include an internal calculation of axial current. We therefore implemented a rectilinear control volume for calculating divergence along the outermost axial and radial borders of the domain as shown in Fig. 7.3b. Finally, the charge utilization (Eq. 3.8) was determined by summing the flux of these species across the integration surface.

To better understand the drivers of trends in anode efficiency with increasing power, we also analyze the output component efficiencies from Hall2De as calculated with the method outlined above. We note that the vertical scale for anode efficiency in these plots ranges from 0.5 to 0.8, while the scale for component efficiencies ranges from 0.7 to 1.0 for ease of comparison. The circular markers in gray shown on the plot of anode efficiency indicate the anode efficiency calculated as the product of the component efficiencies (Eq. 3.6). While the thrust-calculated anode efficiency is scaled by $1/0.92^2$, the product anode efficiency is not scaled.

The behavior of the mass utilization (explored in Sec. 7.4) explains why the efficiency gap closes between the two propellants, but not why the xenon anode efficiency remains nearly constant with increasing current. This latter trend can be explained by analyzing the beam current efficiency. While the beam efficiency of krypton remains approximately the same throughout the range of currents, varying by less than 1.5%, the beam efficiency of xenon monotonically decreases by 4% a total of from 15 to 30 A (Fig. B.1b). Since all the other efficiency modes remain relatively unchanged with increasing current, the slight increase in mass utilization is nullified by the decrease in beam utilization, resulting in an overall anode efficiency for xenon that remains nearly constant with increasing current. The other noteworthy trend seen in our simulated efficiencies that is reflected in experimental data is the different dependence of beam current utilization with increasing current. We saw in our previous work on the H9 MUSCLE thruster that the beam utilization of xenon follows more of a classical scaling with increasing current density, i.e. the beam utilization decreases as current increases (Sec. 6.4.2.2). This classical scaling stems from the fact that with additional collisions, we would expect the electron-ion collisions to increase. However, the behavior of krypton was shown to be much more Bohm-like, i.e. independent of changes in current. We note that the lack of decrease in krypton’s beam utilization is even more notable given the lower magnetic field strength used at the 20 and 30 A conditions (Tab. 7.1). Typically, we would expect that a lower magnetic field means weaker electron confinement and therefore more electrons escaping to the anode, manifesting as a loss in beam utilization. Krypton’s near-constant beam utilization seen in Fig. B.1b is indicative that neither the increasing density nor the lower magnetic field have a strong impact on the behavior of krypton’s beam utilization.

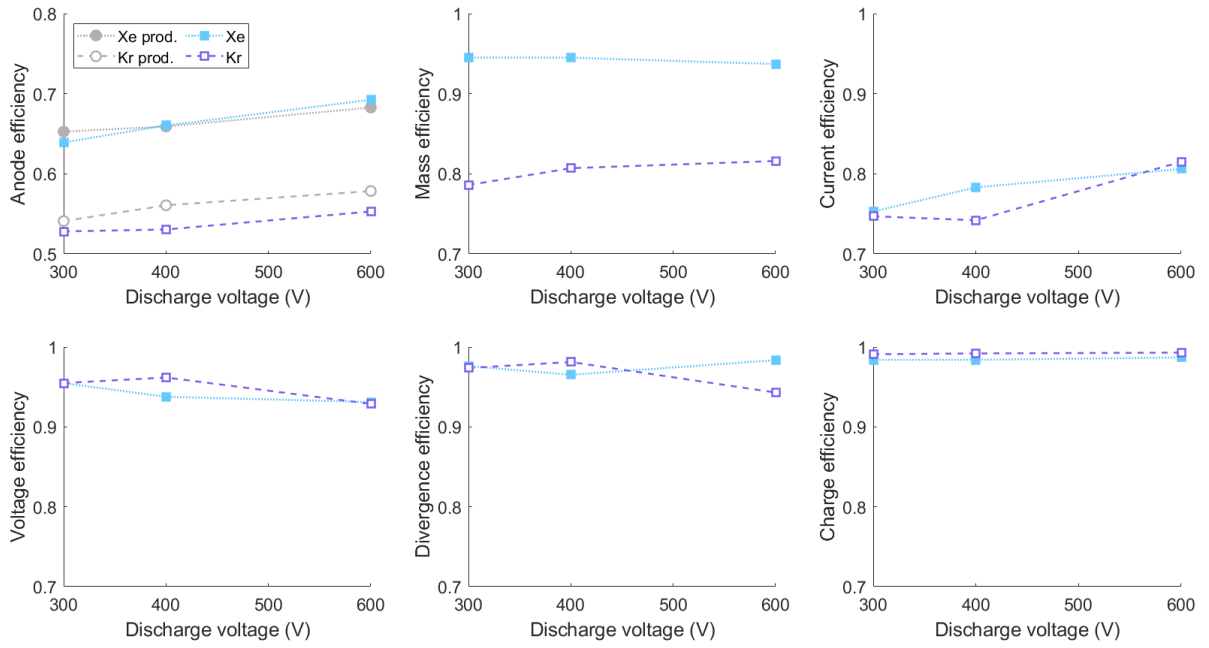
It is apparent that the stronger inverse relationship between beam utilization and current density for xenon is what causes the overall anode efficiency to stay constant from 15 to 30 A—the decrease in beam utilization nullifies the increase in mass utilization. Indeed, at even higher currents, we expect xenon’s efficiency begin to decrease (as it does for the H9 MUSCLE shown in Sec. 6.3.1) as our mass utilization cannot increase past unity. While the physical mechanisms underlying the different responses of beam utilization in each propellant to current remains unknown, this trend is favorable towards closing the gap between xenon and krypton efficiency. We do in fact see the gap between propellants in simulated anode efficiency close from 11% at 15 A to 4% at 30 A.

We present all efficiency components in Fig. B.1. The simulated values of mass utilization, voltage utilization, and charge utilization efficiencies match nearly within uncertainty of simulated data at all conditions. The beam utilization is consistently lower than measured experimentally by 7–13%, and the divergence efficiency is consistently higher than measured

