

## University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Munro-O'Brien, Thomas F. 2026 "Alternative Propellant Performance in Hall Effect Thrusters", University of Southampton, Astronautical Engineering, PhD Thesis, pagination.

Data: Munro-O'Brien, Thomas F. 2026 "Alternative Propellant Performance in Hall Effect Thrusters", DOI: <https://doi.org/10.5258/SOTON/PG/D196>



**UNIVERSITY OF SOUTHAMPTON**

Faculty of Engineering and Physical Sciences  
Department of Aeronautical and Astronautical Engineering

**Alternative Propellant Performance  
in Hall Effect Thrusters**

DOI: [10.5258/soton/pg/t196](https://doi.org/10.5258/soton/pg/t196)

*by*

**Thomas Francis Munro-O'Brien**

MEng Aeronautical & Astronautical Engineering:  
Spacecraft Engineering

ORCID: [0000-0002-8387-3143](https://orcid.org/0000-0002-8387-3143)

*A thesis for the degree of  
Doctor of Philosophy*

May 2026



University of Southampton

Abstract

Faculty of Engineering and Physical Sciences  
Department of Aeronautical and Astronautical Engineering

Doctor of Philosophy

**Alternative Propellant Performance  
in Hall Effect Thrusters**

by Thomas Francis Munro-O'Brien

This thesis investigates the scaling, design, and performance of a novel modular high-voltage Hall effect thruster operating on alternative propellants, with the aim of enhancing understanding of alternative propellant performance and reducing mission costs compared to conventional xenon-propelled systems. As the electric propulsion community seeks to reduce reliance on xenon, a propellant marred by scarcity and volatile price, lighter alternatives such as krypton and argon offer potential cost-saving and advantages in specific impulse without the xenon-associated supply chain constraints. However, in traditional Hall thrusters this potential increase in specific impulse is often offset by lower propellant utilisation, stemming from their smaller ionisation cross-sections and higher ionisation energies. Furthermore, while prior work has established empirical scaling methods for xenon-fed thrusters, the lack of comprehensive performance data for alternative propellants has hindered their optimisation.

To address this gap, two key contributions are presented: the development of an extended semi-empirical scaling methodology that preserves propellant-specific properties, and the experimental characterisation of a modular discharge channel thruster designed to investigate the influence of channel geometry, namely mean channel diameter, and channel width, on alternative propellant performance.

The scaling framework builds upon semi-empirical methodologies by Dannermayer et al. and Kim et al., extending previous work through preservation of propellant-dependent metrics and enabling direct scaling for krypton Hall effect thrusters. To validate this scaling methodology, a novel modular Hall thruster was designed, featuring a discharge channel with adjustable width and mean diameter while maintaining constant magnetic, electrical, diagnostic, and fluidic configurations. This design enabled a systematic comparison of xenon, krypton, and argon across nine distinct channel geometries at discharge voltages up to 600 V. The final thruster, scaled using this new method for krypton at 2.5 kW and 600 V, was designated the

Southampton High-voltage Anode-layer Research Krypton 600 V “SHARK-600V”. An optional thruster-with-anode-layer configuration was also tested as a potential method for enhancing alternative propellant performance, though full characterisation was not possible.

The experimental results from the University of Southampton campaign reveal clear propellant-specific trends: xenon exhibits efficiency degradation at high voltages, krypton shows similar degradation only for certain geometries, and argon demonstrates reduced performance sensitivity to geometry, with greater stability sensitivity to geometry across the range of tested flow rates and voltages. At optimised conditions, krypton achieved a specific impulse of 2809 s and an anode efficiency of 35.8 %, exceeding that of xenon of (2422 s, 34.7 %) but with slightly reduced thrust (109.3 mN vs. 115 mN). Argon attained comparable specific impulse (2537 s) but with reduced thrust (81.7 mN) and anode efficiency (24.8 %).

To further utilise this dataset, a Gaussian process regression model was trained on this experimental dataset, creating a continuous surrogate-model of thruster performance across the input parameter space. The model enabled interpolation between tested geometries and extrapolation to untested conditions, all with quantified uncertainty. Crucially, the optimal geometries identified by the GPR model showed meaningful agreement with the semi-empirical scaling predictions, providing experimental validation of the scaling approach developed for this thesis. For argon, where conventional scaling methods cannot be applied due to insufficient literature data, the GPR model provided a novel experimentally derived optimal channel geometry, offering preliminary design guidance for future argon thruster development.

Cross-facility validation at the University of Michigan on krypton further strengthened these findings through direct thrust measurements and far-field plume diagnostics using a retarding potential analyser,  $E \times B$  probe, and a Faraday probe. Direct thrust measurements for krypton identified a peak performance at 50 sccm and 600 V, with a specific impulse of 3511.2 s, thrust of 107.3 mN, and an anode efficiency of 58.1 %. At 60 sccm and 500 V, direct thrust measurements yielded 2901.5 s, 106.4 mN, and 50.6 % ( $\pm 2.24$  %) efficiency, while probe-derived results closely matched, reporting 50.7 % anode efficiency comprised of 97.1 % charge efficiency, 70.2 % divergence efficiency, 88.4 % current efficiency, 95.0 % voltage efficiency, and 88.5 % mass utilisation efficiency.

Testing at the University of Michigan revealed consistently higher performance in terms of anode efficiency, thrust, and specific impulse. These results from direct thrust measurements are corroborated by the far-field probe measurements, showing good agreement between the direct thrust anode efficiency and the probe-derived anode efficiency. These cross-facility results both validate the core geometry and highlight the high-performance capability of high-voltage Hall thrusters operating on krypton.

Importantly, although absolute performance at Michigan exceeded that at Southampton for the same configuration, both facilities reached peak anode efficiency at the same operating point: 50 sccm, 600 V. The enhanced performance is attributed to the thruster operating at a consistently lower discharge current with a corresponding slight increase in thrust output. This results in a lower power draw for a slightly enhanced performance at the Michigan test site compared to the University of Southampton test site.



## Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

Munro-O'Brien, Thomas F., and Charles N. Ryan. "Performance of a low power Hall effect thruster with several gaseous propellants." *Acta Astronautica* 206 (2023): 257-273.

Munro-O'Brien, Thomas F., and Charles N. Ryan. "Performance Comparison Between Modular SPT and TAL Type Hall Thrusters Operating on Krypton." *International Electric Propulsion Conference, Toulouse, France. 38th Edition. 2024.*

Munro-O'Brien, Thomas, and Charles Ryan. "High-Power Hall effect Thruster Activity at The University of Southampton." *APS Annual Gaseous Electronics Meeting Abstracts. 2024.*

Thomas F Munro-O'Brien, Mohamed Ahmed, Andrea Lucca Fabris, Charles N Ryan. "Inter-Laboratory Characterisation of a Low-Power Channel-Less Hall-Effect Thruster: Performance Comparisons and Lessons Learnt." *Aerospace* 2025, 12(7), 601; <https://doi.org/10.3390/aerospace12070601> . 1 July 2025.

Munro-O'Brien, Thomas, and Charles Ryan. "Effect of channel width on the performance of a modular Hall effect thruster operating on Xenon, Krypton, and

Argon." International Electric Propulsion Conference, London, United Kingdom. 39th Edition. 2025.

Signed:.....

Date:.....

## Acknowledgements

First of all, I would like to thank Professor Charlie Ryan for taking me on as an individual project student as a lowly undergraduate student, which sparked my passion for electric propulsion and for pushing me to grow as a researcher. I am also grateful to him for providing me constant opportunities throughout my time at Southampton. Such as when I was given unprecedented access to the large vacuum chamber laboratory and allowed me to build invaluable skills over the plague summer. It is not an overstatement to say that without his continued and unwavering support, I would not be the researcher I am today. I would also like to thank my fellow researchers from the Southampton laboratory: Szymon Dworski, Chengyu Ma, Vlad-George Trilia, Nazli Turin, and Euan Donovan-Hill, who not only taught me so much about engineering but also patiently listened to me complain about work in the pub. I would also like to thank Yaseen Reza for keeping me company through several houses, putting up with my antics, and always reminding me that while my PhD might be rough at times, at least I was not doing a computational project; for which I am immensely grateful.

I must also thank the many people who have collaborated with me over the years and helped shape numerous aspects of this project. At the University of Surrey, I am grateful to Mohamed Ahmed and Andrea Lucca Fabris for their support in cross-facility collaborations, which were pivotal for validating the thrust balance. At the University of Michigan, I thank the entire PEPL team, who made me feel at home for whilst on the other side of the world. Thank you for letting this little British man drink and entertain you during my time in Ann Arbor. Furthermore, I am indebted to Professor Ben Jorns for hosting me in his laboratory and for fostering such a welcoming team of talented engineers. I would be remiss not to mention Dr Leanne Su, who has been a source of friendship, kindness, and whose work has been of personal inspiration.

Moreover, I would like to thank all those who directly supported the efforts covered in this thesis. In particular, I owe much to the talented machinists within the EDMC. Specifically, to Wes Tucker: thank you for listening to every ambitious request (and occasional terrible idea) I brought to you over the past few years. I would also like to thank two outstanding technicians, Gus Gillam and Simon Klitz, for the support and guidance you have provided over several years. Without your help, and that of the University of Southampton's technicians and engineers as a whole, the achievements of this work would not have been possible.

To those in my life who have supported me throughout this journey. I would like to thank my two amazing sisters, Hannah, and Caitlin for reminding me that while I might be studying rockets, I will never be as cool as them, and for being a constant source of support and care throughout the past four years of this degree. To my

parents, who have been endlessly supportive and forgiving when I forget to call for weeks on end, as my poor time management always seems to leave me in some kind of time crunch, thank you. To my friends who have helped me countless times and number too many to list here, I am forever grateful. To cousin Terry, for inspiring me to be passionate about science from a young age: I look forward to joining the many Dr O'Briens who have come before me.

*“Electrons are barely even real.”*  
– Terrance Francis O’Brien, Ph.D. (2019)



# Contents

<b>Declaration of Authorship</b>	<b>vii</b>
<b>Acknowledgements</b>	<b>ix</b>
<b>List of Figures</b>	<b>xix</b>
<b>List of Tables</b>	<b>xxxi</b>
<b>Abbreviations</b>	<b>xxxv</b>
<b>Nomenclature</b>	<b>xxxix</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Problem Statement . . . . .	3
1.1.1 Alternative propellants . . . . .	4
1.1.2 High Specific Impulse . . . . .	7
1.2 Research Objectives and Contributions . . . . .	9
1.3 Organisation . . . . .	11
<b>2 Background and Fundamentals</b>	<b>13</b>
2.1 Introduction . . . . .	13
2.1.1 Basic Rocket Science . . . . .	13
2.1.2 Chemical Space Propulsion . . . . .	15
2.1.3 Electric Space Propulsion . . . . .	16
2.1.3.1 Gridded Ion Thrusters . . . . .	17
2.1.3.2 Hall Effect Thrusters . . . . .	18
2.2 Physics of Hall Effect Thruster Operation . . . . .	19
2.3 Evaluation of Hall Effect Thruster Performance . . . . .	20
2.3.1 Thrust . . . . .	20
2.3.2 Specific Impulse . . . . .	21
2.3.3 Thruster Efficiency . . . . .	22
2.3.3.1 Anode Efficiency . . . . .	23
2.3.3.2 Current Efficiency . . . . .	23
2.3.3.3 Charge Efficiency . . . . .	23
2.3.3.4 Mass Efficiency . . . . .	24
2.3.3.5 Voltage Efficiency . . . . .	24
2.3.3.6 Divergence Efficiency . . . . .	24
2.3.4 Thrust-to-Power Ratio . . . . .	25

2.4	High Voltage Operation . . . . .	25
2.5	Thruster with Anode-layer Operation . . . . .	26
2.6	Alternative Propellant Hall Effect Thrusters . . . . .	27
2.6.1	Gaseous Alternatives . . . . .	28
2.6.2	Molecular and Condensable Alternatives . . . . .	29
2.6.3	Metallic Alternatives . . . . .	30
2.7	Alternative Propellant Operation . . . . .	31
2.7.1	Mass . . . . .	31
2.7.2	Ionisation . . . . .	34
2.7.3	Cross-Sectional Area of Ionisation . . . . .	36
2.7.4	Reaction Rate Coefficients . . . . .	39
2.8	Sheath-Wall Potential in the Presence of Secondary Electron Emission . .	41
2.9	Summary of Background and Fundamentals . . . . .	46
<b>3</b>	<b>Scaling Methodology</b>	<b>49</b>
3.1	Introduction . . . . .	49
3.2	Review of Existing Scaling Literature . . . . .	49
3.3	Hall Effect Thruster Database . . . . .	53
3.3.1	Channel Width and Mean Diameter Trends . . . . .	54
3.4	Derivation of Scaling Methodology Employed . . . . .	56
3.4.1	Electron Confinement . . . . .	56
3.4.2	Neutral Density . . . . .	58
3.4.3	Channel Length . . . . .	60
3.4.4	Specific Impulse . . . . .	61
3.4.5	Thrust . . . . .	62
3.4.6	Anode Power . . . . .	63
3.4.7	Anode Current . . . . .	67
3.4.8	Channel Width . . . . .	67
3.5	Application of Scaling Method for Alternative Propellant High-Voltage Operation . . . . .	69
3.5.1	Introduction . . . . .	69
3.5.2	Scaling Methodology . . . . .	70
3.6	Scaling Results . . . . .	71
3.7	Summary of Scaling Methodology . . . . .	73
<b>4</b>	<b>SHARK-600V Hall Thruster Design</b>	<b>75</b>
4.1	Modular SHARK-600V Krypton Hall Effect Thruster Concept . . . . .	76
4.2	Thruster with Anode-Layer Design . . . . .	77
4.3	Magnetic Field Design . . . . .	79
4.3.0.1	Magnetic Field Simulation: Prototyping . . . . .	80
4.3.1	Magnetic Field Simulation: Final . . . . .	84
4.3.2	Magnetic Field Measurements . . . . .	86
4.4	Anode Design . . . . .	90
4.4.1	Anode X-Ray Computed Tomography Scans . . . . .	91
4.4.1.1	Acknowledgement . . . . .	94
4.5	Thermal Design Considerations . . . . .	94
4.5.1	Zero-Dimensional Thermal Model . . . . .	95

---

4.5.1.1	Ionisation Power . . . . .	95
4.5.1.2	Plume Kinetic Power . . . . .	96
4.5.1.3	Plume Thermal Power . . . . .	96
4.5.1.4	Wall Power . . . . .	97
4.5.1.5	Radiated and Excited Power . . . . .	99
4.5.1.6	Anode Heating . . . . .	99
4.5.1.7	Total Thermal Flux . . . . .	99
4.5.1.8	Thermal Model Results . . . . .	100
4.5.2	Thermal Design . . . . .	102
4.6	SHARK-600V Design Summary . . . . .	102
<b>5</b>	<b>Southampton Test Campaign</b>	<b>105</b>
5.1	Introduction . . . . .	105
5.2	David Fearn Large Vacuum Chamber Facility . . . . .	105
5.3	Cathode . . . . .	106
5.4	Electrical Setup . . . . .	108
5.4.1	Anode Power Supply . . . . .	108
5.4.1.1	Low-Pass Filter Design . . . . .	109
5.4.2	Cathode Power Supply . . . . .	109
5.4.3	Electromagnet Power Supply . . . . .	110
5.5	Fluidic Setup . . . . .	111
5.6	Thrust Balance Development . . . . .	113
5.6.1	Thrust Balance Variants . . . . .	114
5.6.1.1	Hanging Pendulum . . . . .	114
5.6.1.2	Inverted Pendulum . . . . .	114
5.6.1.3	Torsional Pendulum . . . . .	114
5.6.2	Double Inverted Pendulum Thrust Balance . . . . .	114
5.6.3	Thrust Balance Calibration . . . . .	117
5.6.4	Calibration Results . . . . .	119
5.6.5	Cross-Laboratory Validation of Thrust Balance . . . . .	119
5.6.5.1	Introduction . . . . .	119
5.6.5.2	External Plasma Thruster . . . . .	121
5.6.5.3	Experimental Results . . . . .	121
5.6.5.4	Conclusion . . . . .	121
5.7	Data Acquisition Set-up . . . . .	123
5.7.1	Power Supply Communication . . . . .	123
5.7.2	Flow Controller Communication . . . . .	125
5.7.3	Multi-threading Structure . . . . .	125
5.7.4	Graphical User Interface . . . . .	126
5.7.5	Data logger . . . . .	126
5.7.6	Thrust Balance Measurements . . . . .	126
5.8	Data Analysis . . . . .	128
5.9	Experimental Results . . . . .	130
5.9.1	Introduction . . . . .	130
5.9.1.1	How to Read Grid Plots . . . . .	131
5.9.2	Global Performance Trends . . . . .	131
5.9.3	Xenon Performance Trends . . . . .	143

5.9.3.1	Xenon: Effect of Channel Width . . . . .	143
5.9.3.2	Xenon: Effect of Channel Mean Diameter . . . . .	145
5.9.3.3	Xenon: Max Performance Grid . . . . .	146
5.9.4	Krypton Performance Trends . . . . .	153
5.9.4.1	Krypton: Effect of Channel Width . . . . .	153
5.9.4.2	Krypton: Effect of Channel Mean Diameter . . . . .	154
5.9.4.3	Krypton: Max Performance Grid . . . . .	157
5.9.5	Argon Performance Trends . . . . .	163
5.9.5.1	Argon: Effect of Channel Width . . . . .	164
5.9.5.2	Argon: Effect of Channel Mean Diameter . . . . .	165
5.9.5.3	Argon: Max Performance Grid . . . . .	167
5.9.6	Anomalous Behaviour . . . . .	171
5.9.7	Thruster with Anode-Layer Configuration . . . . .	171
5.9.7.1	TAL: Experimental Results . . . . .	172
5.9.7.2	TAL: Performance . . . . .	173
5.9.7.3	TAL: Notable Events . . . . .	175
5.9.7.4	TAL: Conclusion . . . . .	176
5.10	Experimental Campaign Discussion . . . . .	176
5.11	Gaussian Process Regression Model of SHARK-600V . . . . .	179
5.11.1	Gaussian Process Regression . . . . .	179
5.11.2	Gaussian Process Regression Model Training . . . . .	182
5.11.2.1	Model Training Results . . . . .	182
5.11.3	Gaussian Process Regression Model of SHARK-600V Performance . . . . .	186
5.11.4	Scaling Implications . . . . .	191
5.11.5	Gaussian Process Regression Model Summary . . . . .	196
5.12	Conclusions . . . . .	197
<b>6</b>	<b>Michigan Test Campaign</b>	<b>201</b>
6.1	Introduction . . . . .	201
6.1.1	Funding . . . . .	201
6.2	Experimental Set-up . . . . .	202
6.2.1	Large Vacuum Test Facility . . . . .	202
6.2.2	Thruster Configuration . . . . .	203
6.2.3	Cathode . . . . .	204
6.3	Diagnostics . . . . .	205
6.3.1	Thrust Balance . . . . .	205
6.3.2	E×B Probe . . . . .	205
6.3.2.1	E×B Probe: Charge-Exchange Correction . . . . .	210
6.3.2.2	E×B Probe: Current Fractions . . . . .	212
6.3.3	Faraday Probe . . . . .	212
6.3.3.1	Faraday Probe: Charge-Exchange Correction . . . . .	214
6.3.4	Retarding Potential Analyser . . . . .	216
6.4	Experimental Results . . . . .	217
6.4.1	Probe Analysis . . . . .	218
6.4.2	Charge Efficiency . . . . .	219
6.4.3	Beam Efficiency . . . . .	222
6.4.4	Divergence Efficiency . . . . .	224

---

6.4.5	Voltage Efficiency . . . . .	225
6.4.6	Mass Efficiency . . . . .	227
6.4.7	Anode Efficiency Comparison . . . . .	228
6.4.8	Comparison to Southampton Results . . . . .	229
6.5	Conclusion . . . . .	233
<b>7</b>	<b>Summary of Key Contributions, Conclusion and, Future Work</b>	<b>235</b>
7.1	Summary of Major Contributions . . . . .	235
7.2	Conclusion . . . . .	238
7.3	Future Work . . . . .	239
<b>8</b>	<b>Publication History</b>	<b>241</b>
<b>A</b>	<b>Magnetic Field Simulation results</b>	<b>243</b>
<b>B</b>	<b>Gaussian Process Regression Modelling: Additional Figures</b>	<b>247</b>
B.1	GPR Training plots . . . . .	247
B.2	GPR Model Error . . . . .	249
<b>C</b>	<b>Michigan Campaign: Uncertainty Analysis &amp; Additional Result</b>	<b>253</b>
C.1	Uncertainty Analysis of Far-Field Probes . . . . .	253
C.1.1	$E \times B$ Probe Error . . . . .	254
C.1.2	Faraday Probe Error . . . . .	256
C.1.3	Retarding Potential Analyser Error . . . . .	257
C.2	Additional Results from Michigan Campaign . . . . .	259
C.2.1	$E \times B$ Additional Results . . . . .	259
C.2.2	Faraday Probe Additional Results . . . . .	260
C.2.3	Retarding Potential Analyser Additional Results . . . . .	261



# List of Figures

1.1	The cost of a kilogram of xenon in euros over time. Xenon pricing data prior to 2015 are from Herman and Unfried [13]; pricing between 2015 and 2022 is from Unfried et al. [17]; <b>A</b> : 2022 estimate is from Munro-O'Brien and Ryan [12]; <b>B</b> : 2023 value is from Hansen et al. [18]; <b>C</b> : 2024 estimate is from a receipt dated 2 October 2024 for an order quantity of 50 L; <b>D</b> : 2026 estimate is from a quote dated 2 February 2026 for an order quantity of 50 L. The data were adjusted for inflation using the Consumer Price Index for All Urban Consumers (CPI-U) from the U.S. Bureau of Labor Statistics [Series ID: CUUR0000SA0L1E]. Euro values for data prior to 2015 were converted using an exchange rate of 0.93 EUR to 1 USD accessed on 1 April 2025. Euro values for data post-2015 use year-averaged exchange rates. . . . .	5
1.2	Illustrative side-by-side comparison of a cross-section of stationary plasma thruster and a thruster with anode-layer configuration. This schematic is non-physical and intended for conceptual comparison only, highlighting the primary structural and operational differences between the two Hall thruster types. . . . .	6
1.3	Representation of several popular types of electric propulsion being compared directly with Hall thrusters operating on alternative propellants being separated. The dot-dash iso-lines represent the total efficiency as labelled within the line. Data plotted is taken from [21, 23, 24, 35–66]. . . . .	9
2.1	Schematic diagram showing the cross-section of a traditional gridded ion thruster [73]. . . . .	17
2.2	Illustrative cross-section of a Hall effect thruster featuring a ceramic discharge channel. The three key geometric parameters used throughout this thesis: channel length $L$ , channel width $h$ , and mean channel diameter $d$ are indicated. . . . .	18
2.3	A schematic view of the cross section of a thruster with anode-layer with the hollow anode and the guard rings highlighted. . . . .	26
2.4	Log-log plot showing the relationship between anode power and mean channel diameter for SPT-type and TAL-type Hall thrusters [20–23, 37, 41, 43, 46, 48, 51, 54, 57–60, 64, 92–111]. . . . .	27
2.5	Estimated xenon neutral temperature at the channel centreline over the channel length relative to the exit plane ( $Z = 0$ mm) of the Stanford Hall thruster, measured via laser induced florescence for several discharge voltage levels [139]. The temperature values shown are under the assumption that the neutrals are thermalised. . . . .	33

2.6	Variation of mean free path for xenon, krypton, and argon at anode voltages of 200 V, 300 V, and 600 V, as a function of neutral temperature. Calculated assuming optimal neutral number density [12]. . . . .	34
2.7	An illustration of the ionisation power per ion current assuming a purely mono-ionic ion current from Eq. 2.41. . . . .	37
2.8	Cross-sectional area of ionisation for potential propellants versus bulk plasma temperature in electron volts. Cross-sectional area data: xenon and krypton from [142], argon from [143], and neon and diatomic nitrogen from [144]. . . . .	39
2.9	Computed reaction rate coefficients of potential propellants as a function of bulk plasma temperature in electron volts. . . . .	40
2.10	Schematic of the plasma–wall transition layer and the two Debye sheaths either side are shown [148]. . . . .	42
2.11	Schematic of particle fluxes at the plasma-facing wall, including primary electron flux $\Gamma_{ew}$ , ion flux $\Gamma_{iw}$ , and secondary electron flux $\Gamma_{SEE}$ . Under steady-state conditions, the total current to the wall must be zero. . . . .	43
2.12	Schematic illustrating electron interactions with a boron nitride surface. Primary electrons impact the surface, resulting in energy transfer, absorption, and secondary electron emission depending on the local material structure. . . . .	44
2.13	Schematic illustrating electron interactions with a metallic surface. Primary electrons impact the surface, resulting in energy transfer, absorption, and secondary electron emission depending on the local material structure. . . . .	44
3.1	Histogram of the database anode power values using 150 bins per propellant-thruster type [21–23, 37, 41, 43, 45, 46, 48, 50, 51, 54, 59, 60, 92, 96, 97, 99–101, 103, 104, 107, 109, 111]. . . . .	53
3.2	Channel width plotted against mean channel diameter for SPT and TAL thrusters within the dataset. Least-squares linear fits are shown for both types. . . . .	55
3.3	Box plots of channel width to mean diameter ratio ( $h/d$ ) for various thruster families. The data highlight inter-family variance and the presence of outliers. . . . .	56
3.4	Box and whisker plots of neutral number density values for xenon and krypton propellants across the Hall thruster database. The total population mean neutral density is $\bar{n}_n _{Xe} = 1.66 \times 10^{19} \text{ m}^{-3}$ for $N = 851$ , and $\bar{n}_n _{Kr} = 1.34 \times 10^{19} \text{ m}^{-3}$ for $N = 338$ . . . . .	59
3.5	Database of thrusters with their specific impulse at peak anode efficiency against the normalised scaling value from Eq. 3.20 for xenon and krypton operating in SPT type Hall effect thrusters. $R^2 _{Xenon} = 0.8143$ and $R^2 _{Krypton} = 0.6305$ . . . . .	62
3.6	Normalised thrust plotted against the scaling proportionality function described in Eq. 3.22 for the SPT-type thrusters listed in Tab. 3.1. $R^2 _{Xenon} = 0.9886$ and $R^2 _{Krypton} = 0.9834$ . . . . .	63
3.7	Comparison of anode power scaling fits using different characteristic dimension models. . . . .	66
3.8	The result of Eq. 3.31 for both xenon and krypton SPTs from Tab. 3.1. $R^2 _{Xenon} = 0.9905$ and $R^2 _{Krypton} = 0.9837$ . . . . .	67

3.9	Normalised channel width as a function of the normalised scaling parameter from Eq. 3.35, evaluated at the operating point corresponding to maximum specific impulse. $R^2 _{\text{xenon}} = 0.9197$ and $R^2 _{\text{krypton}} = 0.8511$ . . . . .	69
4.1	A graphical illustration of the channel inserts, showing the three constant mean diameter cases “+0, +0”, “+1, +1” and “+2, +2”. . . . .	77
4.2	Computer renderings of cross sections of the SHARK-600V in all of the thruster with anode-layer configurations designed. The magnetic circuit and anode are illustrative only and do not represent the final design. Here “L” the TAL configuration channel length is shown as reported in Tab. 4.2. . . . .	78
4.3	A graphical rendering of the cross-section and iso-metric views of the SHARK-600V Hall effect thruster with labels highlighting the electromagnet locations. . . . .	79
4.4	Illustration of a single simulated configuration of the FEMM mesh and output. . . . .	81
4.5	Magnetic field simulations for selected geometry sweep 3 showing magnetic field streamlines for 2A inner coil and 2A outer coil currents. . . . .	83
4.6	Comparison of full-thruster magnetic field (shown in units of Gauss) topology at coil currents of 1 A and 2 A. Note: the internal structure has been obscured. . . . .	84
4.7	Comparison of radial magnetic field (shown in units of Gauss) component ( $B_z$ ) at coil currents of 1 A and 2 A. Note: the internal structure has been obscured. . . . .	85
4.8	The SHARK-600V prior to testing mounted on a flat plate for magnetic field mapping. . . . .	86
4.9	Comparison of radial magnetic field measurements of the SHARK-600V at an anti-node with 1 A coil current using two mapping approaches. In Fig. 4.9a, the lines shown are iso-contours bounding regions of constant magnitude within the colour map bins, <i>not</i> magnetic field lines. . . . .	87
4.10	Comparison of radial magnetic field strength measured along the inner diameter (ID), outer diameter (OD), and mean diameter (MD) of the SHARK-600V with coil currents of 1 A. . . . .	89
4.11	Front one image of the SHARK-600V Hall effect thruster in “+2, +2” channel configurations prior to testing where anode injection holes are clearly visible. . . . .	90
4.12	A illustration from the computer rendering of the anode design where the angled propellant injection is visible, highlighted with two red-lines indicating the inner and outer edges of the hole and the centre-line indicating the holes angle relative to the front surface. . . . .	91
4.13	X-ray CT slices of the SHARK-600V anode at different depths, showing internal channels and features. . . . .	92
4.14	Side-on X-ray CT annotated slice of the anode showing the internal propellant pipe and highlighting a manufacturing irregularity in the bottom channel. . . . .	93
4.15	Approximate xenon ion thermal temperature prior to acceleration, estimated from LIF measurements by Hargus et al. using Eq. 4.5 [139]. . . . .	97

4.16	A illustration of the results from Tab. 4.3 for both SPT and TAL type for xenon, krypton, and argon. The $P_{\text{thermal}}$ column is the sum of the thermal powers, the $P_{\text{tot}}$ column is the sum of all powers from the model. . . . .	101
5.1	The David Fearn Electric Propulsion Laboratory at the University of Southampton where all experimental testing presented in Chapter 5 was undertaken.	106
5.2	Illustration of the electrical, fluidic and data connections of the vacuum chamber testing that was used for this campaign. . . . .	107
5.3	The HCES 5000 hollow cathode used during the Southampton testing campaign. . . . .	107
5.4	Low-pass filter implementation showing (a) the circuit diagram and (b) the assembled hardware. . . . .	110
5.5	Piping and Instrumentation Diagram for the multi and mixed propellant delivery system. . . . .	111
5.6	Photograph of the MAXPDS front control board as implemented for the SHARK-600V test campaign. . . . .	112
5.7	Illustration of the double inverted pendulum thrust stand with counter-balance [183]. . . . .	115
5.8	Design views of the inverted pendulum thrust balance. . . . .	116
5.9	The manufactured inverted pendulum thrust balance installed in the Southampton large vacuum chamber, with the SHARK-600V mounted post-test. . . . .	117
5.10	Laser triangulation sensor output for a standard calibration, showing the raw signal, filtered response, and steady-state levels. . . . .	118
5.11	Example output of the calibration process for the thrust balance used in these tests. The corresponding statistical results and error analysis are summarised in Tab. 5.1. . . . .	120
5.12	Operation of the XPT on xenon in the David Fearn Electric Propulsion Laboratory at Southampton (left) and in the Daedalus facility at Surrey (right). . . . .	121
5.13	Thrust measured for the xenon-fed XPT at different anode mass flow rates in the University of Surrey and University of Southampton facilities.	122
5.14	Anode efficiency of the XPT at varying xenon flow rates, measured at Surrey and Southampton. . . . .	122
5.15	High-level connection map showing how instruments were linked to the laboratory computer during operation of the DAQ system. . . . .	124
5.16	Thread-level connection map used for the DAQ system, illustrating parallelisation and the separate communication protocols for each instrument.	126
5.17	Screenshot of the data acquisition front panel showing all connected power supplies and flow controllers. . . . .	127
5.18	Example thrust measurement from the analysis routine, showing laser displacement, anode voltage, anode current, and anode mass flow rate. The red vertical line marks the on-off transition, shaded regions represent the averaging windows, and the horizontal dot-dashed lines indicate the two laser displacement averages used to calculate the thrust-induced deflection. . . . .	128
5.19	A side by side comparison of the thruster operating in the large vacuum chamber at the University of Southampton on xenon (top left), krypton (top right), argon (bottom left), and neon-xenon mix (bottom right). . . .	130

5.20	Illustration of the modular thruster channel configurations arranged in a grid format. Each subplot corresponds to the inner and outer channel dimensions listed in Tab. 5.2. The figure is displayed at a reduced scale of 1:3 relative to the actual geometry for clarity (assuming displayed on a standard A4 page). . . . .	132
5.21	Measured thrust across all SPT configurations and propellants. Both anode power and thrust error bars are shown. Anode power uncertainties are derived from the standard deviation of voltage and current over the averaging range, while thrust uncertainties are calculated from the calibration procedure described in Sec. 5.6.3. . . . .	133
5.22	Thrust versus anode power grouped by anode propellant (xenon, krypton, and argon) for each channel configuration. Each subplot corresponds to one thruster geometry, arranged as in Tab. 5.2. . . . .	135
5.23	thrust-to-power versus anode specific impulse with anode efficiency iso-lines, grouped by propellant (xenon, krypton, argon) and anode voltage, for each channel configuration tested. . . . .	136
5.24	Anode specific impulse versus anode power for the "+0, +0" and "+2, +1" channel configuration, grouped by propellant and discharge voltage. . .	137
5.25	thrust-to-power ratio versus anode specific impulse for the "+0, +0" and "+2, +1" channel configuration, grouped by propellant and discharge voltage. . . . .	138
5.26	thrust-to-power against anode specific impulse for each tested propellant and anode voltage in the configurations of constant mean diameter of 70 mm ("+0, +0", "+1, +1", and "+2, +2"). Here iso-lines have been applied to represent lines of constant anode efficiency. . . . .	139
5.27	Anode current draw against anode discharge voltage for xenon, grouped by anode mass flow rate. . . . .	141
5.28	thrust-to-power ratio against specific impulse for xenon, grouped by discharge voltage. . . . .	142
5.29	The SHARK-600V operating on xenon in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory. . . . .	143
5.30	Comparison of anode specific impulse against anode power for xenon across channel width configurations ("+0, +0" vs. "+1, +1" vs. "+2, +2"), grouped by discharge voltage. . . . .	144
5.31	Comparison of anode efficiency against anode power for xenon across channel width configurations ("+0, +0" vs. "+1, +1" vs. "+2, +2"), grouped by discharge voltage. . . . .	144
5.32	Comparison of anode specific impulse against anode power for xenon across channel mean diameter configurations ("+1, +0" vs. "+0, +1"), grouped by discharge voltage. . . . .	146
5.33	Comparison of anode efficiency against anode power for xenon across channel mean diameter configurations ("+1, +0" vs. "+0, +1"), grouped by discharge voltage. . . . .	146
5.34	Comparison of anode specific impulse against anode power for xenon across channel mean diameter configurations ("+1, +2" vs. "+2, +1"), grouped by discharge voltage. . . . .	147
5.35	Comparison of anode efficiency against anode power for xenon across channel mean diameter configurations ("+1, +2" vs. "+2, +1"), grouped by discharge voltage. . . . .	147

5.36	Performance grid of thrust-to-power ratio (mN/kW) for xenon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. .	149
5.37	Performance grid of anode efficiency (%) for xenon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. . . . .	150
5.38	Anode current draw against anode discharge voltage for krypton, grouped by anode mass flow rate. . . . .	151
5.39	thrust-to-power ratio against specific impulse anode power for krypton, grouped by discharge voltage. . . . .	152
5.40	The SHARK-600V operating on krypton in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory. . . . .	153
5.41	Comparison of anode specific impulse against anode power for krypton across varying channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage. . . . .	154
5.42	Comparison of anode efficiency against anode power for krypton across varying channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage. . . . .	155
5.43	Comparison of anode specific impulse against anode power for krypton across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage. . . . .	156
5.44	Comparison of anode efficiency against anode power for krypton across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage. . . . .	156
5.45	Comparison of anode specific impulse against anode power for krypton across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage. . . . .	157
5.46	Comparison of anode efficiency against anode power for krypton across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage. . . . .	158
5.47	Performance grid of thrust-to-power (mN/kW) ratio for krypton. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. .	159
5.48	Performance grid of anode efficiency (%) for krypton. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. . . . .	160
5.49	Anode current draw against anode discharge voltage for argon, grouped by anode mass flow rate. . . . .	161
5.50	thrust-to-power ratio verses specific impulse for argon, grouped by discharge voltage. . . . .	162

5.51	The SHARK-600V operating on argon in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory. . . . .	163
5.52	Comparison of anode specific impulse against anode power for argon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage. . . . .	164
5.53	Comparison of anode efficiency against anode power for argon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage. . . . .	165
5.54	Comparison of anode specific impulse against anode power for argon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage. . . . .	166
5.55	Comparison of anode efficiency against anode power for argon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage. . . . .	166
5.56	Comparison of anode specific impulse against anode power for argon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage. . . . .	167
5.57	Comparison of anode efficiency against anode power for argon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage. . . . .	168
5.58	Performance grid of thrust-to-power (mN/kW) ratio for argon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. . . . .	169
5.59	Performance grid of anode efficiency (%) for argon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale. . . . .	170
5.60	A photo of the thruster captured after the plasma was extinguished manually due to overheating concerns where the boron nitride channel is emitting in the visible wavelength due to high temperatures. . . . .	171
5.61	The TAL-1 configuration operating at approximately 1.3 kW on krypton within the large vacuum chamber facility at the University of Southampton. . . . .	172
5.62	Anode efficiency for the two TAL configurations tested against anode power grouped by krypton flow rate in sccm and anode configuration. . . . .	173
5.63	Anode specific impulse against anode power for the “TAL-1” configurations grouped by krypton flow rates. . . . .	174
5.64	Thrust against anode power for the “TAL-1” configurations grouped by krypton flow rates. . . . .	174
5.65	Thermal radiation from the anode in the “TAL-1” configuration following 1.5 kW operation. . . . .	175
5.66	“TAL+1” firing at 1.3 kW where excessive thermal loading of the anode caused localised melting of the stainless-steel sheaths. . . . .	175
5.67	The “TAL+1” metallic anode after high thermal load conditions. The image shows the post-operation state and the resultant structural failure after melting, observed once the chamber was returned to atmosphere. . . . .	176

5.68	Mean relative error of the Gaussian process regression model for xenon as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (100 for xenon). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models. . . . .	183
5.69	Mean relative error of the Gaussian process regression model for xenon as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models. . . . .	184
5.70	Geometry map of the channel configurations from Tab. 5.2 shown in the same format as subplots in Figs. 5.71–5.73. . . . .	186
5.71	Array of Gaussian process regression model for xenon, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	188
5.72	Array of Gaussian process regression model for krypton, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	189
5.73	Array of Gaussian process regression model for argon, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	190
6.1	Photo of Alex D. Gallimore Large Vacuum Facility taken at the University of Michigan, Ann Arbor (from [81] with permission). . . . .	202
6.2	The SHARK-600V thruster in “+2, +1” channel configurations mounted in the LVTF at the University of Michigan prior to testing. . . . .	203
6.3	Cathode design used during the Michigan test campaign (image from [168], used with permission). . . . .	204
6.4	Schematic of a Wien filter-style $E \times B$ probe, illustrating the orthogonal electric and magnetic fields and ion trajectories. . . . .	206
6.5	Step-by-step illustration of the procedure for fitting multiple twin-Gaussian distributions to a representative probe signal. The example corresponds to an operating point of 300 V anode voltage and 90 sccm krypton anode mass flow rate. . . . .	209

6.6	Charge exchange ion correction factor versus acceleration voltage for each ion species considered. A background pressure of $5.3 \times 10^{-6}$ mbar is assumed, with the shaded region showing $\pm 25\%$ pressure variation. As can be seen in Eq. 6.7, a higher pressure will result in a lower $(j/j_0)_n$ value. . . . .	211
6.7	An illustration of a Faraday probe with important dimensions required for analysis labelled. . . . .	212
6.8	Faraday probe measurements of the SHARK-600V plume at the University of Michigan for 90 sccm of krypton and an anode voltage of 300 V. Both flat-subtraction and exponential fitting corrections are shown. The plume angle shown is calculated using the exponential fitting correction and Eq. 2.32. . . . .	215
6.9	An illustration of the retarding potential analyser measurement for 300 V and 90 sccm anode voltage and flow respectively and the resulting derivative, $-dC/dV$ , of the signal, with the most probable voltage indicated. . . . .	217
6.10	The SHARK-600V operating at 300 V anode voltage and 60 sccm anode mass flow rate of krypton in the LVTF at the University of Michigan. . . . .	218
6.11	Comparison of individual probe anode efficiency, multiplicative probe anode efficiency, and thrust anode efficiency for 50 sccm of krypton. . . . .	219
6.12	Comparison of each " $E \times B$ " probe trace at 50 sccm for anode voltages ranging from 200 V to 500 V. . . . .	220
6.13	Resultant ion current fractions expressed as a percentage versus anode mass flow rate of krypton from the $E \times B$ probe. . . . .	220
6.14	Average charge of axial krypton ions against anode mass flow rate in the plume of the SHARK-600V, as measured by the $E \times B$ probe and calculated using Eq. 6.16. . . . .	221
6.15	Charge efficiency against anode power grouped by discharge voltage of the SHARK-600V operating on krypton in the LVTF. . . . .	222
6.16	Faraday sweeps for all anode flow rates tested at 300 V. Due to the effect of the cathode, the trace shown is from $-90^\circ$ to $0^\circ$ , mirrored about the centreline to cover $0^\circ$ to $90^\circ$ . . . . .	223
6.17	Current efficiency against anode power grouped by discharge voltage as measured from the Faraday probe. . . . .	223
6.18	Beam divergence angles against anode power grouped by discharge voltage. . . . .	224
6.19	Divergence efficiency against anode power grouped by discharge voltage. . . . .	225
6.20	Voltage efficiency against anode power grouped by discharge voltage. . . . .	226
6.21	Measured voltage loss against discharge voltage grouped by anode flow rate. . . . .	226
6.22	Mass efficiency against anode mass flow rate grouped by discharge voltage. . . . .	227
6.23	A comparison of anode efficiency as measured with direct thrust measurements (blue) and inferred from probe measurements (red) against anode power. . . . .	228
6.24	Comparison of thrust as measured by direct thrust measurements from Michigan testing and Southampton testing. . . . .	229
6.25	Comparison of direct thrust measurements against anode discharge current from the Michigan and Southampton testing. . . . .	230

6.26	Comparison of measured anode current difference between Michigan and Southampton testing as a function of anode voltage. Note: negative values indicate a lower anode current measured at Michigan than at Southampton. . . . .	230
6.27	Comparison of anode efficiency as measured by direct thrust measurements from University of Michigan test campaign and University of Southampton test campaign. . . . .	231
6.28	Comparison of anode specific impulse as measured by direct thrust measurements from Michigan and Southampton testing. . . . .	232
A.1	Magnetic field simulations for selected geometry sweep 1 showing magnetic field streamlines for 1A inner coil and 1A outer coil currents. . . . .	244
A.2	Magnetic field simulations for selected geometry sweep 2 showing magnetic field streamlines for 2A inner coil and 2A outer coil currents. . . . .	245
B.1	Mean relative error of the Gaussian process regression model for krypton as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (109 for krypton). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models. . . . .	247
B.2	Mean relative error of the Gaussian process regression model for argon as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (35 for argon). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models. . . . .	248
B.3	Mean relative error of the Gaussian process regression model for krypton as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models. . . . .	248
B.4	Mean relative error of the Gaussian process regression model for xenon as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models. . . . .	249

B.5	Array of Gaussian process regression model error for xenon from Fig. 5.71, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	250
B.6	Array of Gaussian process regression model error for krypton from Fig. 5.72, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	251
B.7	Array of Gaussian process regression model error for argon from Fig. 5.73, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested. . . . .	252
C.1	Charge exchange ion correction factor versus chamber background pressure for each ion species considered, for a range of discharge voltages. The background neutral temperature is assumed to be 300 K. The dotted vertical line represents $5.33 \times 10^{-6}$ mbar, the measured pressure during testing. . . . .	255
C.2	E×B probe data for krypton at 40 sccm for each tested anode voltage. . .	259
C.3	E×B probe data for krypton at 60 sccm for each tested anode voltage. . .	259
C.4	E×B probe data for krypton at 70 sccm for each tested anode voltage. . .	260
C.5	Faraday probe radial sweep for krypton at 400 V for all tested flow rates.	260
C.6	Faraday probe radial sweep for krypton at 500 V for all tested flow rates.	261
C.7	RPA sweep for krypton at 40 sccm across all tested voltages. . . . .	261
C.8	RPA sweep for krypton at 50 sccm across all tested voltages. . . . .	262
C.9	RPA sweep for krypton at 60 sccm across all tested voltages. . . . .	262
C.10	RPA sweep for krypton at 70 sccm across all tested voltages. . . . .	263
C.11	RPA sweep for krypton at 300 V across all tested flow rates. . . . .	263