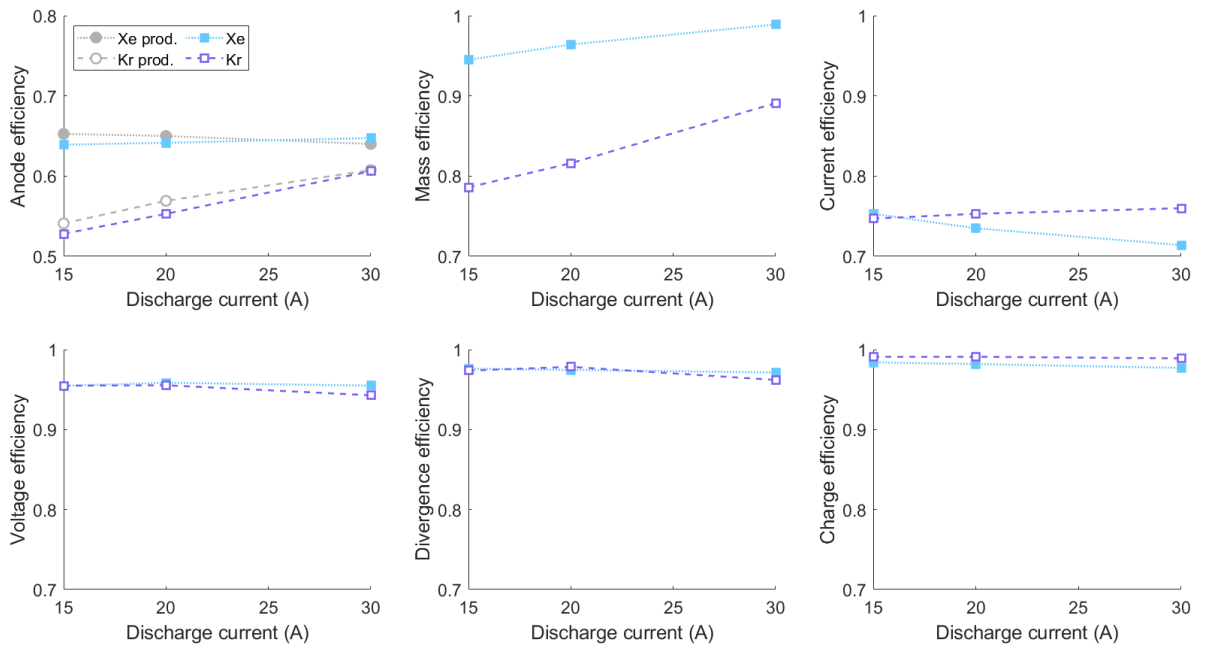
experimentally by 8–17%. The low beam utilization is likely due to Hall2De overestimating the electron current, possibly the result of shortcomings in our ability to model anomalous electron current. The high divergence efficiency—correlated with low divergence angle—may be because Hall2De under-predicts the amount of ions born at low potentials downstream of the exit plane due to the fluid nature of the code, leading to a smaller population of high-radial velocity particles. That said, it is also possible that the experimental data is not entirely accurate; as has been noted in multiple other works, isolating precise values of beam current (a necessary value for evaluating mass efficiency) is difficult due to the effect of charge exchange in testing facilities [120, 122, 123, 130]. Due to these discrepancies, our focus is on comparing the trends in the component efficiencies rather than the values themselves. This is also true for comparing the two methods of calculating anode efficiency—the unscaled anode efficiency as calculated with Eq. 3.6 should be analyzed only for its trends with increasing power rather than the magnitude of the values.

We note that all the simulated trends in all efficiency modes excepting divergence angle match the experimentally observed trends (i.e. increasing, decreasing, or unchanging) between 300 and 400 V with constant discharge current (c.f. Fig. B.1a and Tab. 7.3). While the simulated values of divergence efficiency trend opposite that of experimental, the magnitudes of change for both propellants in both the experimental and simulated cases are under 2%. We therefore cannot draw conclusive trends regarding the change in efficiency between these two conditions. Between 400 and 600 V, we see a slight increase in the overall anode efficiency for both propellants but minimal change in the gap between them.

Similarly, the trends between 15 and 20 A with constant discharge voltage match with what has been observed experimentally with the exception of voltage utilization (c.f. Fig. B.1b and Tab. 7.3). The change in voltage efficiency between these two conditions is less than 0.5% in all cases, again minimizing the significance of these “opposing” trends. The only efficiency mode for which there is a statistically significant change between 15 and 20 A is the mass utilization for krypton, which increases by 4.6% experimentally and 3% in simulation. The overall anode efficiency is relatively constant for xenon (less than 1% variation) and monotonically increases for krypton by a total of $\sim 8\%$ in the 15 to 30 A range.



(a)



(b)

Figure B.1: Simulated component efficiencies of xenon and krypton operation as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

B.3.2 1D Profiles for All Conditions




Figure B.2: Axial profile of electron temperature along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

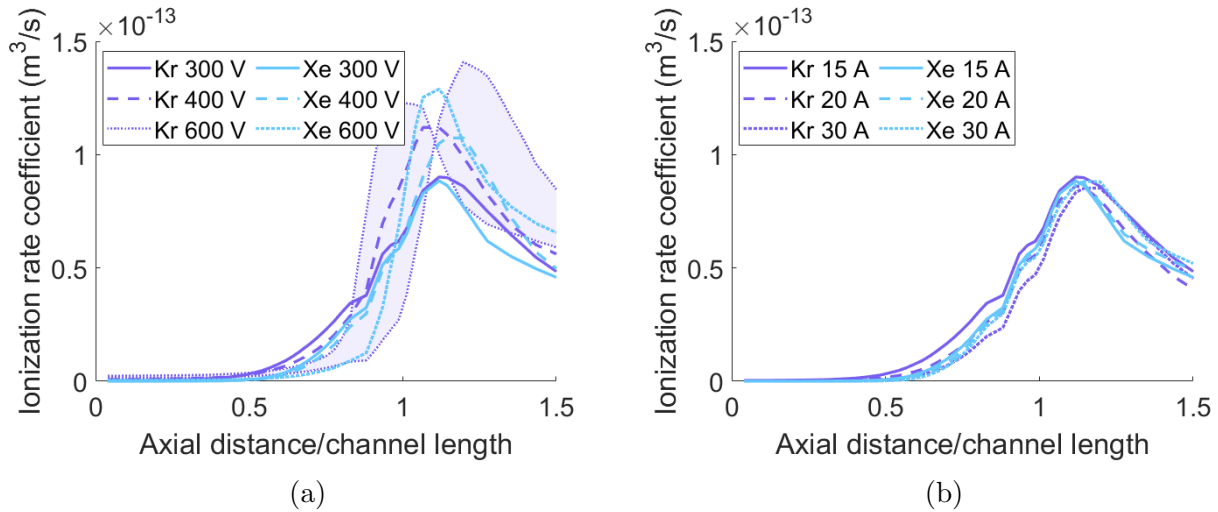


Figure B.3: Axial profile of ionization rate coefficient along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.




Figure B.4: Axial profile of electron density along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.




Figure B.5: Axial profile of neutral velocity along channel centerline for xenon and krypton operation as measured experimentally and as found through simulation for a) increasing discharge voltage with a constant discharge current of 15 A, and b) increasing discharge current with a constant discharge voltage of 300 V.