# List of Tables

2.1	Comparison of the mass, ionisation energies (IE), and natural abundance of several potential gaseous propellants with xenon. Percentages are shown relative to xenon. The ionisation energy for nitrogen refers to dinitrogen ( $N_2$ ); within the plasma both diatomic and atomic states would be present [11, 120]. †: Ionisation energy for $N_2 \rightarrow N_2^+ + e$ . . . . .	29
2.2	Comparison of the properties of several alternative condensable and metallic alternative propellants with xenon [119]. . . . .	30
2.3	Comparison of cross-sectional areas and reaction rate coefficients for candidate propellants at an assumed 300 V discharge voltage, using Eq. 2.45.	40
2.4	Power law fitting parameters for secondary electron emission. Values for boron nitride are taken from Pigeon et al. [149], and for stainless steel from Choueiri [29]. . . . .	43
2.5	Critical electron temperature $T^*$ for sheath reversal in Hall thrusters with boron nitride and stainless-steel walls for selected propellants. . . . .	45
3.1	Thruster geometry database used for scaling. †: Primary anode, secondary anode. . . . .	54
3.2	Summary of neutral number density statistics for xenon and krypton across the dataset for SPT type thrusters and peak performance subsets from Fig. 3.4. $\bar{n}_n$ : Mean number density; $\tilde{n}_n$ : Median number density. The <b>Min.</b> and <b>Max.</b> columns refer to the minimum and maximum values, respectively, within the dataset for the corresponding row. . . . .	58
3.3	Scaling inputs and initial scaling results used to design the SHARK-600V Hall effect thruster for this thesis. . . . .	72
3.4	Scaling estimation for the chosen dimensions for both krypton and xenon. . . . .	73
3.5	Scaling inputs and updated scaling results for the design parameters of the original thruster using the reduced dataset. . . . .	73
4.1	Channel width (mm) and mean channel diameter (mm) for each configuration of the modular thruster. Each cell shows “h mm, d mm” for that inner and outer configuration. . . . .	76
4.2	Table of all the thruster with anode-layer configurations of variable channel lengths as seen in Fig. 4.2. Here negative values represent that the anode extends past thrusters magnetic exit plane. . . . .	78
4.3	Comparison of power terms, in watts, in the thermal model for a thruster operating at 1350 W and 600 V, scaled to input values for the SHARK-600V. . . . .	100

5.1	A table of the results from one of the calibration runs of the thrust balance used for these testing. Individual calibrations are run for each test in-situ.	120
5.2	Channel width (mm) and mean channel diameter (mm) for each configuration of the modular thruster. Each cell shows “h mm, d mm : “+ #, + #” for that inner and outer configuration. . . . .	131
5.3	Dataset sizes used for training and testing the GPR surrogate models for each propellant. . . . .	182
5.4	Statistical results of the Gaussian Process Regression test set for xenon dataset against the trained model. Here $R^2$ is the coefficient of determination, RSME, is the root mean square error between the tested values and the model predictions, $RMSE \%_{mean}$ is the RMSE relative to the mean value of test value, and the CV(RMSE) is Coefficient of Variation of the RMSE and is mathamatically described in Eq. 5.15. . . . .	184
5.5	Statistical results of the Gaussian Process Regression test set for krypton dataset against the trained model. Here $R^2$ is the coefficient of determination, RSME, is the root mean square error between the tested values and the model predictions, $RMSE \%_{mean}$ is the RMSE relative to the mean value of test value, and the CV(RMSE) is Coefficient of Variation of the RMSE and is mathamatically described in Eq. 5.15. . . . .	185
5.6	Statistical results of the Gaussian Process Regression test set for argon dataset against the trained model. Here $R^2$ is the coefficient of determination, RSME, is the root mean square error between the tested values and the model predictions, $RMSE \%_{mean}$ is the RMSE relative to the mean value of test value, and the CV(RMSE) is Coefficient of Variation of the RMSE and is mathamatically described in Eq. 5.15. . . . .	185
5.7	Comparison of GPR-predicted optimal geometry and performance against the semi-empirical scaling methodology as described in Chapter 3 for the updated database for xenon at 600 V and 1.35 kW. Percentage differences are calculated as relative error of the GPR model with respect to the scaling result. No feasible GPR solutions were found for krypton or argon at this operating point. . . . .	192
5.8	GPR-predicted optimal geometries and performance for xenon, krypton, and argon at 300 V and 1.35 kW. . . . .	193
5.9	Results of the semi-empirical scaling methodology for xenon and krypton at 300 V and 1.35 kW, with relative errors calculated with respect to the GPR-predicted optima shown in Tab. 5.8. Percentage differences are expressed as $(GPR - scaling) / scaling \times 100\%$ . . . . .	193
5.10	GPR-predicted optimal geometries and performance for xenon, krypton, and argon at 600 V and 2.5 kW. †: Argon exhibits large uncertainty in the GPR model for this operational point due to the flow rate required being much less than the tested range. . . . .	194
5.11	Comparison of semi-empirical scaling methodology results with GPR-predicted optima for xenon and krypton at 600 V and 2.5 kW. Percentage differences are calculated as the relative error of the GPR model with respect to the scaling result. Argon is excluded from this comparison due to the inability of the scaling method to be applied to argon. Percentage differences are expressed as $(GPR - scaling) / scaling \times 100\%$ . . . . .	195

- 
- 6.1 Channel width, mm and mean channel diameter, mm for the possible SPT configuration. Here the chosen configurations for the University of Michigan campaign can be seen in **bold** and red. . . . . 203
- 6.2 Secondary electron emission (SEE) yields for molybdenum for impinging krypton ions. Values for  $\text{Kr}^+$  and  $\text{Kr}^{2+}$  from [212]. ‡: No data for  $\text{Kr}^{3+}$ ; 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 krypton impinging on molybdenum. [213]. . . . . 213



# Abbreviations

1D	One-Dimensional
2D	Two-Dimensional
3D	Three-Dimensional
A	Ampère
AC	Alternating Current
amu	Atomic Mass Units
Ar	Argon
ASCII	American Standard Code for Information Interchange
Bi	Bismuth
BN	Boron Nitride
CAT-6	Category 6 Cable
Cd	Cadmium
CEX	Charge Exchange
Co.	Company
COM Port	Communication Port
CPI	Consumer Price Index
CPU	Central Processing Unit
CSV	Comma-Separated Values
CV(RMSE)	Coefficient of Variation of the Root Mean Squared Error
DAQ	Data Acquisition
DC	Direct Current
EP	Electric Propulsion
ESA	European Space Agency
etc.	et cetera
eV	Electron Volts
FC	Flow Controller
FEMM	Finite Element Method Magnetics
GIT	Gridded Ion Thrusters
GUI	Graphical User Interface
GP	Gaussian Process
GPR	Gaussian Process Regression
HET	Hall Effect Thruster

I	Iodine
ID	Inner Diameter
IEDF	Ion Energy Distribution Function
i.e.	id est
IE	Ionisation Energy
Inc.	Incorporated
IP Address	Internet Protocol Address
IVDF	Ion Velocity Distribution Function
KE	Kinetic Energy
Kr	Krypton
LaB <sub>6</sub>	Lanthanum Hexaboride
LAN	Local Area Network
LH	Left Hand
LIF	Laser Induced Fluorescence
LVTF	Large Vacuum Test Facility
Max.	Maximum
MAXPDS	Multi And miXed Propellant Delivery System
MD	Mean Diameter
Med.	Medium
MFC	Mass Flow Controller
Mg	Magnesium
mg/s	Milligrams per second
Min.	Minimum
MIT	Massachusetts Institute of Technology
mm	Millimetre
mN	Millinewton
MS	Magnetic Shielding
MPV	Most Probable Voltage
N <sub>2</sub>	Dinitrogen
N5.0	99.999% Purity
NASA	National Aeronautics and Space Administration
Ne	Neon
NEXT	NASA's Evolutionary Xenon Thruster
NN	Neural Network
OD	Outer Diameter
PIC	Particle-in-cell
P&ID	Piping and Instrumentation Diagram
PEPL	Plasmadynamics and Electric Propulsion Laboratory
ppm	Parts Per Million
PS	Power Supply
RH	Right Hand

---

RMSE	Root Mean Square Error
RPA	Retarding Potential Analyser
s	Seconds
sccm	Standard Cubic Centimetre per Minute
SCPI	Standard Commands for Programmable Instruments
SEE	Secondary Electron Emission
SERT-1	Space Electric Rocket Test 1
SHARK-600V	Southampton Hall Anode-layer Research Krypton 600 V
SmCo	Samarium–Cobalt
Soton	University of Southampton
SPT	Stationary Plasma Thrusters
SS	Stainless Steel
STP	Standard Temperature and Pressure
TAL	Thruster with Anode-Layer
TsNIIMash	Central Scientific Research Institute of Machine Building
T2P	Thrust-to-Power Ratio
U.S.	United States of America
USA	United States of America
UK	United Kingdom of Great Britain and Northern Ireland
UKSA	UK Space Agency
USB	Universal Serial Bus
USD	United States Dollar
USSR	Union of Soviet Socialist Republics
V	Volts
W	Watts
X-CT	X-ray Computed Tomography
Xe	Xenon
XPT	eXternal Plasma Thruster
Zn	Zinc



# Nomenclature

$A_c$	Cross-sectional area of thrust discharge channel	$m^2$
$B$	Magnetic field density	T
$C_{I_d}$	Anode current scaling coefficient	–
$C_{I_{sp}}$	Specific impulse scaling coefficient	–
$C_L$	Channel length scaling coefficient	–
$C_P$	Anode power scaling coefficient	–
$C_T$	Thrust scaling coefficient	–
$C_h$	Channel width scaling coefficient	–
$E_i$	Ion energy	J
$\mathbb{E}[\cdot]$	Expected value operator	–
$F_{cal}$	Calibration force	N
$\mathcal{GP}$	A Gaussian process distribution	–
$H_{ion}$	Ion beam width	m
$I_{axial}$	Axial component of measured plume current	A
$I_{EM}$	Electromagnet current	A
$I_b$	Beam current	A
$I_c$	Cathode keeper current	A
$I_c$	Probe collected current	A
$I_d$	Discharge current	A
$I_{sp}$	Specific impulse	s
$J$	Measured current of the ion species	A
$J_0$	Unattenuated current of the ion species at the thruster	A
$L$	Thruster channel length	m
$L_a$	Anode to the magnetic pole plate distance	m
$L_{ionisation}$	Ionisation length	m
$M$	Total system mass	kg
$M_d$	Total dry spacecraft mass	kg
$M_p$	Total propellant mass	kg

$M_{pl}$	Total payload mass	kg
$M_{pp}$	Total power plant mass	kg
$M_{cal}$	Calibration mass	kg
$N_{cal}$	Number of calibration points	–
$\mathcal{N}(0, \sigma_\epsilon^2)$	Normal distribution with mean 0 and variance $\sigma_\epsilon^2$	–
$P$	Power	W
$P_{KE}$	Kinetic power	W
$P_a$	Anode power	W
$P_c$	Cathode power	W
$P_d$	Discharge power	W
$P_{EM}$	Electromagnet power	W
$P_{Plume}$	Plume power	W
$P_{ion}$	Ionisation power	W
$P_y$	Plumb line length	m
$\mathbb{P}(f   \mathbf{X}_t, \mathbf{y}_t)$	Posterior probability distribution over functions	–
$\mathbb{P}(f   \mathbf{X}_t)$	Prior probability distribution over functions	–
$\mathbb{P}(\mathbf{y}_t   f, \mathbf{X}_t)$	Probability	–
$\mathbb{P}(\mathbf{y}_t   \mathbf{X}_t)$	Marginal probability	–
$R$	Faraday probe radius	m
$R^2$	Coefficient of determination	–
$S_{area}$	Plasma surface area	m <sup>2</sup>
$S_{cal}$	Calibration coefficient	m·N <sup>-1</sup>
$T$	Thrust	N
$T2P$	Thrust to power ratio	N·W <sup>-1</sup>
$T^*$	Critical electron temperature	eV
$T_a$	Anode temperature	K
$T_{eV}$	Plasma electron temperature	eV
$U_c$	Cathode keeper voltage	V
$U_b$	Beam voltage	V
$U_{EM}$	Electromagnet voltage	V
$U_d$	Discharge voltage	V
$V_{MPV}$	Most probable voltage	V
$V_c$	Cathode keeper voltage	V
$V_{plasma}$	Plasma volume	m <sup>3</sup>
$Z_n$	Charge of the n <sup>th</sup> state	C
$a, b$	Power fitting coefficients for SEE emission	–
$a_1, b_1, c_1, a_2, b_2, c_2$	Twin-Gaussian fitting parameters	–

$b_{\text{cal}}$	Offset of calibration	m
$c$	Effective propellant exhaust velocity	$\text{m}\cdot\text{s}^{-1}$
$c$	Speed of light in a vacuum	$\text{m}\cdot\text{s}^{-1}$
$\text{cal}_{\text{error}}$	Error of the calibration	%
$d$	Thruster channel mean diameter	m
$e$	Electron charge	C
$f(\mathbf{x})$	Underlying signal function	–
$f_i(u)$	Twin-Gaussian fit	–
$g(w, T_{eV})$	Maxwell–Boltzmann probability distribution for the microscopic electron velocity	–
$g_0$	Gravitational acceleration	$\text{m}\cdot\text{s}^{-2}$
$h$	Thruster channel width	m
$h_{jk}$	Height of the saddle point of the overlap of the $j^{\text{th}}$ and the $k^{\text{th}}$ ion species twin-Gaussians	–
$h_{\text{mag}}$	The radial separation between the inner and outer magnetic poles	m
$h_{\text{thickness}}$	Channel wall thickness	m
$j(\theta_p)$	Current density at sweep angle $\theta_p$	$\text{A}\cdot\text{m}^{-2}$
$k(\mathbf{x}, \mathbf{x}')$	Covariance function of a kernel function	–
$k_B$	Boltzmann constant	$\text{m}^2\cdot\text{kg}\cdot\text{s}^{-2}\cdot\text{K}^{-1}$
$\dot{m}$	Mass flux	$\text{kg}\cdot\text{s}^{-1}$
$\dot{m}_a$	Anode mass flow rate	$\text{kg}\cdot\text{s}^{-1}$ , $\text{mg}\cdot\text{s}^{-1}$ , sccm
$\dot{m}_n$	Neutral mass flow rate	$\text{kg}\cdot\text{s}^{-1}$
$m_e$	Electron mass	kg
$m_i$	Ion mass	kg
$m(\mathbf{x})$	Mean function of a Gaussian Process	–
$m_n$	Neutral mass	kg
$n$	Number of samples	–
$n_0$	Chamber background neutral number density	$\text{m}^{-3}$
$n_e$	Electron number density	$\text{m}^{-3}$
$n_n$	Neutral number density	$\text{m}^{-3}$
$\bar{q}$	Average charge	C
$r_{L_e}$	Electron Larmor radius	m
$r_{L_i}$	Ion Larmor radius	m
$s_\delta$	Variance of the measured displacement	$\text{m}^2$
$s_{jk}$	The saddle point of the overlap of the $j^{\text{th}}$ and the $k^{\text{th}}$ ion species twin-Gaussians	–

$s_{\text{cal}}$	Variance of the measured calibration coefficient	$\text{m}^2 \cdot \text{N}^{-2}$
$t$	Time	s
$t_{\text{life}}$	Thruster lifetime	s
$u_{i_n}$	Species specific ion velocity	$\text{m} \cdot \text{s}^{-1}$
$u_n$	Neutral velocity	$\text{m} \cdot \text{s}^{-1}$
$v_B$	Bohm velocity	$\text{m} \cdot \text{s}^{-1}$
$v_{ex}$	Exhaust velocity	$\text{m} \cdot \text{s}^{-1}$
$v_i$	Ion velocity	$\text{m} \cdot \text{s}^{-1}$
$v_w$	Wien filter velocity	$\text{m} \cdot \text{s}^{-1}$
$v_\epsilon$	Velocity associated with ionisation energy threshold	$\text{m} \cdot \text{s}^{-1}$
$w$	Microscopic electron velocity	$\text{m} \cdot \text{s}^{-1}$
$\mathbf{x}$	Input vector	–
$\mathbf{X}_t$	Matrix of training input vectors	–
$y$	Output value (e.g., thrust, anode current)	–
$y_i$	True output value	–
$\bar{y}$	Mean of the true outputs	–
$\hat{y}_i$	Predicted output value	–
$\mathbf{y}_t$	Vector of training outputs	–
$z$	Distance between the thruster exit plane and the probe	m
$\%_{\text{mean}}$	Percentage relative to the mean	%
$\Gamma$	Euler Gamma function	–
$\Gamma_{\text{SEE}}$	Secondary electron flux from the wall	$\text{s}^{-1}$
$\Gamma_{ew}$	Electron flux to the wall	$\text{s}^{-1}$
$\Gamma_{iw}$	Ion flux to the wall	$\text{s}^{-1}$
$\Delta$	Voltage loss term	V
$\Delta V$	Change in velocity	$\text{m} \cdot \text{s}^{-1}$
$\Delta\phi_{\text{plate}}$	Voltage difference applied across the $\mathbf{E} \times \mathbf{B}$ probe bias plates	V
$\Delta\Phi_{\text{sheath-wall}}$	Plasma potential drop from sheath to the wall	V
$\Phi$	Potential	V
$\Omega_n$	$n^{\text{th}}$ species current ratio	–
$\Omega_{n,\text{raw}}$	The raw uncorrected $n^{\text{th}}$ species current ratio	–
$\Delta\Omega_n$	Uncertainty in the current fraction for the $n^{\text{th}}$ ion species	–
$\alpha_A(\theta_p)$	The angle between the direction of ion flux and the thruster's axial direction	rad

$\gamma_{\text{SEE}}$	Secondary electron emission coefficient	–
$\gamma_n$	The SEE yield of the collector plate material for each ion species n	–
$\delta$	Small distance calibration mass is displaced	m
$\hat{\delta}i$	Thrust balance displacement response to a force	m
$\epsilon$	Noise term	–
$\epsilon_c$	Specific chemical energy of combustion	J
$\epsilon_e$	Effective energy of impinging electron	J
$\epsilon_i^{n+}$	Ionisation energy of the nth ionisation state	J
$\eta$	Efficiency	–
$\eta_{I_d}$	Current efficiency	–
$\eta_N$	Characteristic nozzle efficiency	–
$\eta_V$	Voltage efficiency	–
$\eta_{\theta_d}$	Beam divergence efficiency	–
$\eta_a$	Anode efficiency	–
$\eta_m$	Mass utilisation efficiency	–
$\eta_q$	Charge efficiency	–
$\theta_d$	Effective beam divergence angle	rad
$\theta_p$	Probe sweep angle	rad
$\kappa_A$	Geometric correction factor that accounts for the probe-thruster angle	–
$\kappa_D$	Correction for differences in path length from near and far channel exit points to the probe	–
$\kappa_G$	Geometric correction factor	–
$\kappa_{\text{SEE}}$	Correction factor accounting for secondary electron emission	–
$\lambda$	Length-scale of a kernel hyperparameter	–
$\lambda_i$	Mean free path of ionisation	m
$v_e$	Electron velocity	$\text{m}\cdot\text{s}^{-1}$
$v_i$	Ionisation frequency	$\text{s}^{-1}$
$\pi$	Pi	–
$\sigma_{\text{cal}}$	Standard deviation of the measured calibration coefficient	$\text{m}\cdot\text{N}^{-1}$
$\sigma_F$	Standard deviation of the measured thrust	N
$\sigma_f^2$	Signal variance of a kernel hyperparameter	–

---

$\sigma_i$	Cross-sectional area of electron impact ionisation	$\text{m}^2$
$\langle \sigma_i \nu_e \rangle$	Reaction rate coefficient	$\text{m}^3 \cdot \text{s}^{-1}$
$\sigma_n$	Effective CEX cross-sectional area	$\text{m}^2$
$\sigma_y$	Standard deviation of the true outputs	–
$\sigma_x$	Standard deviation of displacement measurement	$\mu\text{m}$
$\sigma_\delta$	Standard deviation of the measured displacement	$\text{m}$
$\sigma_\epsilon^2$	Noise variance	–
$\chi$	Power-law scaling exponent for mean channel diameter	–
$\psi$	Excitation and in-elastic collision loss coefficient per ionisation	–

# Chapter 1

## Introduction

The concept of space propulsion was first formally described by Konstantin Tsiolkovsky in his 1903 publication: *Exploration of Outer Space by Means of Rocket Devices* [1, 2], where he proposed electrostatic propulsion as a theoretical means of achieving space travel. While visionary for its time, electric propulsion remained largely theoretical through the early 20<sup>th</sup> century and was widely regarded as a technology of the distant future. This perception was altered with the launch of the first electrically propelled spacecraft on July 20<sup>th</sup> 1964, a milestone that marked the practical realisation of Tsiolkovsky's vision and the dawn of a new era in spacecraft propulsion.

Electric propulsion (EP) is a branch of space propulsion that, unlike chemical propulsion, derives the energy for accelerating propellant from an external electrical power source. EP systems are an attractive propulsion choice for the majority of space missions due to their ability to achieve significantly larger changes in velocity,  $\Delta V$ , for the same propellant mass when compared to chemical systems, enabled by specific impulses typically an order of magnitude higher. This enhanced efficiency enables longer mission durations, greater number of attainable orbits, greater payload masses, and increased flexibility in spacecraft design and operations.

One form of electric propulsion is the Hall effect thruster (HET), which is frequently selected for a wide variety of satellite missions due to its ability to deliver high specific impulse while maintaining a higher thrust-to-power ratio than other EP systems of comparable power. The use of Hall effect thrusters in-space has a history of over half a century, with the first use of a HET being the former Union of Soviet Socialist Republics (USSR) Meteor satellite launched in 1972 [2–4]. These performance advantages have led to the widespread adoption of HETs, making them the most utilised form of propulsion in operation today [5, 6].

Hall effect thrusters are often compared to gridded ion thrusters, another form of electric propulsion, where HETs offer several unique advantages. HETs operate with

higher thrust-to-power ratios, achieve higher power densities for a given physical envelope, and feature simpler thruster and system architectures. These characteristics combine to make HETs relatively lower in cost, capable of higher thrust, more reliable, and therefore they can reduce launch costs. Nevertheless, gridded ion thrusters retain advantages for specific applications, offering higher maximum specific impulses and, in some cases, lower erosion rates.

Historically, research in and in-space applications of HETs has been focused around the 1 kW power regime. However, with the recent increased accessibility of space this has in turn increased interest into smaller spacecraft, which require lower power EP systems such as micro-Hall thrusters (10 W - 300 W). Conversely, there is also a desire to develop very high-power EP systems (10 kW - 100 kW) for next generation interplanetary missions, human orientated missions, and defence applications.

Despite their extensive flight heritage, the fundamental physics governing Hall effect thrusters remains insufficiently understood to enable design from first principles. Unlike chemical propulsion systems, where nozzle geometries can be readily defined using isentropic flow relations, Hall thrusters lack an equivalent analytical framework, making rapid high-performance design impractical. While aspects of HET operation can be investigated through numerical modelling, such approaches are often prohibitively time-consuming and computationally expensive and are typically built upon highly specialised codes developed over years specifically for Hall thrusters. A recent three-dimensional particle-in-cell (PIC) simulation of a small volume (13 mm  $\times$  13 mm  $\times$  26 mm) of a Hall effect thruster plasma required  $1.4 \times 10^6$  Central Processing Unit (CPU) hours (approximately 159 years) to simulate merely 10  $\mu$ s [7]. Whilst the real time taken can be much less with the use of large super-computing clusters, on the order of tens of days. However, no existing model fully captures the entire three-dimensional channel domain and strongly coupled plasma dynamics that govern Hall thruster behaviour.

As a result, the geometric design of new Hall thrusters typically relies on first-order scaling relations informed by large databases of flight-proven xenon thrusters. This empirical approach has proven effective for high-power (kilowatt-class) thrusters operating on xenon, where abundant historical data provides a reliable basis for design. However, for alternative propellants such as krypton or argon, relevant performance databases are often limited or entirely absent, even though different geometric configurations may yield improved performance. This scarcity complicates the design process, particularly when novel geometries or operating regimes are considered.

Additionally, the reliance on existing databases can unintentionally constrain innovation, as designs often reproduce legacy configurations and inherit their associated limitations. Without a strong analytical foundation, such empirical

methods risk reinforcing historical design trends without fully understanding their underlying justification or their adaptability to new regimes.

This challenge is further compounded by the space industry's historical preference for conventional xenon-fed stationary plasma thrusters (SPTs), while alternative HET variants such as the thruster with anode-layer (TAL) have received comparatively little attention. Consequently, both the operational behaviour of alternative propellants and the potential advantages of non-SPT geometries remain under-explored. Nevertheless, these configurations may offer valuable pathways to overcoming the limitations associated with high specific impulse operation and alternative propellant compatibility.

## 1.1 Problem Statement

Hall effect thrusters have become a cornerstone of the satellite industry, offering an attractive balance of thrust density and specific impulse making them well suited for a wide range of space missions [6, 8]. However, there is a heavy reliance on xenon as a propellant, which presents significant limitations that threaten the sustainability and capability of future propulsion systems. Xenon's scarcity and rising cost volatility have created urgent pressure to adopt alternative propellants. Furthermore, the performance limitations of conventional HET designs prevent them from meeting the high specific impulse requirements of ambitious deep-space missions.

The transition to alternative propellants like krypton and argon faces substantial theoretical and practical barriers. Current scaling methodologies, developed primarily for xenon, fail to adequately account for the different plasma physics of lighter propellants, particularly their lower ionisation cross-sections and higher ionisation energies. This limitation is compounded by a reduction in available experimental data, especially at the high voltage regimes ( $>300$  V) that may be necessary for alternative propellants to match the performance of xenon.

A fundamental gap exists in understanding how thruster configuration, particularly channel geometry, interacts with alternative propellant characteristics. Moreover, the absence of systematic studies investigating the coupled effects of propellant choice, channel geometry, and operating voltage creates significant uncertainty in designing Hall effect thrusters optimised for these applications.

Also, TAL configurations are known to achieve greater power densities compared to traditional annular-channel Hall thrusters, suggesting potential advantages for high-power applications [9]. Theoretically, TALs could offer advantages for alternative propellants through higher secondary electron emission thresholds. In practice, their

implementation remains poorly understood, as unresolved stability and thermal management challenges continue to limit their development.

The lack of comprehensive diagnostic data for alternative propellants, particularly in the critical 300–600 V range which could be advantageous for alternative propellants, further hinders the development of accurate scaling relationships. Without reliable methods to predict and optimise performance for non-xenon propellants, the electric propulsion community faces significant challenges in transitioning to more fiscally sustainable propellant options while maintaining the mission capabilities currently achievable with xenon-based systems.

### 1.1.1 Alternative propellants

Xenon is the current propellant of choice for electric propulsion systems, although it was not initially the first selected; the first electric propulsion systems were operated on caesium and mercury [6]. The advent of Xenon's use for space propulsion applications was in 1971, when the USSR's Meteor 1 satellite employed a xenon-fuelled Hall thruster [10]. For electric propulsion systems such as Hall thrusters, thrust levels are inherently constrained by the available electrical power. To maximise thrust in an EP system, a more massive propellant is desirable, as this increases the momentum exchange per atom exhausted. Consequently, both caesium and xenon (heavy elements with stable isotopes) were early choices for EP systems.

In addition to their high atomic mass, these elements offer other advantages: their large atomic "radius" provides a greater cross-sectional area for ionising electron collisions, and their low ionisation energies facilitate efficient plasma generation. From a system perspective, the heavy atomic mass also benefits storage density by reducing the size of high-pressure tanks and completely eliminating the need for tanks in the case of solid propellants. Collectively, these characteristics have led to xenon becoming the dominant propellant for electric propulsion missions over the past several decades.

Despite its excellent properties, including chemical inertness, low ionisation energy, and high atomic mass, xenon's supply and cost presents significant challenges for satellite mission designers. Xenon is produced primarily as a by-product of other industrial processes, and its low natural abundance leads to high costs and volatile supply chains [11]. As a result, the xenon price reached a recent peak of approximately 20 k€/kg in the years following the global pandemic [12]. Global xenon production is approximately 53,000 kg per year [13], meaning missions requiring several tonnes of propellant can significantly impact both cost and global supply. For example, Mars-Earth mission analyses predict the need for 20,000–23,000 kg of xenon for a single round trip, while the Lunar Gateway mission requires approximately 2,750 kg

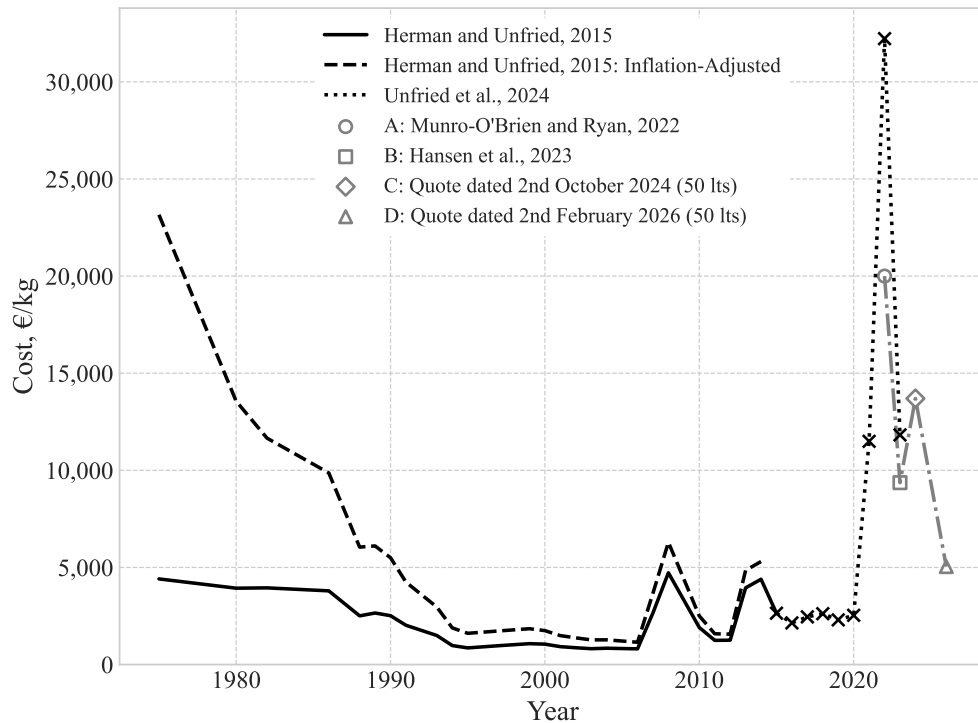


FIGURE 1.1: The cost of a kilogram of xenon in euros over time. Xenon pricing data prior to 2015 are from Herman and Unfried [13]; pricing between 2015 and 2022 is from Unfried et al. [17]; A: 2022 estimate is from Munro-O'Brien and Ryan [12]; B: 2023 value is from Hansen et al. [18]; C: 2024 estimate is from a receipt dated 2 October 2024 for an order quantity of 50 L; D: 2026 estimate is from a quote dated 2 February 2026 for an order quantity of 50 L. The data were adjusted for inflation using the Consumer Price Index for All Urban Consumers (CPI-U) from the U.S. Bureau of Labor Statistics [Series ID: CUUR0000SA0L1E]. Euro values for data prior to 2015 were converted using an exchange rate of 0.93 EUR to 1 USD accessed on 1 April 2025. Euro values for data post-2015 use year-averaged exchange rates.

for orbit transfer and several-tens of kilograms annually for station keeping [14–16]. As shown in Fig. 1.1, xenon prices more than quadrupled between 2015 and 2022, reflecting increasing demand and supply constraints [17]. However, the price as of 2025 has recovered to pre-pandemic levels. Nevertheless, given the projected growth in EP applications, the volatility and cost of xenon are expected to worsen still.

It is important to note that the high cost of xenon alone does not necessarily make it an unsuitable choice for space applications, as the space industry is generally cost-tolerant and recognises that operations in space are inherently expensive. However, it is the volatility in xenon pricing, and the impact this has on long-term planning; that has driven interest in finding alternative propellants. Space missions typically require several years of preparation, and large fluctuations in xenon prices during this period can create significant challenges in budgeting and procurement, making it difficult to ensure mission readiness.

In response to these challenges, krypton has been investigated as an alternative




FIGURE 1.2: Illustrative side-by-side comparison of a cross-section of stationary plasma thruster and a thruster with anode-layer configuration. This schematic is non-physical and intended for conceptual comparison only, highlighting the primary structural and operational differences between the two Hall thruster types.

propellant for nearly two decades. Research has demonstrated promising results at high powers, although progress at lower power levels remains limited [19–23]. However, prior studies have found that krypton operation results in lower anode efficiencies compared to xenon in the same thruster configuration [23–25], primarily due to reduced mass utilisation and current utilisation efficiencies [19]. A lighter alternative, argon, may also be viable, offering similar operational characteristics to krypton and potential additional cost savings. Furthermore, theoretically argon can deliver higher specific impulses for equivalent discharge voltages assuming similar efficiencies.

Parallel to propellant studies, efforts to improve thruster lifetimes have explored novel magnetic field topologies such as magnetic shielding (MS). In addition to reducing channel erosion to approximately zero, MS configurations tend to generate higher plasma temperatures in the ionisation region compared to unshielded configurations [19, 26]. These elevated temperatures can enhance ionisation rates for propellants with higher ionisation energies; however, they also increase radiative losses, presenting a trade-off. Furthermore, magnetic field optimisation beyond purely MS topology can yield improved performance for alternative propellants, an aspect that is currently not encompassed by the design nor scaling methodologies. Moreover, it has been seen that the optimal magnetic field topology for efficient operation is also dependant on discharge voltage and flow rate, suggesting that for alternative propellant where novel discharge voltages are required care in magnetic field design should be taken [27].

In addition to alternative propellants and magnetic topologies, different thruster architectures may offer advantages. The two primary types of Hall thrusters are the stationary plasma thruster (also known as the magnetic layer thruster) and the thruster with anode-layer. These configurations are illustrated in Fig. 1.2. SPTs use ceramic channel walls to confine the plasma, while TAL thrusters employ metallic walls biased to the anode potential. Although TAL thrusters generally experience higher plasma temperatures due to reduced secondary electron emission from metallic walls (a disadvantage for xenon operation), they may prove beneficial for alternative propellants [28, 29]. The hotter plasmas in TALs could increase ionisation rates for harder-to-ionise species like krypton and argon, making them a promising option for broadening the operational envelope of Hall thrusters.

Along with the thruster-specific considerations discussed above, the cathode is of critical importance to the operation of Hall effect thrusters. The cathode supplies the electrons needed for initiating ionisation and for neutralising the exhausted ions. However, there is complex coupling between the bulk thruster plasma and the cathode that directly affects the operation and performance of the thruster, yet this is rarely considered. This oversight largely stems from the fact that cathodes are themselves a separate technology that is not simple to understand; as a result, many electric propulsion researchers prefer to ignore the interplay between thruster and cathode, treating it as a mere integration issue. Regardless, the cathode can be integrated into the thruster system in several ways. The majority opt for an external radial mounting of the cathode due to the simplicity, but the generally higher-performance option is centrally mounting the cathode [30]. Centrally mounted cathodes are placed within the hollowed core of a Hall effect thruster, avoiding the asymmetry of the radially mounted option and generally providing improved performance. The primary drawback is that the thruster must be of sufficient size to house a cathode in its centre, which is not possible for lower-power HETs. Cathodes require propellant to operate; the flow rate is normally a fraction of the thruster flow (1–5%), yet the cathode flow rate relative to thruster operation presents another opportunity for optimisation. Similarly to the other aspects discussed here, the literature on these relationships beyond those empirically derived for xenon is sparse, representing another barrier to alternative propellant HET design.

### 1.1.2 High Specific Impulse

High specific impulses are desirable for missions that require large changes in velocity,  $\Delta V$ ; these are primarily interplanetary or scientific missions. Electric propulsion systems with specific impulses exceeding 3000 s have been developed, such as NASA's Evolutionary Xenon Thruster (NEXT), a 7.3 kW gridded ion thruster [31]. Such high performance has mainly been achieved with gridded ion thruster architectures, which

readily attain high specific impulses but have seen limited application due to high costs and scaling constraints [32, 33]. For example, the T6 gridded ion thruster used by the European Space Agency's (ESA) BepiColombo mission to Mercury achieved specific impulses exceeding 4100 s at a discharge power of 4.5 kW [34].

Fig. 1.3 shows the thrust-to-power ratio of several electric propulsion systems as a function of their operational specific impulse. Each class of EP thruster occupies a distinct envelope of operation. Electrothermal thrusters, such as resistojets and arcjets, exhibit high thrust-to-power ratios but low specific impulses. Hall thrusters occupy an intermediate regime between thrust-to-power ratio and specific impulse. Finally, gridded ion thrusters achieve the highest specific impulses but exhibit relatively low thrust-to-power ratios. The wide range of specific impulses achievable with GITs arises from the decoupling of the ion acceleration mechanism, which allows for the application of very high discharge voltages.

The inverse relationship between thrust-to-power ratio and specific impulse arises from the definition of electric propulsion system efficiency. The kinetic power of the exhaust products compared to the electrical power delivered to the system is given by

$$\eta = \frac{T^2}{2P\dot{m}} = \frac{T}{P} \frac{g_0}{2} I_{sp} \propto \frac{T}{P} I_{sp} , \quad (1.1)$$

where  $\eta$  is the efficiency,  $T$  is the thrust,  $P$  is the electrical power,  $\dot{m}$  is the mass flow rate,  $g_0$  is acceleration due to gravity at sea-level, and  $I_{sp}$  is the specific impulse. From this, it is clear that for constant efficiency, the thrust-to-power ratio and specific impulse are inversely proportional. These relations will be discussed in Chapter 2 in greater detail.

The discussion of Fig. 1.3 highlights a secondary motivation for this thesis. Hall thrusters already demonstrate significant versatility; therefore, if the specific impulse of these systems can be increased without sacrificing efficiency, their applicability would broaden. Although increasing specific impulse at constant efficiency reduces the thrust-to-power ratio, the favourable thrust-to-weight and thrust-to-volume characteristics could be extended into high specific impulse applications. Despite being able to achieve higher specific impulses than HETs, GIT systems are often excluded from many missions due to their inherent complexity, numerous potential failure modes, and poor mass scaling with power compared to HETs [67, 68].

Accordingly, this thesis evaluates the potential of alternative propellants to extend the operational envelope of Hall thrusters towards higher specific impulses. The use of lighter propellants such as krypton and argon, together with operation at elevated discharge voltages, may enable Hall thrusters to enter the specific impulse regime currently occupied by gridded ion thrusters. Achieving specific impulses in excess of 3000 s, even at the expense of reduced thrust-to-power ratios, would represent a

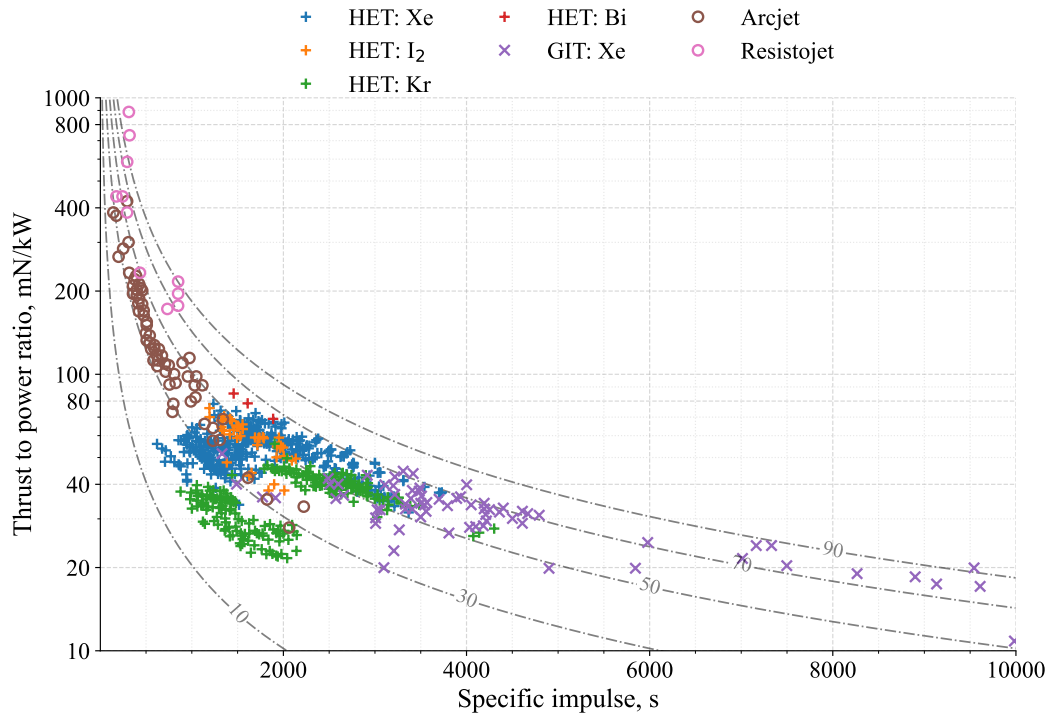


FIGURE 1.3: Representation of several popular types of electric propulsion being compared directly with Hall thrusters operating on alternative propellants being separated. The dot-dash iso-lines represent the total efficiency as labelled within the line. Data plotted is taken from [21, 23, 24, 35–66].

mission-enabling improvement. Such an advancement would leverage the inherently favourable power and thrust scaling characteristics of Hall thrusters compared with gridded ion thrusters, offering significant advantages in system mass, physical envelope, and overall mission cost.

## 1.2 Research Objectives and Contributions

The primary aim of this thesis is to advance the understanding of Hall thruster operation on alternative propellants, with a focus on how channel geometry couples to thruster performance. Furthermore, this work aims to enable the utilisation of alternative, low-mass propellants to attain high specific impulse operation. The investigation is pursued through several methods:

1. **Modification of current scaling methodologies to enable direct scaling of initial Hall thruster geometry for operation with non-xenon propellants.** Extensions to existing scaling methodologies were derived and applied to a database of Hall thrusters from the literature. The updated scaling approach enables direct scaling for krypton using the current database.

**2. Implementation of a thrust measurement system capable of accurately characterising thruster performance.**

An inverted double pendulum thrust balance was designed, manufactured, and validated for direct thrust measurements. To verify the accuracy of the new balance, a collaborative testing campaign with the University of Surrey was conducted, during which a novel external plasma thruster developed at Surrey was tested at both facilities. Thrust measurements obtained with the new balance were successfully compared against data from Surrey's established torsional thrust balance, demonstrating the accuracy and reliability of the system.

**3. Design and manufacture of a Hall thruster capable of interrogating the relationship between channel geometry and thruster performance.**

The "Southampton Hall Anode-layer Research Krypton 600 V" (SHARK-600V) thruster was designed with variable discharge channel inserts that could be exchanged without modifying the thruster testing setup.

**4. Comprehensive testing of alternative gaseous propellants in a variable channel geometry Hall thruster scaled for krypton operation.**

Direct thrust measurements of the SHARK-600V were performed at the University of Southampton for nine distinct channel geometries, operating on xenon, krypton, and argon, yielding a total of 1220 unique data points.

**5. Development and application of a Gaussian process regression surrogate model.**

A Gaussian process regression model was developed from the Southampton experimental dataset of the SHARK-600V Hall thruster, creating a continuous, smooth function describing thruster performance across the input parameter space of anode voltage, anode mass flow rate, channel width, and mean channel diameter. Separate models were trained for each propellant and each output parameter, providing predictions with quantified uncertainty. The optimal geometries identified by the GPR model showed meaningful agreement with the semi-empirical scaling predictions, most notably for krypton at the 2.5 kW, 600 V condition where the predicted mean diameters agreed to within 0.6 mm. This serves as experimental validation of the scaling approach.

**6. Development and application of diagnostic techniques for non-invasive investigation of Hall thruster discharges.**

The SHARK-600V, in a single configuration, was transported to the University of Michigan's Large Vacuum Test Facility (LVTF) at the Plasmadynamics and Electric Propulsion Laboratory (PEPL) for krypton testing. Here, in collaboration with PEPL staff, the thruster was operated over the same parameter envelope as at Southampton, and far-field plume measurements were obtained using a suite of probes, including a retarding potential analyser, a Faraday probe, and an  $E \times B$  probe. These diagnostics were supplemented with direct thrust measurements.

### **7. Investigation of the feasibility and performance characteristics of thruster-with-anode-layer configurations operating on alternative propellants.**

An optional TAL configuration was incorporated into the SHARK-600V thruster architecture to enable preliminary investigations of anode-layer designs for alternative propellant operation. Although full characterisation was constrained by thermal and operational stability challenges, the work provides initial experimental insights into TAL behaviour at elevated voltages and highlights critical areas such as electron confinement and channel wall interactions that must be addressed for successful future TAL development.

This work makes four principal contributions to the field of electric propulsion. First, it demonstrates that high-voltage operation can mitigate the ionisation limitations of lighter propellants, with argon exhibiting particular resilience to voltage-induced efficiency losses. Second, the modular thruster design reveals how channel geometry interacts with propellant choice, offering a large dataset for optimising future alternative-propellant Hall effect thrusters. Third, the extended scaling methodology offers a pathway to design alternative-propellant Hall effect thrusters once sufficient experimental data are collected. Fourth, the development of a Gaussian process regression surrogate model enables continuous interpolation across the design space and provides experimental validation of the scaling methodology using the experimental dataset. Collectively, these advances bridge the gap between conventional xenon-based scaling and the needs of next-generation thrusters, supporting the adoption of krypton and argon in missions where cost and specific impulse are critical. By decoupling geometric and propellant effects, this research lays the groundwork for further exploration of high-voltage, alternative-propellant Hall effect thrusters.

## **1.3 Organisation**

The structure of this thesis is as follows: Chapter 1 introduces the motivation, context, and objectives of the research, with particular focus on the hurdles and drawbacks of operating Hall effect thrusters on alternative propellants.

Chapter 2 provides the historical and theoretical background to electric propulsion, with emphasis on the physics and performance evaluation of Hall effect thrusters, and a review of candidate alternative propellants.

Chapter 3 presents the development of the extended semi-empirical scaling methodology, reviewing existing approaches, introducing a database of Hall thruster dimensions and performance, and deriving a scaling methodology modified for alternative propellants at high-voltage operation used in this work.

Chapter 4 details the design of the Southampton High-voltage Anode-layer Research Krypton 600 V (SHARK-600V) thruster, describing its modular configuration, magnetic field optimisation, anode design, and thermal modelling considerations.

Chapter 5 describes the Southampton experimental campaign, including the facility, diagnostics, and thrust balance development. Results for xenon, krypton, and argon across nine thruster geometries are presented, together with an assessment of anode-layer configurations. A Gaussian process regression model trained from this dataset enables continuous interpolation between tested geometries and comparison with the semi-empirical scaling methodology, providing experimental validation of the scaling approach and offering the first experimentally derived channel geometry guidance for argon operation.

Chapter 6 outlines the inter-laboratory testing campaign conducted at the University of Michigan. This chapter introduces several far-field plume diagnostic techniques and focuses on the cross-validation of the Southampton results through comparison with direct thrust measurements and detailed plume diagnostics.

Chapter 7 summarises the key findings, the conclusion of the thesis and proposed future work. A record of the associated publication history is provided in Chapter 8.

## Chapter 2

# Background and Fundamentals

### 2.1 Introduction

This chapter provides the necessary background to support the motivations, methodologies, and objectives of this thesis. Fundamental plasma physics derivations are omitted to maintain focus, but recommended references are provided for readers seeking a basic introduction to plasma concepts [69, 70]. The broader context of space propulsion is first discussed. Several classes of electric propulsion are introduced and described, alongside an overview of the key performance metrics used to evaluate propulsion.

The chapter then presents a detailed discussion of Hall effect thrusters, covering their structural design, operating principles, and performance evaluation methods. Derivations for thrust, specific impulse, and anode efficiency are developed under idealised assumptions and later corrected for practical considerations. Finally, the role of alternative propellants is examined, with a focus on the limitations of xenon and the motivation for investigating gaseous and solid propellants that could enhance the versatility and reduce the cost of future electric propulsion missions.

#### 2.1.1 Basic Rocket Science

The method of quantifying space propulsion was first described by Konstantin Tsiolkovsky, who introduced the equations governing the change in velocity of a variable-mass system. This foundational result, known as the Tsiolkovsky rocket equation, describes the effect of constant thrust acting on a system with decreasing mass. The force experienced by the spacecraft can be expressed as

$$\text{Force} = M(t) \frac{dv(t)}{dt} , \quad (2.1)$$

where  $M(t)$  is the spacecraft mass and  $v(t)$  is the spacecraft velocity, both as functions of time  $t$ , while the force is applied. The thrust results from the conservation of momentum between the spacecraft and the exhausted propellant. Thus, the thrust experienced by the spacecraft is equal and opposite to the change in momentum of the propellant, expressed as

$$Thrust = -\frac{d}{dt}(M_p c) = -\frac{dM_p}{dt}c = -\dot{M}_p c , \quad (2.2)$$

where  $M_p$  is the propellant mass of the spacecraft and  $c$  is the effective propellant exhaust velocity in the spacecraft's frame of reference. Note that in Eq. 2.2, any contribution from pressure thrust, which would be present for a chemical system operating in space without an infinitely expanding nozzle, is neglected for simplicity and due to its generally low magnitude with respect to thrust produced [68, 71]. In practice, the total mass of the spacecraft is a function of time and can be expressed as

$$M(t) = M_d + M_p(t) , \quad (2.3)$$

where  $M_d$  is the dry mass of the spacecraft, which, for chemically propelled applications, is typically equal to the payload mass. However, for electrically powered propulsion systems,  $M_d = M_{pl} + M_{pp}$ , where  $M_{pl}$  and  $M_{pp}$  represent the payload mass and the power plant mass, respectively.

From Eq. 2.3, it can be seen that the rate of change of the spacecraft mass with time is determined solely by the rate of propellant consumption, such that

$$\frac{dM}{dt} = \frac{dM_p}{dt} . \quad (2.4)$$

Substituting Eq. 2.4 into Eq. 2.1 and Eq. 2.2 yields

$$M \frac{dv}{dt} = -c \frac{dM}{dt} , \quad (2.5)$$

which can also be rearranged as

$$dv = -c \frac{1}{M} dM . \quad (2.6)$$

Solving this for an initial velocity and final velocity,  $v_i$  and  $v_f$ , and an initial and final mass,  $M_i = M_d + M_p$  and  $M_f = M_d$ , via integration such

$$\int_{v_i}^{v_f} dv = -c \int_{M_d+M_p}^{M_d} \frac{1}{M} dM . \quad (2.7)$$

This integrates to

$$v_i - v_f = -\Delta V = c \ln \left[ \frac{M_d}{M_d + M_p} \right] . \quad (2.8)$$

This equation is often expressed in an alternative form by introducing the specific impulse,  $I_{sp}$ , which standardises the representation across both imperial and metric unit systems. Thus, the final form of the rocket equation is

$$\Delta V = \frac{c}{g_0} g_0 \ln \left[ \frac{M_d + M_p}{M_d} \right] = I_{sp} g_0 \ln \left[ \frac{M_d + M_p}{M_d} \right] , \quad (2.9)$$

where  $I_{sp}$  is the specific impulse and  $g_0$  is the standard gravitational acceleration at Earth's surface. Thrust and specific impulse are critical metrics for the evaluation of propulsion systems, as they enable mission planning and spacecraft mass budgeting. From both Eq. 2.2 and Eq. 2.9, the relationship between a propulsion system's operational characteristics and the required propellant mass fraction can be derived.

It is also useful to note that if the thrust and mass flow rate are preserved over the operating lifetime of the thruster, the specific impulse reduces to

$$I_{sp} = \frac{T \int_0^t dt}{g_0 \dot{m}_p \int_0^t dt} = \frac{T}{g_0 \dot{m}_p} = \frac{c}{g_0} . \quad (2.10)$$

Here, Eq. 2.10 represents the more commonly used form when evaluating the specific impulse of a thruster under steady thrust and mass flow rate conditions.

### 2.1.2 Chemical Space Propulsion

Chemical propulsion has the longest heritage in spaceflight and remains a key technology enabling space exploration. By utilising large propellant flow rates and the chemical energy stored within molecular bonds, chemical propulsion systems are capable of producing very high thrust outputs at modest specific impulse values ( $I_{sp} \leq 320$  s).

By utilising the chemical energy stored within the propellant, chemical thrusters can achieve very high kinetic powers, described as

$$P_{KE} = \frac{1}{2} \dot{m} v_{ex}^2 = \eta_N \epsilon_c \dot{m} , \quad (2.11)$$

where  $\dot{m}$  is the propellant mass flow rate,  $v_{ex}$  is the exhaust velocity,  $\eta_N$  is a characteristic nozzle efficiency value and  $\epsilon_c$  is the total mass-specific chemical energy released in combustion. In chemical propulsion systems, the exhaust velocity is determined by the energy released during combustion and the expansion through the nozzle, making it a function of both the propellant chemistry and the nozzle efficiency.

In contrast, for electric propulsion systems, the exhaust velocity can be externally controlled via the applied electrical power. The kinetic power in chemical propulsion

depends solely on the flow rate and chemical energy released in combustion, allowing chemical thrusters to operate at kinetic powers on the order of gigawatts.

However, the use of chemical reactants inherently limits the available energy per unit mass to the combustion energy of the propellant. Although large mass flow rates (analogous to high electrical current) are achievable, making chemical propulsion ideal for launch vehicles and high-thrust applications, the energy imparted per unit mass (analogous to electrical voltage) is fundamentally limited by chemical bond energies. Consequently, chemical propulsion systems achieve relatively low maximum specific impulses compared to electric propulsion, which can impart several orders of magnitude more energy per unit mass of propellant.

This limitation makes chemical propulsion poorly suited for in-space applications where high specific impulse is more critical than high thrust. Nevertheless, the high thrust levels achievable with chemical propulsion maintain its importance for in-space applications requiring impulsive  $\Delta V$  manoeuvres, such as collision avoidance or when time-to-final orbit is mission critical.

### 2.1.3 Electric Space Propulsion

There are several forms of electric propulsion that have been used, categorised by the primary acceleration method employed: electro-thermal, electro-static, electro-magnetic, and magneto-electrostatic. These systems are particularly valued for their ability to achieve significantly higher specific impulses than conventional chemical propulsion, enabling long-duration and fuel-efficient missions.

The majority of the flight heritage in electric propulsion falls within the electro-thermal and electro-static categories. Electro-thermal thrusters, such as resistojets and arcjets, utilise electrical power to thermally heat the propellant, which is then exhausted through a traditional nozzle, making their operation most analogous to chemical propulsion systems. Electro-static thrusters have strong heritage within the United States, notably with the early flights of caesium gridded ion thrusters on the Space Electric Rocket Test 1 (SERT-1) spacecraft in 1964, which achieved approximately 22 mN of thrust [72].

A related class of propulsion is electro-magneto-static propulsion, which is the category in which Hall effect thrusters reside. This term refers to the use of both a static axial electric field and a static radial magnetic field utilised in HETs. Hall effect thrusters, first tested aboard the Soviet Meteor satellite in 1972, operate via a unique acceleration mechanism that cannot be fully classified as purely magneto-static or electro-static [2]. Consequently, they are categorised as magneto-electrostatic thrusters. In HETs, the static electric field does not directly accelerate the ions; rather,




FIGURE 2.1: Schematic diagram showing the cross-section of a traditional gridded ion thruster [73].

it energises the magnetically confined electrons, which in turn create a “virtual” grid that accelerates ions.

A key requirement for electric thrusters is the maintenance of spacecraft charge neutrality. Without it, the spacecraft bus would accumulate negative charge due to the ejection of positive ions; eventually electro-statically attracting the exhausted ions back to the spacecraft and negating net thrust. This is mitigated by the use of an external cathode, which emits cold electrons into the plume, neutralising the ion beam and maintaining spacecraft electrical balance.

### 2.1.3.1 Gridded Ion Thrusters

Gridded ion thrusters are the primary alternative to Hall effect thrusters for high  $\Delta V$  missions, with comparable power demands but significantly higher achievable specific impulses, typically on the order of several thousand seconds. This advantage arises from the decoupling of ionisation and acceleration processes in GITs; in contrast, Hall thrusters couple these processes. While xenon is the most common propellant, alternatives such as krypton, iodine, and caesium have also been investigated [6].

The separation of ionisation and acceleration in GITs allows grid structures to be biased at extremely high discharge voltages without significantly heating the plasma, which would otherwise impose thermal stress on the thruster body. GITs must employ multiple grids for efficient operation, as shown in Fig. 2.1, with components

such as the screen grid used to prevent poorly accelerated ions from eroding structural elements and to aid ionic beam collimation.

The most common configuration, illustrated in Fig. 2.1 and termed the electron ionisation GIT, requires two electron sources: one to ionise the propellant within the discharge chamber and another to neutralise the exhausted ion beam after acceleration. Typically, hollow cathodes have shorter operational lifetimes and higher failure rates than the thrusters themselves. Therefore, using two cathodes can significantly increase the overall failure risk compared to single-cathode systems, such as HETs. Variants of GITs that avoid the need for multiple hollow cathodes, such as radio-frequency ionisation GITs, still require an external cathode for neutralisation.

Additionally, the ion current per grid aperture in a GIT is limited by the space-charge limit. As a result, scaling GITs to produce more beam current, and thus more thrust, at higher powers requires increasing the number of grid apertures. At very high powers, this leads to grids of large physical size, resulting in prohibitively high thruster masses and volumes that may be incompatible with spacecraft payload fairings.

### 2.1.3.2 Hall Effect Thrusters

Hall effect thrusters are an attractive option for spacecraft propulsion due to their unique combination of higher thrust-to-power ratios, good specific impulses (1500–2000 s), and simple physical architecture. As shown in Fig. 2.2, HETs feature an open annular channel structure where plasma is confined by an  $\vec{E} \times \vec{B}$  drift between a radial magnetic field and an axial electric field. This configuration traps electrons, which subsequently ionise the neutral propellant via electron impact ionisation.




FIGURE 2.2: Illustrative cross-section of a Hall effect thruster featuring a ceramic discharge channel. The three key geometric parameters used throughout this thesis: channel length  $L$ , channel width  $h$ , and mean channel diameter  $d$  are indicated.

HETs have a long history of operational use in space and are considered a mature propulsion technology, with more active satellites using Hall thrusters than any other

propulsion system as of 2022 [5]. While these systems are highly reliable, their operational lifetimes are often limited by erosion of the channel walls due to high-energy ions impinging on the ceramic discharge channel. Consequently, HETs have historically been considered less desirable for very long-duration missions compared to gridded ion thrusters, which can operate without direct material erosion. However, recent advances such as magnetic shielding of the channel have demonstrated the potential to extend Hall thruster lifetimes by several orders of magnitude [74].

Hall thrusters demonstrate a broad operational envelope. Extremely low-power HETs have been developed, operating at power levels as low as 20 W [75, 76], while other designs have achieved operation in excess of 100 kW [77, 78]. Future designs targeting 300–600 kW are under development for integration with nuclear fission-powered spacecraft [78, 79].

Moreover, HETs show considerable versatility even within fixed geometries. Standard designs have been successfully operated at currents up to ten times their nominal power without major degradation in performance. Similar operational flexibility has been observed in low-power HETs as well, achieving good performance at multiple times their nominal input power [12, 80, 81].

Collectively, these factors establish HETs as highly promising candidates for a wide range of satellite propulsion applications. The maturity and operational flexibility of Hall thrusters continue to drive research into novel design variants and optimisations for specific mission requirements.

## **2.2 Physics of Hall Effect Thruster Operation**

A typical Hall thruster consists of an annular hollow ceramic discharge channel with a metallic anode located at its base, which also serves as the propellant injection point. The primary thruster structure is composed of a soft magnetic material, which, together with permanent magnets or electromagnets, produces a radial magnetic field orthogonal to the electric field established between the anode and the cathode. Although the magnetic field is nominally radial, its magnitude and curvature vary significantly along both axial and radial directions, forming a complex axisymmetric two-dimensional topology. This topology governs electron confinement, ionisation distribution, and acceleration region placement, and is therefore crucial for efficient thruster design. However, unlike discharge voltage or channel geometry, magnetic topology lacks a scalar description beyond peak magnitude, making direct comparison between thrusters difficult. Consequently, magnetic configuration is often simplified at the scaling stage to categorical distinctions of magnetically shielded

versus unshielded. However, despite this some attempts to provide analytical and computational methods for magnetic topological design have been proposed [82, 83].

A Hall thruster's physical structure can be broadly characterised by three key channel dimensions. The channel, where plasma confinement, ionisation, and acceleration occur, plays a critical role in determining thruster performance and efficiency. The channel length refers to the axial distance between the thruster exit plane and the anode at the base of the channel. The inner and outer radii are typically combined into two parameters: the mean channel diameter and the channel width. The mean diameter defines the average radial position of the discharge, while the channel width is the difference between the inner and outer radii, which together determine the available channel cross-sectional area for plasma processes. Careful selection of these parameters is essential in the design and scaling of Hall thrusters, as they directly influence ionisation efficiency, magnetic field configuration, and operational characteristics.

Hall thrusters utilise the Lorentz force and the  $\vec{E} \times \vec{B}$  drift to accelerate ions that produce thrust. The  $\vec{E} \times \vec{B}$  drift refers to the motion of charged particles in mutually orthogonal electric and magnetic fields, resulting in a spiral trajectory with a net drift perpendicular to both fields.

Within the Hall thruster, thrust is generated through the electrostatic acceleration of ions, while the reaction force is transferred to the thruster body via the restoring Lorentz force acting on the confined electrons. This mechanism defines Hall thrusters as a class of electro-magneto-static propulsion systems [68].

## 2.3 Evaluation of Hall Effect Thruster Performance

### 2.3.1 Thrust

Largely, the force produced by a Hall thruster can be approximated by considering a simple electrostatic force accelerating the ion plume through an effective potential, under the assumption that electrons are well confined. This assumption simplifies analysis, avoiding the need to explicitly account for the complex electron momentum transfer mechanisms to the thruster body. The reaction force on the spacecraft from ion acceleration can be expressed as

$$T = \frac{d}{dt} (m_i v_i) = \dot{m}_i v_i , \quad (2.12)$$

where  $\dot{m}_i$  is the mass flow rate of the exhausted ions, and  $v_i$  is the exhaust velocity.

The mass flow rate of ions can be related to the ion beam current as

$$\dot{m}_i = \frac{I_b m_i}{\bar{q}} , \quad (2.13)$$

where  $I_b$  is the ion beam current,  $m_i$  is the ion mass, and  $\bar{q}$  is the average ion charge.

The ion exhaust velocity can be approximated by considering the acceleration of an ion through the beam potential  $U_b$ , such that

$$v_i = \sqrt{\frac{2\bar{q}U_b}{m_i}} , \quad (2.14)$$

where  $U_b$  represents the bulk plasma potential. In Hall thrusters, this is typically approximated by the discharge voltage minus a voltage loss term, such that

$$U_b = U_d - \Delta , \quad (2.15)$$

where  $\Delta$  represents voltage losses.

Substituting the expressions for ion mass flow rate and ion velocity (given in Eqs. 2.13 and 2.14, respectively) into the general thrust equation from Eq. 2.12, yields

$$T = \frac{I_b m_i}{\bar{q}} \sqrt{\frac{2\bar{q}U_b}{m_i}} . \quad (2.16)$$

This expression contains many values that are not known or directly controlled during the operation of a Hall thruster. For example,  $I_b$  and  $U_b$  are both values that are experimentally obtained. To get the thrust equation in a form that can be used in a predictive manor several efficiency values are required.

### 2.3.2 Specific Impulse

For the calculation of specific impulse in HETs, the definition from Eq. 2.10 is used in a modified form, considering only the mass flux of the neutral propellant such

$$I_{sp} = \frac{T}{g_0 \dot{m}_p} . \quad (2.17)$$

Moreover, substituting in Eq. 2.16 and Eq. 2.13 where  $\dot{m}_p \approx \dot{m}_i$ , specific impulse for Hall effect thrusters can be seen to be

$$I_{sp} = \frac{1}{g_0 \dot{m}_i} \dot{m}_i \sqrt{\frac{2\bar{q}U_b}{m_i}} = \frac{1}{g_0} \sqrt{\frac{2\bar{q}U_b}{m_i}} . \quad (2.18)$$

### 2.3.3 Thruster Efficiency

The efficiency of a thruster can be described as the ratio of the axial kinetic power of the exhausted ions to the total input power supplied to the propulsion system. The total input power for a Hall thruster is given by

$$P_{\text{Total}} = \underbrace{I_d U_d}_{P_a} + \underbrace{I_c U_c}_{P_c} + \underbrace{I_{\text{EM}} U_{\text{EM}}}_{P_{\text{EM}}} , \quad (2.19)$$

where  $I_d$  and  $U_d$  are the anode discharge current and voltage,  $I_c$  and  $U_c$  are the cathode current and voltage, and  $I_{\text{EM}}$  and  $U_{\text{EM}}$  are the coil current and voltage if an electromagnet (EM) is used.

The kinetic power of the ion plume is described as

$$P_{\text{Plume}} = \frac{d}{dt} (KE_{\text{Plume}}) = \frac{1}{2} \frac{d(m_i v_i^2)}{dt} = \frac{1}{2} \dot{m}_i v_i^2 , \quad (2.20)$$

where  $KE_{\text{Plume}}$  is the kinetic energy of the plume,  $m_i$  is the ion mass, and  $v_i$  is the exhaust velocity.

Substituting Eq. 2.12 into Eq. 2.20 yields

$$P_{\text{Plume}} = \frac{1}{2} \dot{m}_i v_i^2 = \frac{1}{2 \dot{m}} \dot{m}^2 v_i^2 = \frac{T^2}{2 \dot{m}} . \quad (2.21)$$

Thus, the efficiency can be expressed as

$$\eta_{\text{total}} = \frac{P_{\text{Plume}}}{P_{\text{total}}} = \frac{T^2}{2(I_d U_d + I_c U_c + I_{\text{EM}} U_{\text{EM}}) \dot{m}} , \quad (2.22)$$

where  $\dot{m}$  is the total mass flow to the thruster, which is the sum of the anode and cathode flow rates.

When discussing thruster design, it is useful to define an anode efficiency,  $\eta_a$ , that only considers the power and mass flow associated with the anode, as the efficiencies of the cathode and magnetic circuit are not the focus. This is defined as

$$\eta_a = \frac{T^2}{2 P_a \dot{m}_a} , \quad (2.23)$$

where  $\dot{m}_a$  is the neutral mass flow rate to the anode and  $P_a = I_d U_d$  is the anode input power.

### 2.3.3.1 Anode Efficiency

The anode efficiency can also be thought of as the product of several other efficiency terms that represent individual processes as

$$\eta_a = \eta_{I_d} \eta_q \eta_m \eta_V \eta_{\theta_d} , \quad (2.24)$$

where  $\eta_{I_d}$  is the current efficiency (also known as beam efficiency),  $\eta_q$  is the charge efficiency,  $\eta_m$  is the mass efficiency (also known as mass utilisation),  $\eta_V$  is the voltage efficiency, and  $\eta_{\theta_d}$  is the beam divergence efficiency.

### 2.3.3.2 Current Efficiency

The current efficiency represents the fraction of discharge current contributing to the ion beam rather than being lost to the walls. It is defined as

$$\eta_{I_d} = \frac{I_b}{I_d} , \quad (2.25)$$

where  $I_b$  is the ion beam current and  $I_d$  is the anode discharge current.

### 2.3.3.3 Charge Efficiency

The ion plume contains a non-zero quantity of multiple charged states and cannot be assumed to be purely monogenetic. Higher ionisation states provide increased thrust for the same potential drop; however, as seen in Eq. 2.14, the velocity, and by extension, the momentum exchange, of the ion scales with the square root of the charge and is therefore less efficient. To quantify this charge efficiency,  $\eta_q$ , it can be seen that

$$\eta_q = \frac{\left( \sum_n \frac{\Omega_n}{\sqrt{Z_n}} \right)^2}{\sum_n \frac{\Omega_n}{Z_n}} , \quad (2.26)$$

where  $Z_n$  is the  $n^{\text{th}}$  charge state, and  $\Omega_n$  is defined as

$$\Omega_n = \frac{I_n}{\sum_{n=1}^N I_n} = \frac{I_n}{I_b} , \quad (2.27)$$

where,  $I_n$  is the total beam current contribution from the  $n^{\text{th}}$  ion species, and  $I_b$  is the total measured beam current. Experimental data shows that the majority of the ion population under nominal conditions is singly charged, with negligible populations with  $Z_n > 3$ . However, in principle, charge states can extend up to the atomic number of the propellant species.

### 2.3.3.4 Mass Efficiency

The mass utilisation efficiency,  $\eta_m$ , accounts for the fraction of the input propellant converted to thrust producing ions, such that

$$\eta_m = \frac{\dot{m}_i}{\dot{m}_n} = \frac{I_b m_n}{\bar{q} \dot{m}_n} = \eta_{I_d} \frac{I_d m_n}{\bar{q} \dot{m}_n}, \quad (2.28)$$

where  $\dot{m}_i$  is the ion mass flow rate,  $\dot{m}_n$  is the neutral propellant mass flow rate into the thruster,  $m_n$  is the atomic mass, and  $\bar{q}$  is the average charge state.

To obtain an accurate estimate of the mass efficiency, both the ion beam current and the ion species population must be measured using a Faraday probe and an “E × B” probe, respectively. However, a rough estimation can be made by using the thrust from Eq. 2.16, and assuming  $m_n \approx m_i$ ,  $\bar{q} \approx e$ , and  $U_b \approx U_d$ , to form

$$\eta_m = \frac{I_b m_n}{\bar{q} \dot{m}_n} \frac{1}{\dot{m}_n} \approx \frac{T}{\dot{m}_n} \sqrt{\frac{m_i}{2\bar{q}U_b}} = \frac{T}{\dot{m}_n} \sqrt{\frac{m_n}{2eU_d}}. \quad (2.29)$$

### 2.3.3.5 Voltage Efficiency

A voltage utilisation efficiency,  $\eta_V$ , is introduced to account for the non-ideal acceleration of ions. The accelerating potential in Hall thrusters is determined by the self-consistent electric field maintained by confined electrons. This results in a complex relationship between the applied anode voltage and the accelerating potential experienced by the ions.

Moreover, because the ionisation and acceleration regions in a HET are not discrete, ions can be produced throughout the plasma and, as a result, may experience different accelerating potential drops. Additionally, it is understood that a portion of the voltage (on the order of tens of volts) is lost in extracting electrons from the cathode. As such

$$\eta_V = \frac{U_b}{U_d} \approx \frac{U_d - \Delta}{U_d}, \quad (2.30)$$

where  $U_b$  is the effective beam voltage and  $\Delta$  is the voltage drop between the anode potential and the plasma acceleration potential.

### 2.3.3.6 Divergence Efficiency

The beam divergence efficiency,  $\eta_{\theta_d}$ , corrects for the fact that ions do not accelerate strictly along the thruster central axis and have a non-axial momentum component. It is defined as

$$\eta_{\theta_d} = \cos^2(\theta_d). \quad (2.31)$$

Here,  $\theta_d$  is the effective beam divergence half-angle, defined as

$$\theta_d = \cos^{-1} (\langle \cos(\theta) \rangle_j) = \cos^{-1} \left( \frac{I_{\text{axial}}}{I_b} \right) . \quad (2.32)$$

In this expression,  $\langle \cos(\theta) \rangle_j$  represents the current-weighted average cosine of the beam half-angle,  $I_{\text{axial}}$  is the axial component of the measured plume current, and  $I_b$  is the total collected ion beam current. Calculations of  $I_{\text{axial}}$  and  $I_b$  from Faraday probe measurements are described in Sec. 6.3.3.

### 2.3.4 Thrust-to-Power Ratio

The thrust-to-power ratio is a critical performance metric in electric propulsion, directly linking the propulsive force generated to the electrical power consumed. It provides an important measure of system efficiency and operational capability, especially when comparing different thruster types or scaling designs for specific mission requirements. Higher thrust-to-power ratios are generally desirable for minimising transit times or maximising payload delivery for a given onboard power budget.

Thrust-to-power ratio is simply defined as

$$T2P = \frac{T}{P_{\text{Total}}} , \quad (2.33)$$

where  $T2P$  is the thrust-to-power ratio,  $T$  is the thrust, and  $P_{\text{Total}}$  is the total electrical power input to the thruster system.

## 2.4 High Voltage Operation

Numerous experimental campaigns have explored the operation of xenon Hall thrusters at elevated discharge voltages (typically exceeding 300 V), with varying objectives and generally favourable outcomes. These investigations have frequently been conducted using thrusters originally optimised for nominal operation around 300 V, with the intent of characterising their extended operational envelope rather than specifically maximising performance at higher voltages [84].

Despite this, such studies have contributed substantially to the understanding of high-voltage discharge dynamics. One of the primary limitations identified at elevated voltages is the significant increase in heat deposition to the channel walls, which adversely impacts thruster efficiency. Notably, this issue has been shown to be partially mitigated through the use of modified magnetic field configurations. For example, configurations that shift the magnetic field peak outside the channel have




FIGURE 2.3: A schematic view of the cross section of a thruster with anode-layer with the hollow anode and the guard rings highlighted.

demonstrated improved performance by ensuring that the majority of the potential drop occurs downstream of the ionisation region [49, 85]. In one such configuration, peak efficiency was achieved at a discharge voltage of 850 V. Furthermore, recent testing of the H10 Hall thruster, a conducting wall magnetically shielded thruster with novel heat pipes, was efficiently operated up to 800 V [86].

Further research has examined high-voltage operation using alternative propellants, such as krypton. Jakubczak et al. [87] demonstrated that high specific impulses could be achieved using krypton at elevated voltages. However, the study also revealed significant limitations, including pronounced channel erosion. This was attributed to the use of a conventional magnetic field topology that did not incorporate magnetic shielding, thereby allowing direct plasma–wall interaction at high energies.

Although most modern efforts have focused on SPT-type Hall thrusters, early investigations into TAL configurations exhibited particularly promising results at discharge voltages exceeding 1 kV. Studies of both single-stage and dual-stage (i.e., secondary anode) TALs reported anode efficiencies greater than 50 % under high-voltage conditions [88–90]. These investigations also noted that the combination of low mass flow rate and high voltage produced the highest thermal loads on the thruster, a critical design consideration for long-duration operation [89].

## 2.5 Thruster with Anode-layer Operation

Hall thrusters with an anode-layer operate on the same fundamental principles as conventional Hall effect thrusters, but differ structurally by eliminating the ceramic discharge channel, illustrated in Fig. 2.3. Instead, TAL designs extend the anode structure itself in a hollow annular configuration to the thruster exit plane. This modification reduces the financial cost associated with ceramic materials, however historically, TALs have exhibited marginally lower performance compared to conventional boron nitride channelled HETs.

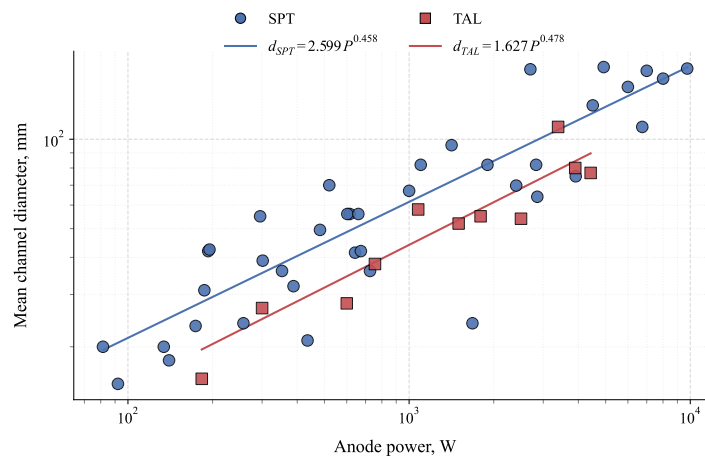


FIGURE 2.4: Log-log plot showing the relationship between anode power and mean channel diameter for SPT-type and TAL-type Hall thrusters [20–23, 37, 41, 43, 46, 48, 51, 54, 57–60, 64, 92–111].

TAL thrusters offer significantly higher power densities relative to conventional SPT-type Hall thrusters, this relationship is clearly illustrated in Fig. 2.4. For the same discharge power, TALs typically exhibit much smaller mean channel diameters [9]. Their compact geometry, featuring a short acceleration zone surrounded by metallic guard rings biased to cathode potential, reduces ion bombardment of internal surfaces, which can potentially improve lifetime. Propellant is supplied through a hollow anode, maintaining plasma generation and ion acceleration. Compared to SPT designs, TALs achieve higher thrust densities and demonstrate particular advantages at high-power operating regimes. However, despite these promising characteristics, limited experimental research into long-term erosion and stability effects in TAL thrusters means that the evidence for improved operational lifetime remains inconclusive [91].

## 2.6 Alternative Propellant Hall Effect Thrusters

Research into alternative propellants for HETs has been an ongoing endeavour for several decades, and several different classes of alternatives have been considered. The easiest to integrate into current experimental and operational architectures are inert monoatomic gaseous alternatives since they are compatible with the same fluidic and diagnostic systems used to characterise xenon.

Molecular and condensable propellants present a range of behaviours depending on their chemistry. Polyatomic species introduce additional excitation, vibrational, and dissociation modes that can reduce efficiency. Some dissociation products, such as reactive oxygen species, introduce material compatibility challenges. Liquid or solid propellants (for example, iodine or hydrocarbon-based propellants) require dedicated

delivery systems to produce the gaseous feed required for ionisation. This inherently increases system complexity and power consumption, as electric heaters are typically required to facilitate the phase change. These alternatives also raise additional compatibility issues: iodine is highly corrosive, and hydrocarbons can produce “sooty” carbon deposits that clog channels and shorten system lifetime through “caking” and electrical shorting, from the resulting conductive carbon layer.

Metallic propellants are a further alternative. Metallic propellants are not new to electric propulsion; magnetohydrodynamic thrusters [112–114], pulsed plasma thrusters [115], and vacuum arc thrusters [116–118] all have metallic-fuelled variants. For Hall effect thrusters, metallic propellants are attractive because they offer substantially higher storage densities than gaseous alternatives, remove the need for heavy pressurised propellant tanks, and can have favourable ionisation characteristics [119]. However, they suffer from similar integration issues as condensable propellants; novel propellant delivery systems are typically required, high temperatures for phase change may be necessary, and system power and thermal management requirements increase.

Moreover, a compatibility hurdle that is easily overlooked when developing alternative propellant Hall effect thrusters is cathode compatibility. Several cathode types, specifically  $\text{LaB}_6$  and barium oxide cathodes, can be sensitive to the operational environment and can experience “poisoning”. Although not exhaustive,  $\text{LaB}_6$  and barium oxide are both common cathode types for electric propulsion and have compatibility issues with several alternative propellants. Cathode research is ongoing and not the focus here; however, it remains an important consideration.

### 2.6.1 Gaseous Alternatives

Xenon remains an excellent propellant for electric propulsion owing to its combination of high storage density, chemical inertness, large atomic mass, low first ionisation energy and a large ionisation cross-section. Several gaseous alternatives have been compared to xenon in Tab. 2.1, where key physical properties relevant to thruster performance are listed.

Prior work has shown krypton, when operated in the same thruster geometry, typically exhibits lower anode and overall efficiency than xenon [23–25]; this reduction has been attributed primarily to decreases in mass utilisation and current efficiency for krypton. Novel magnetic field topologies, such as magnetic shielding, have been proposed to improve the viability of harder-to-ionise propellants; magnetic shielding configurations often produce higher plasma temperatures in the ionisation region compared to unshielded designs [19], which can increase ionisation rates at the cost of greater radiative losses. The H9-MUSCLE results by Su [81] showed krypton

outperforming xenon at high current densities when the dominant ion species for xenon became  $\text{Xe}^{2+}$ , whereas krypton maintained high charge efficiency under those conditions.

TABLE 2.1: Comparison of the mass, ionisation energies (IE), and natural abundance of several potential gaseous propellants with xenon. Percentages are shown relative to xenon. The ionisation energy for nitrogen refers to dinitrogen ( $\text{N}_2$ ); within the plasma both diatomic and atomic states would be present [11, 120]. †: Ionisation energy for  $\text{N}_2 \rightarrow \text{N}_2^+ + e$ .

		Xe	Kr	Ar	Ne	$\text{N}_2$
Mass,	amu	131.3	83.8	39.9	20.2	28.0
1 <sup>st</sup> IE,	eV	12.1	14.0	15.8	21.6	15.58 <sup>†</sup>
2 <sup>nd</sup> IE,	eV	21.0	24.4	27.6	41.0	-
Abundance,	ppm	0.087	1.14	9340	18.21	780840
Mass,	% <sub>Xe</sub>	-	63.8 %	30.4 %	15.4 %	21.3 %
1 <sup>st</sup> IE,	% <sub>Xe</sub>	-	115.7 %	130.6 %	178.5 %	128.8 % <sup>†</sup>
2 <sup>nd</sup> IE,	% <sub>Xe</sub>	-	116.2 %	131.4 %	195.2 %	-

## 2.6.2 Molecular and Condensable Alternatives

Molecular propellants span gaseous diatomic or polyatomic species as well as condensable solids and liquids that can be sublimated or vapourised for feed. Polyatomic gases introduce additional internal degrees of freedom that increase the number of collisional energy dissipation mechanisms; this tends to lower ionisation efficiency compared to noble gases of similar mass [121]. Most condensable propellants, such as iodine, offer higher storage densities and, in some cases, lower ionisation energies than xenon and other gaseous alternatives. However, for hydrocarbon-based molecular propellants, the ionisation energy is not strictly less than that of xenon. A comparison of selected condensable propellants is given in Tab. 2.2.

Sublimation or vaporisation incurs an energy cost that scales with mass flux and reduces overall system efficiency. The feed system must ensure the propellant remains gaseous from tank to thruster, which may require tank and feed heating and constrain integration options. Some condensable materials, such as bismuth, require non-trivial temperatures to achieve sufficient vapour pressure for practical mass flow rates, which increases heater power demand and complicates thermal integration with the spacecraft bus [119].

Other molecular-fuelled Hall thrusters have been operated on water-vapour, dinitrogen, carbon dioxide, iodine, and hydrocarbon-based propellants. Interest in water-vapour propulsion has increased because of in-situ resource utilisation concepts, where ice could be harvested in-space for use as propellant. Although

TABLE 2.2: Comparison of the properties of several alternative condensable and metallic alternative propellants with xenon [119].

	Xe	I	Mg	Zn	Bi	Cd
Mass, amu	131.3	126.9	24.3	65.4	209	112.4
1 <sup>st</sup> IE, eV	12.1	10.5	7.6	9.4	7.3	8.9
Density STP, g cm <sup>-3</sup>	1.6	4.9	1.7	7.1	9.8	8.7
Melting point, °C	-112	113.7	650	420	271	321
Vapour pressure, Pa	-	$2.34 \times 10^4$	384.01	21.49	$7.50 \times 10^{-3}$	15.34
Toxicity	-	Med.	Med.	Low	Low	High

realisation is likely decades away, systems that can operate on water are of research interest. Several Hall thrusters have been tested on water vapour with mixed success, producing relatively low anode efficiencies in the power regimes tested [122–124]. This low anode efficiency results primarily from very low mass utilisation, which highlights the difficulty of ionising water vapour.

Dinitrogen and carbon dioxide have also been tested. Motivations include low cost and potential air-breathing operation on Earth atmosphere or Martian atmosphere. Carbon dioxide experiments, similar to water vapour, show poor anode efficiency due to low mass utilisation [125, 126]. Dinitrogen testing targets very low Earth orbit air-breathing concepts or cost reduction; in some cases, performance has been comparable to low performing argon operation [12, 127].

Hydrocarbon propellants offer large ion masses and therefore potential for high thrust, with higher storage densities and lower cost compared with xenon or krypton. Examples tested include Buckminsterfullerene (C<sub>60</sub>) [128, 129], naphthalene (C<sub>10</sub>H<sub>8</sub>) [130], and adamantane (C<sub>10</sub>H<sub>16</sub>) [131]. Whilst this list is not exhaustive, there is a range of hydrocarbon-based research for Hall effect thrusters.

Iodine Hall effect thrusters have seen greater interest because iodine is close to xenon on the periodic table, yielding similar atomic mass and comparable ionisation cross-sections. Iodine introduces compatibility challenges related to toxicity and corrosive interactions. Despite these issues, several Hall thrusters have been operated successfully on iodine [46, 132, 133].

### 2.6.3 Metallic Alternatives

Metallic propellants combine high storage density with generally favourable ionisation properties [119]. Integration challenges are significant and varied; they include feed design to deliver metal vapour or liquid to the anode. This can lead to deposition within the thruster channel and propellant feed lines, which may cause failure if a conductive layer forms and shorts the anode to the thruster body or satellite

bus. In the case of bismuth, the solid phase has a lower density than the liquid phase, a 3.3 % volume expansion on freezing; solidification within feed lines can therefore cause rupture [134]. Prior studies across different electric propulsion architectures indicate both promise and practical challenges for metallic propellants [112, 115, 116].

Despite these challenges, several Hall effect thrusters have been operated on metallic propellants in experimental settings. The BHT-Bi-1500-V was operated successfully on bismuth between 350 W and 880 W, achieving high anode efficiencies (60 %) [60]. The HA-Zn-HET was operated on zinc between 140 W and 240 W and produced thrusts up to 5 mN; whilst lower performing, it was a significant development for operation in this power regime [135]. Zinc and magnesium have also been operated at higher powers with good performance [46].

## 2.7 Alternative Propellant Operation

In this section we describe and discuss the operational changes to a Hall thruster discharge when operating on alternative gaseous monatomic propellants. These noble-gas alternatives have been chosen as the focus of this work because of the ease with which they can be integrated into the existing testing framework and to avoid the system-level integration issues present with other alternatives.

### 2.7.1 Mass

Propellant selection affects several key aspects of Hall effect thruster operation, particularly the ionisation dynamics of neutral atoms and plasma properties. The propellant's neutral atomic mass is a key driver in both thrust and specific impulse outputs. By rearranging  $\dot{m}_i \approx \eta_m \dot{m}_n$ , from Eq. 2.28 and substituting in, the thrust and specific impulse can be expressed as

$$T = \frac{d(m_i v_i)}{dt} = \dot{m}_i v_i = \frac{\overbrace{\dot{m}_i}^{\dot{m}_i}}{\bar{q}} \underbrace{\eta_m \eta_{I_d} m_n I_d}_{v_i} \eta_{\theta_d} \sqrt{\frac{2 \bar{q} U_b}{m_n}} = \eta_m \eta_{I_d} \eta_{\theta_d} I_d \sqrt{\frac{2 m_n \eta_V U_d}{\bar{q}}} \propto \sqrt{m_n} . \quad (2.34)$$

Thus, the thrust scales with the square root of atomic mass, reflecting the dependence of momentum transfer on particle mass.

For specific impulse

$$I_{sp} = \frac{T}{g_0 \dot{m}_n} = \frac{\dot{m}_i v_i}{\dot{m}_n g_0} = \eta_m \eta_{\theta_d} \frac{1}{g_0} \sqrt{\frac{2qU_b}{m_n}} \propto \frac{1}{\sqrt{m_n}}. \quad (2.35)$$

Specific impulse, therefore, scales inversely with the square root of atomic mass, emphasising the trade-off between thrust and exhaust velocity.

From Tab. 2.1, it is evident that low atomic mass propellants such as neon offer the potential for extremely high specific impulse at the same discharge voltage compared to xenon. However, this gain is typically offset by a reduction in thruster efficiency, due to lower mass utilisation efficiency, leading to the high specific impulse not being achieved. It is a secondary motivation of this research to explore methods of closing this gap, enabling the use of lighter propellants to extend Hall thruster high specific impulse capabilities.

Furthermore, atomic mass also significantly influences neutral gas dynamics within a Hall thruster. Assuming a thermalised neutral velocity based on the anode temperature,  $T_a$ , yields

$$u_n = \sqrt{\frac{8k_B T_a}{\pi m_n}}, \quad (2.36)$$

where,  $u_n$  is the neutral velocity, and  $k_B$  is the Boltzmann constant. The thermal velocity of neutrals therefore scales inversely with the square root of their atomic mass. For example, at the same anode temperature, krypton neutrals move at approximately 125 % the speed of xenon neutrals. Affecting both the ionisation probability and as a result increasing the mean free path of ionisation. This effect becomes more pronounced with lighter propellants; for instance, neon neutrals travel at roughly 255 % the speed of xenon neutrals under equivalent conditions.

Furthermore, the literature reports a wide range of assumed and measured values for neutral temperature, with stated values ranging from 800–1300 K [136–138]. However, as shown in Fig. 2.5, there is a demonstrable correlation between increasing anode voltage and neutral temperature. Which is not an unexpected result, as higher discharge voltages will result in higher powers and consequentially result in larger thermal loads to the thruster, heating the anode. Furthermore, the neutral temperature is clearly not constant along the thruster's axis which is a common assumption in Hall thruster scaling. However, the values presented in Fig. 2.5 were obtained via laser induced fluorescence (LIF) measurements, a non-invasive plasma diagnostic method that can have significant uncertainty (shown as shaded regions on the figure), for a single laboratory thruster. As such, this data is not used to define a temperature model but is included here to illustrate the broader uncertainty surrounding this common assumption.