BIBLIOGRAPHY

- [1] Space Nuclear Propulsion for Human Mars Exploration. Technical report, National Academies of Sciences, Engineering, and Medicine, Washington, DC, 2021.
- [2] Next Space Technologies for Exploration Partnerships NNH15ZCQ001K. Technical report, NASA, Nov 2014.
- [3] R. Myers and C. Carpenter. High power solar electric propulsion for human space exploration architectures. In *32nd International Electric Propulsion Conference*, Wiesbaden, Germany, 2011.
- [4] R.M. Myers, C.R. Joyner, R.J. Cassady, S. Overton, T. Kokan, H. Horton, and W.A. Hoskins. Affordable exploration architectures using the space launch system and high power solar electric propulsion. In *34th International Electric Propulsion Conference*, Kobe, Japan, 2015.
- [5] How rare-gas supply adapted to Russia’s war, Mar 2023. URL <https://www.economist.com/finance-and-economics/2023/03/30/how-rare-gas-supply-adapted-to-russias-war>.
- [6] A. I. Bugrova, A. M. Bishaev, A. V. Desyatskov, M. V. Kozintseva, A. S. Lipatov, and M. Dudeck. Experimental investigations of a krypton stationary plasma thruster. *Int. J. Aerospace Eng.*, page 686132, 2013. ISSN 16875966. doi: 10.1155/2013/686132.
- [7] Jacek Kurzyna, Maciej Jakubczak, Agnieszka Szelecka, and Käthe Dannenmayer. Performance tests of IPPLM’s krypton Hall thruster. *Laser Part. Beams*, 36(1):105–114, Mar 2018. ISSN 1469803X. doi: 10.1017/S0263034618000046.
- [8] Richard R Hofer, Peter Y Peterson, David T Jacobson, and David M Manzella. Factors Affecting the Efficiency of Krypton Hall Thrusters. In *Meeting of the APS Division of Plasma Physics*, Savannah, GA, 2004.
- [9] Jesse A. Linnell and Alec D. Gallimore. Efficiency analysis of a Hall thruster operating with krypton and xenon. *J. Propul. Power*, 22(6):1402–1418, Nov 2006. doi: 10.2514/1.19613.
- [10] Hani Kamhawi, Thomas Haag, David Jacobson, and David Manzella. Performance Evaluation of the NASA-300M 20 kW Hall Thruster. In *Joint Propulsion Conference and Exhibit*, San Diego, CA, 2011. doi: 10.2514/6.2011-5521.
- [11] Peter Peterson, David Jacobson, David Manzella, and Jeremy John. The Performance and Wear Characterization of a High-Power High-Isp NASA Hall Thruster. In *Joint Propulsion Conference and Exhibit*, Tucson, AZ, 2005. doi: 10.2514/6.2005-4243.
- [12] Colleen Marrese, Alec D. Gallimore, James Haas, John Foster, Brad King, Sang Wook Kim, and Sergei Khartov. An investigation of stationary plasma thruster performance with krypton propellant. In *Joint Propulsion Conference and Exhibit*, San Diego, CA,

1995. ISBN 9780000000002. doi: 10.2514/6.1995-2932.
- [13] David T. Jacobson and David H. Manzella. 50 kW Class Krypton Hall Thruster Performance. In *Joint Propulsion Conference and Exhibit*, Huntsville, Alabama, 2003.
- [14] Ioannis G. Mikellides, Ira Katz, Richard R. Hofer, and Dan M. Goebel. Magnetic shielding of a laboratory Hall thruster. I. theory and validation. *J. Appl. Phys.*, 115(4):043303, 2014. doi: 10.1063/1.4862313.
- [15] Richard R. Hofer, Dan M. Goebel, Ioannis G. Mikellides, and Ira Katz. Magnetic shielding of a laboratory Hall thruster. II. experiments. *J. Appl. Phys.*, 115(4):043304, 2014. doi: 10.1063/1.4862314.
- [16] Randall Munroe. Space mission hearing. <https://xkcd.com/2124/>.
- [17] Francis F Chen. *Introduction to Plasma Physics and Controlled Fusion*, volume 1. Springer, New York, NY, 1984. doi: 10.1007/978-3-319-22309-4.
- [18] Dan M Goebel and Ira Katz. *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*. Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 2008.
- [19] Thomas A. Marks. *Modeling anomalous electron transport in a fluid Hall thruster code*. PhD thesis, University of Michigan, 2023.
- [20] D Patrascu. Gen2 Starlink Internet Satellites to Hit Orbit Powered by New Argon Hall Thrusters. <https://www.autoevolution.com/news/gen2-starlink-internet-satellites-to-hit-orbit-powered-by-new-argon-hall-thrusters-211005.html>, Feb 2023.
- [21] David Y Oh, Steve Collins, Tracy Drain, William Hart, Travis Imken, Kristina Larson, Danielle Marsh, Dhack Muthulingam, John Steven Snyder, Denis Trofimov, et al. Development of the Psyche mission for NASA’s discovery program. In *International Electric Propulsion Conference*, number 192, Vienna, Austria, 2019. Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space
- [22] John S Snyder, Vernon H Chaplin, Dan M Goebel, Richard R Hofer, Alejandro Lopez Ortega, Ioannis G Mikellides, Taylor Kerl, Giovanni Lenguito, Faraz Aghazadeh, and Ian Johnson. Electric propulsion for the psyche mission: Development activities and status. In *AIAA Propulsion and Energy 2020 Forum*, page 3607, 2020.
- [23] 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.
- [24] Jean Pierre Boeuf. Tutorial: Physics and modeling of Hall thrusters. *J. Appl. Phys.*, 121(1):011101, Jan 2017. ISSN 10897550. doi: 10.1063/1.4972269.
- [25] Jason D. Frieman, Thomas M. Liu, and Mitchell L. R. Walker. Background flow model of Hall thruster neutral ingestion. *J. Propul. Power*, 33(5):1087–1101, Jan 2017. ISSN 0748-4658. doi: 10.2514/1.b36269.
- [26] Scott J. Hall. *Characterization of a 100-kW Class Nested-Channel Hall Thruster*. PhD thesis, University of Michigan, 2018.
- [27] Richard R Hofer and Alec D Gallimore. High-specific impulse Hall thrusters, part 2: Efficiency analysis. *J. Propul. Power*, 22(4):732–740, 2006. doi: 10.2514/1.15954.
- [28] Richard Hofer, Ira Katz, Dan Goebel, Kristina Jameson, Regina Sullivan, Lee Johnson, and Ioannis Mikellides. Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations. In *44th Joint Propulsion Conference and Exhibit*, page 4924,

- Hartford, CT, 2008. AIAA. doi: 10.2514/6.2008-4924.
- [29] Richard R Hofer. *Development and Characterization of High-Efficiency, High-Specific Impulse Xenon Hall Thrusters*. PhD thesis, University of Michigan, 2004.
- [30] Ira Katz and Ioannis G Mikellides. Neutral gas free molecular flow algorithm including ionization and walls for use in plasma simulations. *J. Comput. Phys.*, 230(4):1454–1464, 2011. doi: 10.1016/j.jcp.2010.11.013.
- [31] Ruojian Pan, Junxue Ren, Renfan Mao, and Haibin Tang. Practical analysis of different neutral algorithms for particle simulation of Hall thruster. *Plasma Sources Sci. T.*, 32(3):034005, 2023. doi: 10.1088/1361-6595/acc134.
- [32] Horatiu C Dragnea, Alejandro Lopez Ortega, Hani Kamhawi, and Iain D Boyd. Simulation of a Hall effect thruster using krypton propellant. *Journal of Propulsion and Power*, 36(3):335–345, 2020.
- [33] R Rejoub, BG Lindsay, and RF Stebbings. Determination of the absolute partial and total cross sections for electron-impact ionization of the rare gases. *Physical Review A*, 65(4):042713, 2002.
- [34] Robert C Wetzell, Frank A Baiocchi, Todd R Hayes, and Robert S Freund. Absolute cross sections for electron-impact ionization of the rare-gas atoms by the fast-neutral-beam method. *Physical Review A*, 35(2):559, 1987.
- [35] AI Morozov, A Ya Kislov, and IP Zubkov. A high-current plasma accelerator with closed electron drift. *Pis' ma Zh. Tekh. Fiz.*, 7:224, 1968.
- [36] EY Choueiri. Fundamental difference between the two hall thruster variants. *Physics of Plasmas*, 8(11):5025–5033, 2001.
- [37] S Mazouffre, S Tsikata, and J Vaudolon. Development and experimental characterization of a wall-less Hall thruster. *Journal of Applied Physics*, 116(24), 2014.
- [38] Jacob Simmonds and Yevgeny Raitses. Ion acceleration in a wall-less Hall thruster. *Journal of Applied Physics*, 130(9), 2021.
- [39] Käthe Dannenmayer and Stéphane Mazouffre. Elementary scaling relations for Hall effect thrusters. *J. Propul. Power*, 27(1):236–245, 2011. doi: 10.2514/1.48382.
- [40] Kristi de Grys, Alex Mathers, Ben Welander, and Vadim Khayms. Demonstration of 10,400 hours of operation on 4.5 kW qualification model Hall thruster. In *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 6698, 2010.
- [41] Ioannis G Mikellides, Ira Katz, Richard R Hofer, and Dan M Goebel. Magnetic shielding of walls from the unmagnetized ion beam in a Hall thruster. *Applied Physics Letters*, 102(2), 2013.
- [42] Richard R Hofer, Sarah E Cusson, Robert B Lobbia, and Alec D Gallimore. The H9 magnetically shielded Hall thruster. In *35th International Electric Propulsion Conference*, number 232, Atlanta, GA, 2017. Electric Rocket Propulsion Society.
- [43] Sarah E Cusson, Richard R Hofer, Robert Lobbia, Benjamin A Jorns, and Alec D Gallimore. Performance of the H9 magnetically shielded Hall thrusters. In *35th International Electric Propulsion Conference*, Atlanta, Georgia, 2017. Electric Rocket Propulsion Society.
- [44] 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.
- [45] Jerry Jackson, Joseph Cassady, May Allen, Roger Myers, Todd A Tofil, Daniel A