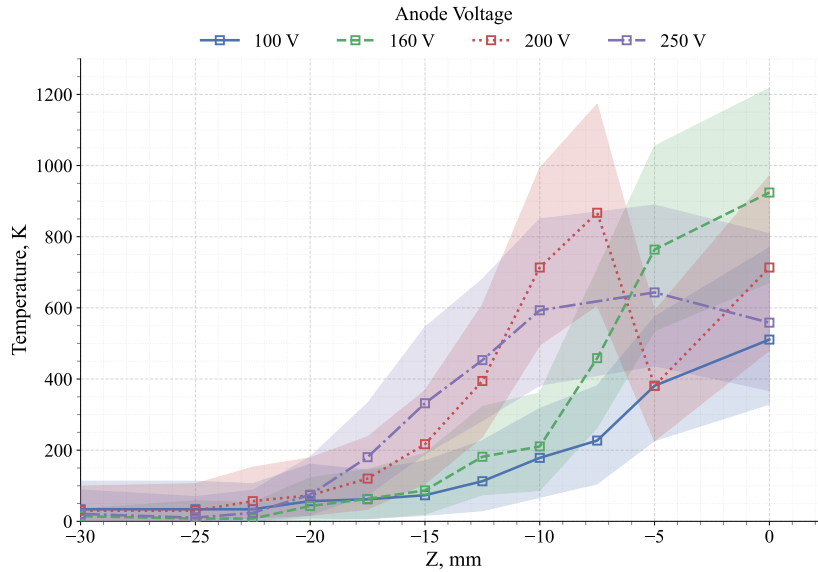


FIGURE 2.5: Estimated xenon neutral temperature at the channel centreline over the channel length relative to the exit plane ( $Z = 0$  mm) of the Stanford Hall thruster, measured via laser induced fluorescence for several discharge voltage levels [139]. The temperature values shown are under the assumption that the neutrals are thermalised.

Neutral velocity plays a significant role in the ionisation dynamics of a Hall thruster. One way to quantify this is through the ionisation mean free path,  $\lambda_i$ , defined as

$$\lambda_i = \frac{u_n}{v_i} = \frac{u_n}{n_e \langle \sigma_i v_e \rangle}, \quad (2.37)$$

where  $v_i$  is the ionisation frequency,  $n_e$  is the electron number density, and  $\langle \sigma_i v_e \rangle$  is the reaction rate coefficient, which will be discussed in detail in Sec. 2.7.4. Experimentally, it has been observed that the electron number density is approximately 10 % of the neutral number density, leading to the approximation  $n_e \approx 0.1n_n$  [136].

The variation of ionisation mean free path with neutral temperature for xenon, krypton and argon is shown in Fig. 2.6. The assumed plasma temperature estimate used for the mean free path calculations shown in Fig. 2.6 was empirically derived for xenon and is assumed propellant-agnostic due to limited data for alternative propellants [136]. The calculation of the reaction rate coefficient required for the mean free path calculation is described in greater detail in Sec. 2.7.4. Lighter propellants and lower voltages exhibit greater sensitivity to neutral temperature, underscoring the need to account for mass and temperature effects in scaling laws.

A key observation of Fig. 2.6 is that the mean free path can vary by an order of magnitude, depending on propellant choice and neutral temperature. This variation is substantial and makes quantitative evaluation difficult. Given these observations, it is clear that both propellant mass and neutral temperature should be explicitly

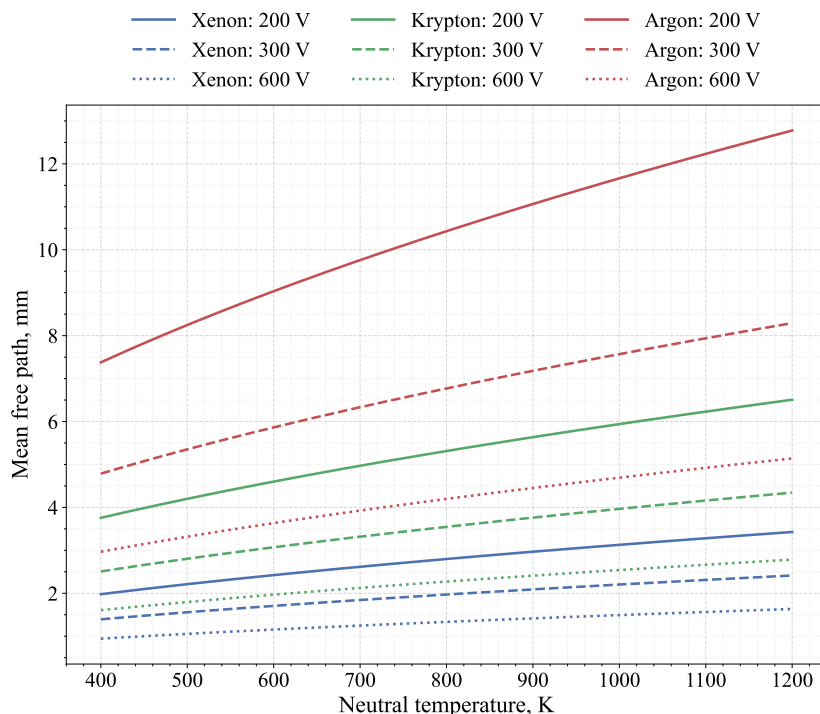


FIGURE 2.6: Variation of mean free path for xenon, krypton, and argon at anode voltages of 200 V, 300 V, and 600 V, as a function of neutral temperature. Calculated assuming optimal neutral number density [12].

incorporated into modern Hall thruster scaling approaches if non-xenon propellants are to be considered.

## 2.7.2 Ionisation

Propellant choice also changes the ionisation mechanics due to each propellant having differing electron-orbital structure, ionisation energy, and cross-sectional area of electron-impact ionisation. The ionisation mechanics of Hall thrusters critically dictate plasma properties. Efficient operation requires satisfying several ionisation-related criteria, including sufficient ionisation frequency and appropriate energy management.

To evaluate the mean free path of ionisation from Eq. 2.37 the ionisation frequency,  $\nu_i$ , needs to be determined. This can be achieved by via

$$\nu_i = n_e \langle \sigma_i v_e \rangle, \quad (2.38)$$

where  $\sigma_i$  is the ionisation cross-section, and  $v_e$  is the electron velocity. The  $\langle \sigma_i v_e \rangle$ , sometimes referred to as the reaction rate coefficient, accounts for the Maxwellian-Boltzmann distribution of electron velocities and their contribution to ionising collisions, as discussed in Sec. 2.7.4.

The neutral number density,  $n_n$ , can be determined from the mass flow rate,  $\dot{m}_n$ , and channel geometry, under the assumption of uniform propellant distribution, via

$$n_n = \frac{\dot{m}_n}{A_c m_n u_n} , \quad (2.39)$$

where  $A_c$  is the channel cross-sectional area. For annular discharge channels it can easily be seen that the cross-sectional area can be described by the two characteristic dimensions of a Hall thruster, as given in Fig. 2.2, such

$$A_c = \pi h d . \quad (2.40)$$

As shown in Tab. 2.1, xenon possesses the lowest first ionisation energy among the considered gases, making it highly favourable for efficient plasma generation. Lighter noble gases, such as krypton and argon, require higher electron energies to have equivalent ionisation rates, increasing the energy demands on the discharge. This trend arises because, for noble gases, decreasing atomic number results in the valence electrons being more tightly bound to the nucleus due to reduced electron shielding and increased proximity to nuclear charge.

Another consideration for Hall thruster ionisation dynamics is the power required to ionise the propellant. Within a HET ions are primarily produced via inelastic electron impact ionisation of neutral propellant atoms. For each ion created, there is a corresponding energy cost equal to the ionisation energy,  $\epsilon_i$ , that is propellant-specific, typically measured in electron volts. Assuming quasi-neutrality, the ion production rate is proportional to the discharge current, and the power consumed by ionisation,  $P_{\text{ion}}$ , can be expressed generally as

$$P_{\text{ion}} = \sum_{n=1}^N \left[ e \epsilon_i^{n+} \frac{I^{n+}}{en} \right] , \quad (2.41)$$

where  $I^{n+}$  is the species specific ion current,  $\epsilon_i^{n+}$  is the ionisation energy associated with the  $n^{\text{th}}$  ionisation state,  $N$  is the atomic number of the propellant, and  $I^{n+}/en$  is the number of ions in the  $n^{\text{th}}$  charge state per second.

For a constant mass flow rate, switching from xenon to krypton increases the ionisation power requirement. This is primarily due to krypton's higher first ionisation energy and the greater number of neutral atoms per unit mass, given its lower atomic mass.

Assuming singly charged ions, equal thruster geometry, and neutral temperature, from Eq. 2.41 the ratio of the ionisation power required for krypton relative to xenon

simplifies to

$$\frac{P_{\text{ion}}|_{\text{Kr}}}{P_{\text{ion}}|_{\text{Xe}}} \propto \left[ \frac{\epsilon_i^{+1}}{\sqrt{m_n}} \right]_{\text{Kr}} / \left[ \frac{\epsilon_i^{+1}}{\sqrt{m_n}} \right]_{\text{Xe}}, \quad (2.42)$$

where  $P_{\text{ion}}|_P$  is the propellant specific ionisation power. Using the ionisation energies and atomic masses listed in Tab. 2.1, the ratio of ionisation power required for krypton relative to xenon is

$$\frac{P_{\text{ion}}|_{\text{Kr}}}{P_{\text{ion}}|_{\text{Xe}}} = 1.448. \quad (2.43)$$

This indicates that krypton requires approximately 44.8 % more ionisation-power than xenon to ionise the same mass flux under equivalent conditions.

Applying the same methodology to argon yields

$$\frac{P_{\text{ion}}|_{\text{Ar}}}{P_{\text{ion}}|_{\text{Xe}}} = 2.369. \quad (2.44)$$

Thus, ionising an equivalent mass flux of argon requires approximately 136.9 % more ionisation-power than xenon. These results highlight the increasing ionisation cost associated with lighter propellants, reinforcing the need for tailored thruster designs when considering non-xenon operation.

This could, in principle, lead to colder plasmas and lower ionisation efficiencies as more energy is being extracted from the plasma to ionise the propellant. However, the absolute magnitude of this ionisation power is relatively small compared to the discharge power, as shown in Fig. 2.7. For a 1000 W, 300 V discharge (discharge current of 3.3 A), the resulting ionisation power required is 40.4 W, 46.7 W and 52.5 W for xenon, krypton, and argon, respectively. This comparison assumes constant discharge current rather than constant mass flux. Conversely, for constant mass flux, ion currents scale with the square root of mass (Eq. 2.34); as a result, the ionisation power required would scale accordingly.

Thus, propellant choice directly affects ionisation energy requirements and overall thruster efficiency, and must be carefully considered in any effort to optimise Hall thruster performance for alternative propellants.

### 2.7.3 Cross-Sectional Area of Ionisation

The cross-sectional area of ionisation,  $\sigma_i$ , represents the effective target area a neutral atom presents to an incident ionising electron. This parameter is a function of the electron temperature and depends on the atomic or molecular species, as well as the ionisation state, thus providing a probabilistic measure of ionisation likelihood.

Plasma electron temperature,  $T_{eV}$ , is a key parameter in calculating ionisation rates and is often related empirically to the discharge voltage. One widely accepted

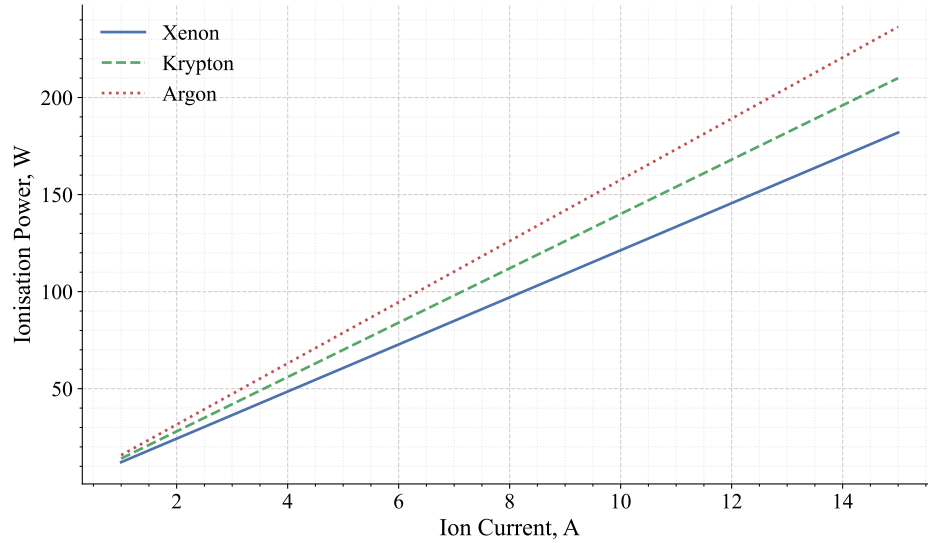


FIGURE 2.7: An illustration of the ionisation power per ion current assuming a purely mono-ionic ion current from Eq. 2.41.

relation, determined experimentally by Dannenmayer et al. [136], is given by

$$T_{eV} = 0.12U_d . \quad (2.45)$$

This linear relation has been validated for xenon discharges in standard annular HETs and has been adopted by multiple studies due to its simplicity and robustness [26, 68, 136]. Whilst, the result of Eq. 2.45 has not been verified as accurate for alternative propellant discharges it is adopted for this thesis due to a lack of a viable alternative. Some results have shown that xenon and krypton discharges have similar plasma temperatures at the same voltage within the NASA-173Mv1 Hall thruster [61]. This suggests that Eq. 2.45 is valid for krypton, however the NASA-173Mv1 is a very large magnetically shielded Hall thruster and not very representative of a “traditional” Hall thruster.

Alternatively, other formulations have also been proposed for bulk electron temperature within a HET discharge. For instance, Azziz [140] derived a power-law relationship from plume measurements of the BHT-1500 thruster is

$$T_{eV} = 1.828U_d^{0.518} . \quad (2.46)$$

While such non-linear models may offer improved fidelity for specific configurations, they can exhibit instability when extrapolated beyond the validated range. Consequently, the linear formulation of Dannenmayer and Mazouffre [136] will be used throughout this thesis for consistency and general applicability. However, there is no consensus that this relationship holds for alternative propellants.

The cross-sectional area of electron-neutral impact ionisation plays a critical role in

determining thruster performance. Specifically, it strongly influences the mean free path for ionisation events, a key factor in satisfying the Melikov–Morozov criterion, which requires that the mean free path be significantly smaller than the characteristic length scale of the discharge channel [68, 98, 136]

$$\lambda_i \ll L , \quad (2.47)$$

where  $L$  is the discharge channel length.

It should be noted, however, that the Melikov–Morozov criterion is a broad generalisation of the ionisation dynamics in Hall thrusters. In reality, ionisation predominantly occurs in a limited region near the channel exit where the magnetic field strength peaks. Experimental observations suggest that the effective ionisation region has a characteristic size on the order of the channel width,  $h$ , implying that  $h \gg \lambda_i$  should also be satisfied for efficient operation [141].

In practice, HETs typically exhibit  $L/h$  ratios between 1 and 2. For instance, the SPT-100 Hall thruster has  $L/h \approx 1.23$ , while the MaSMi-60 demonstrates  $L/h \approx 2.0$  [51, 98]. This suggests that the Melikov–Morozov criterion is reasonably satisfied whether using the channel length or width as the characteristic length.

The variation in cross-sectional area as a function of electron energy for several potential propellants is shown in Fig. 2.8. Xenon exhibits the largest cross-sectional area, followed by krypton, with argon and nitrogen presenting comparable values, and neon showing significantly lower ionisation cross-sections.

It is also noteworthy that in the case of diatomic nitrogen, the dominant ionisation pathway is  $\text{N}_2 \rightarrow \text{N}_2^+ + e$ , rather than dissociative ionisation into atomic nitrogen ions [145, 146]. The cross-sectional area for  $\text{N}_2^+$  formation is more than an order of magnitude greater than the combined cross-sections for  $\text{N}^+$  and  $\text{N}_2^{2+}$  formation across the whole range of electron temperatures tested [146]. Additionally, nitrogen's high dissociation energy of 9.8 eV means that dissociation processes are energetically less favourable compared to direct ionisation, further supporting the assumption that the majority of ions exhausted from a nitrogen-fuelled Hall thruster will remain diatomic.

By contrast, for iodine ( $\text{I}_2$ ), another diatomic species considered for alternative propellants, the dissociation energy is only 1.54 eV, leading to near-complete dissociation in thruster plasmas [147]. Energy spent on dissociation does not contribute to thrust production, instead manifesting as an efficiency loss. Thus, it is reasonable to treat  $\text{N}_2$  ions as diatomic when modelling nitrogen-based HET discharges.

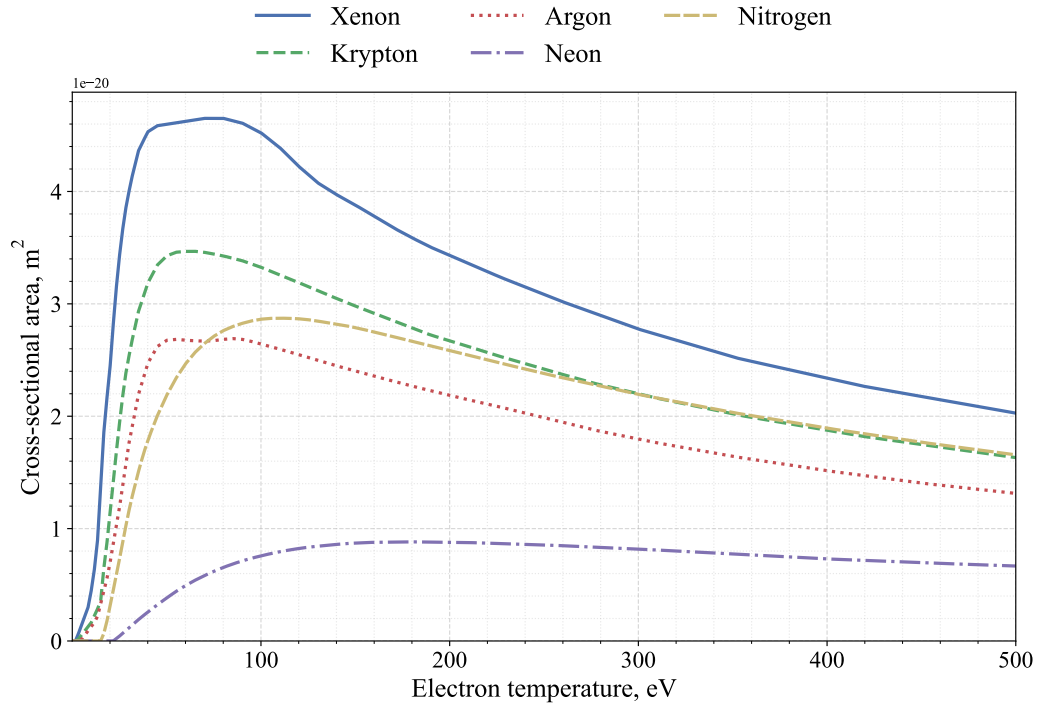


FIGURE 2.8: Cross-sectional area of ionisation for potential propellants versus bulk plasma temperature in electron volts. Cross-sectional area data: xenon and krypton from [142], argon from [143], and neon and diatomic nitrogen from [144].

#### 2.7.4 Reaction Rate Coefficients

To accurately estimate ionisation rates, the reaction rate coefficient  $\langle \sigma_i v_e \rangle$  must be evaluated. This quantity is the product of the cross-sectional area and the electron velocity, averaged over the assumed Maxwellian-Boltzmann velocity distribution of the electrons.

The probability distribution for the microscopic electron velocity  $w$  at a bulk electron temperature  $T_{eV}$  is given by the Maxwell-Boltzmann distribution such

$$g(w, T_{eV}) = \sqrt{\frac{2}{\pi}} \left( \frac{m_e}{eT_{eV}} \right)^{3/2} w^2 \exp\left(-\frac{m_e w^2}{2eT_{eV}}\right), \quad (2.48)$$

where  $m_e$  is the electron mass, and  $e$  is the electron charge and is needed to convert the temperature from electron volts to joules.

The reaction rate coefficient is then expressed as

$$\langle \sigma_i v_e \rangle = \int_{v_e}^c w \sigma_i(T_{eV}(w)) g(w, T_{eV}) dw, \quad (2.49)$$

where  $v_e$  corresponds to the velocity associated with the ionisation energy threshold, and  $c$  is the speed of light in a vacuum.

Eq. 2.49 was solved numerically for the candidate propellants, with the results presented in Fig. 2.9.

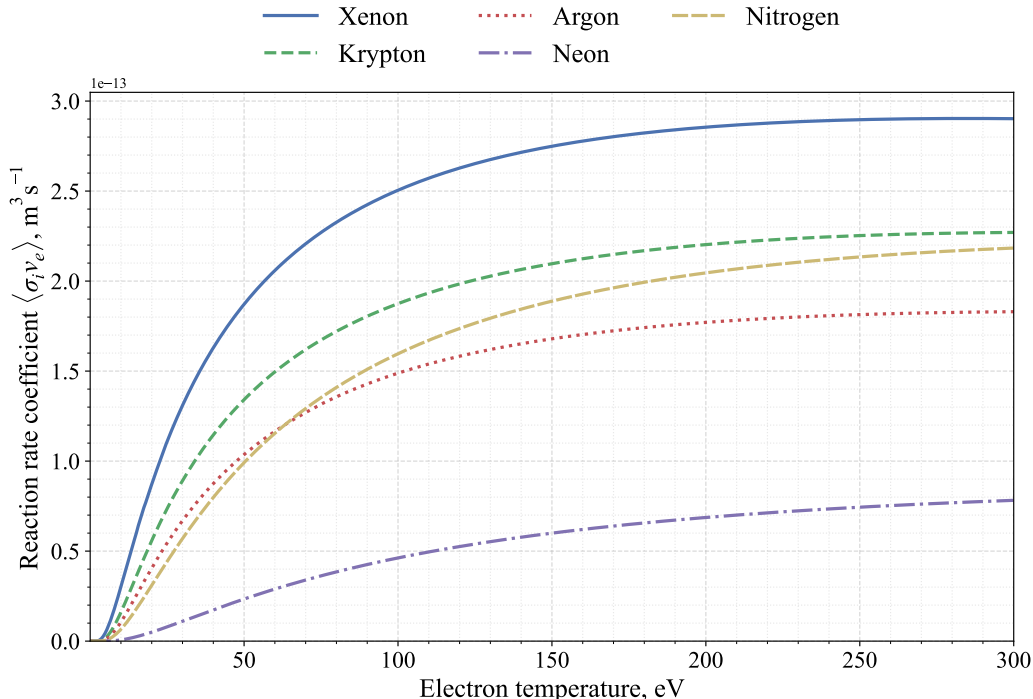


FIGURE 2.9: Computed reaction rate coefficients of potential propellants as a function of bulk plasma temperature in electron volts.

Tab. 2.3 summarises the cross-sectional areas and reaction rate coefficients at an assumed discharge voltage of 300 V, corresponding approximately to a plasma electron temperature of 36 eV, according to Eq. 2.45 [136].

TABLE 2.3: Comparison of cross-sectional areas and reaction rate coefficients for candidate propellants at an assumed 300 V discharge voltage, using Eq. 2.45.

	Xe	Kr	Ar	Ne	N <sub>2</sub>
$\sigma_i, 10^{-20} m^2$	4.01	2.55	1.84	0.12	1.17
$\langle \sigma_i v_e \rangle, 10^{-14} m^3 s^{-1}$	13.6	8.85	6.33	0.96	5.16
$\sigma_i, \%_{Xe}$	-	63.6 %	45.9 %	2.9 %	29.2 %
$\langle \sigma_i v_e \rangle, \%_{Xe}$	-	65.0 %	46.5 %	7.1 %	37.9 %

As shown in Tab. 2.3 and Figs. 2.8 and Fig. 2.9, xenon exhibits the highest reaction rate coefficient and cross-sectional area, followed by krypton. Argon and nitrogen demonstrate comparable values, while neon trails significantly behind.

Given the similar atomic mass, ionisation energy and reaction rate profiles of argon and nitrogen, it is reasonable to expect comparable thruster performance for these propellants under similar mass flow conditions, a conclusion supported by preliminary experimental results of a 100 W HET operating on argon and dinitrogen [12].

## 2.8 Sheath-Wall Potential in the Presence of Secondary Electron Emission

One of the key distinctions between traditional stationary plasma thrusters and thrusters with an anode-layer lies in the secondary electron emission (SEE) characteristics of their channel walls. These differences primarily arise from how electrons interact with the microstructure of the wall material. When a sufficiently energetic electron strikes the channel wall, it may release a secondary electron via the SEE process. SEE effect from impinging ions is not considered here as it is assumed they are pre-acceleration and have only low thermal energies, below the threshold for SEE to occur.

The emitted secondary electrons are typically low in energy, compared to the incident electron, and migrate into the bulk plasma, effectively cooling it. Meanwhile, the original high-energy primary electron is absorbed by the wall. For this exchange to occur, the primary electron must traverse the plasma wall-sheath, the idealised structure of which is shown in Fig. 2.10.

The structure depicted in Fig. 2.10 represents an idealised plasma of arbitrary potential bounded by two infinitely large plates at zero potential [148]. While simplified, this model supports several important insights into sheath behaviour.

To maintain charge neutrality and zero net potential at the wall, there must be no charge accumulation. Under steady-state conditions, this requires zero net current to the wall. Therefore, the balance of the positively charged ion flux and the negatively charged electron and secondary electron fluxes must satisfy current continuity. Expressing this balance in terms of particle fluxes, the current continuity condition at the wall can be described analytically as

$$\Gamma_{iw} = \Gamma_{ew} - \Gamma_{SEE} = \Gamma_{ew} (1 - \gamma_{SEE}) \quad , \quad (2.50)$$

where  $\Gamma_{iw}$  is the ion flux to the wall,  $\Gamma_{ew}$  is the primary electron flux,  $\Gamma_{SEE}$  is the secondary electron flux, and  $\gamma_{SEE}$  is the secondary electron emission coefficient. This equation is illustrated schematically in Fig. 2.11.




FIGURE 2.10: Schematic of the plasma–wall transition layer and the two Debye sheaths either side are shown [148].

Rearranging Eq. 2.50 as an expression for the required primary electron flux that satisfies current continuity yields

$$\Gamma_{ew} = \frac{\Gamma_{iw}}{1 - \gamma_{SEE}} . \quad (2.51)$$

This expression shows that as the SEE coefficient increases, a proportionally larger primary electron flux is required to maintain net current balance at the wall. Therefore, for any non-zero value of  $\gamma_{SEE}$ , the electron flux to the wall exceeds the ion flux.

The SEE coefficient itself can be approximated using a power-law function of the incident electron energy via [29]

$$\gamma_{SEE} = a\epsilon_e^b . \quad (2.52)$$

Here,  $a$  and  $b$  are material-specific fitting coefficients, and  $\epsilon_e$  denotes the effective energy of the incoming electron, typically on the order of several to tens of electron volts. Representative values for  $a$  and  $b$  for common Hall thruster wall materials are provided in Tab. 2.4.

The difference in SEE characteristics between ceramics and metals stems from how electrons are confined within the material microstructure. In ceramics, such as boron nitride, electrons are not free to move within the crystalline structure and are strongly localised. When a primary electron strikes the surface, the energy is distributed over a small number of electrons, allowing sufficient energy transfer for secondary emission

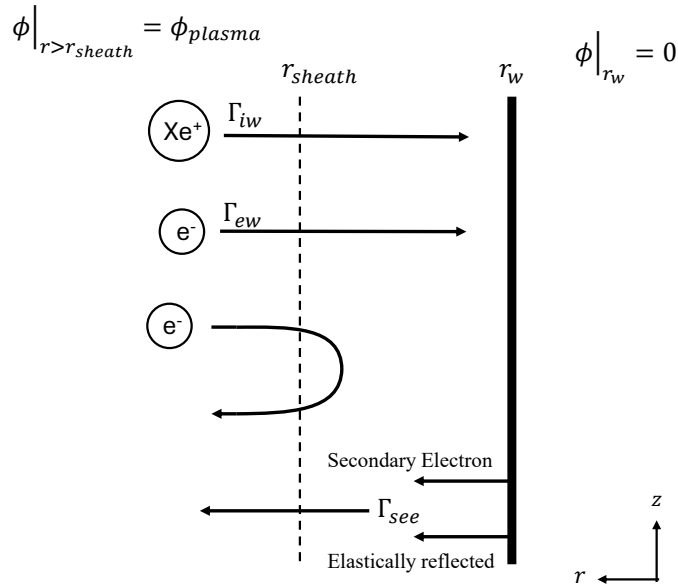


FIGURE 2.11: Schematic of particle fluxes at the plasma-facing wall, including primary electron flux  $\Gamma_{ew}$ , ion flux  $\Gamma_{iw}$ , and secondary electron flux  $\Gamma_{SEE}$ . Under steady-state conditions, the total current to the wall must be zero.

TABLE 2.4: Power law fitting parameters for secondary electron emission. Values for boron nitride are taken from Pigeon et al. [149], and for stainless steel from Choueiri [29].

Material	$a$	$b$
Boron nitride	0.0819	0.66
Stainless steel	0.0400	0.61

to occur relatively easily. This process is illustrated in Fig. 2.12, where a highly simplified example shows a high-energy electron interacting with boron nitride (BN); the figure highlights the highly localised nature of the energy distribution resulting from the primary electron impact.

In contrast, in metallic materials such as stainless steel, conduction electrons are free to move throughout the material and readily interact with one another. When a primary electron impacts a metal surface, the deposited energy is rapidly shared among many electrons, reducing the energy per electron and raising the threshold energy required for secondary emission. As a result, metallic surfaces typically exhibit lower SEE yields compared to ceramics for a given primary electron energy. This process is illustrated in Fig. 2.13 where a very simplified example of an high energy electron interacts with a metallic crystal structure.

However, it can be quickly seen that Eq. 2.51 suggests a potential issue as  $\gamma_{SEE} \rightarrow 1$ , where the electron current would rapidly increase, resulting in a “short-circuit” of electrons to the wall. In this regime, the increasing electron current would further




FIGURE 2.12: Schematic illustrating electron interactions with a boron nitride surface. Primary electrons impact the surface, resulting in energy transfer, absorption, and secondary electron emission depending on the local material structure.

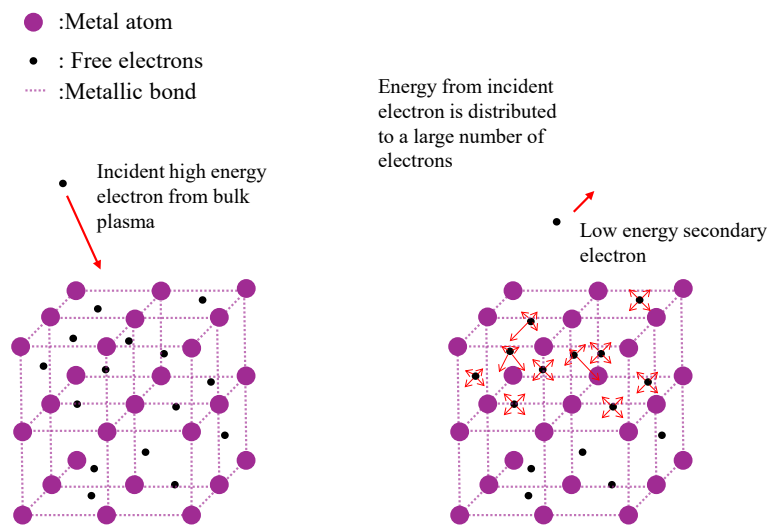


FIGURE 2.13: Schematic illustrating electron interactions with a metallic surface. Primary electrons impact the surface, resulting in energy transfer, absorption, and secondary electron emission depending on the local material structure.

enhance secondary electron emission, in a run-away-current increase. This occurs because one of the primary assumptions of the plasma sheath breaks down, as the potential will reverse, and the sheath will become an electron attracting sheath. Before the sheath reversal occurs, it is accepted that the secondary electrons will begin to be re-absorbed by the wall preventing this sheath reversal.

The presence of these secondary electrons has a non-negligible effect on the plasma-wall sheath potential required for quantifying the ion flux to the wall. With the ions assumed cold and at the Bohm velocity when reaching the sheath's edge the energy of the ions before the potential drop is  $E_i = 1/2 m_i v_B^2$ . As described by Hobbs and Wesson [150], the plasma potential drop that the ion experiences in the presence of SEE is

$$\Delta\phi_{\text{sheath-wall}} = T_{eV} \ln \left( \frac{1 - \gamma_{\text{SEE}}}{2\pi \frac{m_e}{m_i}} \right). \quad (2.53)$$

However, this balance only holds for  $\gamma_{\text{SEE}} < \gamma^*$ , where  $\gamma^*$  is the critical secondary electron emission coefficient at which the wall sheath undergoes a potential reversal: from electron-repelling to electron-attracting. This critical value  $\gamma^*$  occurs when the energy of the primary electrons reaches a threshold, defined by Choueiri [29] as

$$T^* = \left[ \frac{1 - \left( \frac{2\pi m_e}{e m_i} \right)^{0.5}}{a \Gamma(2 + b)} \right]^{1/b} \quad (2.54)$$

where  $\Gamma(x)$  is the Euler Gamma function. This method of estimating  $\gamma_{\text{SEE}}$  has been widely used in multidimensional Hall thruster simulations [151].

Using the values in Tab. 2.4, where boron nitride represents typical stationary plasma thruster walls and stainless steel (SS) represents typical thruster with anode-layer walls, the critical electron temperatures for xenon, krypton, and argon are listed in Tab. 2.5.

TABLE 2.5: Critical electron temperature  $T^*$  for sheath reversal in Hall thrusters with boron nitride and stainless-steel walls for selected propellants.

Propellant	Boron nitride, eV	Stainless steel, eV
Xenon	23.87	106.72
Krypton	23.82	106.49
Argon	23.72	105.99

As the wall electron temperature  $T_{eV}$  approaches  $T^*$ , the emission of cold secondary electrons increases. If  $T_{eV}$  were to closely approach or exceed  $T^*$ , a shallow potential well would form near the wall, producing a so-called "double sheath". This structure reflects cold emitted electrons back toward the wall, effectively capping the energy of

primary electrons near  $T^*$ . The formation of the double sheath thus serves to preserve current continuity without requiring an unphysical increase in hot electron flux.

Understanding the plasma-wall interaction, particularly the role of secondary electron emission, is essential for predicting and optimising Hall thruster performance when operating on alternative propellants and in high-voltage regimes. Since TAL-type thrusters use metallic walls with different SEE behaviour compared to traditional ceramic channels, the sheath characteristics and plasma confinement properties are fundamentally altered. These effects directly influence efficiency, lifetime, and stability; key factors when designing Hall thrusters for non-xenon propellants and high-specific impulse missions. Thus, accurately accounting for SEE processes is a critical part of advancing the capability of Hall effect thrusters to meet the goals of this thesis: enabling higher specific impulse operation and broadening the range of viable propellant options for future space missions.

## 2.9 Summary of Background and Fundamentals

This chapter has introduced the key physical and technological concepts underpinning the research presented in this thesis. The discussion first established the growing importance of electric propulsion for spacecraft missions where high efficiency and long operational lifetimes are required, contrasting this with the performance limitations of conventional chemical propulsion systems.

Several electric propulsion technologies were then introduced, highlighting the emergence of Hall effect thrusters as one of the dominant propulsion system selected for mission today. The widespread adoption of Hall thrusters has historically relied on xenon propellant, which benefits from favourable ionisation characteristics and extensive operational heritage. However, increasing demand within the space industry and price-volatility has motivated the investigation of alternative propellants, particularly krypton and argon.

These propellants are of particular interest as they can be integrated into existing propellant storage and feed architectures without requiring fundamentally new infrastructure, however their lower atomic mass and higher ionisation energies generally lead to reduced thruster performance when compared with xenon. As a result, several approaches have been proposed to mitigate these performance penalties, including operation at elevated anode discharge voltages and the exploration of alternative thruster architectures such as the thruster with anode-layer configuration.

Finally, the role of secondary electron emission was introduced as an important factor influencing electron temperature balance and ionisation processes within Hall thruster

discharges. Differences in secondary electron emission behaviour between ceramic-walled magnetic-layer thrusters and metallic-walled anode-layer thrusters may therefore influence the ionisation efficiency of alternative propellant discharges.

These considerations highlight the need for a more systematic understanding of how thruster design parameters interact with alternative propellant properties, motivating the investigation presented in the following chapters.



## Chapter 3

# Scaling Methodology

### 3.1 Introduction

Hall effect thrusters have been employed as a primary propulsion of spacecraft for several decades and are widely regarded as a mature and reliable form of electric propulsion. Despite their extensive testing and research since the 1970s, a consensus on how the physical dimensions of the thrusters themselves directly contribute to operation has still eluded the community as a whole. Consequently, the design of new Hall thrusters is not typically based on purely analytical methods but rather relies on empirical or semi-empirical trends derived from previously well characterised HETs.

To address this gap in understanding, empirical methods often employ scaling approaches that use databases of thruster geometries and their operational parameters to infer design guidelines. In this section, a brief overview of relevant scaling methodologies from the literature is first presented, followed by the formulation and justification of a modified set of scaling relationships developed specifically for the thruster investigated in this thesis. The resulting scaling predictions are then applied to the target propellant, krypton, with corresponding operating parameters also derived for the same thruster geometry operating on xenon.

### 3.2 Review of Existing Scaling Literature

Hall thrusters have been the subject of sustained research for several decades, with numerous designs having been developed, tested, and flown in-space. Despite this long history, a unified consensus on how to analytically size a Hall thruster to meet a mission's performance requirements remains elusive. Currently, there is no purely analytical first-principles method that allows for direct thruster sizing based on desired performance metrics. As a result, various scaling strategies ranging from

theoretical to semi-empirical and fully empirical have been developed and applied across the literature.

The earliest systematic scaling efforts can be traced back to the Soviet Union, where both stationary plasma thrusters and thrusters with anode-layers were developed in parallel. SPT development was led by A.I. Morozov at the Kurchatov Institute, while TAL thruster research was conducted under A.V. Zharinov at the Central Scientific Research Institute of Machine Building (TsNIIMash<sup>1</sup>) [152]. One of the earliest proposed sets of geometric relations for SPT and TAL thrusters sought to distinguish these two classes of Hall thrusters based on their relative physical dimensions: for SPT thrusters

$$\frac{L}{h} < 1 \quad , \quad \frac{L}{L_{\text{ionisation}}} \geq 1 \quad , \quad \frac{H_{\text{ion}}}{h} = 1 \quad , \quad \frac{h - H_{\text{ion}}}{H_{\text{ion}}} = 0 \quad , \quad (3.1)$$

and for TAL thrusters

$$\frac{L}{h} > 1 \quad , \quad \frac{L}{L_{\text{ionisation}}} < 1 \quad , \quad \frac{H_{\text{ion}}}{h} < 1 \quad , \quad \frac{h - H_{\text{ion}}}{H_{\text{ion}}} = 1 \quad , \quad (3.2)$$

where  $L$  is the channel length,  $h$  is the channel width,  $L_{\text{ionisation}}$  is the characteristic length over which the majority of ionisation occurs, and  $H_{\text{ion}}$  is the width of the ion beam [152]. These relations reflect early insights into the physical and operational differences between the two thruster types. For example, TAL thrusters were observed to feature shorter channels with ionisation and acceleration occurring predominantly outside the physical channel, as opposed to the more distributed processes seen in SPTs.

Subsequent research at the Kurchatov Institute advanced these early efforts by developing a purely empirical set of geometric scaling laws for SPT thrusters [153, 154]. These relationships connected various geometric parameters but did not inherently allow scaling to a desired power level or thrust. Additional geometric terms were introduced, such as  $h_{\text{mag}}$ , the radial separation between the inner and outer magnetic poles, and  $L_a$ , the distance from the anode to the magnetic pole plate. The empirical relationships were given as

$$h_{\text{mag}} = 0.3d \quad , \quad (3.3)$$

$$h = 6 + 0.375h_{\text{mag}} \quad , \quad (3.4)$$

$$L = 0.32h_{\text{mag}} \quad , \quad (3.5)$$

$$L_a = 2L \quad , \quad (3.6)$$

$$L \geq 1.1L_a \quad . \quad (3.7)$$

---

<sup>1</sup>This acronym is derived from the original Cyrillic and not from the English translation.

From these relations, several important geometric design constraints are inferred. For instance, Eq. 3.3 implies a linear relationship between the mean channel diameter  $d$  and the magnetic pole spacing. Eq. 3.7 suggests that the ceramic discharge channel should extend at least 10 % beyond the axial location of the magnetic pole plate (assumed to coincide with the peak magnetic field), thereby ensuring full containment of the acceleration region within the insulating channel.

Furthermore, early American studies of Hall thrusters established a useful empirical relationship between thrust and channel diameter, reported by Gulczinski III [154] as

$$d^2 \propto \dot{m}_a \propto T , \quad (3.8)$$

where  $\dot{m}_a$  is the anode mass flow rate and  $T$  is the thrust. These results were found by comparing the SPT-70, SPT-100, SPT-140, SPT-200, SPT-290, T-100E, and the T-160E. The "T" designated thrusters are Soviet TAL thrusters. This proportionality has been widely adopted as a method for scaling Hall thrusters to different power or thrust levels using experimentally derived coefficients.

Although these historical methods offer valuable insights, they remain primarily descriptive and often lack predictive capability for novel designs or alternative operating conditions (e.g., non-xenon propellants or, high-voltage operation).

Modern work, though similar in structure, has adopted semi-empirical methods that utilise analytical governing equations for Hall thruster operation to form a set of performance-geometry scaling relations incorporating empirically determined coefficients. These approaches offer practical utility in thruster design due to their relative simplicity and reliance on measurable quantities. However, their predictive capability is constrained by the underlying data: the accuracy and generalisability of these models are directly tied to the composition and diversity of the database from which scaling coefficients are derived. For instance, Dannenmayer and Mazouffre [136] used a database of 33 thrusters spanning 10 W to 50 kW, while Lee et al. [98] employed a smaller set of 17 sub-kilowatt devices, leading to subtle but meaningful variations in the derived scaling laws. Both approaches aim to derive thruster channel geometry based on input quantities such as anode power and voltage, although other combinations such as thrust and specific impulse can also serve as input variables with minor modification to the framework.

It is noteworthy that both Dannenmayer et al. and Kim et al. propose the same geometric scaling between channel width and mean diameter of

$$h = 0.242d . \quad (3.9)$$

While Dannenmayer et al. do not explicitly state this coefficient, their data reveals a strong linear trend that closely matches Eqs. 3.3 and 3.9. They suggest that this trend

likely emerges from the inclusion of thrusters that were themselves scaled with these early SPT scaling methods. However, as will be discussed in Sec. 3.3.1, with our database, the robustness of this scaling relation diminishes when incorporating a broader diversity of thruster architectures and operating regimes [12].

A distinct methodology, known as photographic scaling, was introduced by Khayms and Martinez-Sanchez [155]. This approach applies linear scaling to all thruster dimensions (mean diameter, channel width, and length) under the assumption that ionisation conditions can be preserved across scales. To maintain ionisation at reduced power levels, the plasma density must increase, which in turn requires higher magnetic field strengths. This requirement introduces potential material limitations, including magnetic saturation and increased channel erosion. Ahedo and Gallardo [156] subsequently investigated this issue and concluded that linear downscaling reduces both efficiency and lifetime, primarily due to higher wall thermal loads and constrained magnetic design options.

Ashkenazy, Shitrit, and Appelbaum [157] also examined low-power scaling, focusing on preserving propellant utilisation while improving thruster lifetime. They proposed that thruster lifetime scales with the ratio of wall thickness of the channel and cross-sectional area to the propellant mass flow rate, such that

$$t_{\text{life}} \propto \frac{h_{\text{thickness}} A_c}{\dot{m}_n} . \quad (3.10)$$

In their approach, thruster lifetime improves through reduced propellant flow while maintaining fixed geometry. To counteract the decrease in ionisation, they advocate lengthening the channel. However, subsequent work has shown that this only enhances ionisation efficiency within bounds of an optimal channel length [141, 158–160].

More recently, machine learning techniques have been explored as a means of bypassing explicit scaling laws. For example, Plyashkov et al. [161] employed neural networks (NNs) trained on large thruster databases to infer geometry from performance requirements. These models show promise in interpolating within known data but lack predictive capability outside of the training data and are highly sensitive to biases within the training data. While NN methods have begun to incorporate non-standard propellants, their applicability remains limited.

In contrast to purely empirical or machine learning-based approaches, recent work by Lafleur and Chabert [162] presents a major step forward in analytical modelling. Their one-dimensional steady-state model of a Hall thruster plasma yields a fully closed-form solution, capturing the profiles of plasma density, potential, and ionisation rate as a function of position. The model introduces three key similarity parameters: a normalised discharge voltage, a normalised discharge current, and a

combined magnetic-geometry scaling factor. Remarkably, this formulation enables prediction of thrust, specific impulse, and efficiency in agreement with experimental results for the SPT-100. By offering insight into how physical dimensions, discharge parameters, and field profiles interact, this model provides a useful analytical foundation for both design and code verification.

### 3.3 Hall Effect Thruster Database

The database used for the scaling analysis in this thesis is compiled from publicly available data and publications. The final derivation of scaling coefficients used for thruster sizing, only includes thrusters with well-documented and reliable dimensional data.