- Herman, and Eric J Pencil. Development of high power Hall thruster systems to enable the NASA exploration vision. Technical report, 2018.
- [46] 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. In *Joint Propulsion Conference and Exhibit*, 2016. doi: 10.2514/6.2016-5027.
- [47] George J Williams, Hani Kamhawi, Maria Choi, Thomas W Haag, Wensheng Huang, Daniel A Herman, James H Gilland, and Peter Y Peterson. Wear trends of the HER-MeS thruster as a function of throttle point. In *International Electric Propulsion Conference*, Atlanta, GA, 2019.
- [48] Nasa, aerojet rocketdyne put gateway thruster system to the test. <https://www.nasa.gov/feature/glenn/2023/nasa-aerojet-rocketdyne-put-gateway-thruster-system-to-the-test>, July 2023.
- [49] Richard Hofer, Dan Goebel, Ioannis Mikellides, and Ira Katz. Design of a Laboratory Hall Thruster with Magnetically Shielded Channel Walls, Phase II: Experiments. In *48th Joint Propulsion Conference and Exhibit*, Atlanta, GA, Jul 2012. American Institute of Aeronautics and Astronautics.
- [50] Richard R Hofer, James E Polk, Michael J Sekerak, Ioannis G Mikellides, Hani Kamhawi, Timothy R Sarver-Verhey, Daniel A Herman, and George Williams. The 12.5 kw hall effect rocket with magnetic shielding (hermes) for the asteroid redirect robotic mission. In *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, page 4825, Salt Lake City, UT, 2016.
- [51] Richard R Hofer, Benjamin A Jorns, Ira Katz, and John R Brophy. Hall effect thruster electrical configuration, Nov 2019. US Patent US-10480493-B2.
- [52] Dan M. Goebel, Richard R. Hofer, Ioannis G. Mikellides, Ira Katz, James E. Polk, and Brandon N. Dotson. Conducting wall Hall thrusters. *IEEE T. Plasma Sci.*, 43(1): 118–126, Jan 2015. doi: 10.1109/TPS.2014.2321110.
- [53] Lou Grimaud and Stéphane Mazouffre. Conducting wall Hall thrusters in magnetic shielding and standard configurations. *Journal of Applied Physics*, 122(3):033305, 2017. doi: 10.1063/1.4995285.
- [54] Richard R Hofer, Robert Lobbia, and Steven Arestie. Performance of a conducting wall, magnetically shielded Hall thruster at 3000-s specific impulse. In *37th International Electric Propulsion Conference*, number 401, Boston, MA, 2022. Electric Rocket Propulsion Society.
- [55] Nicolas Gascon, Scott Crawford, Ronald Corey, and Mark Cappelli. Coaxial Hall thruster with diamond inner channel wall. In *42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 4995, 2006.
- [56] James Szabo, Bruce Pote, Surjeet Paintal, Mike Robin, Adam Hillier, Richard D Branam, and Richard E Huffmann. Performance evaluation of an iodine-vapor Hall thruster. *J. Propul. Power*, 28(4):848–857, 2012. doi: 10.2514/1.B34291.
- [57] David R Lide. *CRC handbook of chemistry and physics*, volume 85. CRC press, 2004.
- [58] Richard P Welle. Xenon and krypton availability for electric propulsion - An updated assessment. In *29th Joint Propulsion Conference and Exhibit*, El Segundo, CA, Jun 1993. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.1993-2401.
- [59] Timothy Kokan and C. Russell Joyner. Mission comparison of Hall effect and gridded

- ion thrusters utilizing various propellant options. In *AIAA SPACE Conference and Exposition*, 2012. doi: 10.2514/6.2012-5237.
- [60] Daniel A Herman and Kenneth G Unfried. Xenon Acquisition Strategies for High Power Electric Propulsion NASA Missions. In *Joint Army Navy NASA Air Force*, Nashville, TN, 2015.
- [61] V. Giannetti, T. Andreussi, A. Leporini, S. Gregucci, M. Saravia, A. Rossodivita, M. Andrenucci, D. Estublier, and C. Edwards. Electric Propulsion System Trade-Off Analysis Based on Alternative Propellant Selection. In *Space Propulsion Conference*, number 1, Rome, Italy, 2016.
- [62] Adam Shabshelowitz, Alec D Gallimore, and Peter Y Peterson. Performance of a helicon Hall thruster operating with xenon, argon, and nitrogen. *Journal of Propulsion and Power*, 30(3):664–671, 2014.
- [63] Leonid Pekker and Michael Keidar. Analysis of airbreathing Hall-effect thrusters. *Journal of propulsion and power*, 28(6):1399–1405, 2012.
- [64] Francesco Marchioni and Mark A Cappelli. Extended channel Hall thruster for air-breathing electric propulsion. *J. Appl. Phys.*, 130(5):053306, 2021. doi: 10.1063/5.0048283.
- [65] Anita Sengupta, Colleen Marrese-Reading, Mark Capelli, David Scharfe, Sergey Tverdokhlebov, Sasha Semenkin, Oleg Tverdokhlebov, Ian Boyd, Michael Keidar, Azer Yalin, Vijaya Surla, Tom Markusic, and Kurt Polzin. An overview of the VHITAL program: a two-stage bismuth fed very high specific impulse thruster with anode layer. In *29th International Electric Propulsion Conference*, number 238, Princeton, NJ, 2005. Electric Rocket Propulsion Society.
- [66] S D Grishin, V S Erofeev, A V Zharinov, V P Naumkin, and I V Safronov. Characteristics of a two-stage ion accelerator with an anode layer. *J. Appl. Mech. Tech. Phy.*, 19(2):166–173, 1978.
- [67] James Szabo, Mike Robin, and Vlad Hruby. Bismuth vapor Hall effect thruster performance and plume experiments. In *Proceedings of the the 35th International Electric Propulsion Conference*, pages 2017–25, 2017.
- [68] Dean Massey, Lyon King, and Jason Makela. Development of a direct evaporation bismuth Hall thruster. In *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 4520, 2008.
- [69] John Dankanich, Kurt A Polzin, and Hani Kamhawi. Iodine Hall thruster demonstration mission concept and development. In *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, page 3910, 2014.
- [70] Wired. May 30, 1898: Krypton discovered, decades before superman arrives. <https://www.wired.com/2008/05/may-30-1898-krypton-discovered-decades-before-superman-arrives/>.
- [71] Leanne L Su and Benjamin A Jorns. Performance comparison of a 9-kW magnetically shielded Hall thruster operating on xenon and krypton. *J. Appl. Phys.*, 130(16):163306, 2021. doi: 10.1063/5.0066849.
- [72] Leanne L Su, Parker J Roberts, Tate Gill, William Hurley, Thomas A Marks, Christopher L Sercel, Madison Allen, Collin B Whittaker, Eric Viges, and Benjamin A. Jorns. High current density performance of a magnetically shielded Hall thruster. *J. Propul. Power*, 2023. In review.

- [73] Y Raitses, A Smirnov, D Staack, and NJ Fisch. Measurements of secondary electron emission effects in the Hall thruster discharge. *Physics of Plasmas*, 13(1), 2006. doi: 10.1063/1.2162809.
- [74] Michael Patterson. A new space enterprise of exploration. Technical report, NASA, 2010.
- [75] Mitchell L R Walker, Dan Lev, Maryam Saeedifard, Benjamin Jorns, John Foster, Alec D Gallimore, Alex Gorodetsky, Joshua L Rovey, Huck Beng Chew, Deborah Levin, John D Williams, Azer Yalin, Richard Wirz, Jaime Marian, Iain Boyd, Kentaro Hara, and Kristina Lemmer. Overview of the Joint Advanced Propulsion Institute (JANUS). In *International Electric Propulsion Conference*, number 156, Boston, MA, 2022.
- [76] Kevin D Diamant, Raymond Liang, and Ronald L Corey. The effect of background pressure on SPT-100 Hall thruster performance. In *50th Joint Propulsion Conference and Exhibit*, page 3710, Cleveland, OH, 2014. AIAA. doi: 10.2514/6.2014-3710.
- [77] 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. *J. Propul. Power*, 33(3):539–555, 2017. doi: 10.2514/1.B35564.
- [78] John W. Dankanich, Bob Vondra, and Andrew V. Ilin. Fast transits to Mars using electric propulsion. In *46th Joint Propulsion Conference and Exhibit*, page 6771, Nashville, TN, 2010. AIAA. doi: 10.2514/6.2010-6771.
- [79] Matthew E Duchek, William Machemer, Christopher Harnack, Matteo A Clark, Alejandro R Pensado, Kelsa B Palomares, Kurt A Polzin, Adam Martin, Frank Curran, Roger Myers, Mitchell Rodriguez, Dasari Rao, Rodger Dyson, and Robert Scheidegger. Key performance parameters for MW-class NEP elements and their interfaces. In *ASCEND 2022*, number 4273, page 4273, Las Vegas, NV, Oct 2022. AIAA. doi: 10.2514/6.2022-4273.
- [80] Robert Jankovsky, Sergey Tverdokhlebov, and David Manzella. High power Hall thrusters. In *Joint Propulsion Conference and Exhibit*, number 2949, Los Angeles, CA, 1999. doi: 10.2514/6.1999-2949.
- [81] B Pote, V Hruby, and J Monheiser. Performance of an 8 kW Hall thruster. Technical report, Busek, Natick, MA, Jan 2000.
- [82] Hani Kamhawi, Thomas W Haag, Timothy Smith, Daniel Herman, Wensheng Huang, Rohit Shastry, Peter Peterson, and Alex Mathers. Performance characterization of the air force transformational satellite 12 kW Hall thruster. In *JANNAF*, Colorado Springs, CO, 2013.
- [83] James Szabo, Bruce Pote, Vald Hruby, Lawrence Byrne, Rachel Tadrake, George Kolencik, Hani Kamhawi, and Thomas Haag. A commercial one newton Hall effect thruster for high power in-space missions. In *Joint Propulsion Conference and Exhibit*, number 6152, San Diego, CA, 2011. doi: 10.2514/6.2011-6152.
- [84] George Soulas, Thomas Haag, Daniel Herman, Wensheng Huang, Hani Kamhawi, and Rohit Shastry. Performance test results of the NASA-457M v2 Hall thruster. In *Joint Propulsion Conference and Exhibit*, number 3940, Atlanta, GA, 2012. doi: 10.2514/6.2012-3940.
- [85] Lee Mason, Robert Jankovsky, and David Manzella. 1000 hours of testing on a 10 kilowatt Hall effect thruster. In *Joint Propulsion Conference and Exhibit*, number

- 3773, 2001. doi: 10.2514/6.2001-3773.
- [86] Scott J Hall, Benjamin A Jorns, Sarah E Cusson, Alec D Gallimore, Hani Kamhawi, Peter Y Peterson, Thomas W Haag, Jonathan A Mackey, Matthew J Baird, and James H Gilland. Performance and high-speed characterization of a 100-kW nested hall thruster. *J. Propul. Power*, 38(1):40–50, 2022. doi: 10.2514/1.B38080.
- [87] David H Manzella, Charles Sarmiento, John Sankovich, and Tom Haag. Performance evaluation of the SPT-140. In *International Electric Propulsion Conference*, number 059, Cleveland, OH, 1997.
- [88] A I Morozov and V V Savelyev. Fundamentals of stationary plasma thruster theory. *R. Plasma Phys.*, 21:203–391, Mar Mar 2000. doi: {10.1007/978-1-4615-4309-1.2}.
- [89] Jacob Simmonds, Yevgeny Raitses, and Andrei Smolyakov. A theoretical thrust density limit for Hall thrusters. *J. Elec. Propul.*, 2(12), 2023. doi: 10.1007/s44205-023-00048-9.
- [90] Mitchell L.R. Walker and Alec D. Gallimore. Hall thruster cluster operation with a shared cathode. *Journal of Propulsion and Power*, 23(3):528–536, may 2007.
- [91] Mitchell L.R. Walker and Alec D. Gallimore. Performance characteristics of a cluster of 5-kW laboratory hall thrusters. *Journal of Propulsion and Power*, 23(1):35–43, jan 2007.
- [92] Sarah E Cusson, Richard R Hofer, Dan M Goebel, Marcel P Georgin, Alexander R Vazsonyi, Benjamin A Jorns, Alec D Gallimore, and Iain D Boyd. Development of a 30-kW Class Magnetically Shielded Nested Hall Thruster. In *36th International Electric Propulsion Conference*, Vienna, Austria, 2019.
- [93] Raymond Liang. *The Combination of Two Concentric Discharge Channels into a Nested Hall-Effect Thruster*. PhD thesis, University of Michigan, Ann Arbor, MI, 2013.
- [94] S.J. Hall, R.E. Florenz, A.D. Gallimore, H. Kamhawi, D.L. Brown, J.E. Polk, D.M. Goebel, and R.R. Hofer. Implementation and initial validation of a 100-kW class nested-channel hall thruster. In *Joint Propulsion Conference, AIAA-2014-3815*, Cleveland, OH, 2014. American Institute of Aeronautics and Astronautics.
- [95] Scott J. Hall, Benjamin Jorns, Alec Gallimore, and Richard R. Hofer. Expanded Thruster Mass Model Incorporating Nested Hall Thrusters. In *Joint Propulsion Conference and Exhibit*, number 4729, Atlanta, GA, Jul 2017. doi: 10.2514/6.2017-4729.
- [96] Sarah Cusson. *Impact of neutral density on the operation of high-power magnetically shielded Hall thrusters*. PhD thesis, University of Michigan, 2019.
- [97] R. Florenz, A.D. Gallimore, and P.Y. Peterson. Developmental status of a 100-kW class laboratory nested channel hall thruster. In *32nd International Electric Propulsion Conference*, pages 2011–246, Wiesbaden, Germany, 2011.
- [98] S. J. Hall, B. A. Jorns, A. D. Gallimore, H. Kamhawi, T. W. Haag, J. A. Mackey, J. H. Gilland, P. Y. Peterson, and M. J. Baird. High-power performance of a 100-kW class nested Hall thruster. In *35th International Electric Propulsion Conference*, pages 2017–228, Atlanta, GA, 2017.
- [99] S.E. Cusson, S.J. Hall, R.R. Hofer, B.A. Jorns, and A.D. Gallimore. The impact of magnetic field coupling between channels in a nested Hall thruster. In *35th International Electric Propulsion Conference*, pages 2017–507, Atlanta, Georgia, 2017.
- [100] C. Moore, E. Pencil, R. Hardy, K. Bollweg, and M. Ching. NASA’s NextSTEP advanced electric propulsion activities. In *AIAA Propulsion and Energy Forum*, july