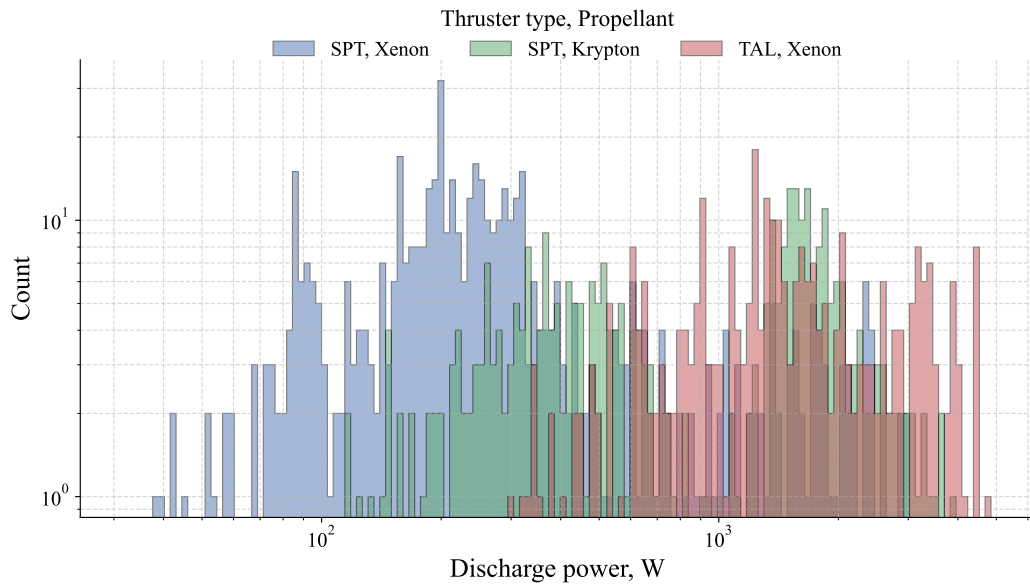


FIGURE 3.1: Histogram of the database anode power values using 150 bins per propellant-thruster type [21–23, 37, 41, 43, 45, 46, 48, 50, 51, 54, 59, 60, 92, 96, 97, 99–101, 103, 104, 107, 109, 111].

The performance data for these thrusters were collated across a range of input powers, discharge voltages, mass flow rates, and propellants. The database comprises approximately 1190 unique data points for 25 thrusters; 19 of which are SPT-type Hall effect thrusters with fully documented dimensions, and 6 of which are TAL-type thrusters. The distribution of operational discharge power for xenon and krypton for SPT and TAL type thruster within the database can be seen in Fig. 3.1. The thrusters contained within the database are tabulated in Tab. 3.1 with the used dimension shown.

TABLE 3.1: Thruster geometry database used for scaling. †: Primary anode, secondary anode.

Thruster name	$d$ , mm	$h$ , mm	$L$ , mm	Type	MS	Source
85W	20	6	8.4	SPT	No	[37]
BHT-200	21	5.6	-	SPT	No	[41]
BHT-600	56	16	10	SPT	No	[43, 105]
BHT-1000	69.8	10	10	SPT	No	[46]
BHT-1500	85	12	-	SPT	No	[92, 163]
CAMILA	43	12	47.35	SPT	No	[96]
CSU Hall thruster	87	17	32	SPT	No	[59]
D-55	55	17	4	TAL	No	[109]
D-80	80	† 5, 3	18.3	TAL	No	[97]
HIKHET	42	8	15	SPT	No	[100]
ISCT200	37	10	-	SPT	No	[101]
KLIMT	42	8	12-35	SPT	No	[23]
KM-32	32	6	16	SPT	No	[50]
MaSMi-40	36	8	-	SPT	Yes	[164]
MaSMi-60	60	9.42	19	SPT	Yes	[51, 105]
RAIJIN-66	54.1	11.9	2.1	TAL	No	[110]
RAIJIN-94	77	17	3	TAL	No	[111]
SPT-20	15	5	32	SPT	No	[105]
SPT-25	20	5	10	SPT	No	[105]
SPT-30	22.5	7.5	-	SPT	No	[48]
SPT-50	40	10	25	SPT	No	[54]
SPT-70	56	14	25	SPT	No	[105]
SPT-100	85	15	21	SPT	No	[21]
TALT-2	58	4	35	TAL	No	[109]
TAL-110	110	40	-	TAL	No	[109]

### 3.3.1 Channel Width and Mean Diameter Trends

The relationship between the channel width,  $h$ , and the mean channel diameter,  $d$ , plays a critical role in Hall thruster design and are two of the primary desired outputs from scaling methods. Several scaling models in the literature assume a direct linear correlation between these parameters, often derived from empirical fits to existing thruster datasets. This simplification is attractive for design purposes but may obscure important dependencies that vary across thruster families and operational regimes.

For example, Lee et al. [98], based on a dataset of 17 sub-kilowatt thrusters, proposed a linear relationship of the form  $h = 0.242d$  (Eq. 3.9). In contrast, a least-squares fit to the this more diverse database compiled for this thesis yields  $h = 0.200d$  for the 19 SPT-type thrusters. The reduced linearity,  $R^2|_{\text{SPT}} = 0.4609$ , as seen in Fig. 3.2 suggests increased sensitivity to additional design parameters.

Further insight is provided in Fig. 3.3, which presents box plots of  $h/d$  for various thruster families. The BHT series, developed by Busek Co. Inc. in the United States

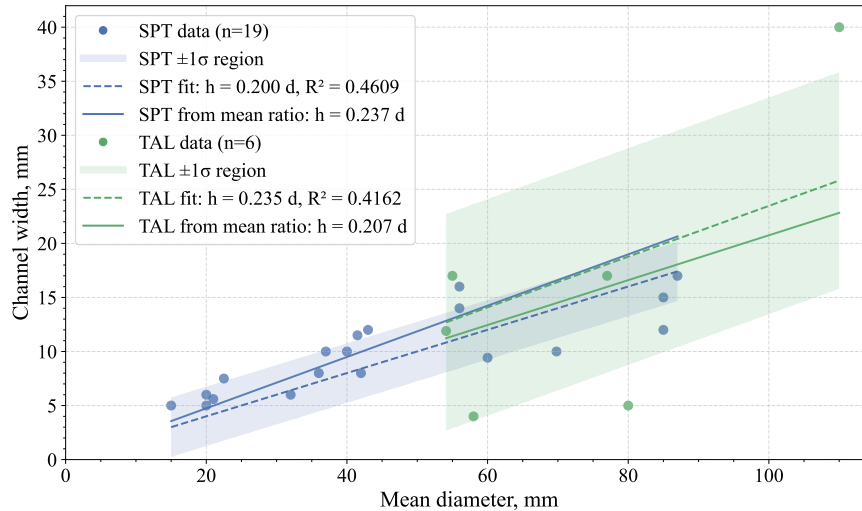


FIGURE 3.2: Channel width plotted against mean channel diameter for SPT and TAL thrusters within the dataset. Least-squares linear fits are shown for both types.

and recognised for its high performance, exhibits a wide distribution and deviates from the typical SPT trend. The BHT family of thrusters has  $(h/\bar{d}) = 0.205$ ,  $(h/\bar{d}) = 0.209$ ,  $h/d|_{\min} = 0.142$ , and  $h/d|_{\max} = 0.286$ . This supports the conclusion that simple geometric scaling relations are insufficient to capture the range of valid design practices across the field.

The SPT family, which refers to the thrusters in Tab. 3.1 with “SPT” in their designation, also deviates from the overall group in terms of distribution, even though it shares a similar mean value. The SPT family has  $(h/\bar{d}) = 0.25$ ,  $(h/\bar{d}) = 0.266$ ,  $h/d|_{\min} = 0.177$ , and  $h/d|_{\max} = 0.333$ .

Finally, the RAIJIN family, a set of Japanese TAL laboratory thrusters, was photographically scaled from the 5 kW-class RAIJIN-94 to the 2 kW-class RAIJIN-66. As a result of this photographic scaling, the RAIJIN family maintains an identical value of  $h/d = 0.22$ .

Interestingly, the TAL thruster-type exhibits a broader range of  $h/d$  ratios, with a median lower than that of the SPT thruster-type which can be seen in Fig. 3.3. However, correctly analysing TAL thrusters is challenging due to inconsistent reporting in the literature; some sources report the anode width, while others refer to the spacing between guard rings. Without direct inspection of the physical hardware, these discrepancies are difficult to reconcile quantitatively.

Although a nominal  $h/d \approx 0.2$  serves as a reasonable first-order approximation, the significant variability observed, particularly across different propellants and power and thruster families, indicates that this ratio alone is insufficient for reliable scaling. The diversity in design approaches and operational regimes suggests that more




FIGURE 3.3: Box plots of channel width to mean diameter ratio ( $h/d$ ) for various thruster families. The data highlight inter-family variance and the presence of outliers.

flexible, performance-oriented scaling strategies are required for robust thruster development.

## 3.4 Derivation of Scaling Methodology Employed

This section presents the scaling relations used to size the Hall thruster employed in this thesis. The approach builds upon methodologies established by Dannenmayer and Mazouffre [136] and Lee et al. [98], with extensions to preserve propellant-specific metrics and incorporate semi-empirical correlations. In addition, several novel relations, unique to this work and informed by a broader and updated dataset, are introduced.

### 3.4.1 Electron Confinement

Electron confinement and by extension, plasma confinement is a key design driver for Hall effect thrusters. As previously discussed in Sec. 2.3, the confined electrons are responsible for sustaining the ionisation process as well as the momentum transferal mechanism responsible for producing thrust. Therefore, the magnetic circuit of a Hall effect thruster must be carefully designed to ensure effective electron confinement.

In a typical annular Hall thruster, electrons are confined by the applied radial magnetic field. The magnetic field strength is selected such that electrons are magnetised, while the much heavier ions remain unmagnetised or only weakly magnetised. Classically, thermal electrons moving in a magnetic field gyrate in circular orbits with a radius defined by the electron Larmor radius, defined as

$$r_{L_e} = \frac{m_e v_e(T_{eV})}{eB} , \quad (3.11)$$

where  $r_{L_e}$  is the electron Larmor radius,  $v_e(T_{eV})$  is the electron thermal velocity, and  $B$  is the magnetic field magnitude.

For electrons to be effectively magnetised [165], the following criterion must be satisfied

$$r_{L_e} \ll [h, L]_{\min} . \quad (3.12)$$

That is, the electron Larmor radius must be much smaller than the smallest characteristic dimension of the discharge channel. To preserve this condition across thruster scales, the magnetic field strength must scale as

$$B \propto \frac{v_e(T_{eV})}{[h, L]_{\min}} . \quad (3.13)$$

However, this classical confinement criterion can typically be satisfied using a moderate magnetic field. Garcia, Tang, and Ren [166] proposed that a more stringent and physically relevant constraint is to maintain a constant ratio between the electron gyro-period and the mean time between ionising collisions. This is in order to make sure that the electrons are confined for sufficient time to ionise the neutral propellant flux. From this consideration, a new magnetic field scaling criterion is derived as

$$B \propto \sqrt{\frac{1}{n_n h L_i^2}} , \quad (3.14)$$

where  $n_n$  is the neutral number density and  $L_i$  is the characteristic ionisation region length. This result draws an inverse square root relationship between magnetic field strength and number density. The result of this is that for higher number density (higher mass flow rate) a lower magnetic field should be chosen. This result directly contradicts the suggested requirement that the electron gyro-period should be less than the time between ionising electron-neutral collisions which yields

$$B \propto \langle \sigma_i v_e \rangle n_n , \quad (3.15)$$

which suggests for higher mass flow rates will require greater magnetic field strengths to prevent a reduction in electron confinement [167].

### 3.4.2 Neutral Density

Neutral number density,  $n_n$ , defines the quantity of neutral propellant atoms or molecules per unit volume within a thruster's discharge channel. As discussed in Sec. 2.6, this parameter plays a critical role in ionisation dynamics and electron confinement in Hall effect thrusters. As shown in Eqs. 2.39 and 2.40, the neutral number density can be expressed as

$$n_n = \frac{\dot{m}_n}{\pi h d m_n u_n}. \quad (3.16)$$

As previously discussed in Sec. 2.7.1,  $u_n$  is also a function of propellant atomic mass,  $m_n$ , and assumed anode temperature (also called assumed neutral temperature),  $T_a$ . Whilst a range of value of  $T_a$  exist in literature, for the analysis undertaken here an assumed value of  $T_a \approx 800$  K was chosen and assumed propellant agnostic for simplicity in scaling.

Previous work suggests that there exists an optimal  $n_n$  for efficient HET operation. For xenon, this critical value has been empirically identified as  $n_n \approx 1.2 \times 10^{19} \text{ m}^{-3}$  [68, 136]. Maintaining a roughly constant neutral density across scaling regimes helps mitigate thermal loads and erosion. If  $n_n$  drops too low, the time between ionising collisions becomes too long to sustain effective ionisation. Conversely, excessive  $n_n$  increases electron-neutral collisions, degrading electron confinement due to frequent momentum transfer [136].

TABLE 3.2: Summary of neutral number density statistics for xenon and krypton across the dataset for SPT type thrusters and peak performance subsets from Fig. 3.4.  $\bar{n}_n$ : Mean number density;  $\tilde{n}_n$ : Median number density. The **Min.** and **Max.** columns refer to the minimum and maximum values, respectively, within the dataset for the corresponding row.

Propellant		$N$	$\bar{n}_n$ ( $\times 10^{19} \text{ m}^{-3}$ )	$\tilde{n}_n$ ( $\times 10^{19} \text{ m}^{-3}$ )	<b>Min.</b> ( $\times 10^{19} \text{ m}^{-3}$ )	<b>Max.</b> ( $\times 10^{19} \text{ m}^{-3}$ )
Xenon	All data	851	1.66	1.48	0.45	10.3
	Peak Thrust	19	2.41	1.72	0.91	8.07
	Peak $I_{sp}$	19	2.00	1.60	0.76	5.24
	Peak $\eta_a$	20	2.08	1.51	0.76	8.07
Krypton	All data	338	1.34	1.43	0.64	2.11
	Peak Thrust	4	1.68	1.77	1.06	2.11
	Peak $I_{sp}$	4	1.54	1.60	0.85	2.11
	Peak $\eta_a$	3	1.54	1.46	1.06	2.11




FIGURE 3.4: Box and whisker plots of neutral number density values for xenon and krypton propellants across the Hall thruster database. The total population mean neutral density is  $\bar{n}_n|_{\text{Xe}} = 1.66 \times 10^{19} \text{ m}^{-3}$  for  $N = 851$ , and  $\bar{n}_n|_{\text{Kr}} = 1.34 \times 10^{19} \text{ m}^{-3}$  for  $N = 338$ .

It remains uncertain whether this critical value holds for alternative propellants. Using the thruster geometries and operating data from Tab. 3.1, the neutral number densities for xenon and krypton were computed. The results are illustrated in Fig. 3.4 and summarised in Tab. 3.2. The peak values in Tab. 3.2 represent the operational points at which the respective maxima were achieved for each thruster.

A few observations emerge: krypton appears more sensitive to neutral number density than xenon, as the peak performance values each occur around a similar neutral density. However, this may be influenced by the limited krypton data available. Furthermore, the mean and median values for both xenon and krypton exceed the critical value of  $n_n \approx 1.2 \times 10^{19} \text{ m}^{-3}$ , with  $\bar{n}_n = 1.48 \times 10^{19}$  and  $\bar{n}_n = 1.66 \times 10^{19}$  for xenon, and  $\bar{n}_n = 1.43 \times 10^{19}$  and  $\bar{n}_n = 1.34 \times 10^{19}$  for krypton. This suggests that such neutral density targets must be carefully tailored to specific applications.

For instance, Su [81] demonstrated successful operation at high current and neutral number densities of a modified H9 Hall thruster. As the H9's dimensions are not stated explicitly, an exact value is not known; however, the nominal maximum flow rate for the H9 in literature has been 20 mg/s of xenon, with the H9-MUSCLE being tested up to 68.1 mg/s of xenon, corresponding to 340.5 % of the maximum flow rate [81, 168].

In some cases, peak performance is observed at significantly higher neutral densities than the established critical value (max column in Tab. 3.2). Consequently, performance values reported in the literature may not correspond to thermally sustainable operation, an aspect not captured in this analysis. This may reflect a

limitation of the database approach, as it does not implicitly account for thermal loading.

Complications arise from the fact that most thrusters tested on krypton were originally designed and optimised for xenon operation. These thrusters are then operated on krypton, sometimes without design modifications. Consequently, the observed trends for krypton may not reflect an ideal design point but rather emulate the performance of xenon-optimised configurations operating with krypton as the propellant.

### 3.4.3 Channel Length

Empirically scaling the channel length presents a challenge, as this parameter is not reported as frequently as others, such as power or thrust. This stems from the relatively weak direct correlation between channel length and key operational parameters. Instead, channel lengths are typically designed to satisfy the Melikov–Morozov criterion, discussed in Sec. 2.7.3, which states

$$\lambda_i \ll L . \quad (2.47)$$

To apply this criterion, an estimate of the ionisation mean free path is required. The mean free path is defined as

$$\lambda_i = \frac{u_n}{n_e \langle \sigma_i v_e \rangle} . \quad (2.37)$$

As discussed in Sec. 2.7.4, the ionisation “reaction rate” is used here as

$$\langle \sigma_i v_e \rangle = \int_{v_e}^c w \sigma_i(T_{eV}(w)) g(w, T_{eV}) dw , \quad (2.49)$$

where  $g(w, T_{eV})$  is defined in Eq. 2.48.

To obtain an estimate of the reaction rate, an estimate of the plasma electron temperature is required. Following the method outlined in Sec. 2.7.3, a widely used empirical relation between the bulk plasma temperature (in eV) and the anode discharge voltage for xenon discharges is given by

$$T_{eV} = 0.12U_d . \quad (2.45)$$

Although this empirical relation has been obtained for xenon discharges, it is assumed in this work to remain approximately independent of propellant species. This assumption is adopted in the absence of comprehensive alternative-propellant datasets and is considered reasonable provided the dominant physical mechanisms governing the relation remain unchanged.

Adhering to the Melikov-Morozov criterion allows a scaling coefficient to be defined from

$$L \propto \lambda_i . \quad (3.17)$$

This leads to a simple scaling form

$$L = C_L \lambda_i , \quad (3.18)$$

where  $C_L$  is a dimensionless channel length scaling coefficient.

### 3.4.4 Specific Impulse

Specific impulse is an important consideration when designing a thruster, as different missions or applications may require different ranges of specific impulse. As discussed in Sec. 2.1.1, it is a measure of how efficiently propellant mass is used in the momentum exchange process for rocket propulsion. The definition of specific impulse from Sec. 2.3.2 is defined as

$$I_{sp} = \frac{T}{g_0 \dot{m}_n} . \quad (2.17)$$

By substituting thrust from Eq. 2.34, specific impulse can also be expressed as

$$I_{sp} = \eta_m \frac{1}{g_0} \sqrt{\frac{2\bar{q}U_b}{m_n}} . \quad (3.19)$$

By applying several simplifying assumptions, proportional relationships between controlled parameters can be derived. Including the beam is assumed mono-ionic  $\bar{q} = e$ ,  $U_b \approx U_d$ , implying  $\eta_V \approx 1$ , and similarly mass utilisation is assumed close to unity,  $\eta_m \approx 1$ .

Although these assumptions are non-trivial, database scaling methodology combines results from many thrusters tested across different facilities, where detailed diagnostic measurements are often unavailable or reported inconsistently. Consequently, adopting simplified relationships enables a consistent comparison between thrusters while retaining the dominant physical trends. Since database scaling is intended to provide a first-order approximation rather than a predictive design model, it is expected that the resulting empirical scaling coefficients will implicitly describe the aggregate effects of these simplifying assumptions.

From Eq. 3.19 it can be seen that specific impulse appears to be independent to thruster geometry and purely a function of propellant and discharge voltage. As a




FIGURE 3.5: Database of thrusters with their specific impulse at peak anode efficiency against the normalised scaling value from Eq. 3.20 for xenon and krypton operating in SPT type Hall effect thrusters.  $R^2|_{\text{Xenon}} = 0.8143$  and  $R^2|_{\text{Krypton}} = 0.6305$ .

result, we can define a specific impulse scaling coefficient as

$$I_{sp} = C_{I_{sp}} \sqrt{\frac{U_d}{m_n}}, \quad (3.20)$$

where  $C_{I_{sp}}$  a propellant specific fitting coefficient. The result of this scaling coefficient is shown in Fig. 3.5 where the trend of specific impulse for discharge voltage and propellant is shown.

### 3.4.5 Thrust

Thrust is the fundamental performance parameter for spacecraft propulsion. It is defined as the momentum exchange with the exhaust which yields

$$T = \dot{m}_i v_i. \quad (2.12)$$

This relationship demonstrates that the thrust of a Hall effect thruster is proportional to both the ion mass flow rate and the ion velocity. By applying the same assumptions used in Sec. 3.4.4, the exhaust velocity can be seen as proportional to the square root of the discharge voltage divided by the propellant mass, and the ion mass flow rate as proportional to the neutral mass flow rate delivered to the thruster as

$$T \propto \dot{m}_n \sqrt{\frac{U_d}{m_n}}. \quad (3.21)$$

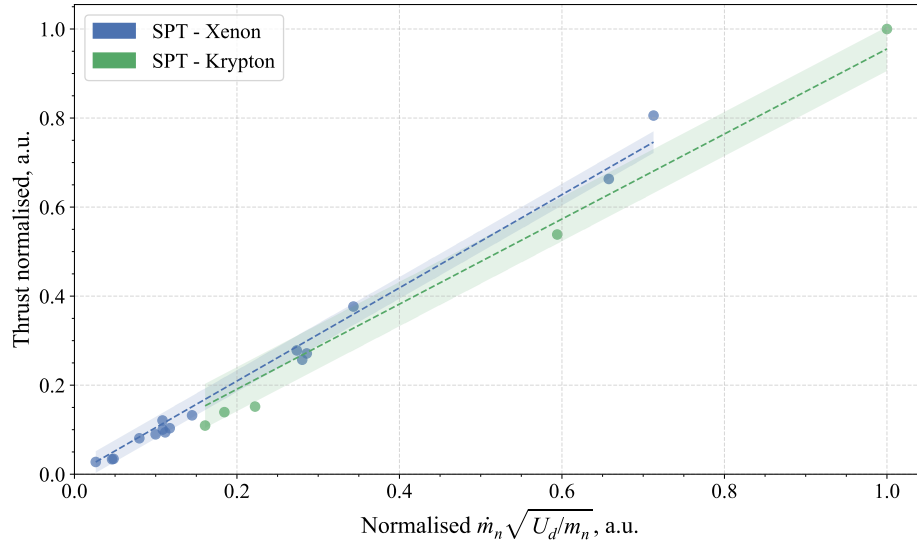


FIGURE 3.6: Normalised thrust plotted against the scaling proportionality function described in Eq. 3.22 for the SPT-type thrusters listed in Tab. 3.1.  $R^2|_{\text{Xenon}} = 0.9886$  and  $R^2|_{\text{Krypton}} = 0.9834$ .

This proportionality can be directly used to define a simple scaling law as

$$T = C_T \dot{m}_n \sqrt{\frac{U_d}{m_n}}, \quad (3.22)$$

where  $C_T$  is a propellant specific thrust scaling coefficient derived from database analysis. The result of Eq. 3.22 is illustrated in Fig. 3.6.

### 3.4.6 Anode Power

The primary contributor to the total power supplied to a Hall thruster is the anode discharge power, which arises from the applied discharge voltage and the resulting electron current through the magnetic confinement in the plasma discharge. The anode discharge power is given by

$$P = I_d U_d. \quad (3.23)$$

The anode discharge current  $I_d$  is the sum of the currents from all ionic species in the plume such as

$$I_d = \sum_{n=1}^N I^{n+}. \quad (3.24)$$

In Hall thrusters, this current can also be expressed using the mass flow rate and ionisation efficiency as

$$I_d \approx \eta_m \frac{\bar{q} \dot{m}_n}{m_n}. \quad (3.25)$$

By substituting in the mass flux from rearranging Eq. 2.39 and cross-sectional area from Eq. 2.40, it follows that

$$I_d \propto u_n h d . \quad (3.26)$$

Consequently, the anode discharge power becomes

$$P \propto u_n h d U_d . \quad (3.27)$$

This relationship is illustrated in Fig. 3.7a which demonstrates a reasonably linear trend.

Assuming a constant neutral density, the input power can be scaled in several ways. One common approach, suggested in prior scaling models, assumes  $h \propto d$ , which allows the simplification of Eq. 3.27 into

$$P \propto u_n U_d d^2 . \quad (3.28)$$

The results of this relationship are shown in Fig. 3.7b, where a linear fit is applied.

Furthermore, as illustrated in Fig. 2.4, a power-law relationship between the mean channel diameter and anode power is observed. To investigate this, a generalised power-law fit is applied under the assumption of

$$P \propto U_d u_n d^\chi . \quad (3.29)$$

The corresponding fit is shown in Fig. 3.7c. Although this formulation produces the best fit in terms of ( $R^2$ ) among the tested models, such behaviour is expected because power-law relations possess considerable flexibility when fitted to limited datasets. Furthermore, these fits are known to be sensitive to both the number of available data points and their distribution. In the present case, the krypton dataset is relatively small, preventing the determination of a statistically robust scaling exponent.

This introduces a methodological dilemma regarding the most appropriate formulation for power scaling. The relation given in Eq. 3.27, which retains both channel height  $h$  and diameter  $d$ , preserves more geometric information about the thruster and is therefore physically attractive. However, this formulation complicates the scaling procedure, since it prevents  $h$  and  $d$  from being treated as independent parameters, which is a key requirement of the scaling methodology adopted in this work.

Alternatively, the power-law expression in Eq. 3.29 provides a more compact representation and yields a better empirical fit to the available data. Nevertheless, this improvement likely reflects over-fitting due to the limited number of krypton

thrusters within the dataset. Indeed, the fitted exponent  $\chi$  varies between approximately 1.7 for xenon and 2.2 for krypton, suggesting that the value is sensitive to dataset composition and may lack universality.

To balance physical interpretability with empirical robustness, a representative value of  $\chi = 2$  is adopted. This choice is consistent with the geometric scaling of the thruster cross-sectional area and lies within the range obtained from the empirical fits. The resulting power scaling relation therefore becomes

$$P = C_p u_n U_d^2 , \quad (3.30)$$

where  $C_p$  is a propellant-specific power scaling coefficient.

This formulation provides a sufficiently accurate representation of the available data ( $R^2 > 0.9$ ) while maintaining a simple and generalisable scaling relation.



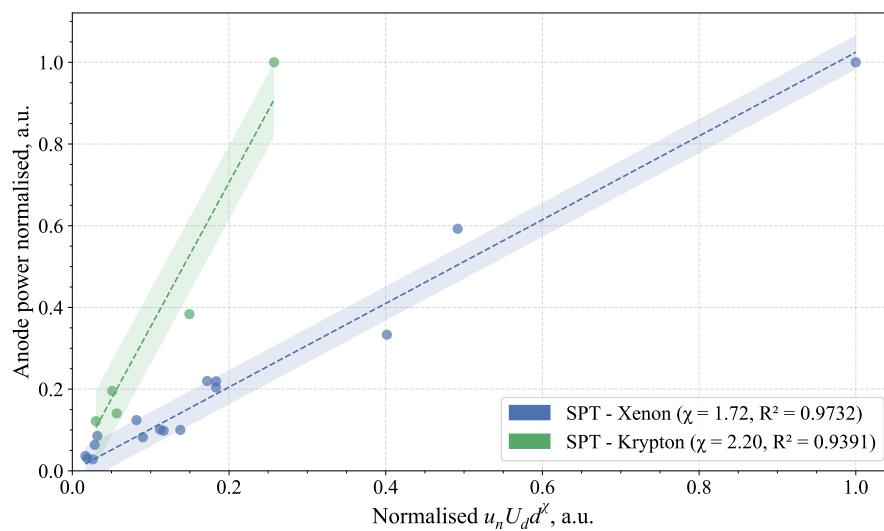
A. Power scaling with  $hd$  fit.B. Power scaling with  $d^2$  fit.C. Power scaling with  $\chi$  fit.

FIGURE 3.7: Comparison of anode power scaling fits using different characteristic dimension models.




FIGURE 3.8: The result of Eq. 3.31 for both xenon and krypton SPTs from Tab. 3.1.  $R^2|_{\text{Xenon}} = 0.9905$  and  $R^2|_{\text{Krypton}} = 0.9837$ .

### 3.4.7 Anode Current

In Hall effect thrusters, the ion beam current is directly related to the discharge current drawn by the anode as a result of quasi-neutrality within the plasma, making it a critical parameter for thruster operation. As a result, it is useful to draw scaling relations between the anode current based on fundamental thruster operation. A simplified relationship can be derived from Eq. 3.25, resulting in

$$I_d = C_{I_d} \frac{\dot{m}_n}{m_n}, \quad (3.31)$$

where  $C_{I_d}$  is the current scaling coefficient. The correlation from Eq. 3.31 is illustrated in Fig. 3.8, showing a strong linear relationship for both xenon and krypton propellants.

### 3.4.8 Channel Width

To move beyond a purely geometric correlation for channel width with respect to mean channel diameter, an alternative approach was considered based on thruster anode efficiency. One of the key sources of inefficiency in Hall thrusters is energy loss due to plasma-wall interactions. As thrusters are miniaturised, the discharge volume decreases faster than the surface area, leading to proportionally greater losses. This highlights the importance of the volume-to-surface-area ratio in determining thruster performance.

Assuming a toroidal discharge plasma volume, this ratio can be approximated as

$$\frac{V_{\text{plasma}}}{S_{\text{area}}} = \frac{\frac{1}{2}\pi^2 d h^2}{2\pi^2 d h} \propto h, \quad (3.32)$$

where  $V_{\text{plasma}}$  and  $S_{\text{area}}$  denote the plasma volume and surface area, respectively. This implies that increasing the channel width improves the volume-to-surface area ratio, which may in turn improve efficiency, subject to limits imposed by the current density required to sustain an efficient plasma and the physical constraint  $h < d$ .

As a result, it is suggested that the anode efficiency can be used to develop a scaling method for the channel width. From Eq. 2.22, it is seen that anode efficiency can be described in terms of more useful values for scaling as

$$\eta_a = \frac{T^2}{2U_d I_d \dot{m}_n}. \quad (3.33)$$

By inserting the scaling expression for thrust from Eq. 3.21 and mass flow rate from Eq. 2.39, and the expanded form of mass flow rate from Eq. 2.39, it can be shown that

$$\eta_a \propto \frac{n_n u_n \pi h d}{2I_d}. \quad (3.34)$$

Solving for the channel width and collecting anode efficiency within the scaling coefficient yields

$$h = C_h \frac{I_d}{n_n u_n d}. \quad (3.35)$$

This version is shown in Fig. 3.9. This formulation still reveals strong clustering by propellant type, indicating that this simplified relation retains predictive utility. This result is not a new concept with Dannenmayer [167] highlighting the relationship between larger channel widths and efficiency.

In summary, while historical scaling approaches often assume a fixed linear relation between channel width and mean diameter, analysis of a broader dataset reveals substantial variability, particularly across different thruster families and power classes. The conventional  $h/d$  approximation, while useful for initial estimates, is insufficient for capturing the nuances of modern Hall thruster design.

These findings reinforce the need for more nuanced, performance-informed scaling strategies when designing Hall thrusters, especially in regimes involving alternative propellants, high voltages, or miniaturised platforms. The proposed methodology offers a more flexible foundation for such efforts.

Using the scaling coefficient described in Eq. 3.35 achieves significantly improved  $R^2$  values compared to a direct link between the mean channel diameter and the channel




FIGURE 3.9: Normalised channel width as a function of the normalised scaling parameter from Eq. 3.35, evaluated at the operating point corresponding to maximum specific impulse.  $R^2|_{\text{xenon}} = 0.9197$  and  $R^2|_{\text{krypton}} = 0.8511$ .

width. The  $R^2$  values shown in Fig. 3.9 are  $R^2|_{\text{xenon}} = 0.9197$  and  $R^2|_{\text{krypton}} = 0.8511$ , compared to Fig. 3.2, which has  $R^2|_{\text{SPT}} = 0.4609$  using the same dataset.

Furthermore, the difference between the krypton and xenon values again highlights the need for a method that accounts for alternative propellants at the scaling stage.

## 3.5 Application of Scaling Method for Alternative Propellant High-Voltage Operation

### 3.5.1 Introduction

One of the core objectives of this thesis was to develop a high-voltage, krypton-optimised Hall effect thruster capable of supporting an anode-layer configuration. However, during the scaling analysis, it became clear that both the quantity and fidelity of available data were insufficient to directly design a krypton-optimised thruster; nor more broadly, any thruster using non-xenon propellants without introducing significant uncertainty.

As a result, the thruster designed in this work used a first-order estimation approach for krypton operation but was intentionally constructed with a fully modular discharge channel. This modularity enables systematic variation of channel geometry and supports future experimental studies that can build a more robust krypton scaling database in the 1–5 kW power regime.

A discharge power of 1.35 kW was selected for the scaling analysis. This corresponds to the design power of the SPT-100 thruster, a widely studied and validated device in the literature. A discharge voltage of 600 V was chosen for two main reasons. First, one objective of this study is to investigate high-voltage operation and its potential to enhance ionisation efficiency when using alternative propellants. Second, for a given power level, operating at a higher voltage reduces the discharge current, and thus the required propellant mass flow rate which helps reduce the load on pumps during vacuum chamber testing.

### 3.5.2 Scaling Methodology

To apply the scaling process outlined in this chapter, the following methodology was implemented:

1. A target performance parameter was selected (thrust, specific impulse, or anode efficiency). This parameter determined the optimisation objective and was used to extract a reference condition from the database.
2. The thruster database was filtered to include only SPT-type thrusters operating on krypton.
3. The operational condition corresponding to the maximum value of the selected performance parameter was identified. The neutral number density,  $n_n$ , was then calculated from the anode flow rate and the reported channel geometry.
4. A coefficient-fitting routine was applied to derive empirical scaling constants. For each scaling relation, the relevant independent and dependent variables were extracted from the dataset.
5. The following theoretical relations were used to define the form of each empirical coefficient:
  - $C_P$ : Anode power coefficient, fitted using Eq. 3.30.
  - $C_{I_d}$ : Anode current coefficient, fitted using Eq. 3.31 to determine the mass flow rate.
  - $C_T$ : Thrust coefficient, fitted using Eq. 3.22 to estimate thrust.
  - $C_h$ : Channel width coefficient, fitted using Eq. 3.35 to determine channel width.
  - $C_L$ : Channel length coefficient, derived from the ionisation mean free path estimate in Eq. 3.18.
  - $C_{I_{sp}}$ : Specific impulse coefficient, fitted independently using Eq. 3.20.
6. The derived coefficients were compiled into a coefficient table for krypton. These coefficients were then input into a predictive scaling function, which accepts discharge power and voltage as inputs and returns estimates of the geometry and

performance of a krypton-optimised Hall effect thruster. The predicted outputs are:

- Channel diameter,  $d$
- Channel width,  $h$
- Channel length,  $L$
- Thrust,  $T$
- Specific impulse,  $I_{sp}$
- Anode mass flow rate,  $\dot{m}_a$
- Anode efficiency,  $\eta_a$

7. The predictive scaling function was evaluated at a discharge power of 1.35 kW and a discharge voltage of 600 V. The resulting outputs were assessed for consistency with established scaling trends and physical feasibility.

### 3.6 Scaling Results

To scale the thruster for this thesis, the maximum specific impulse value within the database for each thruster operating on krypton was considered. This was motivated by one of the aims of this work: to investigate high specific impulse Hall effect thrusters. A secondary aim was to investigate high-voltage operation as a method to enhance the ionisation characteristics of alternative propellants. As a result, an anode discharge voltage of 600 V was chosen as the scaling target. This also provided the secondary benefit of reducing the mass flow rate for a given power, improving the likelihood that the vacuum chamber could achieve a lower base pressure during operation.

Initial scaling was performed with targets of 1.35 kW anode power and 600 V anode voltage. The 1.35 kW value was chosen to match the SPT-100's discharge power, as this thruster has been extensively tested in the literature. The results of this initial scaling can be seen in Tab. 3.3. However, this scaling was conducted using a larger database that contained some poorly performing thrusters, as well as several thrusters with incomplete geometric data.

Because the initial thruster scaling analysis was conducted using a larger database, the resulting target parameters did not align well with the intended design objective of a 1.35 kW thruster operating at a discharge voltage of 600 V. The expanded database contained several research thrusters with relatively low performance, thrusters with incomplete geometric information, and, in some cases, channel dimensions that had to be inferred from secondary sources. Such limitations introduce significant uncertainty and potential error into the scaling relations used to design the thruster.

TABLE 3.3: Scaling inputs and initial scaling results used to design the SHARK-600V Hall effect thruster for this thesis.

	Units	Value
Scaling inputs		
Anode power	W	1350
Anode voltage	V	600
Propellant		Krypton
Scaling outputs		
Channel width	mm	8.51
Channel mean diameter	mm	72.1
Channel length	mm	10.2
Chosen design		
Channel width	mm	8
Channel mean diameter	mm	70
Channel length	mm	15

To mitigate these issues, a refined thruster database was constructed, as shown in Tab. 3.1. This reduced dataset includes only thrusters with well-documented geometric parameters and demonstrated high-performance operation. This improved database minimise uncertainties and errors associated with poorly characterised and under-reported thrusters. The refined database was therefore used to re-evaluate the expected operating conditions for the designed thruster.

Because the thruster geometry had already been selected based on the initial scaling results, the resulting design does not correspond exactly to the original design objective of a 1.35 kW thruster operating at 600 V when using krypton propellant. Consequently, the true nominal operating conditions of the thruster are expected to differ from those originally targeted.

To determine the appropriate scaled operating point, the refined database was used to evaluate the design space corresponding to the selected channel geometry, specifically a channel width of 8 mm and a mean channel diameter of 70 mm. The resulting predicted operating conditions indicate an anode power of approximately 2.5 kW, with discharge voltages of approximately 580 V for krypton and 870 V for xenon. The full set of corrected scaling predictions for each propellant is presented in Tab. 3.4.

Argon was initially considered as a candidate propellant for scaling; however, due to the lack of available literature, this was not possible.

Furthermore, using the reduced dataset, the scaling results for a 1.35 kW, 600 V krypton Hall effect thruster differ from those obtained previously. Applying the method described in Sec. 3.5.2 results in a different geometry, as shown in Tab. 3.5.

TABLE 3.4: Scaling estimation for the chosen dimensions for both krypton and xenon.

	Units	Value
Target dimensions		
Mean channel diameter	mm	70
Channel width	mm	8
Results		
Propellant	Krypton	
Anode power	W	2553.2
Anode voltage	V	586.1
Thrust	mN	90.33
Specific impulse	s	2415.6
Anode mass flow rate	mg/s	3.42
Propellant	Xenon	
Anode power	W	2453.65
Anode voltage	V	874
Thrust	mN	82.55
Specific impulse	s	2850
Anode mass flow rate	mg/s	2.65

TABLE 3.5: Scaling inputs and updated scaling results for the design parameters of the original thruster using the reduced dataset.

	Units	Value
Scaling inputs		
Anode power	W	1350
Anode voltage	V	600
Propellant		Krypton
Scaling outputs		
Channel width	mm	11.73
Channel mean diameter	mm	51.43
Channel length	mm	7.64
Anode mass flow rate	mg/s	1.84

### 3.7 Summary of Scaling Methodology

This chapter introduced the historical development of scaling methodologies for Hall effect thrusters, ranging from the earliest published scaling relations to more recent data-driven approaches such as neural network models. Building on this background, the semi-empirical database scaling method was selected as the primary methodology which would build on prior methods and be extended for alternative propellants for

this work. This approach utilises a database of existing thrusters and their reported performance, combined with sets of governing relations derived from the underlying discharge physics. These relations aim to capture the complex plasma processes within empirical scaling coefficients that link key performance parameters, such as anode power, thrust, and discharge voltage, to the geometric dimensions of the thruster discharge channel.

However, the relative scarcity of well-characterised performance data for alternative propellants within literature was identified as a limitation. This restricts the direct application of database scaling techniques when designing thrusters intended to operate with propellants other than xenon. Despite this limitation, the methodology was applied to a large thruster database to determine an initial design corresponding to a target operating point of 1.35 kW at a discharge voltage of 600 V using krypton propellant. The resulting scaling analysis suggested channel dimensions of approximately 8 mm channel width and 70 mm mean channel diameter.

Given the limited number of well-characterised krypton thrusters available for inclusion in the database, the confidence in these scaling predictions is inherently limited. This uncertainty informed the design choice adopted in this thesis, whereby the thruster incorporates a modular discharge channel. This configuration enables systematic investigation of the influence of channel dimensions on thruster performance across multiple propellants.

Subsequently, the thruster database was refined to provide a more reliable estimate of the nominal operating conditions for the selected geometry. The reduced dataset included only high-performance thrusters with fully reported geometric parameters, thereby improving the robustness of the scaling relations. Application of this refined database suggested that the expected design power for the selected channel dimensions is closer to approximately 2.5 kW for xenon or krypton operation.

## Chapter 4

# SHARK-600V Hall Thruster Design

The design process for a Hall effect thruster involves several challenges in achieving an efficiently operating device, even after scaling has been completed. For the SHARK-600V, additional difficulties arose from the need for modularity and the requirement to operate with three different propellants. Furthermore, the thruster incorporated an optional anode-layer configuration, which had to be integrated into a design that also employed the same gas distributor for magnetic-layer operation.

This chapter describes the process taken to design the Southampton High-voltage Anode-layer Research Krypton 600 V Hall effect thruster, termed the "SHARK-600V" thruster. The design started with implementation of scaling outputs as described in Sec. 3.5, at a nominal operating point corresponded to a discharge power of 1.35 kW and an anode voltage of 600 V. However, because of the modular discharge channel, this value was treated as a general design guideline rather than an absolute target.

A range of channel widths was selected around the mean channel diameter of 70 mm. Once the range of channel widths of interest had been established for the selected mean diameter, the magnetic field design process was undertaken. This initially employed two-dimensional axisymmetric numerical methods and was subsequently refined using three-dimensional finite element simulations, as described in Sec. 4.3.

In parallel, the propellant delivery system and anode assembly were developed, as detailed in Sec. 4.4 and Sec. 5.5. The results of the thermal modelling were also incorporated to guide the overall design process, as described in Sec. 4.5.1.

## 4.1 Modular SHARK-600V Krypton Hall Effect Thruster Concept

Once the channel dimensions had been estimated using the scaling methodology described in Sec. 3.4, an initial design was obtained with a mean channel diameter of 72.1 mm, a channel width of 8.51 mm, and a channel length of 6.7 mm. For practical design and manufacturing considerations, these values were rounded to nominal dimensions of 70 mm for the mean channel diameter, 8 mm for the channel width, and 10 mm for the channel length.

Starting from this nominal configuration, a modular channel architecture was adopted to enable systematic investigation of the sensitivity of thruster performance to channel geometry. A range of channel widths was selected while maintaining a constant mean diameter of 70 mm, specifically 8 mm, 12 mm, and 17 mm. These widths were chosen such that each successive constant mean diameter configuration represented approximately a 50% increase in channel cross-sectional area relative to the previous configuration. These configurations were implemented through interchangeable inner and outer channel inserts, allowing the effective channel width to be varied without altering the overall thruster structure. The smallest channel width corresponds to the scaling output and was selected due to the high aspect ratio required by the high-voltage design. For tests at standard voltages while maintaining the same power, wider channel dimensions are preferable, providing flexibility to explore a wide range of operating conditions.

TABLE 4.1: Channel width (mm) and mean channel diameter (mm) for each configuration of the modular thruster. Each cell shows “h mm, d mm” for that inner and outer configuration.

Outer	Inner		
	SPT +0	SPT +1	SPT +2
SPT +0	8, 70	10, 68	12.5, 65.5
SPT +1	10, 72	12, 70	14.5, 67.5
SPT +2	12.5, 74.5	14.5, 72.5	17, 70

The selected width increments were deliberately non-uniform to capture performance gradients across the range of geometries tested. Each insert extends to a depth of 10 mm, corresponding to the channel length suggested from the scaling required for sufficient ionisation efficiency. With the exception of the widest configuration, where the channel body already matches the required inner and outer diameters. The use of the inserts is illustrated in Fig. 4.1 for the three constant mean diameter cases. By systematically combining inner and outer inserts, both the channel width and mean channel diameter can be varied, allowing multiple channel configurations to be realised. This approach expands the range of testable geometries within a single thruster platform requiring minimal modification between tests. Since the inserts are




FIGURE 4.1: A graphical illustration of the channel inserts, showing the three constant mean diameter cases “+0, +0”, “+1, +1” and “+2, +2”.

mounted at the front of the thruster, only the retaining ring must be removed before interchanging the inserts to achieve any of the channel configurations listed in Tab. 4.1, allowing for fast changing of the thruster geometry.

## 4.2 Thruster with Anode-Layer Design

Some literature has suggested that a Hall thruster operating in an anode-layer configuration could achieve higher plasma temperatures [29] and given the importance of electron temperature to ionisation of alternative propellants an anode-layer variant was designed. This effect arises from the secondary electron emission yield of conducting walls, which reduces the number of cold secondary electrons introduced into the plasma, as described in Sec. 2.8. Whilst this behaviour is generally considered detrimental for xenon operation, where the propellant is already readily ionisable in magnetic-layer thrusters, alternative propellants with lower ionisation cross-sections may benefit from the higher plasma temperatures of TAL configurations.