- 2018.
- [101] B.A. Jorns, A.D. Gallimore, S.J. Hall, P.Y. Peterson, J.E. Gilland, D.M. Goebel, R.R. Hofer, and I. Mikellides. Update on the nested Hall thruster subsystem for the NextSTEP XR-100 program. In *Joint Propulsion Conference, AIAA-2018-4418*, Cincinnati, OH, 2018. American Institute of Aeronautics and Astronautics.
 - [102] Sarah W.H. Shark, Scott J. Hall, Benjamin A. Jorns, Richard R. Hofer, and Dan M. Goebel. High power demonstration of a 100 kW nested Hall thruster system. In *AIAA Propulsion and Energy 2019 Forum*, Indianapolis, IN, 2019.
 - [103] Leanne L Su, Scott J Hall, Sarah E Cusson, and Benjamin A Jorns. Model for the increase in thruster efficiency from cross-channel coupling in nested hall thrusters. In *International Electric Propulsion Conference*, number IEPC-2019-204, 2019.
 - [104] Leanne L. Su, Alexander R. Vazsonyi, and Benjamin Jorns. Performance of a 9-kw magnetically-shielded Hall thruster with krypton. In *AIAA Propulsion and Energy 2020 Forum*. American Institute of Aeronautics and Astronautics, Aug 2020.
 - [105] 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.
 - [106] IP Zubkov, A Ya Kislov, and AI Morozov. Experimental study of a two-lens accelerator. *Sov. Phys.-Tech. Phys.(Engl. Transl.)*, 15(11):1796–1800, 1971.
 - [107] R. R. Hofer, D. M. Goebel, and R. M. Walker. Compact Lab6 Hollow Cathode for a 6-kW Hall Thruster. In *54th JANNAF Propulsion Meeting*, pages 14–17, 2018.
 - [108] Leanne L Su and Benjamin A Jorns. Performance at High Current Densities of a Magnetically-Shielded Hall Thruster. In *AIAA Propulsion and Energy 2021 Forum*, page 3405, Virtual, 2021. AIAA. doi: 10.2514/6.2021-3405.
 - [109] Dan M Goebel, Ron M Watkins, and Kristina K Jameson. LaB6 hollow cathodes for ion and Hall thrusters. *Journal of Propulsion and Power*, 23(3):552–558, 2007.
 - [110] Marcel P Georgin, Benjamin A Jorns, and Alec D Gallimore. Transient non-classical transport in the hollow cathode plume I: measurements of time-varying electron collision frequency. *Plasma Sources Sci. Technol.*, 29:13, 2020. doi: 10.1088/1361-6595/ab b0ce.
 - [111] Eric A Vigés, Benjamin A Jorns, Alec D Gallimore, and J P Sheehan. University of Michigan’s upgraded Large Vacuum Test Facility. In *36th International Electric Propulsion Conference*, Vienna, Austria, 2019. Electric Rocket Propulsion Society.
 - [112] Chuck Garner et al. Methods for cryopumping xenon. In *Joint Propulsion Conference*, number 3206, 1996.
 - [113] 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. *J. Propul. Power*, 33(3):668–680, 2017. doi: 10.2514/1.B35478.
 - [114] Chih-Shun Lu and Owen Lewis. Investigation of film-thickness determination by oscillating quartz resonators with large mass load. *Journal of Applied Physics*, 43(11): 4385–4390, 1972. doi: 10.1063/1.1660931.
 - [115] Kunning G. Xu and Mitchell L.R. Walker. High-Power, Null-Type, Inverted Pendulum Thrust Stand. *Rev. Sci. Instrum.*, 80(5):055103, 2009. ISSN 00346748. doi: 10.1063/1.3125626.

- [116] Jon Mackey, Scott J Hall, Thomas Haag, Peter Y Peterson, and Hani Kamhawi. Uncertainty in inverted pendulum thrust measurements. In *2018 Joint Propulsion Conference*, page 4516, 2018.
- [117] Robert B. Lobbia and Brian E. Beal. Recommended practice for use of Langmuir probes in electric propulsion testing. *J. Propul. Power*, 33(3):566–581, 2017. ISSN 15333876. doi: 10.2514/1.B35531.
- [118] J. P. Sheehan, Yevgeny Raitses, Noah Hershkowitz, and Michael McDonald. Recommended practice for use of emissive probes in electric propulsion testing. *J. Propul. Power*, 33(3):614–637, 2017.
- [119] Wensheng Huang, Rohit Shastry, George Soulas, and Hani Kamhawi. Farfield plume measurement and analysis on the NASA-300M and NASA-300MS. In *33rd International Electric Propulsion Conference*, Washington, D.C., 2013. Electric Rocket Propulsion Society.
- [120] 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. *J. Propul. Power*, 33(3):582–613, 2017. doi: 10.2514/1.B35696.
- [121] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. CRC Press, Boca Raton, FL, 1994.
- [122] Rohit Shastry, Richard R. Hofer, Bryan M. Reid, and Alec D. Gallimore. Method for analyzing $E \times B$ probe spectra from Hall thruster plumes. *Rev. Sci. Instrum.*, 80(6):063502, 2009. doi: 10.1063/1.3152218.
- [123] Wensheng Huang and Rohit Shastry. Analysis of Wien Filter Spectra from Hall Thruster Plumes. *Review of Scientific Instrumentation*, 86(7):073502, Jul 2015. ISSN 10897623. doi: 10.1063/1.4923282.
- [124] J. Scott Miller, Steve H. Pullins, Dale J. Levandier, Yu Hui Chiu, and Rainer A. Dressler. Xenon charge exchange cross sections for electrostatic thruster models. *J. Appl. Phys.*, 91(3):984–991, Feb 2002. ISSN 00218979. doi: 10.1063/1.1426246.
- [125] Michael L. Hause, Benjamin D. Prince, and Raymond J. Bemish. Krypton charge exchange cross sections for Hall effect thruster models. *J. Appl. Phys.*, 113(16):16330, Apr 2013. ISSN 00218979. doi: 10.1063/1.4802432.
- [126] Daniel L. Brown and Alec D. Gallimore. Evaluation of facility effects on ion migration in a Hall thruster plume. *J. Propul. Power*, 27(3):573–585, Apr 2011. doi: 10.2514/1.B34068.
- [127] Homer D Hagstrum. Auger ejection of electrons from molybdenum by noble gas ions. *Phys. Rev.*, 104(3):672, Jul 1956. doi: 10.1103/PhysRev.104.672.
- [128] Homer D Hagstrum. Auger ejection of electrons from tungsten by noble gas ions. *Phys. Rev.*, 96(2):325, Oct 1954. doi: 10.1103/PhysRev.96.325.
- [129] James M Haas. *Low-Perturbation Interrogation of the Internal and Near-Field Plasma Structure of a Hall Thruster Using a High-Speed Probe Positioning System*. PhD thesis, University of Michigan, Ann Arbor, 2001.
- [130] Bryan Michael Reid. *The Influence of Neutral Flow Rate in the Operation of Hall Thrusters*. PhD thesis, University of Michigan, 2009.
- [131] Richard R Hofer and Alec D Gallimore. Recent Results from Internal and Very-Near-Field Plasma Diagnostics of a High Specific Impulse Hall Thruster. In *28th International Electric Propulsion Conference*, Toulouse, France, 2003.