However, due to the reduced research interest in TAL-type thrusters after the early 2000s, comprehensive scaling laws and established design philosophies for TALs are limited. Consequently, the TAL design for SHARK-600V was based on the scaling outputs used for the magnetic-layer configuration, to maintain consistency and enable straightforward interchangeability between configurations. The design used the same channel width of 8 mm and mean channel diameter of 70 mm from the scaling results in Sec. 3.5, whereas here the channel width is defined as the distance between the inner-diameter and outer-diameter anode sheaths. As with the magnetic-layer thruster the TAL design also incorporated modularity, ensuring that multiple TAL variants could be tested without extensive reassembly of the thruster. The modularity for the TAL configurations focuses on the variable anode length as interactions with the bulk plasma are likely a key driver in TAL performance.

The resulting TAL anode cross-sections are shown in Fig. 4.2, where the variation in anode-layer length is evident (shown in red). The lengths of the anode-layers are




FIGURE 4.2: Computer renderings of cross sections of the SHARK-600V in all of the thruster with anode-layer configurations designed. The magnetic circuit and anode are illustrative only and do not represent the final design. Here “L” the TAL configuration channel length is shown as reported in Tab. 4.2.

TABLE 4.2: Table of all the thruster with anode-layer configurations of variable channel lengths as seen in Fig. 4.2. Here negative values represent that the anode extends past thrusters magnetic exit plane.

Anode-layer configuration	Distance from magnetic exit plane
TAL	2.5 mm
TAL+1	-0.5 mm
TAL-1	5.5 mm
TAL-2	8.5 mm

defined relative to the magnetic front of the thruster, as listed in Tab. 4.2. The four configurations span a total range of approximately 9 mm between the longest and shortest anode tips.

Although in the nominal case the anode-layer tip remains relatively close to the exit plane, 2.5 mm, compared to what is typical for magnetic-layer thrusters, the set of equations in Eq. 3.2 suggest that the anode exit plane separation in a TAL should be on the order of only a few mean free paths. This constraint produces a much smaller region for ionisation and acceleration in TAL operation. A similar conclusion can also be reached analytically using models that account for secondary electron emission and the quenching temperature of electrons impinging on the walls [29].

### 4.3 Magnetic Field Design

The target magnetic-field topology for the SHARK-600V was an unshielded design. Magnetic shielding was not incorporated because such magnetic topologies are typically tailored to a specific discharge-channel geometry, and the modular nature of the SHARK-600V prevented implementation of a bespoke topology that could be varied as channel geometry changed. In addition, the magnetic topology of anode-layer thrusters has historically been based on unshielded designs [138, 169]. Given this thesis' dual mandate to investigating channel geometry and enabling an optional anode-layer configuration, a traditional unshielded magnetic design was therefore selected.

The target magnetic-field strength at the channel exit was 150 to 300 G, which satisfies the magnetic-field requirements discussed in Sec. 3.4.1. The experimental nature of the SHARK-600V motivated the use of electromagnets for the magnetic circuit. Electromagnets allow control of the magnetic-field magnitude within the channel by adjusting coil current; this not only permits parametric investigations of the field but also assists ignition. Specifically, a glow discharge can be established in the absence of a magnetic field, after which the field is introduced to confine electrons, to form a high plasma potential and accelerate the ions.




FIGURE 4.3: A graphical rendering of the cross-section and iso-metric views of the SHARK-600V Hall effect thruster with labels highlighting the electromagnet locations.

The electromagnet coil design was strongly influenced by the laboratory power supplies available, which typically provided a maximum of 3 A of current. Consequently, the coils required a large number of turns to achieve the necessary field intensity. This requirement, combined with the expected thermal load from the thruster, motivated the decision to position the electromagnets behind the channel at the rear of the thruster, as can be seen in Fig. 4.3. This placement increased the thermal

path between the plasma-facing surfaces and the coils, thereby reducing the thermal load and lowering the effective coil resistance. Since resistance is temperature dependent, this arrangement reduced the voltage required to drive the coils. Additionally, positioning the electromagnets at the rear provided a larger physical volume for the coils, permitting more turns without interfering with the discharge channel geometry. Three outer coils and one inner coil was chosen to minimise the number of electromagnets needed whilst providing a sufficient number that a uniform field within the channel could be established.

The overall size of the thruster also necessitated the use of discrete electromagnets rather than continuous electromagnets for the outer coils. Constructing two large coaxial continuous coils would have represented a significant expense. This was undesirable not only due to the initial cost but also because it introduced a higher risk of single-point failure: the loss of one large continuous coil would have required costly replacement.

#### 4.3.0.1 Magnetic Field Simulation: Prototyping

The magnetic field was simulated at several stages of the design process. Initially, rapid two-dimensional axisymmetric simulations were performed using Finite Element Method Magnetics (FEMM). These FEMM simulations were used to confirm that high enough magnetic fields were possible with the electromagnets moved to the rear of the thruster. FEMM employs a finite element approach that is restricted to two dimensions. While useful for preliminary studies, this method cannot accurately capture the field distribution in the SHARK-600V thruster due to the non-axisymmetric nature of the thruster. Nevertheless, these simulations provided a fast and effective means of iteratively refining the larger, non-critical elements of the magnetic circuit. An example of an early FEMM simulation is shown in Fig. 4.4, where the material assignment of “1010 Steel” and the electromagnet wire diameter can be seen. Although these simulations were low fidelity and unable to represent the discrete outer coils, they could be completed in a matter of seconds, enabling rapid design iterations.

Once the initial magnetic circuit was defined, higher-fidelity models were developed. Three-dimensional simulations were performed using COMSOL Multiphysics®<sup>1</sup>. These simulations enabled detailed investigation of the full thruster geometry. In particular, the focus was on parametrically modifying the “magnetic screen”: a ferromagnetic annulus surrounding the ceramic channel, illustrated in Fig. 4.1. This component plays a central role in preventing high magnetic fields near the anode and strongly influences the magnetic topology across the thruster exit plane. Accordingly, a large portion of the optimisation effort concentrated on refining the design of this

<sup>1</sup>COMSOL Multiphysics® v. 6.3. [www.comsol.com](http://www.comsol.com). COMSOL AB, Stockholm, Sweden.

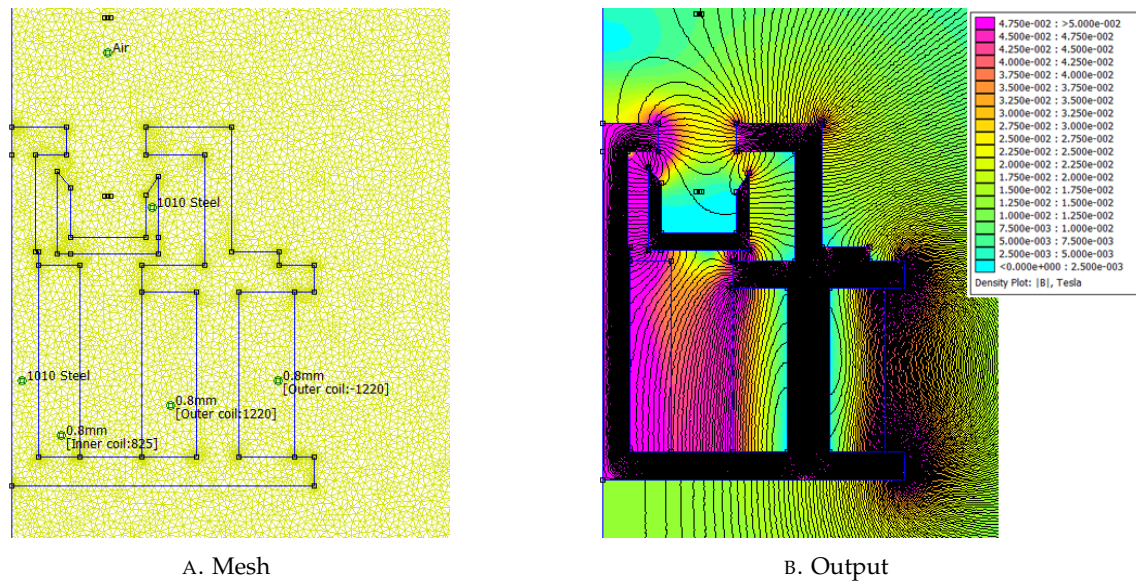


FIGURE 4.4: Illustration of a single simulated configuration of the FEMM mesh and output.

element. All simulation within COMSOL used non-linear B-H curve models for Ampère's Law within solids.

This process was carried out through a comprehensive set of parametric simulations, systematically varying the geometry of the magnetic screen to identify designs that achieved the desired field topology. Three sequential geometry sweeps were performed, each increasing in complexity of the magnetic screen once the previous sweep produced acceptable topologies. To enable the large simulation throughput required by this parametric approach, lower fidelity meshes were employed, and workloads were offloaded to a supercomputing cluster<sup>2</sup>. Access to this facility was essential in enabling multiple iterations of the magnetic field design.

An example of some of the tested configurations is provided in Appendix A and Figs. 4.5a–4.5d. In excess of 500 simulations were undertaken due to the highly parametric nature of the thruster design space. For all simulations shown, the same number of coil turns was used: 800 for the inner coil and 1200 for the three outer coils (Fig. 4.3). Initial simulations varied the number of turns for both the inner and outer electromagnets, with these values ultimately chosen to achieve the correct field with the same current, allowing the coils to be wired in series. These simulations were performed prior to coil assembly for the final thruster, and some variation was anticipated in the assembled device.

For each sweep, the full set of all possible permutations of screen geometry was not simulated; instead, a random selection was employed to keep the number of simulations manageable. This approach ensured that a sufficiently broad range of


<sup>2</sup>The authors acknowledge the use of the IRIDIS High Performance Computing Facility and the associated support services at the University of Southampton in the completion of this work.

possible outcomes was captured. In sweep one, 28 configurations were simulated with varying magnetic screen dimensions, while sweeps two and three each involved 16 configurations. Results from sweep one included both poor- and good-performing cases. Example simulation results from sweeps one and two are provided in Appendix A. Poor-performing screens exhibited asymmetry (Fig. A.1b), large non-radial magnetic field components (Fig. A.1c), or significant field strength at the base of the anode (Fig. A.1a). Identification of good-performing configurations allowed tighter parameter bounds to be applied in subsequent sweeps.


A key observation was that reducing the magnetic field strength at the anode required thick screen walls, but this also suppressed the exit-plane field. This limitation was resolved by introducing two discrete thicknesses for both the inner and outer screen walls: a thick base section to suppress the anode field and a thinner exit section to maintain the high exit-plane field necessary for electron confinement, whilst preserving the required topology. While effective, this dual-thickness approach expanded the parameter space considerably, as the base and exit thicknesses, step locations, and tip geometry all became independent variables. Sweep two was conducted to size these screen elements, but the random application of parameters did not yield any consistently well-performing configurations.

As a result, sweep three was conducted with closer supervision of the parameter choices. The observations from sweeps one and two that guided sweep three were: (i) the inner screen needed to be positioned close to the central pole, (ii) the outer screen should be shorter than the inner screen, and (iii) the axial height of the thick base section had little effect on exit-plane field strength provided it remained short relative to the thinner exit sections. Sweep three therefore focused on variations in the thickness of the thin sections, their axial heights, and the tip angles of both screens. A subset of the 16 simulated configurations is shown in Fig. 4.5. From these results, configuration 10 was chosen as the final design because of its symmetry about the mean diameter line and its favourable magnetic gradient along the channel walls. This design was then further refined, and higher-fidelity simulations were performed.


Furthermore, after the three geometry sweeps there was some concern that there could be some material saturation. As a result, additional simulations were run with the magnetic screen material changed to MuMetal®; however, these resulted in negligible difference allowing the screen to be manufactured out of the much cheaper low carbon steel “1010 Steel”.




A. Parametric configuration 3.



B. Parametric configuration 6.



C. Parametric configuration 10.



D. Parametric configuration 13.

FIGURE 4.5: Magnetic field simulations for selected geometry sweep 3 showing magnetic field streamlines for 2A inner coil and 2A outer coil currents.

### 4.3.1 Magnetic Field Simulation: Final

The final magnetic screen design identified from the low-fidelity parametric study is shown in Fig. 4.5c. This configuration, while effective in simulation, included non-physical features such as zero-radius corners. Although such sharp features could, in principle, be approximated by machining with extremely small radii ( $\ll 0.1$  mm), doing so would significantly increase manufacturing complexity, cost, and time. To address this, small but finite radii were introduced at all critical edges in the design to ensure practical manufacturability without compromising functionality.

Following this modification, the geometry was re-simulated using higher fidelity meshes to capture fine details of the field topology more accurately. These final simulations provided the baseline magnetic field distributions used for thruster assembly and subsequent experimental validation. The results of this high-fidelity simulation for coil currents of 1 A and 2 A applied to both sets of electromagnets are shown in Fig. 4.6.

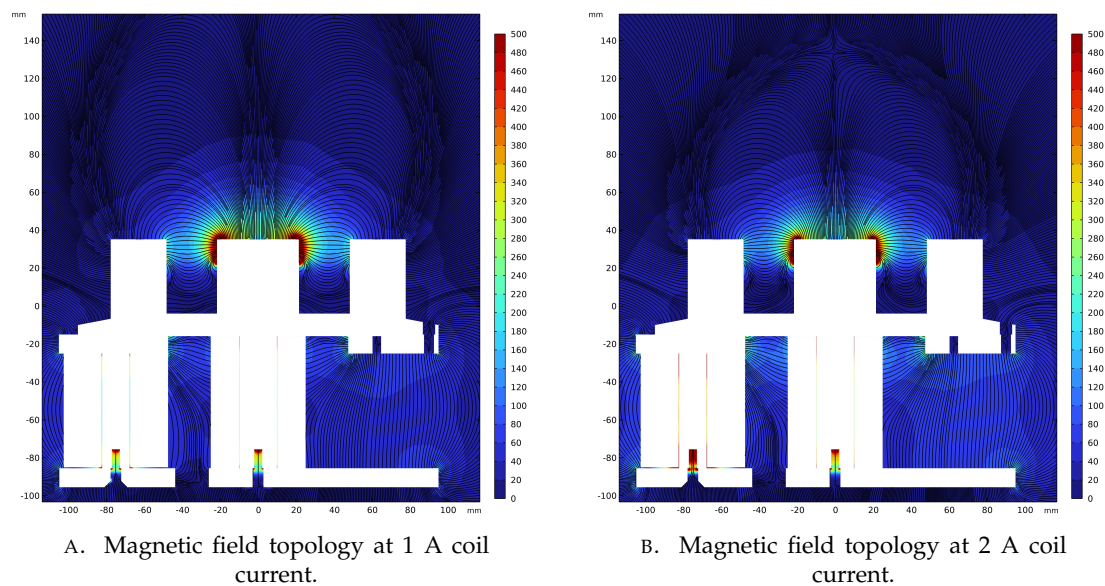



FIGURE 4.6: Comparison of full-thruster magnetic field (shown in units of Gauss) topology at coil currents of 1 A and 2 A. Note: the internal structure has been obscured.


Fig. 4.6 shows the normal magnetic field strength with respect to the magnetic field vector. However, the gauss meter available in the laboratory can only measure the normal magnetic field strength relative to the probe orientation. While it is theoretically possible to reconstruct the magnetic field vector by sweeping the probe angle at the point of interest, this approach would be cumbersome in practice. A three-axis probe could achieve this directly, but such probes typically have a larger sampling volume, which reduces spatial resolution.

For this reason, it is more practical to consider the radial component of the magnetic field, which corresponds to the z-component in the simulations. The z-axis represents

the radial as the simulation cross-section was taken in the  $y$ - $z$  plane, where  $y$  is the vertical axis (thruster axial). The simulation results for the radial component of the magnetic field at coil currents of 1 A and 2 A applied to both the inner and outer electromagnets are shown in Fig. 4.7. These results will allow comparison with the measured spacial map of the radial magnetic field in Sec. 4.3.2.



A. Radial magnetic field (shown in units of Gauss) component ( $B_z$ ) at 1 A coil current.



B. Radial magnetic field (shown in units of Gauss) component ( $B_z$ ) at 2 A coil current.

FIGURE 4.7: Comparison of radial magnetic field (shown in units of Gauss) component ( $B_z$ ) at coil currents of 1 A and 2 A. Note: the internal structure has been obscured.

From Figs. 4.6 and 4.7, it is evident that there is negligible variation in the magnetic field within the channel between the left-hand side (an outer coil node) and the right-hand side (an outer coil antinode). This demonstrates that the discrete nature of the outer coils does not adversely affect the overall magnetic topology.

### 4.3.2 Magnetic Field Measurements

Once the thruster was assembled, the magnetic field was measured to validate the simulation results and to ensure that the magnetic circuit had been machined from the correct material. A Gauss probe mounted on two linear stages was used to obtain measurements in two physical dimensions, enabling the formation of a heat map of the radial magnetic field strength. The experimental setup for these measurements is shown in Fig. 4.8. The Gauss probe used was the unidirectional GM08 Gauss Meter (Hirst Magnetic Instruments, UK), and the linear stages were MTS50-Z8 (Thorlabs Inc., USA), each with a 50 mm travel range. During measurement, the Gauss probe readings were continuously recorded at 1 Hz while the linear stages incremented through a 1 mm grid, pausing for 3 s at each point within the channel to save data. After the full range was measured, a Python script averaged the probe data to form a three-dimensional dataset for plotting.




FIGURE 4.8: The SHARK-600V prior to testing mounted on a flat plate for magnetic field mapping.

The Gauss probe measures only the field component normal to its sensor axis; as a result, only the radial (z-direction in the simulation domain) magnetic field density was measured at each axial and radial location. To verify that there was no azimuthal variation in magnetic field density around the channel, the measurement was repeated at both a coil node and a coil anti-node by offsetting the vertical linear stage to the top and bottom of the thruster channel, as illustrated in Fig. 4.8.

Due to the limitations of the available bench power supplies, the electromagnets were supplied with only 1 A to both the inner and outer coils. During the field mapping experiments, the outer electromagnets were operated in series, independently of the inner electromagnet, whereas during thruster operation all electromagnets are powered in series from a single supply.

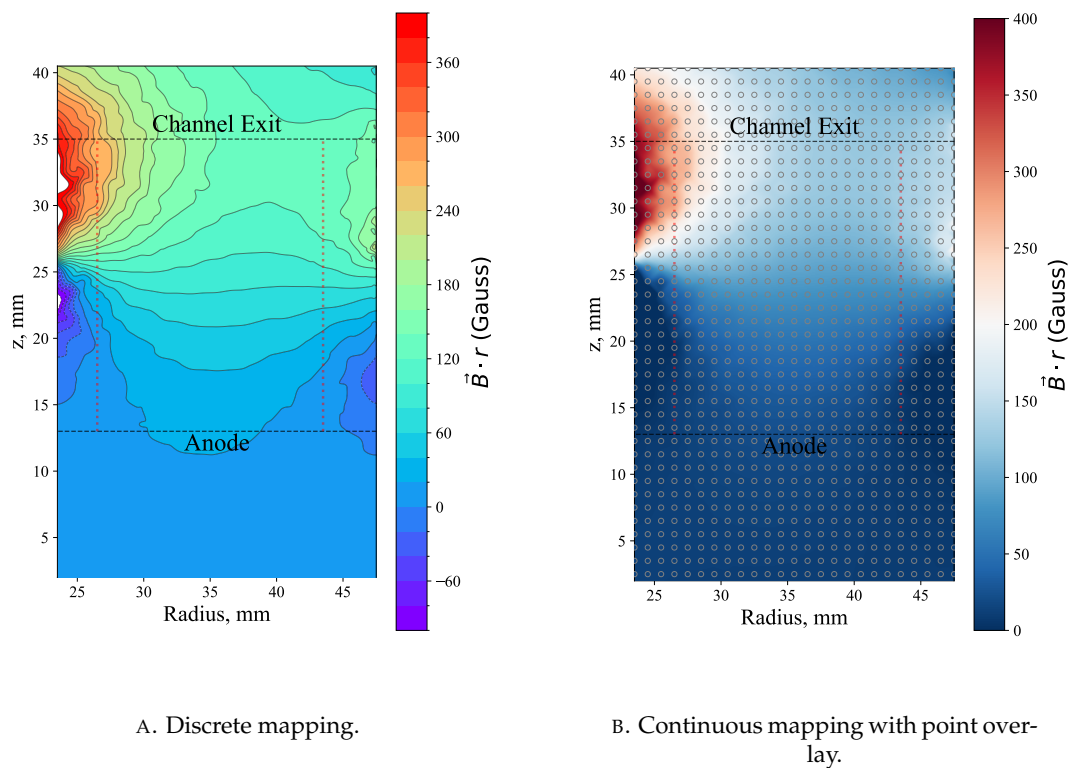


FIGURE 4.9: Comparison of radial magnetic field measurements of the SHARK-600V at an anti-node with 1 A coil current using two mapping approaches. In Fig. 4.9a, the lines shown are iso-contours bounding regions of constant magnitude within the colour map bins, *not* magnetic field lines.

The results of the field mapping are shown in Fig. 4.9, where two mapping approaches are compared. Note: Fig. 4.9a and Fig. 4.9b are plotting the same dataset from the same measurement, however, are showing discrete and continuous interpolating from the discrete spacial radial-magnetic field density measurement. In Fig. 4.9a, the magnetic field is presented as a discrete heat map ranging from 400 G to  $-100$  G, with iso-contours bounding regions of constant value within 20 G increments. The two horizontal lines in both Figs. 4.9a and 4.9b indicate the top of the anode at 13 mm and the thruster exit plane at 35 mm. The two vertical lines mark the inner and outer radii for the widest configuration (“+2, +2”). Fig. 4.9b further illustrates how the data in both figures are interpolated from a set of discrete measurement points, shown as dotted markers.

Because the thruster was designed to be modular, accommodating a wide range of channel diameters and widths, it produces a correspondingly wide range of possible magnetic field strengths along the inner diameter (ID), outer diameter (OD), and geometric mean diameter (MD). Magnetic field values for each geometry can be interpolated from the measurement dataset, and the results of this comparison are shown in Fig. 4.10. A particularly large variation is observed in the peak magnetic

field along the ID, with measured values of 182.2 G, 217.3 G, and 281.5 G for the “+0,+ #”, “+1,+ #” and “+2,+ #” configurations, respectively.

While there is significant variation in the magnetic field along the channel walls across these configurations, the plasma will form slightly above the wall surface. This offset is due to the formation of a plasma sheath between the zero-potential ceramic wall and the plasma, which is near the anode potential. Nevertheless, this field variation has important implications for the  $\mathbf{E} \times \mathbf{B}$  formation within the channel. As described in Sec. 3.4.1, the increased field strength near the channel walls could enhance electron confinement by reflecting electrons back into the bulk plasma via magnetic mirroring, rather than allowing them to impinge upon the channel walls.




FIGURE 4.10: Comparison of radial magnetic field strength measured along the inner diameter (ID), outer diameter (OD), and mean diameter (MD) of the SHARK-600V with coil currents of 1 A.




FIGURE 4.11: Front one image of the SHARK-600V Hall effect thruster in “+2, +2” channel configurations prior to testing where anode injection holes are clearly visible.

#### 4.4 Anode Design

The anode of the SHARK-600V thruster is required to perform several functions within the Hall thruster discharge. Firstly, the anode serves as the biased point for the plasma and therefore must be electrically conductive while sustaining operation under the thermal loads associated with the electron current. Secondly, the anode acts as the distributor for neutral propellant. Uniform distribution of neutral gas into the discharge chamber is a key requirement for efficient operation [170–172]. Lastly, for the SHARK-600V specifically, the anode must be designed such that it can be augmented into an anode-layer configuration without modification.

Building on successful prior designs and considering ease of manufacturing, the anode was fabricated as a single part using selective laser sintering, an additive manufacturing technique suitable for manufacturing from metal [12]. This method allows complex internal channels to be incorporated, enabling delivery of the single propellant inlet evenly to the discharge channel. The anode design incorporates three consecutive parallel channels connected by an increasing number of holes, culminating in the final channel where the gas is expected to be evenly distributed. The anode was additively manufactured from stainless steel 316L, a standard non-magnetic steel.

The directionality of propellant injection has a known influence on thruster performance [170, 171]. Although radial injection can yield higher performance in some thrusters, this approach was not feasible here because the insertable sheaths required for the anode-layer configuration would obstruct such injection.




FIGURE 4.12: A illustration from the computer rendering of the anode design where the angled propellant injection is visible, highlighted with two red-lines indicating the inner and outer edges of the hole and the centre-line indicating the holes angle relative to the front surface.


Consequently, axial propellant injection was adopted, with the injection holes visible on the anode in Figs. 4.11–4.12. To increase neutral residence time in the channel, the injection holes are angled at  $25^\circ$  relative to the front plane. This design choice was motivated by the intent to implement a “swirl” injection approach, which has been shown to enhance performance in some thrusters while remaining compatible with the anode-layer configuration of the SHARK-600V design [49, 173].

Direct measurements of the flow distribution within the channel were not conducted during this test campaign due to the lack of the required supporting infrastructure, such as three-dimensional translation stages and accurate low-pressure pitot tubes and appropriate pressure sensors. Furthermore, rarefied flow simulations were also not undertaken for the design, as the combination of several flow regimes would render such simulations speculative at best. It was therefore decided to iteratively update the design if experimental observations indicated poor neutral distribution from the anode, which was not observed.


#### 4.4.1 Anode X-Ray Computed Tomography Scans

After the manufacturing of the anode, it was of interest to inspect its internal geometry by non-destructive means. To achieve this, the anode was imaged at the  $\mu$ -VIS X-ray Imaging Centre, University of Southampton.


X-ray Computed Tomography (X-CT) utilises high-energy X-rays to acquire a large number of radiographs of the sample from different angles [174]. These 2D projections are then reconstructed into a 3D volumetric dataset using tomographic algorithms, allowing the internal structure to be visualised and analysed without physically




A. Top slice of anode cross-section.




B. Top channel slice.




C. Middle diffuser slice.



D. Middle channel slice.



E. Bottom diffuser slice.



F. Anode pipe slice.

FIGURE 4.13: X-ray CT slices of the SHARK-600V anode at different depths, showing internal channels and features.

sectioning the component. The method is particularly suited for additively manufactured components, such as the SHARK-600V's anode, where complex internal channels and small-scale features are otherwise inaccessible to direct inspection. Due to scheduling access to the X-CT equipment the imaging was undertaken after the test campaign had concluded. However, there is no apparent deformation or alteration from the structure as a result of the testing.

The results of the X-CT scanning can be seen in Figs. 4.13 and 4.14. It is important to note that in X-CT scans of metallic objects, artefacts and anomalies commonly arise due to diffraction and low-energy scattering. This effect is evident in all slices, particularly in the centre of the images in Fig. 4.13, where the inner wall appears "fuzzy." A similar artefact can also be seen in Fig. 4.14, where "ghosting" occurs on the right-hand side of the image.




FIGURE 4.14: Side-on X-ray CT annotated slice of the anode showing the internal propellant pipe and highlighting a manufacturing irregularity in the bottom channel.

The X-CT scans were acquired with the X-ray beam oriented from right to left in these images. As a result, the left-hand side (the "9 o'clock" position) is expected to exhibit significant artefacts caused by scattered X-rays.

There are several notable features in Fig. 4.13. In Fig. 4.13a, the anode injection holes appear oval due to their angle relative to the front plate, also illustrated in Fig. 4.12. Two larger holes are also visible; these do not extend through to the top channel but instead serve as mounting points for anode-layer sheaths. In Fig. 4.13b, the inner wall appears to have a lower density than the outer wall. Although some porosity in this wall was identified previously during water pressure testing, it is unclear whether the apparent lower density in this scan reflects true porosity or arises from artefacts. The region in the "9 o'clock" position is expected to be particularly prone to artifacting;

however, the trend of reduced density in the inner channel wall is consistently observed. The impact of this feature on flow rate is difficult to quantify.

The slices in Figs. 4.13c–4.13e show similar behaviour. Propellant was delivered to the anode through the back of the thruster via a single pipe that was manufactured separately (due to printer bed size constraints) and subsequently press-fitted into the anode body. The notched top of this pipe can be seen clearly in Fig. 4.13f, this was done to prevent sealing the top of the pipe with the ceiling of the bottom channel. The anode pipe is also visible in Fig. 4.14, where a side-on orientation highlights the versatility of X-CT imaging. In this view, the wall between the base channel and the middle channel appears thinner than desired, with noticeable porosity. This raises the possibility that propellant could bypass the first diffusion stage and flow directly into the middle channel, although confirmation would require higher-resolution imaging. As a result of these findings, subsequent anodes were manufactured with thicker separating walls between channels however were not tested within this campaign.

#### 4.4.1.1 Acknowledgement

The authors acknowledge the  $\mu$ -VIS X-ray Imaging Centre (<https://muvis.org>), part of the National Facility for laboratory-based X-ray CT ([nxct.ac.uk](http://nxct.ac.uk) – EPSRC: EP/T02593X/1), at the University of Southampton for provision of tomographic imaging facilities.

## 4.5 Thermal Design Considerations

The thermal design of a Hall thruster is a critical factor in determining both lifetime and operational limits. In TAL-type thrusters, thermal loading has been identified as the key constraint on maximum discharge voltage, which in turn limits achievable specific impulse [89]. This is due to the measured increase in thruster channel temperature with increasing voltage (as well as power) seen in literature [175–177]. Thermal effects influence several aspects of thruster performance. Firstly, the spacecraft bus typically imposes an upper limit on the thermal load from the propulsion subsystem. Secondly, electromagnet coils require increased voltage to maintain constant current at elevated temperatures due to the rise in electrical resistivity. If permanent magnets are employed, their Curie temperature poses a further limit, as exceeding it risks demagnetisation. Finally, high anode temperatures increase the injected propellant's thermal velocity (see Eq. 2.36), reducing residence time in the discharge channel and thereby decreasing ionisation efficiency. While this is not an exhaustive list, it illustrates the central importance of thermal management.

Thermal loads within the thruster can be mitigated in several ways. Reducing plasma–wall interaction (e.g., by adopting a magnetic shielding topology), lowering plasma temperature by operating at reduced discharge voltage, or active and passive cooling mechanisms are common approaches. One such passive method recently employed, novel embedded heat pipes for passive thermal transport have been demonstrated in high-power Hall thrusters [86]. As a result, a method for estimating the thermal load to the thruster during operation is required.

#### 4.5.1 Zero-Dimensional Thermal Model

A useful method for estimating the thermal load is to consider the plasma processes that deposit energy into the thruster walls. This is done by formulating simplified expressions for the dominant power sinks, neglecting detailed geometry to yield a zero-dimensional approximation of thermal fluxes during operation. The model distinguishes between SPT and TAL configurations by incorporating differences in secondary electron emission coefficients, as discussed in Sec. 2.8, and accounts for the influence of SEE on the sheath potential drop at the walls.

To describe how discharge power is partitioned, a control-volume perspective is adopted in which power leaving the system is expressed as the sum of plume, radiation, and thermal loads, such that

$$P_d = \sum_{n=1}^N P_n \approx P_{\text{plume}} + P_{\text{radiated}} + P_{\text{thermal}} , \quad (4.1)$$

where  $P_{\text{plume}}$  represents kinetic energy carried away by the ion beam,  $P_{\text{radiated}}$  is electromagnetic emission from the plasma to the environment, and  $P_{\text{thermal}}$  corresponds to power deposited into the thruster structure itself. For modelling purposes, it is often most convenient to begin by describing each of these terms in terms of the contributing plasma species.

##### 4.5.1.1 Ionisation Power

As described in Sec. 2.7.2, ions are produced via inelastic electron-neutral impact ionisation of the propellant. For each ion produced there is a corresponding energy loss equal to the ionisation energy,  $\epsilon_i^{n+}$ . Assuming quasi-neutrality of the plasma, the discharge current equals the ion production rate, such that the ionisation power flux can be written as

$$P_{\text{ion}} = \sum_{n=1}^N \left[ e\epsilon_i^{n+} \frac{I^{n+}}{en} \right] , \quad (2.41)$$

where  $P_{\text{ion}}$  is the power consumed in ionising the propellant,  $n$  is the  $n^{\text{th}}$  ionisation state, and  $I^{n+}$  is the corresponding  $n^{\text{th}}$  ion species current. In this model only the first

ionisation state is considered. For each propellant used, the first and second ionisation energies are listed in Tab. 2.1.

#### 4.5.1.2 Plume Kinetic Power

The ion kinetic power in the plume was introduced in Sec. 2.3.3 and defined in Eq. 2.20 as  $P_{plume}$ . For completeness, it is repeated here with chapter consistent notation

$$P_{i_{KE}} = \frac{1}{2} m_i v_i^2. \quad (2.20)$$

Whilst the ion kinetic power described by Eq. 2.20 is the only contribution considered in Sec. 2.3.3, as it directly produces thrust, in addition there is electron kinetic power carried away in the plume which must be accounted for in a thermal model.

For the electron kinetic power, only the neutralising cathode electrons are considered since electrons inside the discharge channel are assumed to be sufficiently confined by the applied magnetic field. These neutralising electrons have a velocity set by the cathode-to-plasma potential drop. Their number equals that of the ions in the beam to maintain quasi-neutrality, giving

$$P_{e_{KE}} = e \Delta \frac{I_b}{e}, \quad (4.2)$$

where  $\Delta$  is the cathode potential relative to the plasma and  $I_b$  is the beam current, this potential is normally on the order of tens of volts [136].

The total plume kinetic power,  $P_{p_{KE}}$ , is therefore the sum of ion and electron contributions yielding

$$P_{p_{KE}} = P_{i_{KE}} + P_{e_{KE}}. \quad (4.3)$$

#### 4.5.1.3 Plume Thermal Power

The plume thermal power represents the enthalpy flux associated with the random thermal motion of particles in the exhaust. It consists of contributions from the hot, slow electrons released at the cathode and the residual thermal energy of the ions, described as

$$P_{p_{thermal}} = \frac{5}{2} e T_{ip} \frac{I_b}{e} + \frac{5}{2} e T_{ep} \frac{I_b}{e}, \quad (4.4)$$

where  $T_{ip}$  and  $T_{ep}$  are the ion and electron temperatures in electron volts, respectively, and  $I_b/e$  gives the equivalent particle flux in the plume. The factor of  $\frac{5}{2}$  arises from including both kinetic and pressure terms in the thermal enthalpy.




FIGURE 4.15: Approximate xenon ion thermal temperature prior to acceleration, estimated from LIF measurements by Hargus et al. using Eq. 4.5 [139].

Following Azziz [140], typical values of  $T_{ip} = 1$  eV and  $T_{ep} = 3$  eV are assumed. These values were measured on a relatively high-power Hall thruster (BHT-1500) operating on xenon. Owing to their small magnitude relative to the total discharge power, the variation of  $T_{ip}$  and  $T_{ep}$  across propellants and thrusters is neglected, and constant values are used in this model. Another measurement of ion temperature prior to acceleration can be seen in Fig. 4.15 where the low magnitude of the measured ion temperature can be seen, broadly matching the estimate used here of 1 eV. This result comes from converting the LIF velocity measurements made to a temperature in eV by

$$T_{ip} = v_i^2 \frac{\pi m_i}{8e}, \quad (4.5)$$

where  $v_i$  is the LIF measured velocity.

#### 4.5.1.4 Wall Power

One of the largest sources of loss, and a major lifetime-limiting factor in Hall thrusters, is the flux of electrons and high energy ions deposited onto the channel walls. In this thermal model it is assumed that all ion current not present in the beam is lost to the walls, such that the total wall power is the sum of ion and electron contributions.

The framework for describing plasma–wall interactions, including the role of secondary electron emission (SEE), was outlined in Sec. 2.8. There, expressions were introduced for current continuity (Eq. 2.50), the scaling of electron wall current with SEE (Eq. 2.51), and the effect of SEE on the sheath potential drop. These results are applied here to quantify the corresponding thermal fluxes.

Ions impinging the wall deliver three components of energy: a thermal component associated with their pre-acceleration temperature, a directed component corresponding to the Bohm velocity at the sheath edge, and the energy gained falling through the sheath potential. Combining these contributions gives

$$P_{iw} = \left[ \frac{5}{2} e T_{iw} + \frac{1}{2} m_i v_B^2 + e \Delta \phi_{\text{sheath-wall}} \right] \frac{I_{iw}}{e} , \quad (4.6)$$

where  $T_{iw}$  is the ion temperature prior to acceleration; as a result  $T_{iw} = T_{ip}$ , taken as  $\approx 1$  eV based on LIF data [139],  $v_B$  is the Bohm velocity,  $\Delta \phi_{\text{sheath-wall}}$  is the sheath potential drop defined in Sec. 2.8, and  $I_{iw}$  is the ion current to the wall. For simplicity, the ion beam current fraction is approximated as  $\eta_{I_d} = I_b / I_d = 0.7$ , giving

$$I_{iw} = (1 - \eta_{I_d}) I_d . \quad (4.7)$$

Electrons are attracted to the walls at a rate sufficient to maintain current continuity. As discussed in Sec. 2.8, the SEE process increases the required electron wall current above the ion wall current. The effective energy of these impacting electrons, however, is limited by the formation of a double sheath as the SEE coefficient approaches its critical value. Following the model of Choueiri [29], the effective wall electron temperature is expressed as

$$T_{ew} = T^* \left( 1 - \exp \left[ - \frac{e \Delta \phi_{\text{sheath-wall}}}{T^*} \right] \right) , \quad (4.8)$$

where  $T^*$  is the critical electron temperature for sheath reversal (Tab. 2.5). Using this effective temperature in the SEE power-law relation (Eq. 2.52) gives a value for the SEE coefficient in the wall current balance.

The resulting electron thermal flux is then

$$P_{ew} = 2e T_{ew} \frac{I_{ew}}{e} , \quad (4.9)$$

where  $I_{ew}$  is obtained from the continuity condition of

$$I_{ew} = \frac{I_{iw}}{1 - \gamma_{SEE}} . \quad (4.10)$$

The total power deposited into the walls is therefore the sum of ion and electron contributions, which yields

$$P_w = P_{iw} + P_{ew} . \quad (4.11)$$

This formulation links the plasma sheath behaviour described in Sec. 2.8 directly to the thermal model developed in this section, providing a self-consistent treatment of wall losses under the influence of secondary electron emission.

#### 4.5.1.5 Radiated and Excited Power

In a Hall thruster discharge, many electron–neutral collisions do not lead to ionisation but still contribute to energy loss through processes such as elastic scattering, excitation, ion–ion collisions, and neutral–ion collisions. A significant portion of this lost energy is radiated in the thermal or visible spectrum. Because the plasma is optically thin, these emissions propagate without reabsorption.

To approximate this effect, it is assumed that each ionisation event is accompanied by a proportional energy loss due to radiation or excitation

$$P_{\text{rad}} = \left[ e \epsilon_i^{1+} \psi \right] \frac{I_d}{e} , \quad (4.12)$$

where  $\psi$  denotes the loss coefficient per ionisation. For xenon, simulations indicate a typical value of  $\psi \approx 2.2$  [178], for simplicity and due to the lack of a suitable alternative the xenon value will function as a surrogate for krypton and argon as well.

#### 4.5.1.6 Anode Heating

Electrons streaming to the anode to sustain the discharge also impose a thermal load. This effect can be modelled analogously to electron flux to the channel walls, but with a distinct “electron–anode temperature”

$$P_{ea} = [2eT_{ea}] \frac{I_d}{e} . \quad (4.13)$$

The electron–anode temperature,  $T_{ea}$ , differs from the bulk plasma electron temperature because of energy losses during transit from the  $\vec{E} \times \vec{B}$  state to the anode. Reported estimates of  $T_{ea}$  vary widely from 1.8–15 eV [179–182], with some studies suggesting values up to 20 eV [138]. For the present analysis,  $T_{ea} \approx 10$  eV is assumed.

#### 4.5.1.7 Total Thermal Flux

Summing the major dissipation terms gives the total system power for a given input power and voltage

$$P_{\text{tot}} = P_{\text{ion}} + P_{p_{\text{KE}}} + P_{p_{\text{thermal}}} + P_{i_w} + P_{e_w} + P_{\text{rad}} + P_{ea} , \quad (4.14)$$

where  $P_{\text{tot}}$  is the sum of each individual power component. Only some components, namely  $P_{i_w}$ ,  $P_{e_w}$ , and  $P_{ea}$ , contribute directly to thruster heating. Radiated power  $P_{\text{rad}}$  incident on the thruster is also considered in simulations.

#### 4.5.1.8 Thermal Model Results

The dimensionless thermal model outlined above can be applied to a range of scenarios, including both SPT and TAL thrusters with different propellants. Results are given in Tab. 4.3, where nominal operating values of the SHARK-600V are used for xenon, krypton, and argon, and SEE coefficients for BN and SS materials are taken from Tab. 2.4.

TABLE 4.3: Comparison of power terms, in watts, in the thermal model for a thruster operating at 1350 W and 600 V, scaled to input values for the SHARK-600V.

	SPT			TAL		
	Xe	Kr	Ar	Xe	Kr	Ar
$P_{\text{ion}}$	27.23	31.50	35.55	27.23	31.50	35.55
$P_{i_{KE}}$	648.62	648.62	648.62	648.62	648.62	648.62
$P_{e_{KE}}$	47.25	47.25	47.25	47.25	47.25	47.25
$P_{i_{\text{thermal, plume}}}$	3.94	3.94	3.94	3.94	3.94	3.94
$P_{e_{\text{thermal, plume}}}$	11.81	11.81	11.81	11.81	11.81	11.81
$P_{iw}$	229.39	218.60	200.84	229.29	219.05	202.32
$P_{ew}$	96.09	95.66	94.70	404.89	394.85	375.67
$P_{ea}$	45.00	45.00	45.00	45.00	45.00	45.00
$P_{\text{rad}}$	59.90	69.30	78.21	59.90	69.30	78.21
$P_{\text{thermal}}$	430.38	428.56	418.75	739.08	728.20	701.20
$P_{\text{tot}}$	1169.22	1171.68	1165.92	1477.93	1471.32	1448.37
$P_{\text{input}}$	1350.00	1350.00	1350.00	1350.00	1350.00	1350.00

Tab. 4.3 quantifies each component and the resulting total power predicted by the model. It is important to note that this is a constant-current model, not a constant mass-flux model: less massive propellants therefore correspond to lower mass flow rates,  $\dot{m}_n \approx m_n I_d / e$ .

This model provides initial estimates of thermal fluxes for use in thruster design, enabling early-stage assessment of thermal management prior to manufacturing. It is not intended to replace higher order 1D or 2D simulations, but rather to serve as a simplified tool. As shown in Tab. 4.3 and Fig. 4.16, many of the values remain constant across thruster type and propellant, which is not physically expected. This highlights the simplicity and limitations of the model. However, since it is not intended to be strongly predictive and instead relies on a large number of simplifying assumptions, such behaviour is acceptable for its intended purpose.

As shown in Tab. 4.3 and Fig. 4.16, the approximate kinetic power suggests a thruster efficiency of  $\approx 48\%$ , which while on the high end for alternative propellants is consistent with high-performing Hall thrusters. However, several limitations should be noted. At an input power of 1350 W, the sum of the power components ( $P_{\text{tot}}$ ) increases when moving from SPT to TAL configurations. This behaviour reflects the




FIGURE 4.16: A illustration of the results from Tab. 4.3 for both SPT and TAL type for xenon, krypton, and argon. The  $P_{thermal}$  column is the sum of the thermal powers, the  $P_{tot}$  column is the sum of all powers from the model.

simplicity of the assumptions applied, since the only modified term between SPT and TAL is the electron current to the wall.

This is particularly interesting because, given the lower SEE of metals compared to ceramics (see Sec. 2.8), one would expect a much lower electron current to the walls in order to maintain current continuity. However, due to the much higher  $T^*$  (see Tab. 2.5) for metallic walls, electron energies reaching the surface are not retarded to the same degree as in ceramic-walled thrusters. This results in greater power deposition despite a lower electron current.

A significant number of simplifying assumptions were made in forming this model, such as mono-ionic plume, the relationship between electron temperature and voltage being assumed propellant-agnostic, and excitation losses treated as propellant-agnostic. Moreover, several assumptions are not adjusted between propellant types or between magnetic-layer (also referred to as SPT) and anode-layer thrusters. In practice, the empirical data underpinning many of these assumptions are available primarily for medium-power xenon SPT discharges. This severely limits the predictive ability of the model and results in relatively small differences between propellants.

This discrepancy is even more pronounced when comparing SPT and TAL thrusters, since empirical constants are derived mainly from 1–5 kW xenon SPT thrusters. Additionally, the SEE effect included in wall flux terms should also influence plasma temperature, as high SEE in SPT thrusters is known to lower electron temperature by

replacing hot primaries with cold secondary electrons. The effect of SEE on plasma electron temperature is neglected here to preserve linearity in the equations used.

In summary, the model provides approximate thermal flux values that are useful for design considerations in SPT and TAL thrusters operating with alternative propellants. Despite its limitations the estimates remain valuable, and in the worst case represent an overestimation of thermal load.

#### 4.5.2 Thermal Design

The results from Tab. 4.3 were used to inform simulations of the steady-state temperature of the thruster. The thruster was designed to operate with a range of channel configurations and at high voltages. As a result, it is expected to experience high thermal loads when operating in a low-performance configuration at elevated voltages. This expectation informed several major design decisions.

Firstly, the electromagnets were positioned far from the discharge channel to minimise thermal load and reduce the risk of damage during testing; the resulting estimated coil temperatures during testing are described in Sec. 5.4.3. Their distance from the channel also introduces thermal lag; in the event of a high thermal load, the thruster could be deactivated and allowed to cool as the thermal energy dissipated through the thruster body. Secondly, the thruster was designed with a large thermal mass to limit damage in the event of a sudden increase in thermal load. A larger thermal mass slows the rate of temperature rise, since any deposited thermal energy must heat a substantial mass of metal. However, this choice resulted in increasing the total thruster mass.

### 4.6 SHARK-600V Design Summary

This chapter has presented an overview of the design process for the SHARK-600V, a modular high-voltage Hall effect thruster capable of operating with multiple propellants and in both magnetic-layer (SPT) and anode-layer (TAL) configurations. The design methodology applied analytical scaling laws, numerical magnetic topology simulations, and experimental validation to produce a novel research platform for investigating alternative propellant performance at elevated voltages.

The thruster design began with the application of the scaling methodology described in Chapter 3, which established a nominal operating point and baseline channel geometry. From these scaling outputs, a modular insert system was developed, enabling systematic variation of channel width and mean diameter through interchangeable inner and outer inserts. This approach allows rapid reconfiguration

between tests, facilitating parametric investigation of channel geometry effects without requiring thruster disassembly.

A novel contribution of this work is the integration of an optional anode-layer configuration within the same thruster platform. Multiple TAL variants were designed with varying anode lengths relative to the magnetic exit plane. This flexibility enables direct comparison between magnetic-layer and anode-layer operation under otherwise like-for-like conditions, addressing a notable gap in the literature.

The magnetic circuit design represented a substantial aspect of the development process due to the importance of magnetic topology to efficient Hall thruster operation. Numerous simulations were performed across multiple parametric sweeps using several modelling methods and software packages. Validation of the magnetic design through experimental field mapping confirmed the simulation results and demonstrated negligible azimuthal variation despite the discrete nature of the outer electromagnets.

The anode design utilised additive manufacturing to create the complex internal structure of the gas distributor, with angled injection holes to promote swirl injection. X-ray CT imaging provided non-destructive verification of the internal geometry and revealed manufacturing irregularities, which informed subsequent design iterations.

A zero-dimensional thermal model was developed to estimate power partition and thermal loads prior to manufacturing. The model distinguishes between magnetic-layer and anode-layer configurations by incorporating differences in secondary electron emission coefficients and their effect on sheath potential drops. The analysis suggested that anode-layer configurations experience higher thermal loads despite lower electron currents to the walls, due to the higher critical electron temperature for metallic walls. These thermal estimates directly informed key design decisions, including the rear placement of electromagnets to increase thermal path length and the incorporation of a substantial thermal mass to slow temperature rise during testing. Subsequent operational experience confirmed the adequacy of this thermal management approach.

Several limitations of the thermal model should be acknowledged. The model relies heavily on empirical constants derived primarily from medium-power xenon magnetic-layer thrusters, and many assumptions are necessarily treated as propellant-agnostic due to limited available data. Consequently, the model is best interpreted as providing conservative upper-bound estimates for design purposes rather than precise predictions.



## Chapter 5

# Southampton Test Campaign

### 5.1 Introduction

This chapter outlines the facilities, supporting hardware, and diagnostics employed in the experimental characterisation of the SHARK-600V Hall effect thruster undertaken at the University of Southampton, together with the resulting performance measurements. As described in Chapter 4, The SHARK-600V features a modular discharge channel architecture, enabling systematic variation of both channel width and mean channel diameter without requiring changes to the magnetic circuit, electrical system, or fluidic supply. This flexibility allows the thruster to be operated across a wide parametric envelope while maintaining controlled and repeatable experimental conditions. From the resulting dataset, a Gaussian process regression model is developed to enable continuous interpolation between tested geometries and comparison with the semi-empirical scaling methodology presented in Chapter 3.

### 5.2 David Fearn Large Vacuum Chamber Facility

All experimental testing at the University of Southampton was conducted in the large vacuum chamber within the David Fearn Electric Propulsion Laboratory, shown in Fig. 5.1. This vacuum chamber has an internal length of 4.5 m and a diameter of 2.0 m, providing sufficient volume for plume expansion during thruster operation.

The pumping system consists of a single Oerlikon Leybold LV140C roughing pump with a pumping capacity of 145 m<sup>3</sup>/h, two Coolpower 140T cryo compressors coupled to two Coolpack 6000H 20 K cold heads, and two MAG W 2200 iP magnetically levitated turbopumps. Together, this system provides a combined nominal xenon pumping speed of approximately 26,000 L/s (21,800 L/s from cryogenic pumping and

4,200 L/s from turbopumping). This pumping configuration was able to achieve an average argon pumping speed of  $\approx 59,000$  L/s across a 0–300 sccm<sub>Ar</sub> cold flow range.

In this configuration the facility achieves a base pressure of  $< 6 \times 10^{-7}$  mbar and an operational pressure of  $< 5 \times 10^{-5}$  mbar with 60 sccm of xenon injected. Chamber pressure is monitored using a Pfeiffer (Scotland, UK) D-35614 Assar Pirani gauge and the cold cathode type gauge located at the upstream end of the facility, approximately 0.8 m from the thruster mounting location.



FIGURE 5.1: The David Fearn Electric Propulsion Laboratory at the University of Southampton where all experimental testing presented in Chapter 5 was undertaken.

A high-level diagram of the electrical, fluidic, and data acquisition set-up used for the Southampton experimental campaign is shown in Fig. 5.2. The power supplies are illustrated in the floating configuration employed for this campaign, which is described in greater detail in Sec. 5.4. The fluidic distribution system regulating gas flow to the anode and cathode is also shown and is discussed further in Sec. 5.5. The thrust balance, including the laser triangulation sensor, is likewise indicated; its design, calibration, and validation are described in Sec. 5.6. Finally, the control and data acquisition connections to the laboratory computer are shown, from which power, flow, thrust measurements, and calibration procedures are managed, which is covered in greater detail in Sec. 5.7.

### 5.3 Cathode

The experimental campaign employed an off the shelf Model 5000 Hollow Cathode Electron Source from Intlvac Thin Film Corporation (Georgetown, Canada), which utilises a tungsten filament for electron emission. The cathode used for this test campaign is shown in Fig. 5.3.

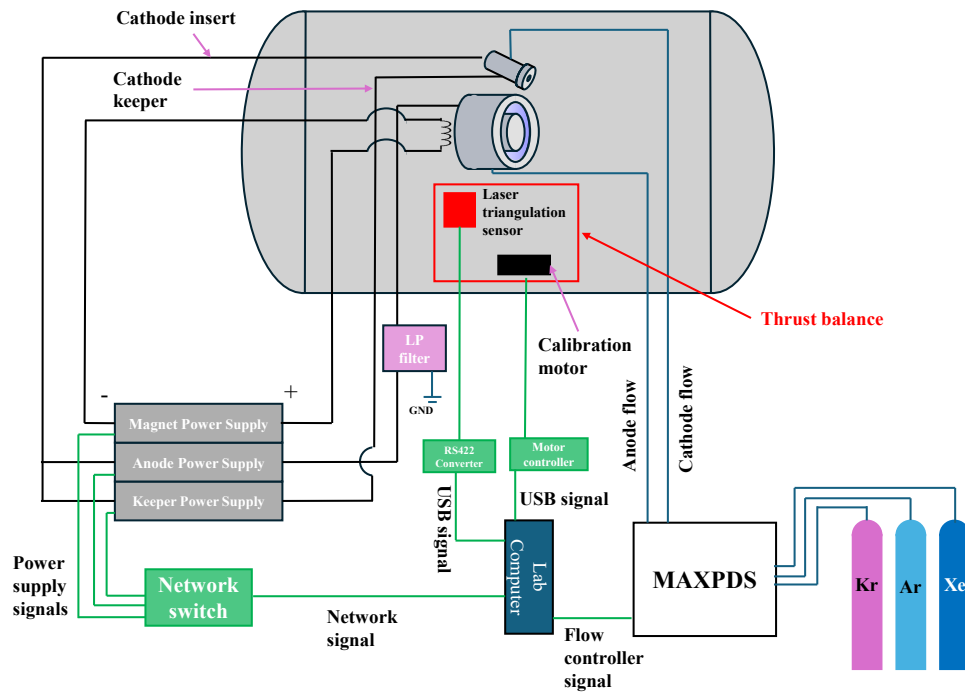


FIGURE 5.2: Illustration of the electrical, fluidic and data connections of the vacuum chamber testing that was used for this campaign.




FIGURE 5.3: The HCES 5000 hollow cathode used during the Southampton testing campaign.

Space-grade hollow cathodes typically operate at comparatively low flow rates relative to the anode, nominally a few percent of the anode flow. However, the operation of the HCES 5000 hollow cathode required a non-negligible propellant flow rate on the order of several sccm. At low anode flow rates, this could introduce the possibility of the thruster ingesting of neutrals and ions from the cathode plume, potentially leading to an artificial enhancement of measured thruster performance.

To evaluate this effect, dedicated tests were previously conducted with a 100 W class thruster and this cathode, where variations in cathode mass flow rate were shown to have a minor impact on overall thruster performance; corresponding to approximately a 2% anode efficiency change across the tested range [58]. As the SHARK-600V is more than an order of magnitude more powerful, with a correspondingly higher anode mass flow rate, it is assumed that cathode neutral flow effects will be negligible. The cathode was operated at a fixed flow rate of 5 sccm of krypton (N5.0 purity) for all tests conducted within this chapter.

Cathode placement can play a significant role in thruster operation. For all tests within Chap. 5, cathode position was kept constant, at 45° orientation relative to the thruster exit plan, cathode orifice approximately 30 mm downstream of the thruster, with the cathode angle down such the injection point is tangential to the near side of the thruster channel. This location was chosen due to early test providing ease of ignition. Furthermore, this configuration allowed for the cathode orifice to be visible from the chamber window such ignition could be visually confirmed.

## 5.4 Electrical Setup

The electrical configuration for thruster testing at the University of Southampton is shown schematically in Fig. 5.2. The negative connections of the anode power supply and cathode keeper power supply were electrically tied to the cathode insert, which was operated in a floating configuration for all tests. This approach was chosen to better replicate in-space conditions and to ensure current continuity between the plasma discharge and the cathode.

The electromagnets were powered on an independent circuit, simplifying the electrical connections and isolating their operation from the plasma discharge loop. The thruster body was left floating but connected to ground through a 100 M $\Omega$  resistor. This allowed the body to reach a nominal floating potential while limiting charge accumulation that could otherwise perturb the plasma or pose a risk to the thruster and diagnostics.

For comparison, additional tests were performed with the thruster body directly grounded, as well as fully floating without the resistor. In both cases, the difference in thruster behaviour was negligible relative to the nominal configuration.

### 5.4.1 Anode Power Supply

The SHARK-600V anode was powered by a LAB-AUTO 1000-30 programmable DC supply (1000 V, 30 A, 10 kW). All tests were conducted in voltage control mode, with

typical discharge voltages in the range 200–600 V and discharge currents up to 15 A. The power supplies positive output connected to the anode was routed through the low-pass filter described in Sec. 5.4.1.1 to protect the supply against plasma-induced transients.

The supply was connected via Local area network (LAN) to the laboratory control computer through a dedicated network switch and operated remotely throughout all tests. This connection also provided data logging of voltage and current with resolutions of 0.1 V and 1 mA at a rate of 20 Hz, which were integrated into the campaign data acquisition system (see Sec. 5.7). Software limits were applied to prevent accidental over-voltage or over-current conditions during testing.

#### 5.4.1.1 Low-Pass Filter Design


To protect the anode power supply from electromagnetic transients during thruster ignition, a low-pass filter was integrated into the supply line. Without this protection, the supply's over-current safeguard was frequently triggered during ignition as the plasma was established during ignition, requiring a time-consuming manual restart of the remotely operated system. The filter could be bypassed by selecting alternative shielded-banana plug sockets; in the optional unfiltered bypass path, high-current diodes are still present to block reverse current into the supply terminals, so protection was retained even when the filter circuit was not.

The chosen design for the circuit is shown in Fig. 5.4a, with input connections from the power supply (left) and the output terminals (right). A dedicated ground point is also included, which should remain connected even when the negative terminals are floated, as it acts as the sink for rejected high-frequency noise and can optionally serve as the negative terminal ground.

The manufactured filter box, used throughout the SHARK-600V test campaign, is shown in Fig. 5.4b. The final design incorporated *Vishay* 1200 V, 80 A rectifier diodes, a *KEMET C4AQ* polypropylene capacitor rated at 1.5 kV DC and 8  $\mu$ F, and a 600 W, wire-wound, 100  $\Omega$  resistor.

#### 5.4.2 Cathode Power Supply

Cathode operation was supported by a LAB-DSP 600-005 programmable DC source (600 V, 5 A, 3 kW). The supply was used to provide the high ignition voltage required by the hollow cathode as well as steady-state keeper current during operation. For ignition, the supply was operated in voltage control mode at up to 500 V. Once discharge current was established, the supply would switch to current control mode to maintain a 2.5 A keeper current. Under these conditions, the keeper voltage was



A. Circuit diagram of the low-pass filter installed between the anode power supply and the thruster.



B. Photograph of the assembled low-pass filter box.

FIGURE 5.4: Low-pass filter implementation showing (a) the circuit diagram and (b) the assembled hardware.

observed to be 30–50 V without an active plasma and 15–25 V once plasma coupling was established with the thruster. The higher keeper voltages would be seen after several hours of operation.

The cathode supply was connected and integrated into the control system using the same LAN-based method as the anode supply, enabling full remote operation. Voltage, current, and protection status were logged continuously during testing as part of the data acquisition system described in Sec. 5.7.

### 5.4.3 Electromagnet Power Supply

The thruster electromagnets were powered using a Kikusui PWX750MHF programmable DC supply (230 V, 10 A, 750 W). The supply was integrated into the facility control system via a LAN connection, enabling remote operation and continuous monitoring during testing.

Although the electromagnets are ideally driven in current control mode to maintain a constant field strength, operation in this mode led to significant oscillations in the delivered current. For stability, the supply was therefore operated in voltage control mode. In practice, the operator incrementally increased the voltage set point throughout testing to compensate for coil resistance growth caused by heating. A nominal current of 2 A was maintained in this manner, with the required supply voltage rising from approximately 26 V at start-up to as high as 52 V under sustained thermal loading. If the required voltage exceeded 52 V, testing was suspended until the coils had cooled sufficiently to return the operating point to within the acceptable range. Assuming the coils were at approximately 23°C initially and using the thermal

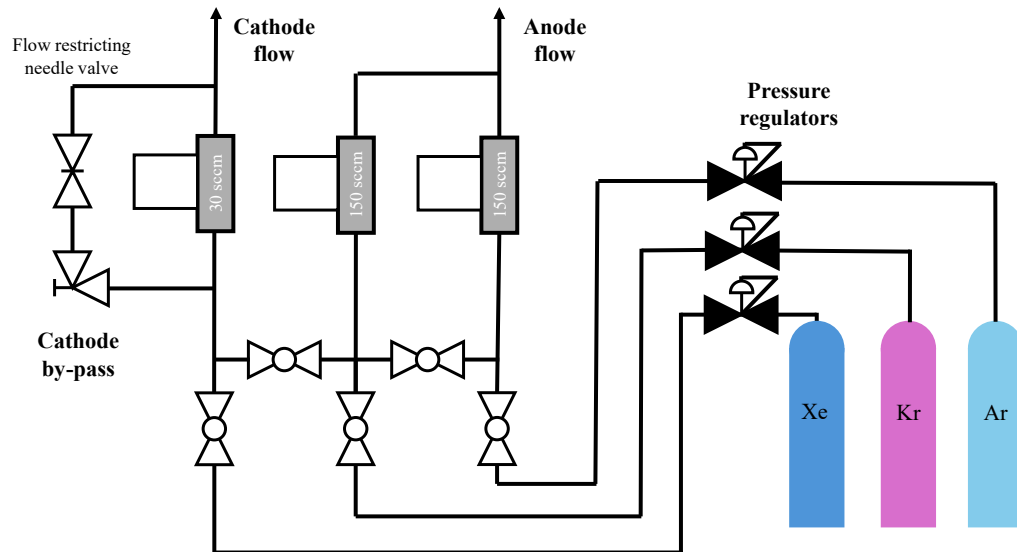


FIGURE 5.5: Piping and Instrumentation Diagram for the multi and mixed propellant delivery system.

coefficient of copper of  $3.9 \times 10^{-3} \text{ } ^\circ\text{C}^{-1}$ , this increase in resistance corresponds to an average coil temperature of approximately  $258^\circ\text{C}$ . This is below the maximum rated temperature of  $300^\circ\text{C}$ , if by a small margin.

This voltage rise is unsurprising given that the test durations extended over several hours with anode discharge powers up to 4 kW. Furthermore, the coils featured a high turn count, making them more sensitive to resistance increases under thermal loading. The choice of a high turn count, as described in Sec. 4.3, was made to achieve the required magnetic field strength at relatively low electromagnet currents compatible with the available power supplies.

## 5.5 Fluidic Setup

To support testing of the SHARK-600V on multiple propellants and at the flow rates required, a dedicated propellant delivery board was developed: the **Multi And miXed Propellant Delivery System (MAXPDS)**. The system was designed to allow rapid switching between gases, independent cathode and anode feeds, and a controlled cathode ignition bypass. The Piping and Instrumentation Diagram (P&ID) of the setup is shown in Fig. 5.5.

The MAXPDS employed three Bronkhorst EL-FLOW Select mass flow controllers (MFCs): one dedicated to the cathode line and two for the anode. The cathode was fitted with an F-201CV-050-RAD-22-V controller with a flow range of 0.6–30 sccm Xe and an accuracy of  $\pm 0.5\%$  of the reading  $\pm 0.1\%$  of the max flow. The anode lines were each fitted with F-201CV-200-RAD-22-V controllers with a wider flow range of



FIGURE 5.6: Photograph of the MAXPDS front control board as implemented for the SHARK-600V test campaign.

3–150 sccm Xe and the same accuracy. All controllers were stainless steel (SS 316L) with Viton seals, and were pre-programmed with Xe, Kr, Ar, Ne, N<sub>2</sub>, and air. The FLOW-BUS which connected to the individual flow controllers itself was operated remotely via a Python script, which continuously logged the mass flow rates during testing, as described further in Sec. 5.7.

During the campaign, the system was operated with xenon, krypton, argon, neon (mixed with xenon), as well as N<sub>2</sub> all at N5.0 purity. Feed pressure was regulated to 1–2 bar from the storage cylinders via the pressure regulators. This aided in rapid switching between propellants since only the small dead volume in the feed lines (1–5 minutes of flow at most) needed to be exhausted, making changing between gases straightforward and fast.

The cathode line also incorporated a dedicated bypass to aid ignition, see Fig. 5.5. Instead of the manually actuated needle valve used in earlier systems, the new design fixed the needle valve partially open and added a manual button valve for short pulses of gas. In practice, only one brief “burp” was typically required to achieve ignition. This improvement, along with the integrated MFC control, provided a higher level of repeatability compared with the previous single-line manual setup, which had required both manual valve actuation and manual flow logging.

The final system provided reliable multi-propellant delivery and ignition support and greatly improved both the efficiency and accuracy of testing compared with previous approaches.

## 5.6 Thrust Balance Development<sup>1</sup>

This section is adapted from “Inter-Laboratory Characterisation of a Low-Power Channel-Less Hall-Effect Thruster: Performance Comparisons and Lessons Learnt” Munro-O’Brien et al. [183], where the thrust balance developed for this thesis was validated in a cross-laboratory test campaign. The thrust balance was also presented previously in “Validation of an adjustable sensitivity inverted pendulum thrust stand with a novel external plasma thruster” Munro-O’Brien et al. [184], where its adjustable sensitivity and application to an external plasma thruster were first introduced.

Thrust measurement is a key diagnostic for evaluating Hall thruster performance, as both specific impulse and anode efficiency can be derived from a measured thrust together with known input parameters such as anode mass flux and discharge power (see Eq. 2.10 and Eq. 2.22). Direct thrust measurements thus provide a non-invasive and reliable means of assessing overall propulsion performance.

Accurate thrust data are particularly critical for electric propulsion systems, which often operate at the milli-Newton level where strain gauges and force transducers are unsuitable. In addition, thrusters impose non-negligible thermal loads on diagnostics, making robust measurement techniques essential. The most widely used ground-based method involves measuring the displacement of a suspended platform or spring-mounted arm under thrust loading [185]. While such balances are conceptually straightforward, they are sensitive to thermal drift due to the temperature dependence of spring constants, which complicates calibration.

Pendulum-type thrust balances are especially well suited for Hall thruster applications due to their sensitivity and relative simplicity. They can operate as static devices for time-averaged measurements or, in dynamic mode, to characterise pulsed operation. This flexibility, together with straightforward static calibration, makes pendulum balances the most practical and widely adopted approach for experimental Hall thruster testing.

Alongside thrust, controlled parameters such as anode mass flux and discharge power were recorded, enabling the calculation of key performance metrics including specific impulse and efficiency.

---

<sup>1</sup>Adapted from “Munro-O’Brien, T.F., Ahmed, M., Lucca Fabris, A. and Ryan, C.N., 2025. Inter-Laboratory Characterisation of a Low-Power Channel-Less Hall-Effect Thruster: Performance Comparisons and Lessons Learnt. *Aerospace*, 12(7), p.601.”, Munro-O’Brien et al. [183]

## 5.6.1 Thrust Balance Variants

There are several forms these thrust balances can take, each with its own advantages and disadvantages. Here, a non-exhaustive selection is described to provide additional context on thrust measurement techniques.

### 5.6.1.1 Hanging Pendulum

Hanging pendulum thrust balances consist of a suspended arm with the thruster mounted at the bottom. Their mechanically simple, self-stabilising design requires minimal components and no active control since gravity restores the system to equilibrium. These balances are relatively robust to thermal drift, as the restoring force is dominated by gravity rather than the thermally sensitive flexures. Sensitivity can be limited compared to other designs. However, nano-newton accuracy hanging pendulum balances have been tested to success for low mass thrusters [186].

### 5.6.1.2 Inverted Pendulum

Inverted pendulum balances are used when higher sensitivity is required or when thruster mass or envelope is large. The thruster mass counteracts the flexure stiffness, significantly increasing responsiveness. Sensitivity can be tuned with differing stiffness' of the restoring springs, but this makes the design more susceptible to thermal drift. To mitigate this, such systems often employ thermal shielding or active cooling. An example of a high-performance null-displacing inverted pendulum stand can be found in Xu and Walker [187].

### 5.6.1.3 Torsional Pendulum

Torsional balances rotate about an axis parallel to gravity, making their sensitivity largely independent of thruster mass. This characteristic makes them suitable for testing thrusters across a wide range of sizes. Through careful design, torsional balances can achieve very high resolution, down to the hundreds of nano-Newton range [188, 189]. They have also demonstrated good accuracy for pulsed thrust measurements [190].

## 5.6.2 Double Inverted Pendulum Thrust Balance

The thrust balance previously used in the David Fearn Electric Propulsion Laboratory was a hanging pendulum design. This balance was employed for testing a range of

sub-kilowatt thrusters [12, 58, 135, 191]. While effective for low-power testing, the balance was ill-suited for heavier thrusters. Furthermore, its sensitivity was relatively low ( $2.6082 \mu\text{m}/\text{mN}$ , Tirila [135]). These limitations motivated the development of a more robust and tuneable solution capable of accommodating a wider range of thruster masses and operating conditions.




FIGURE 5.7: Illustration of the double inverted pendulum thrust stand with counterbalance [183].

A new thrust balance was therefore designed and manufactured to extend the laboratory's testing capability to higher-voltage and higher-mass thrusters. To achieve this, a double inverted pendulum with a counterbalance configuration was selected. The operational principle is illustrated in Fig. 5.7, which provides compactness, tuneable sensitivity, and flexibility for future upgrades. Incorporating aspects of both hanging pendulum and inverted pendulum designs, this configuration allows the thruster to be mounted on top without requiring a superstructure to enclose it, thereby mitigating issues associated with plasma formation within the balance structure.

The thrust balance design drew inspiration from several existing stands. In particular, the thrust stand developed at Boğaziçi University in Turkey by Kokal, Saridede, and Celik [192], itself based on the thrust stand originally designed at the Massachusetts Institute of Technology (MIT) in the United States by Tartler [193], provided a foundation for the concept. Although the pivoting mechanisms differ between these stands, the underlying mechanics of the double inverted pendulum remain consistent. However, while the cited designs were optimised for low-power, low-mass thrusters, the balance developed here represents a significantly scaled-up version of the concept. It incorporates alternative flexure designs, displacement sensing methods, and passive damping systems to accommodate the higher power levels and larger thruster masses targeted in this work.

The double inverted pendulum thrust balance was chosen for several key reasons:

- Compact geometry removes restrictions on thruster size and layout.
- Sensitivity can be tuned via counterweights and adjustable flexure lengths.
- Straightforward upgrade path to a null-displacement configuration.
- Top-mounted thruster reduces plasma–balance interactions observed in earlier tests.

The completed design (Fig. 5.9) incorporates cable management, cathode mounting, and diagnostic access. The frame dimensions (570 mm × 450 mm × 665 mm) and height (715 mm to thruster mount, 1460 mm overall).




FIGURE 5.8: Design views of the inverted pendulum thrust balance.

The thrust balance was successfully manufactured and tested in the University of Southampton's large vacuum chamber facility, as shown in Fig. 5.9.




FIGURE 5.9: The manufactured inverted pendulum thrust balance installed in the Southampton large vacuum chamber, with the SHARK-600V mounted post-test.

### 5.6.3 Thrust Balance Calibration

Calibration of the University of Southampton's thrust balance was performed prior to each test to quantify sensitivity, force–displacement response, and systematic errors. Connecting a thruster to the balance inevitably introduces external stiffness contributions from propellant lines, power cables, and diagnostic connections, all of which must be accounted for when relating measured displacement to thrust. Furthermore, during operation thermal loads from the plasma discharge to the thruster, the balance and the electrical and fluidic lines can cause drift effects. The thermal expansion of cables and fluidic lines especially introduce phantom forces. These can be accounted for by repeating calibration several times throughout testing.

The calibration procedure followed the methodology of Polk et al. [185]. A known calibration mass,  $M_{\text{cal}}$ , was suspended from the balance on a plumb line of length  $P_y$ , of fine wire (with a diameter of approximately  $70 \mu\text{m}$ ). This mass is also attached via another wire to a Thorlabs' (New Jersey, USA) MTS50(/M)-Z8 motorised translation stage. To apply a calibration force that is parallel to the thrust axis, the mass is displaced tens of millimetres by the motion stage such that the horizontal component of the tension in the line applies a force to the thrust stand. By displacing the mass through a small distance  $\delta$ , a restoring tension is applied to the balance corresponding

to a calibration force, such that

$$F_{\text{cal}} = \frac{\delta M_{\text{cal}} g_0}{\sqrt{P_y^2 - \delta^2}} , \quad (5.1)$$

where  $g_0$  is the gravitational acceleration. The resulting displacement of the balance was measured using a Micro-Epsilon ILD1750-10 laser triangulation sensor (resolution  $0.1 \mu\text{m}$ , sampling frequency 5 kHz). A Butterworth filter was applied to the raw signal, and steady-state levels before and after the applied force were averaged to determine the net displacement. An example calibration trace is shown in Fig. 5.10. The same calibration setup and mass was used for all tests presented within this thesis.




FIGURE 5.10: Laser triangulation sensor output for a standard calibration, showing the raw signal, filtered response, and steady-state levels.

The force–displacement relationship was then obtained via least-squares regression, giving

$$\hat{\delta}_i = S_{\text{cal}} F_i + b_{\text{cal}} , \quad (5.2)$$

where  $S_{\text{cal}}$  is the calibration coefficient (m/N) and  $b_{\text{cal}}$  is the offset (m).

For  $N_{\text{cal}}$  calibration points, the variance of the residuals was estimated as

$$s_{\delta}^2 = \frac{\sum_i^{N_{\text{cal}}} (\delta_i - b_{\text{cal}} - S_{\text{cal}} F_i)^2}{N_{\text{cal}} - 2} . \quad (5.3)$$

This gives an estimate of the balance variance,  $\sigma_{\delta}^2 \approx s_{\delta}^2$ , which can be expressed in force units as

$$\sigma_F \approx \frac{s_{\delta}}{S_{\text{cal}}} . \quad (5.4)$$

The uncertainty in the calibration coefficient is given by

$$s_{\text{cal}} = \frac{s_{\delta}}{\sqrt{\sum_i^{N_{\text{cal}}} (F_i - \bar{F}_{\text{cal}})^2}} , \quad (5.5)$$

where  $\bar{F}_{\text{cal}}$  is the mean applied force. For sufficiently large  $N_{\text{cal}}$ ,  $s_{\text{cal}} \approx \sigma_{\text{cal}}$ .

Finally, the calibration error is defined as

$$\text{cal}_{\text{error}} = \frac{\sigma_{S_{\text{cal}}}}{S_{\text{cal}}} , \quad (5.6)$$

and the total fractional uncertainty in a thrust measurement  $T$  is

$$\frac{\sigma_F}{T} = \sqrt{\left(\frac{\sigma_{\delta}}{\delta}\right)^2 + \left(\frac{\sigma_{S_{\text{cal}}}}{S_{\text{cal}}}\right)^2} . \quad (5.7)$$

For each test, a new in-situ calibration was performed, and the resulting coefficients were used to calculate thrust and associated uncertainties. These values are reflected in the error bars presented throughout this work.

## 5.6.4 Calibration Results

The results of the calibration, including the fitted coefficient and calibration offset, are shown in Fig. 5.11. For each applied force, several repeat measurements (nominally five) were taken. These repeats are difficult to distinguish in Fig. 5.11 due to the small variance between them. The corresponding statistical variance and its contribution to the overall thrust measurement error are summarised in Tab. 5.1. As shown, the calibration achieved a high level of precision. For each test campaign, a new in-situ calibration was performed, and the results of that calibration were directly propagated into the error bars of the performance data presented in this thesis.

## 5.6.5 Cross-Laboratory Validation of Thrust Balance<sup>2</sup>

### 5.6.5.1 Introduction

Before commencing SHARK-600V testing, the newly developed double inverted pendulum thrust balance at the University of Southampton was validated against an

<sup>2</sup>Adapted from "Inter-laboratory characterisation of a low power channel-less Hall effect thruster: performance comparisons and lessons learnt", Munro-O'Brien et al. [183]. The original publication lists contributions as follows: conceptualisation, methodology, validation, formal analysis, investigation, data curation, visualisation, and writing by Thomas F. Munro-O'Brien and Mohamed Ahmed; resources, supervision, and project administration by Andrea Lucca Fabris and Charles N. Ryan; and funding acquisition by Mohamed Ahmed.




FIGURE 5.11: Example output of the calibration process for the thrust balance used in these tests. The corresponding statistical results and error analysis are summarised in Tab. 5.1.

TABLE 5.1: A table of the results from one of the calibration runs of the thrust balance used for these testing. Individual calibrations are run for each test in-situ.

Variable	Value	Unit
Average force	43.611	mN
Calibration coefficient	7.0613	$\mu\text{m} / \text{mN}$
Calibration Offset	-6.8217	$\mu\text{m}$
No. of calibration steps	36	#
$\sigma_x$	5.1234	( $\pm$ ) $\mu\text{m}$
$\sigma_F$	725.5699	( $\pm$ ) $\mu\text{N}$
$\sigma_{S_{\text{cal}}}$	0.0304	( $\pm$ ) $\mu\text{m} / \text{mN}$
Calibration error	0.4304	%
Error on a 10 mN measurement	8.043	%
Error on a 50 mN measurement	1.541	%
Error on a 100 mN measurement	0.850	%

independent system. Since no previously characterised thrusters were available locally, collaboration was established with the University of Surrey. Their torsional thrust balance had been used to characterise a novel channel-less Hall effect thruster, termed the eXternal Plasma Thruster (XPT), providing an opportunity for direct comparison with the Southampton balance. The details of this joint campaign are reported by Munro-O'Brien et al. [183].

### 5.6.5.2 External Plasma Thruster

The XPT developed at the University of Surrey is an external plasma thruster that operates on the same principles as a conventional HET but without a discharge channel. The magnetic circuit, consisting of two hollow cylindrical samarium–cobalt (SmCo) permanent magnets, is positioned behind the anode. The anode extends beyond the thruster body to allow direct propellant injection and improved thermal management. A side-by-side comparison of the XPT operating in both facilities is shown in Fig. 5.12.




FIGURE 5.12: Operation of the XPT on xenon in the David Fearn Electric Propulsion Laboratory at Southampton (left) and in the Daedalus facility at Surrey (right).

### 5.6.5.3 Experimental Results

To compare the two thrust balances, xenon-fed XPT performance was measured at both facilities over a range of anode mass flow rates and discharge powers. Fig. 5.13 summarise the results. At low flow rates of 5 to 6 sccm, Southampton consistently measured higher thrust. For example, at 5 sccm xenon and 250 V, Surrey reported 2 mN at 102 W, while Southampton recorded 3 mN at 121 W.

At higher xenon flow rates of 8 to 10 sccm, this trend did not persist. Despite higher anode powers at Southampton, measured thrust was lower, particularly at 10 sccm. At lower anode voltages of 100 and 150 V, results from both facilities overlapped. As the voltage rose to 200 and 250 V, discrepancies of up to 25 % emerged. These differences correlated with variations in magnetic field strength and vacuum conditions specific to each facility.

### 5.6.5.4 Conclusion

The experimental campaign demonstrated strong agreement between facilities in the low-power regime when characterising the miniaturised external plasma thruster. In




FIGURE 5.13: Thrust measured for the xenon-fed XPT at different anode mass flow rates in the University of Surrey and University of Southampton facilities.




FIGURE 5.14: Anode efficiency of the XPT at varying xenon flow rates, measured at Surrey and Southampton.

particular, for the 5 sccm case at both facilities, the results showed close agreement, with the Southampton measurements lying within the error bars of the Surrey data.

At higher powers, however, notable disparities in thrust were observed. These differences likely arose from a combination of thermal effects on magnetic components and the influence of background pressure on thruster behaviour. The divergence in performance between the two facilities became increasingly pronounced at elevated anode mass fluxes and power levels.

Overall, the cross-laboratory tests demonstrated that, while absolute performance values can diverge due to facility-dependent effects, the Southampton thrust balance produced consistent and repeatable measurements. The improved accuracy of this balance, even at low thrust levels, relative to the torsional balance at the University of Surrey is evident in Fig. 5.14, where the Southampton data exhibit smaller error bars. Throughout testing, the uncertainty in thrust measurements at Southampton ranged from 0.07 mN to 0.23 mN, whereas at Surrey the range was 0.26 mN to 0.46 mN. This benchmarking exercise validated the accuracy and reliability of the Southampton balance, providing confidence in its application to the SHARK-600V campaign. Full details of the facility-dependent effects are presented in [183].

## 5.7 Data Acquisition Set-up

The data acquisition (DAQ) system used for these tests represented a major improvement over the previous approach of manually recording each data point, which was prone to human error and rounding inaccuracies. Automated logging enabled continuous recording of voltage and current from all power supplies at a frequency far beyond human capability, facilitating the construction of meaningful error bars for the measured values.

To implement this DAQ system, the power supplies and flow controllers were all remotely operated via a Python-based control script. The script employed multi-threading to handle each instrument as an independent process, ensuring reliable and synchronous communication. Each power supply described in Secs. 5.4.1–5.4.3 was connected to the laboratory control computer through a network switch and addressed using a unique Internet Protocol (IP) address. The flow controllers were operated via the Bronkhorst FLOW-BUS interface, which was connected to the DAQ system through a USB-A 3.0 link. A high-level connection map of the instruments used in this test campaign is shown in Fig. 5.15.

The voltage levels reported by the power supplies were validated with manual multi-meter measurements and found to be accurate; the current values as reported by the power supplies were therefore adopted without independent verification. The flow controllers were calibrated by Bronkhorst prior to their integration into the set-up, and no independent verification was undertaken.

### 5.7.1 Power Supply Communication

Instrument communication for the power supplies was achieved using the Standard Commands for Programmable Instruments (SCPI) protocol, an ASCII-based standard

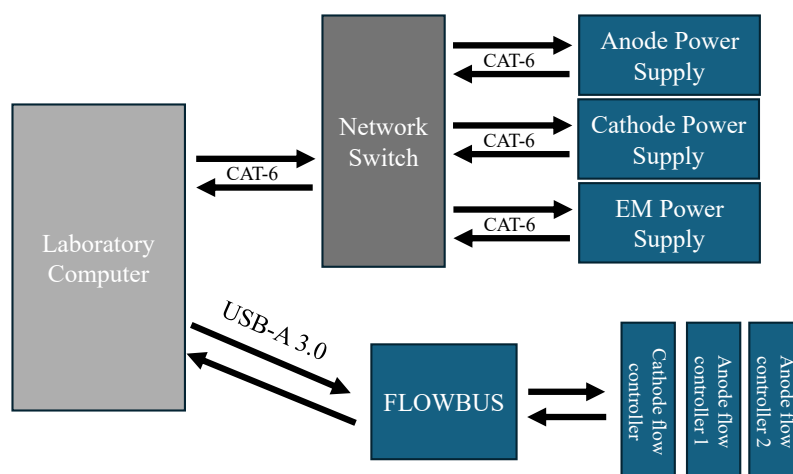


FIGURE 5.15: High-level connection map showing how instruments were linked to the laboratory computer during operation of the DAQ system.

for sending instructions and receiving responses from laboratory equipment. Although SCPI provides a general command framework, each supply used in this work required device-specific syntax, necessitating the development of bespoke Python functions for each instrument. These were implemented as class objects within the Python environment, allowing each power supply to be controlled and queried individually during operation. An example showing only the initialisation section of the LAB-AUTO power supply driver class is given in List. 5.1. While each supply type requires its own bespoke structure, all operate on the same principle and share very similar implementations.

```

1 import pyvisa
2 import threading
3 import time
4
5 class LABAUTO_PowerSupply:
6     def __init__(self, ip_address, name):
7         self.ip_address = ip_address
8         self.name = name
9         self.resource_string = f'TCPIP::{self.ip_address}::5025::SOCKET'
10        self.rm = pyvisa.ResourceManager()
11        self.power_supply = None
12        self.connected = False
13        self.lock = threading.Lock()
14        self.set_voltage_value = 0
15        self.set_current_value = 0
  
```

LISTING 5.1: Initialisation of the LAB-AUTO power supply driver (VISA socket, thread-safety, and cached setpoints).

An example of the query of the output values for the LAB-AUTO 1000-30 power supply used for logging can be seen in List. 5.2. The ASCII request in this case is the MEAS:VOLT?;\*WAI portion, where the ;\*WAI portion instructs the supply to wait for the response before executing the next command. This ensures that requests do not overlap, preventing the response to a voltage query from being misattributed to the current query.

```

1 def query_output(self):
2     with self.lock:
3         if not self.connected:
4             self.connect()
5         try:
6             output_voltage = self.power_supply.query('MEAS:VOLT?;*WAI')
7             time.sleep(0.005)
8             output_current = self.power_supply.query('MEAS:CURR?;*WAI')
9             return float(output_voltage), float(output_current)
10        except Exception as e:
11            print(f"Error querying output: {e}")
12            return None, None

```

LISTING 5.2: Python method used to query output voltage and current from the power supply.

### 5.7.2 Flow Controller Communication

As the flow controllers do not use SCPI as the communication language, their communication method differs from that of the power supplies. Commands are instead issued by writing to numerical parameters within the FLOW-BUS protocol. For example, setting parameter=9 updates the flow set-point of the controller. The Python method used to write values to a specified parameter is shown in List. 5.3.

```

1 def write_value(self, parameter, value):
2     try:
3         self.instrument.writeParameter(parameter, value)
4         print(f"Successfully wrote value {value} to parameter {parameter} "
5             f"on node {self.node_address}")
6     except Exception as e:
7         print(f"Error writing value {value} to parameter {parameter} "
8             f"on node {self.node_address}: {e}")

```

LISTING 5.3: Method for writing a value to a specified parameter on the FLOW-BUS instrument.

### 5.7.3 Multi-threading Structure

To enable simultaneous operation of each instrument without building up large command queues, multi-threading was employed so that each device could be




FIGURE 5.16: Thread-level connection map used for the DAQ system, illustrating parallelisation and the separate communication protocols for each instrument.

accessed independently. A schematic of the thread-level connections between the laboratory computer and the instruments is shown in Fig. 5.16.

#### 5.7.4 Graphical User Interface

Once all instruments were connected and capable of being queried and commanded, they were integrated into the control panel developed for the test campaign. The control panel, implemented using the `tkinter` package (the standard Python interface to the Tcl/Tk graphical user interface toolkit), is shown in Fig. 5.17.

#### 5.7.5 Data logger

The GUI included a “Start Logging All” function, which, when activated, queried each instrument at 20 Hz and stored the results in comma-separated values (CSV) files, with one file generated per thread. This logger again is its own class which is handed the other class objects “devices” and logs them separately into a file.

#### 5.7.6 Thrust Balance Measurements

In addition to logging power supply values and flow controller outputs, the displacement of the thrust balance was continuously monitored. The primary measurement was provided by a Micro-Epsilon ILD1750-10 laser triangulation sensor, a high-speed digital sensor operating at 5 kHz. The laser displacement sensor was mounted to the underside of the thruster mounting plate and measured the

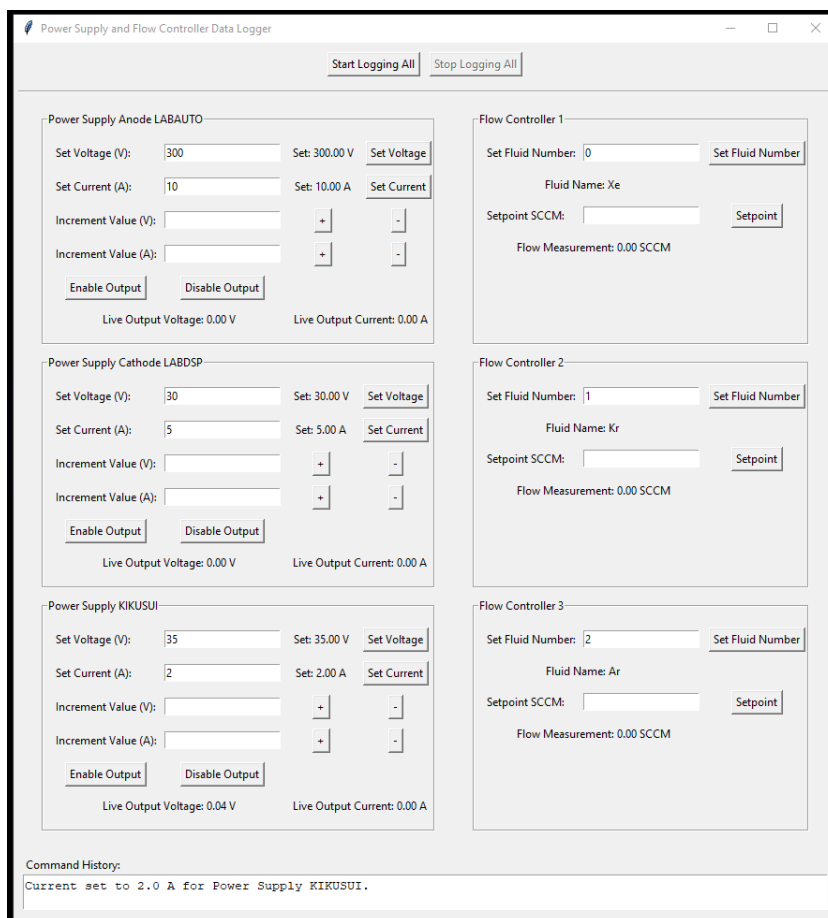


FIGURE 5.17: Screenshot of the data acquisition front panel showing all connected power supplies and flow controllers.

displacement between itself and the bottom plate of the thrust balance. This configuration provides  $2\times$  mechanical amplification, which enhances the sensitivity of the balance. The thrust balance is mounted within the chamber on rubber feet to isolate it from the low-frequency noise generated by the cryo-panel cold heads. However, the signal from the laser required several filters to make it usable for thrust measurements.

During each test, the laser displacement data were logged continuously to a CSV file. To manage file sizes during long-duration tests, a 50-point moving mean was applied in real time. This down-sampling yielded an effective temporal resolution of 100 Hz, while still capturing the thrust balance dynamics with sufficient fidelity. This resolution was adequate to capture the step-function response of a thrust measurement, as illustrated in Fig. 5.18.