- [132] Wensheng Huang, Hani Kamhawi, and Thomas Haag. Facility Effect Characterization Test of NASA’s HERMeS Hall Thruster. In *52nd Joint Propulsion Conference and Exhibit*, page 4828, Salt Lake City, UT, 2016. doi: 10.2514/6.2016-4828.
- [133] Stéphane Mazouffre. Laser-induced fluorescence diagnostics of the cross-field discharge of Hall thrusters. *Plasma Sources Science and Technology*, 22(1):013001, 2012.
- [134] Ethan Dale. *Investigation of the Hall thruster breathing mode*. PhD thesis, University of Michigan, 2020.
- [135] R J Cedolin, W A Hargus, P V Storm, R K Hanson, and M A Cappelli. Laser-induced fluorescence study of a xenon Hall thruster. *Applied Physics B: Lasers and Optics*, 65: 459–469, 1997.
- [136] W. A. Hargus and M. A. Cappelli. Laser-induced fluorescence measurements of velocity within a Hall discharge. *Applied Physics B: Lasers and Optics*, 72(8):961–969, 2001. doi: 10.1007/s003400100589.
- [137] CJ Durot, AD Gallimore, and TB Smith. Validation and evaluation of a novel time-resolved laser-induced fluorescence technique. *Review of Scientific Instruments*, 85(1), 2014.
- [138] Vernon H. Chaplin, Benjamin A. Jorns, Alejandro Lopez Ortega, Ioannis G. Mikelides, Ryan W. Conversano, Robert B. Lobbia, and Richard R. Hofer. Laser-induced fluorescence measurements of acceleration zone scaling in the 12.5 kW HERMeS Hall thruster. *Journal of Applied Physics*, 124(18), nov 2018. doi: 10.1063/1.5040388.
- [139] A. Kramida, Yu. Ralchenko, J. Reader, and and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.9), [Online]. Available: <https://physics.nist.gov/asd> [2022, April 5]. National Institute of Standards and Technology, Gaithersburg, MD., 2021.
- [140] Leanne L. Su, Thomas A. Marks, and Benjamin A. Jorns. Investigation into the efficiency gap between krypton and xenon operation on a magnetically shielded Hall thruster. Boston, MA, 2022. Electric Rocket Propulsion Society.
- [141] William A Hargus, Gregory M Azarnia, and Michael R Nakles. Demonstration of Laser-Induced Fluorescence on a Krypton Hall Effect Thruster. In *32nd International Electric Propulsion Conference*, Wiesbaden, Germany, 2011.
- [142] Timothy B Smith. *Deconvolution of ion velocity distributions from laser-induced fluorescence spectra of xenon electrostatic thruster plumes*. University of Michigan, 2003.
- [143] Wensheng Huang, Brittany Drenkow, and Alec Gallimore. Laser-induced fluorescence of singly-charged xenon inside a 6-kw hall thruster. In *45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 5355, 2009.
- [144] Benjamin Jorns, Christopher A Dodson, John R Anderson, Dan M Goebel, Richard R Hofer, Michael J Sekerak, Alejandro Lopez Ortega, and Ioannis G Mikellides. Mechanisms for pole piece erosion in a 6-kw magnetically-shielded hall thruster. In *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, page 4839, 2016.
- [145] Wensheng Huang, Alec D Gallimore, and Richard R Hofer. Neutral flow evolution in a six-kilowatt hall thruster. *Journal of Propulsion and Power*, 27(3):553–563, 2011.
- [146] Ethan T Dale and Benjamin A Jorns. Two-zone Hall thruster breathing mode mechanism, Part II: Experiment. In *36th International Electric Propulsion Conference*, Vienna, Austria, 2019.
- [147] R. R. Hofer and A. D. Gallimore. Efficiency analysis of a high-specific impulse Hall thruster. In *40th Joint Propulsion Conference*, Ft. Lauderdale, FL, 2004. American

- Institute of Aeronautics and Astronautics. doi: 10.2514/6.2004-3602.
- [148] Benjamin A Jorns, Matthew Byrne, Parker Roberts, Leanne Su, Ethan Dale, and Richard R Hofer. Prediction and mitigation of the mode transition in a magnetically shielded Hall thruster at high-specific impulse and low power. In *International Electric Propulsion Conference*, Boston, MA, 2022. Electric Rocket Propulsion Society.
 - [149] Daniel L. Brown, C. William Larson, Brian E. Beal, and Alec D. Gallimore. Methodology and historical perspective of a Hall thruster efficiency analysis. *J. Propul. Power*, 25(6):1163–1177, Nov 2009. doi: 10.2514/1.38092.
 - [150] M. C. Bordage, S. F. Biagi, L. L. Alves, K. Bartschat, S. Chowdhury, L. C. Pitchford, G. J.M. Hagelaar, W. L. Morgan, V. Puech, and O. Zatsarinny. Comparisons of sets of electron-neutral scattering cross sections and swarm parameters in noble gases: III. Krypton and xenon. *J. Phys. D*, 46(33):334003, Aug 2013. doi: 10.1088/0022-3727/46/33/334003.
 - [151] Sarah E. Cusson, Ethan T. Dale, Benjamin A. Jorns, and Alec D. Gallimore. Acceleration region dynamics in a magnetically shielded Hall thruster. *Phys. Plasmas*, 26(2): 023506, Feb 2019. ISSN 10897674. doi: 10.1063/1.5079414.
 - [152] Jesse A. Linnell and Alec D. Gallimore. Internal plasma potential measurements of a Hall thruster using xenon and krypton propellant. *Physics of Plasmas*, 13(9), 2006. doi: 10.1063/1.2335820.
 - [153] James A R Samson and R B Cairns. Ionization Potential of Molecular Xenon and Krypton. *J. Opt. Soc. Am.*, 56(8):1140–1141, 1966.
 - [154] 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. ISSN 15333876.
 - [155] Sarah E. Cusson, Matthew P. Byrne, Benjamin A. Jorns, and Alec D. Gallimore. Investigation into the use of cathode flow fraction to mitigate pressure-related facility effects on a magnetically shielded Hall thruster. In *AIAA Propulsion and Energy 2019 Forum*. American Institute of Aeronautics and Astronautics, 2019. ISBN 9781624105906.
 - [156] Leanne L Su, Parker J Roberts, Tate Gill, William Hurley, Thomas A Marks, Christopher L Sercel, Madison Allen, Collin B Whittaker, Matthew P. Byrne, Zachariah B. Brown, Eric Vigas, and Benjamin A. Jorns. Operation and performance of a magnetically shielded Hall thruster at ultrahigh current densities on xenon and krypton. In *AIAA SciTech 2023 Forum*, page 0842, National Harbor, MD, 2023. AIAA. doi: 10.2514/6.2023-0842.
 - [157] Benjamin A Jorns and Matthew P Byrne. Model for the dependence of cathode voltage in a Hall thruster on facility pressure. *Plasma Sources Science and Technology*, 30(1), 2021. doi: 10.1088/1361-6595/abd3b6.
 - [158] Vladimir Kim. Main physical features and processes determining the performance of stationary plasma thrusters. *J. Propul. Power*, 14(5):736–743, 1998. doi: 10.2514/2.5335.
 - [159] Burak Karadag, Shinatora Cho, and Ikkoh Funaki. Thrust performance, propellant ionization, and thruster erosion of an external discharge plasma thruster. *J. Appl. Phys.*, 123(15):153302, 2018. doi: 10.1063/1.5023829.
 - [160] Kristina K Jameson, Dan M Goebel, Richard R Hofer, and Ron M Watkins. Cathode coupling in hall thrusters. In *International Electric Propulsion Conference*, number

- 278, Florence, Italy, 2007.
- [161] T Randolph, V Kim, H Kaufman, K Kozubsky, V Zhurin, and M Day. Facility effects on stationary plasma thruster testing. In *International Electric Propulsion Conference*, number 093, Seattle, WA, 1993.
 - [162] Holak Kim, Youbong Lim, Wonho Choe, Sangho Park, and Jongho Seon. Effect of magnetic field configuration on the multiply charged ion and plume characteristics in Hall thruster plasmas. *Appl. Phys. Lett.*, 106(15):154103, 2015. doi: 10.1063/1.4918654.
 - [163] AV Zharinov. Acceleration of plasma by a closed Hall current. *Sov. Phys.-Tech. Phys.*, 12:208–211, 1967.
 - [164] Y Raitses, D Staack, A Smirnov, and NJ Fisch. Space charge saturated sheath regime and electron temperature saturation in Hall thrusters. *Phys. Plasmas*, 12(7):073507, 2005. doi: 10.1063/1.1944328.
 - [165] D Staack, Y Raitses, and NJ Fisch. Temperature gradient in Hall thrusters. *Appl. Phys. Lett.*, 84(16):3028–3030, 2004. doi: 10.1063/1.1710732.
 - [166] Bryan Reid and Alec Gallimore. Plasma potential measurements in the discharge channel of a 6-kW Hall thruster. In *44th Joint Propulsion Conference and Exhibit*, page 5185. AIAA, 2008. doi: 10.2514/6.2008-5185.
 - [167] Thomas A Marks and Benjamin A Jorns. Challenges with the self-consistent implementation of closure models for anomalous electron transport in fluid simulations of Hall thrusters. *Plasma Sources Sci. T.*, 32(4):045016, 2023. doi: 10.1088/1361-6595/accd18.
 - [168] Michael R Nakles and William A Hargus Jr. Background pressure effects on ion velocity distribution within a medium-power Hall thruster. *J. Propul. Power*, 27(4):737–743, 2011. doi: 10.2514/1.48027.
 - [169] Natalia MacDonald-Tenenbaum, Quinn Pratt, Michael Nakles, Nickolas Pilgram, Michael Holmes, and William Hargus Jr. Background pressure effects on ion velocity distributions in an SPT-100 Hall thruster. *J. Propul. Power*, 35(2):403–412, 2019. doi: 10.2514/1.B37133.
 - [170] Electric Propulsion and Plasma Dynamics Laboratory. AF-MPD Thruster Database. <https://alfven.princeton.edu/tools/afmpdt-database/>, 2022.
 - [171] Richard R. Hofer and Thomas M. Randolph. Mass and cost model for selecting thruster size in electric propulsion systems. *J. Propul. and Power*, 29(1):166–177, 2013. doi: 10.2514/1.B34525.
 - [172] 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*, 2023. In review.
 - [173] Ioannis G. Mikellides and Ira Katz. Numerical simulations of Hall-effect plasma accelerators on a magnetic-field-aligned mesh. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 86(4), Oct 2012.
 - [174] Benjamin A. Jorns, Dan M. Goebel, and Richard R. Hofer. Plasma perturbations in high-speed probing of Hall thruster discharge chambers: Quantification and mitigation. In *51st AIAA/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics Inc, AIAA, 2015. ISBN 9781624103216. doi: 10.2514/6.2015-4006.
 - [175] Ioannis G. Mikellides and Alejandro Lopez Ortega. Challenges in the development

- and verification of first-principles models in Hall-effect thruster simulations that are based on anomalous resistivity and generalized Ohm's law. *Plasma Sources Science and Technology*, 28(1), jan 2019. ISSN 13616595. doi: 10.1088/1361-6595/aae63b.
- [176] Alejandro Lopez Ortega and Ioannis G Mikellides. The importance of the cathode plume and its interactions with the ion beam in numerical simulations of Hall thrusters. *Physics of Plasmas*, 23(4), 2016.
- [177] D. Gawron, S. Mazouffre, N. Sadeghi, and A. Héron. Influence of magnetic field and discharge voltage on the acceleration layer features in a Hall effect thruster. *Plasma Sources Science and Technology*, 17(2):0–10, 2008. doi: 10.1088/0963-0252/17/2/025001.
- [178] 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*, 204:389–401, 2023.
- [179] T Andreussi, MM Saravia, and M Andrenucci. Plasma characterization in hall thrusters by langmuir probes. *Journal of Instrumentation*, 14(05):C05011, 2019.
- [180] Michael J Sekerak, Richard R Hofer, James E Polk, Benjamin A Jorns, and Ioannis G Mikellides. Wear testing of a magnetically shielded hall thruster at 2000 s specific impulse. In *International Electric Propulsion Conference*, 2015.
- [181] Alejandro Lopez Ortega, Ioannis G Mikellides, Michael J Sekerak, and Benjamin A Jorns. Plasma simulations in 2-d (rz) geometry for the assessment of pole erosion in a magnetically shielded hall thruster. *Journal of Applied Physics*, 125(3), 2019.
- [182] W. J. Hurley, L. L. Su, P. J. Roberts, M. G. Allen, T. M. Gill, C. L. Sercel, T. A. Marks, C. B. Whittaker, E. A. Vigas, and B. A. Jorns. Efficiency analysis of a magnetically shielded Hall thruster with increasing current density and multiple propellants. 2023. In review.