FIGURE 5.18: Example thrust measurement from the analysis routine, showing laser displacement, anode voltage, anode current, and anode mass flow rate. The red vertical line marks the on–off transition, shaded regions represent the averaging windows, and the horizontal dot-dashed lines indicate the two laser displacement averages used to calculate the thrust-induced deflection.

## 5.8 Data Analysis

Data collection was designed to be continuous; consequently, careful attention to data readback was required to ensure accuracy and reliability in the performance measurements. To perform a thrust measurement, the thruster was brought to the operating point of interest and held until the discharge current stabilised. The stabilisation time varied between operating points; therefore, to ensure consistency, the thruster was held at each operating point for a minimum of 30 s before a measurement was taken. Once equilibrium was reached, the anode power was shut off, producing a step response in the thrust balance laser sensor. This displacement was converted to thrust using the calibration coefficient determined as described in Sec. 5.6.3. The balance was calibrated before testing and periodically throughout to ensure that no significant drift occurred that could affect the measurements.

After each test, the CSV logs from the power supplies, flow controllers, and laser sensor were imported into pandas DataFrames. The calibration file was also loaded, and the thruster configuration was input manually. All datasets were standardised by converting the timestamps to a uniform datetime format.

The datasets were then merged using the laser data as the master timeline, aligning rows on the nearest timestamp. The power supply and flow controller logging rate was 20 Hz, whereas the laser was sampled at 100 Hz. This disparity was resolved

within the merged dataset by duplicating rows from the lower-resolution files to align with the nearest timestamps in the laser signal. To keep signals distinct, all non-timestamp columns were prefixed with their source. The result was a single merged DataFrame containing synchronised measurements from all instruments, ready for subsequent analysis.

Using this new master test dataset, to identify when the anode discharge was terminated, the high-to-low transitions were identified in the merged dataset. Specifically, the anode current column was analysed to find points where the current was high ( $> 0.5$  A) and transition to low ( $< 0.5$  A), with a moving average filter (window size of 20 points) to smooth noise. The corresponding timestamps of these events were returned as a list, marking the points at which the thruster was switched off.

Subsequently, the array of event timestamps is passed to a dedicated analysis function, which processes the thrust balance data and evaluates performance metrics before and after each transition. For each detected event, mean and standard deviation values are calculated for the laser displacement, the anode, cathode, and electromagnet voltages and currents, as well as both anode and cathode mass flow rates, within defined pre- and post-transition windows. The regions immediately before and after the transition are excluded, as they correspond to transient behaviour and are not representative of steady-state operation. Only the laser displacement values from the post-transition window are used to calculate the deflection between on- and off-states; the other signals require only their steady-state on-values for analysis. An example of this analysis method applied to a single thrust measurement, with the regions of interest highlighted, is shown in Fig. 5.18.

Due to the lower temporal resolution of the power supply signal, the thrust measurement appears to pre-date the power supply switching off; this is merely an artifact of the lower resolution and is corrected by the exclusion window.

Propellant flow rates and corresponding propellant names are extracted from the flow controller logs. Using a conversion dictionary, the values are converted from sccm to mass flow in mg/s. This anode mass flow, together with the thrust measurement, is then used to calculate specific impulse and, efficiency. Calibration constants and errors from Sec. 5.6.3 are applied to quantify the thrust magnitude and its associated uncertainty. The final results are compiled into a DataFrame, where each row corresponds to a detected transition and includes all computed parameters (such as thrust, specific impulse, efficiencies, and propellant flow and gas) as well as the aggregated instrument measurements and the derived laser displacement.

## 5.9 Experimental Results




FIGURE 5.19: A side by side comparison of the thruster operating in the large vacuum chamber at the University of Southampton on xenon (top left), krypton (top right), argon (bottom left), and neon-xenon mix (bottom right).

### 5.9.1 Introduction

This section presents and discusses the experimental results of the SHARK-600V Hall effect thruster operated in nine discrete channel configurations in magnetic layer mode, as well as in an optional anode-layer configuration. Each magnetic-layer configuration was fully characterised on xenon, krypton, and argon without any modification to the test set-up. The thruster was successfully operated up to 600 V on xenon and krypton, with a maximum voltage of 500 V achieved on argon. To the author's knowledge, this represents the first instance of a thruster being characterised on xenon, krypton and, argon without alteration to thruster setup or configuration. Furthermore, the range of channel configurations investigated provides an additional insight into how discharge channel dimensions influence Hall effect thruster operation with alternative propellants.

The thruster has also been nominally operated on a neon-xenon mixture as well as dinitrogen. However, the performance was not recorded for these tests as it was primarily a demonstration of the flexibility of Hall effect thruster operation.

### 5.9.1.1 How to Read Grid Plots

Due to the modular nature of the thruster and the testing on three propellants, a large quantity of data is generated that can be difficult to compare when overlaid. To address this, grids of subplots are used to represent discrete thruster configurations. Each subplot corresponds to a specific set of channel dimensions or configuration names. These configurations are denoted as “+0”, “+1”, or “+2” and correspond to the inner and outer channel inserts listed in Tab. 5.2. In-text, these will always be referred to in the format (“inner, outer”).

The grid arrangement of the plots mirrors the structure of Tab. 5.2, enabling direct visual comparison between data and thruster geometry. In Tab. 5.2, the diagonal from the top left to the bottom right represents the constant mean diameter case for increasing channel width. This also holds for the minor diagonals: from the top middle to the left middle, and from the right middle to the bottom middle, each representing a constant channel width case for increasing mean diameter. Furthermore, moving from top to bottom within each column increases both channel width and mean channel diameter. An example of such a grid is shown in Fig. 5.20, where the different channel configurations are illustrated. The dimensions plotted correspond to those listed in Tab. 5.2, and the figure is displayed at a reduced scale of 1:3 relative to the actual geometry for clarity (assuming a standard A4 page).

TABLE 5.2: Channel width (mm) and mean channel diameter (mm) for each configuration of the modular thruster. Each cell shows “h mm, d mm : “+ #, + #” for that inner and outer configuration.

Outer	Inner		
	SPT +0	SPT +1	SPT +2
SPT +0	08.0, 70.0 : “+0, +0”	10.0, 68.0 : “+1, +0”	12.5, 65.5 : “+2, +0”
SPT +1	10.0, 72.0 : “+0, +1”	12.0, 70.0 : “+1, +1”	14.5, 67.5 : “+2, +1”
SPT +2	12.5, 74.5 : “+0, +2”	14.5, 72.5 : “+1, +2”	17.0, 70.0 : “+2, +2”

### 5.9.2 Global Performance Trends

The thruster was operated across a range of discharge voltages and anode flow rates for each propellant. For xenon, the anode flow rate was varied between 20–60 sccm over discharge voltages of 200–600 V. For krypton, the anode flow rate was varied from 40–90 sccm across the same voltage range of 200–600 V. Due to difficulties achieving reliable ignition below 90 sccm, argon testing was restricted to higher flow rates: for one channel configuration the range was 90–150 sccm, while all other configurations were tested between 110–150 sccm. In the case of argon, the discharge voltage range was limited to 200–500 V because instability or discharge powers reaching the 4.2 kW test limit prevented operation at higher voltages.

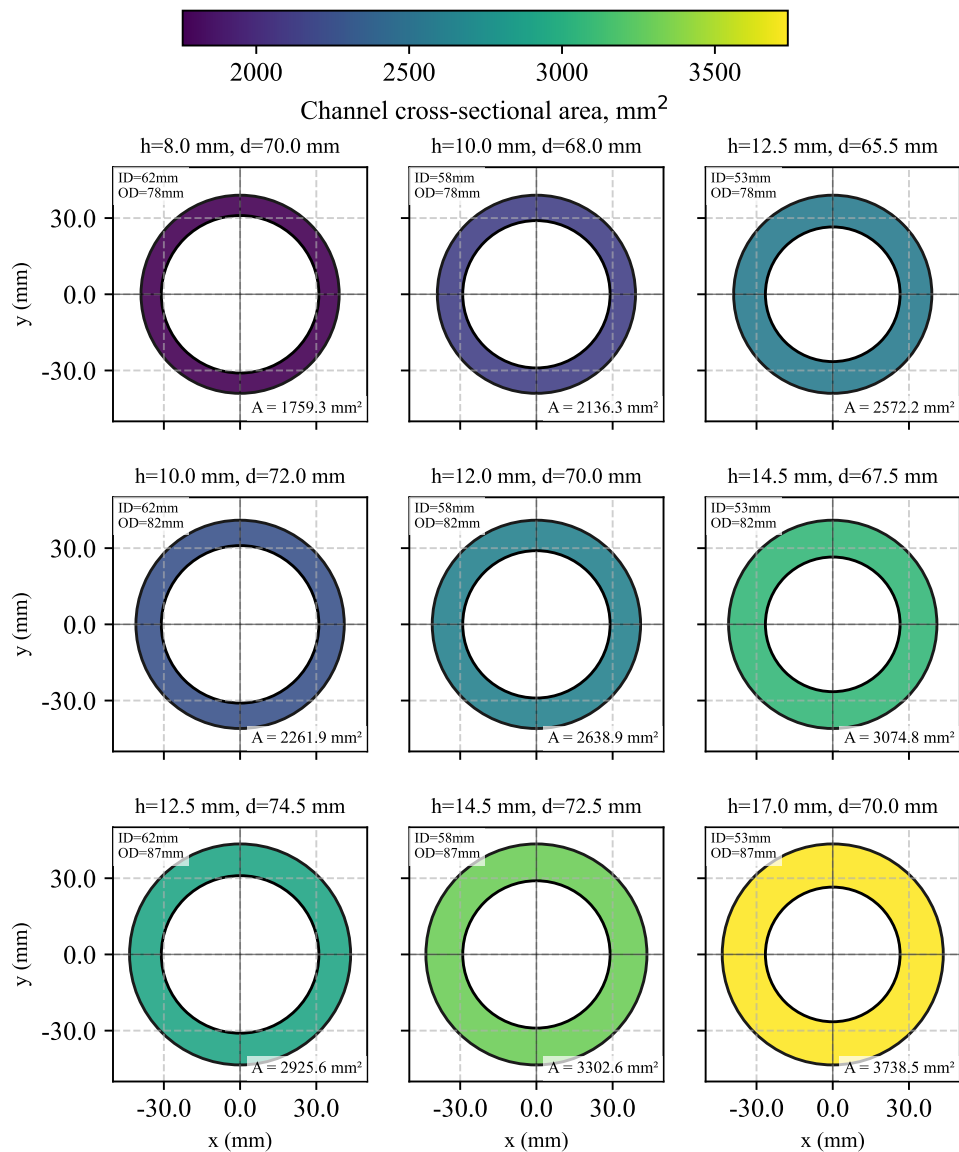


FIGURE 5.20: Illustration of the modular thruster channel configurations arranged in a grid format. Each subplot corresponds to the inner and outer channel dimensions listed in Tab. 5.2. The figure is displayed at a reduced scale of 1:3 relative to the actual geometry for clarity (assuming displayed on a standard A4 page).




FIGURE 5.21: Measured thrust across all SPT configurations and propellants. Both anode power and thrust error bars are shown. Anode power uncertainties are derived from the standard deviation of voltage and current over the averaging range, while thrust uncertainties are calculated from the calibration procedure described in Sec. 5.6.3.

The operational envelope for testing was constrained by several factors. Firstly, at low mass flow rates, the envelope was limited by the ability to ignite and sustain a plasma. Secondly, operation at high voltages was also limited by plasma stability. Thirdly, the maximum input power was capped at 4.2 kW due to concerns over excessive thermal loading. This power limit was rarely reached except during argon operation, where higher number densities (and therefore higher discharge currents) were required to ignite and sustain the discharge. Finally, at certain high-voltage conditions, experienced in several channel configurations, excessive asymmetric heating of the discharge channel was observed. This phenomenon will be discussed in greater detail in Sec. 5.9.6. When such heating was detected, testing was suspended until the thruster channel was no longer emitting in the visible spectrum. Although this heating did not appear to significantly alter performance, its root cause remains unknown.

Top-level performance values for xenon are as follows: The maximum thrust was 115.25 mN ( $\pm 0.59$  mN), recorded at 400 V, 60 sccm, and 3.7 kW in the “+2, +1” channel configuration. The peak specific impulse was 2422.8 s ( $\pm 18.3$  s), achieved at 600 V, 30 sccm, and 2.7 kW in the same “+2, +1” configuration. The maximum anode efficiency reached was 34.71 % ( $\pm 0.63$  %), observed at 600 V, 30 sccm, and 2.7 kW in the “+2, +0” channel configuration.

Top-level performance values for krypton are as follows: The maximum thrust was 109.4 mN ( $\pm 0.58$  mN), recorded at 400 V, 80 sccm, and 3.7 kW in the “+2, +1” channel

configuration. The peak specific impulse was 2809.1 s ( $\pm 17.8$  s), achieved at 600 V, 50 sccm, and 3.3 kW in the same “+2, +1” configuration. The maximum anode efficiency reached was 35.8 % ( $\pm 0.46$  %), observed at the same condition and configuration as the peak specific impulse.

Top-level performance values for argon are as follows: The maximum thrust was 81.75 mN ( $\pm 0.54$  mN), recorded at 300 V, 130 sccm, and 3.7 kW in the “+2, +1” channel configuration. The peak specific impulse was 2537.2 s ( $\pm 24.7$  s), achieved at 500 V, 90 sccm, and 3.3 kW in the “+1, +2” configuration. The maximum anode efficiency reached was 24.8 % ( $\pm 0.31$  %), observed at 400 V, 110 sccm, and 3.7 kW in the “+2, +0” channel configuration.

The thrust versus power results for all tested configurations, grouped by anode propellant, are shown in Fig. 5.21. This figure illustrates the overall performance envelope across all propellants, as well as the imposed 4.2 kW power ceiling. While this plot provides a useful overview, it does not highlight the modular nature of the thruster, where each channel configuration exhibits distinct performance trends and discharge characteristics. To illustrate these differences, Fig. 5.22 presents thrust versus anode power for each channel configuration as separate subplots. The arrangement of these subplots is described in the thruster geometry grid in Fig. 5.20 and described in Sec. 5.9.1.1. It is clear from Fig. 5.22 that different channel geometries yield distinct thrust and power ranges for each propellant. From Fig. 5.22 it can be seen how some configurations, “+2, +0” and “+2, +1”, achieve higher thrusts and could be operated to higher powers.

To highlight the small amount of uncertainty and error within the thrust and power measurement the error bars can be seen added to Figs. 5.21–5.22. It is hard to distinguish as the error bars are smaller than the markers used in most cases. As a result of this low error all other plots within this section error bars are omitted for the sake of readability.

As seen in Fig. 5.23, the overall efficiency of the thruster in the large vacuum facility at the University of Southampton was relatively low. Without access to plasma diagnostics, it is difficult to identify the source of the specific mechanism of the efficiency loss, though facility-dependent factors are likely. Nevertheless, comparison between the three propellants and nine channel geometries remains meaningful and forms the focus of the discussion in this section.

One of the key objectives of this thesis was to enhance alternative propellant performance at high discharge voltages and to investigate the relationship between channel geometries and overall thruster performance. Illustrated in Fig. 5.23, the variation of thrust-to-power ratio against anode specific impulse is shown for each propellant and discharge voltage (propellants share colour, and voltages share markers). Whilst both derived values, from measured thrust and delivery power due

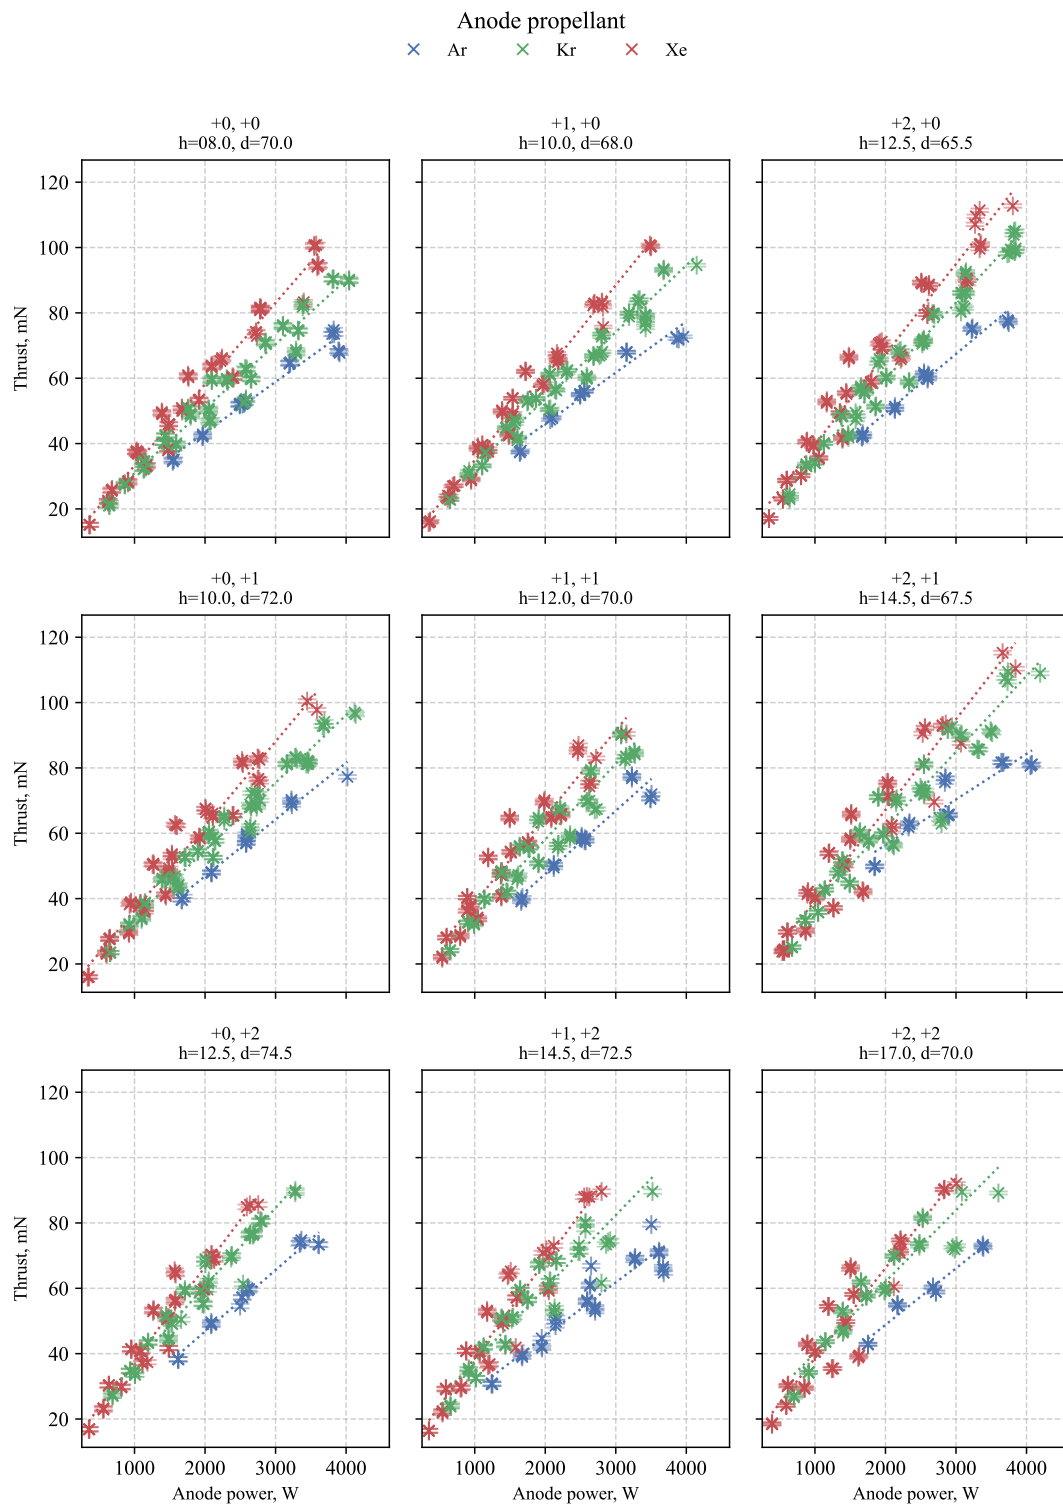


FIGURE 5.22: Thrust versus anode power grouped by anode propellant (xenon, krypton, and argon) for each channel configuration. Each subplot corresponds to one thruster geometry, arranged as in Tab. 5.2.

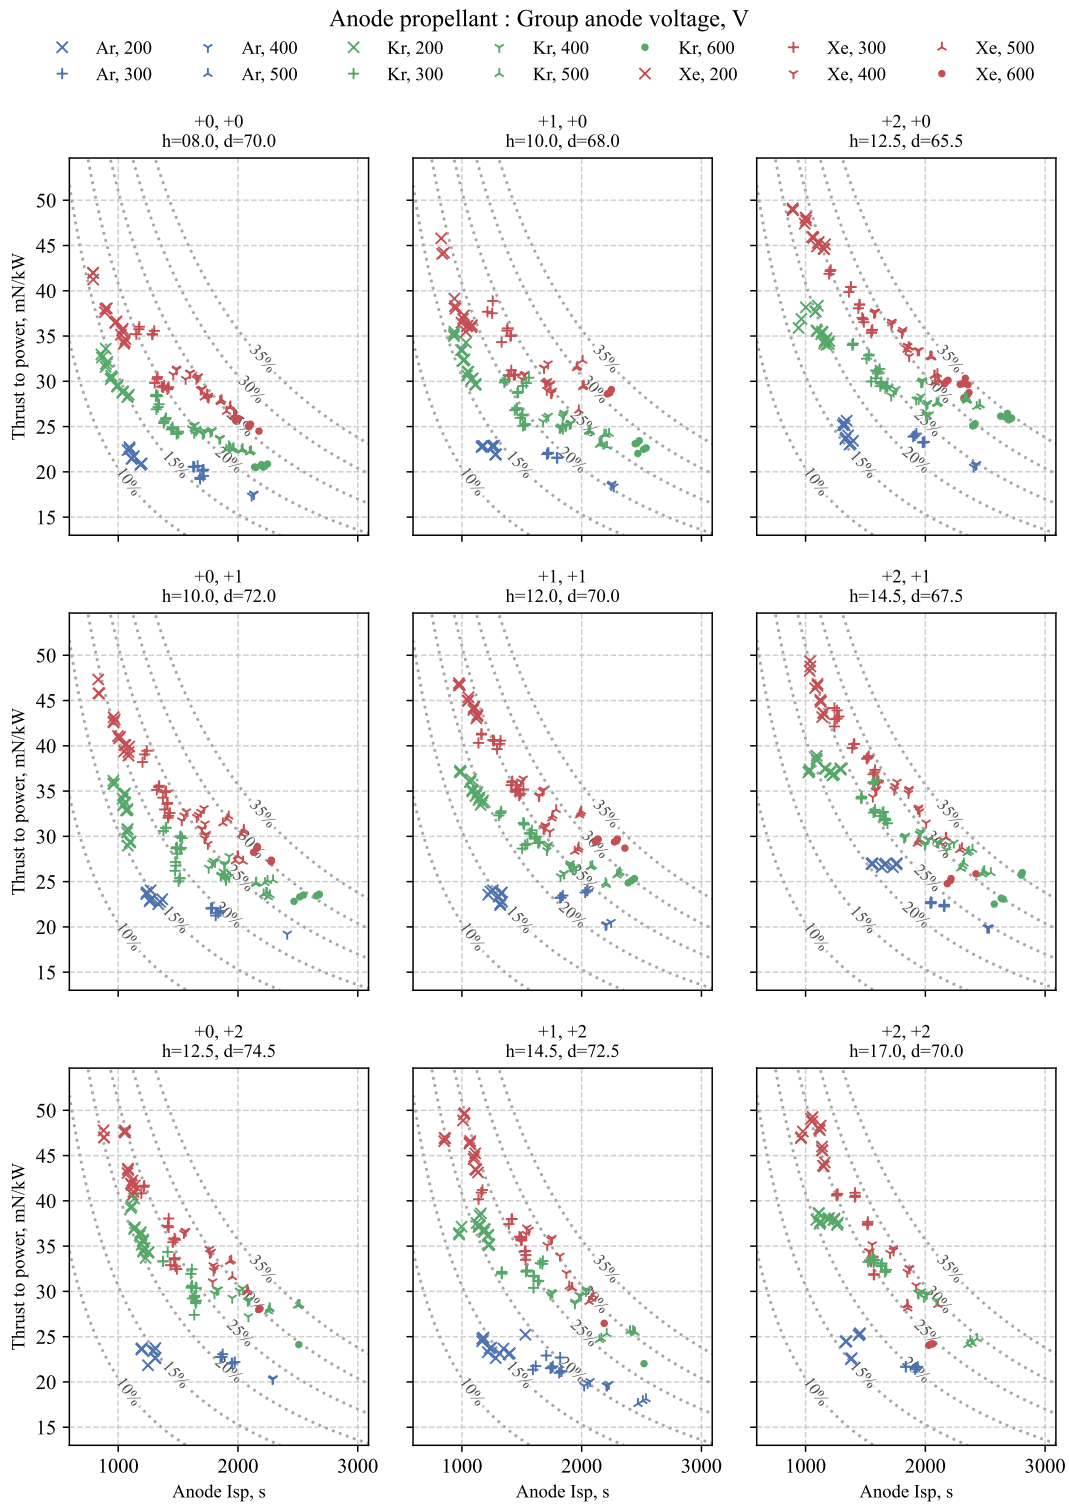


FIGURE 5.23: thrust-to-power versus anode specific impulse with anode efficiency iso-lines, grouped by propellant (xenon, krypton, argon) and anode voltage, for each channel configuration tested.

to the relationship between thrust-to-power, specific impulse and anode efficiency, Eq. 1.1, it is possible to draw lines of constant efficiency on Fig. 5.23. These factors make this type of performance plot very insightful to understand performance and operational envelope.

As can be seen from Fig. 5.23, it is evident that at higher discharge voltages, and for certain channel geometries, krypton performance can be significantly enhanced, in some cases even outperforming xenon under the same conditions. In particular, the “+2, +0” and “+2, +1” channel configurations demonstrated the best krypton performance and are therefore of particular interest.




FIGURE 5.24: Anode specific impulse versus anode power for the “+0, +0” and “+2, +1” channel configuration, grouped by propellant and discharge voltage.

Increasing discharge voltage was also found to increase the specific impulse across all channel configurations. This behaviour is expected, as higher discharge voltages correspond to larger ion acceleration potentials in Eq. 2.18.

It is also noteworthy that argon was able to achieve specific impulses comparable to xenon and krypton at lower discharge voltages. This effect is most evident in the “+0, +0” configuration, where argon achieved specific impulse values at 400 V equivalent to those of krypton and xenon operating at 600 V. However, while the “+0, +0” configuration yielded the lowest overall performance, higher efficiency configurations did not exhibit the same degree of enhancement for argon as observed for krypton. Nevertheless, argon was still able to provide comparable specific impulse at approximately 100 V lower discharge voltage in other configurations.

To better illustrate the effect of channel geometry, the “+2, +1” and “+0, +0” channel configurations are highlighted in Fig. 5.24, where specific impulse is shown against anode power for each propellant and discharge voltage. The “+2, +1” configuration corresponds to a mean channel diameter of 67.5 mm with a channel width of 14.5 mm,

while the “+0, +0” configuration corresponds to a mean channel diameter of 70 mm and a channel width of 8 mm. Comparing these cases therefore isolates the impact of increasing channel width and reducing mean diameter under otherwise identical voltage and flow conditions. This geometric change represents a 78 % increase in channel cross-sectional area, which has a direct influence on the observed performance trends for xenon, krypton, and argon.

Through comparing the two configurations in Fig. 5.24, it can be seen that for argon the “+2, +1” configuration exhibits higher anode power and specific impulse than the “+0, +0” case under the same voltage and flow conditions. This indicates that a larger fraction of the propellant is being ionised, leading to higher discharge current and greater thrust production. The increase in power draw is on the order of 200–300 W, while the corresponding rise in specific impulse is approximately 400–500 s. This improvement is also reflected in anode efficiency which can be seen in Fig. 5.23 but these two configurations are isolated in Fig. 5.25, where the anode efficiency for argon increases by roughly 10 % between the two configurations.




FIGURE 5.25: thrust-to-power ratio versus anode specific impulse for the “+0, +0” and “+2, +1” channel configuration, grouped by propellant and discharge voltage.

For krypton, the power draw at lower voltages (200 V–400 V) remains largely unchanged between the two configurations, while the specific impulse increases for “+2, +1”. This suggests that, for the same ionisation rate, ions in the “+2, +1” channel geometry are either accelerated more axially or accelerated to higher velocities. Unfortunately, it is not possible to identify the specific source of the increased specific impulse without plume diagnostics. Furthermore, at higher voltages (400–600 V), both power and specific impulse increase, indicating enhanced ionisation in addition to improved acceleration efficiency. Xenon follows a similar trend to krypton, though the performance differences between the two channel geometries are less pronounced.

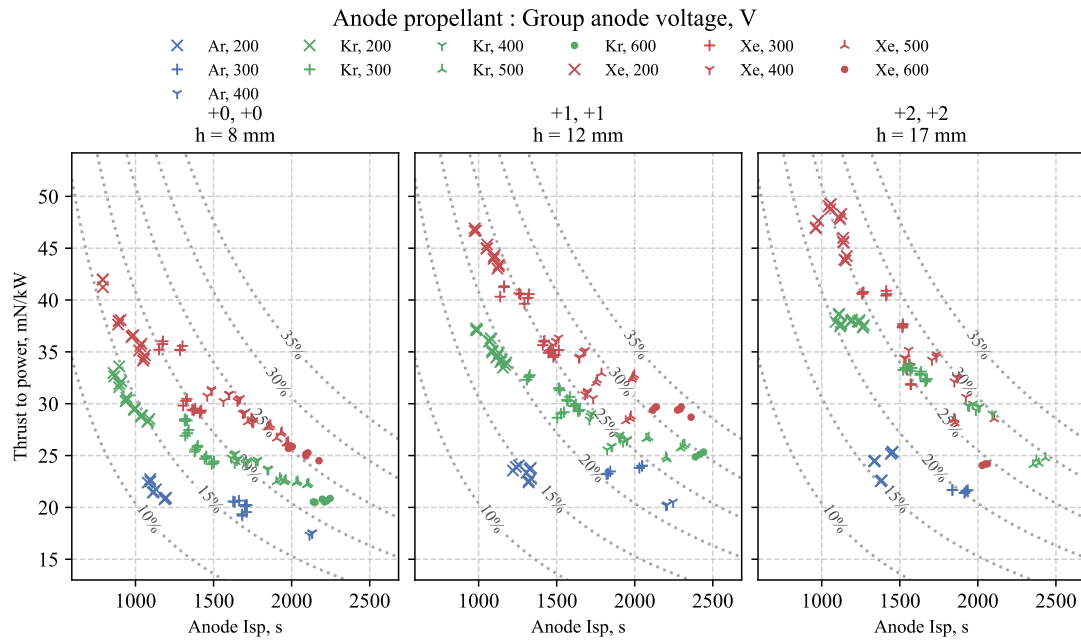


FIGURE 5.26: thrust-to-power against anode specific impulse for each tested propellant and anode voltage in the configurations of constant mean diameter of 70 mm (“+0, +0”, “+1, +1”, and “+2, +2”). Here iso-lines have been applied to represent lines of constant anode efficiency.

Interestingly, in both configurations xenon and krypton exhibit very similar behaviour, with krypton providing a higher specific impulse at elevated voltages. One possible explanation is that at high voltages xenon, owing to its lower ionisation energy and larger ionisation cross-section, produces a greater proportion of multiply charged ions, which depreciates performance. Moreover, as shown in Fig. 5.23, krypton achieves equivalent or even higher anode efficiencies than xenon at greater specific impulses in the “+0, +2”, “+2, +0”, “+2, +1” and “+2, +2” configurations. These observations are consistently associated with high anode voltages in the range of 500–600 V.

Across the three configurations of constant mean diameter (“+0, +0”, “+1, +1”, “+2, +2”) xenon achieved a maximum thrust of 100.8 mN, a peak specific impulse of 2360 s, and a maximum anode efficiency of 33.5 %. Krypton produced a maximum thrust of 90.7 mN, a peak specific impulse of 2440 s, and an efficiency of 30.3 %, in some cases matching or exceeding xenon performance at elevated voltages. Argon, while constrained by higher ignition thresholds and lower efficiency, reached a maximum thrust of 77.6 mN, a peak specific impulse of 2245 s, and a maximum efficiency of 24.1 %. These results highlight the potential of krypton as a competitive alternative to xenon, particularly at high discharge voltages, while also demonstrating that argon can deliver comparable specific impulse at reduced voltages despite lower overall efficiency.

The effect of channel width was evident across all propellants. The narrowest configuration “+0, +0” (channel width of 8 mm) produced the lowest performance but

allowed argon to achieve specific impulse values at 400 V comparable to xenon and krypton at 600 V. The intermediate “+1, +1” (channel width of 12 mm) configuration represented the peak, of this subset, thrust and efficiency for krypton and xenon, while the widest “+2, +2” configuration (channel width of 17 mm) whilst still enhanced performance compared to the “+0, +0” configuration shows a reduced operational envelope and lower peak anode efficiency for xenon and krypton. However, the “+2, +2” configuration showed a high general increased krypton performance across its range, with the reduced envelope limiting the peak achievable. In general, xenon and krypton exhibited similar behaviour, though krypton benefited more strongly from increased channel width, while argon showed less sensitivity to geometry but remained competitive in specific impulse at reduced voltages.

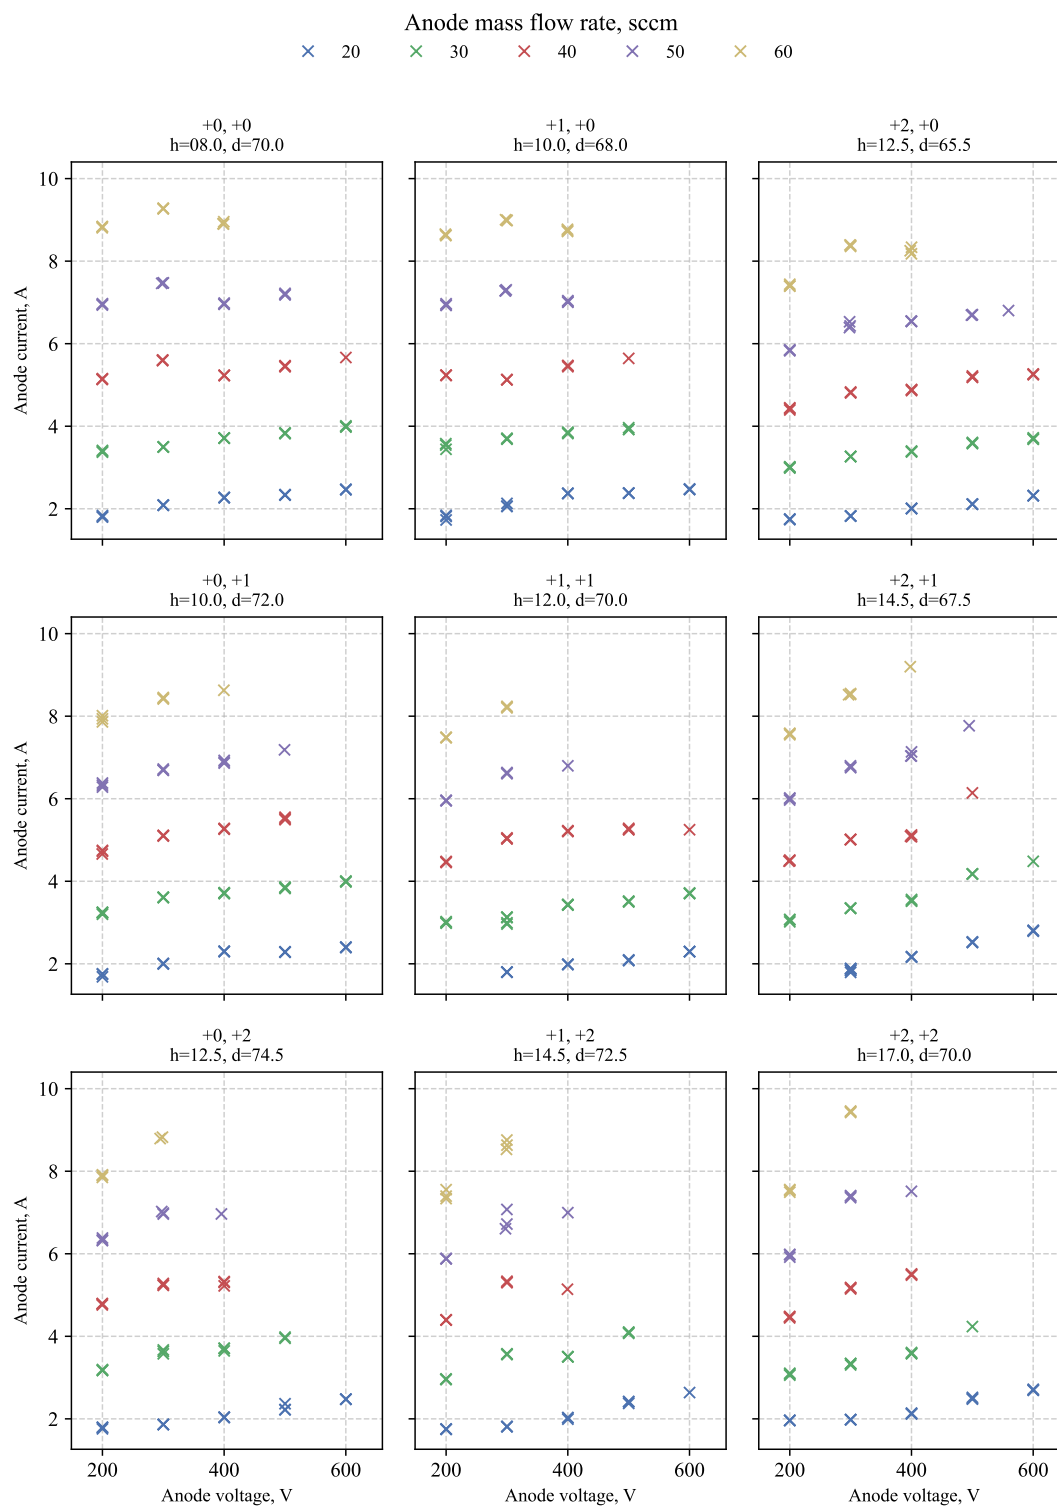


FIGURE 5.27: Anode current draw against anode discharge voltage for xenon, grouped by anode mass flow rate.

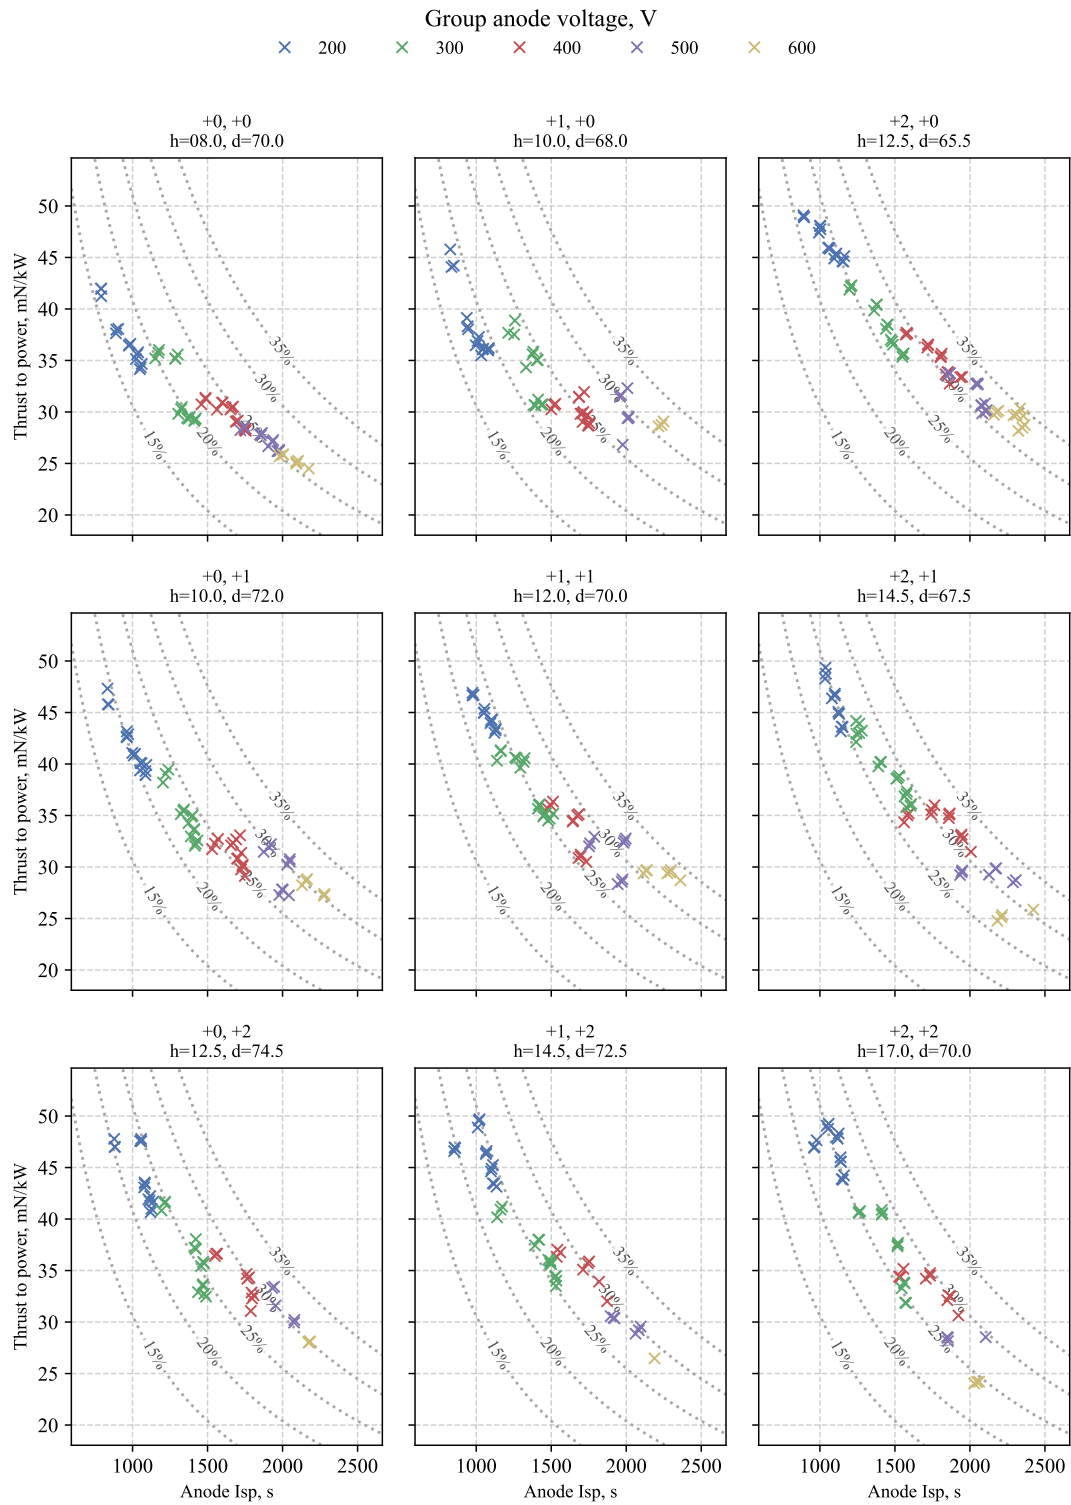


FIGURE 5.28: thrust-to-power ratio against specific impulse for xenon, grouped by discharge voltage.

### 5.9.3 Xenon Performance Trends




FIGURE 5.29: The SHARK-600V operating on xenon in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory.

Xenon operation within the SHARK-600V thruster provided a broad operating envelope, with stable discharges sustained even at relatively low anode mass flow rates. Ignition on xenon was achieved with relative ease. The I–V characteristics for each configuration are shown in Fig. 5.27. At higher discharge voltages, the operating envelope narrows for configurations that combine increased channel width with larger mean diameter (lower left of the grid), and these cases exhibit steep rises in anode current with voltage, particularly at high flow rates. By contrast, configurations with smaller mean diameter (upper right of the grid) maintain a wider high-voltage envelope.

Although the “+0, +0” configuration spans a wide I–V range, its performance is comparatively poor: it delivers the lowest specific impulse and efficiency at each voltage across the power range (see Fig. 5.28). In general, performance improves with increasing channel width and decreasing mean diameter, culminating in the “+2, +0” configuration, which provides the highest anode efficiency among the xenon cases.

#### 5.9.3.1 Xenon: Effect of Channel Width

By isolating the primary diagonal (top left to bottom right) of the configuration grid, the effect of altering the channel width for a constant mean diameter can be assessed. The configurations “+0, +0”, “+1, +1”, and “+2, +2” all have a mean diameter of 70 mm, with channel widths of 8 mm, 12 mm, and 17 mm, respectively. The specific




FIGURE 5.30: Comparison of anode specific impulse against anode power for xenon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

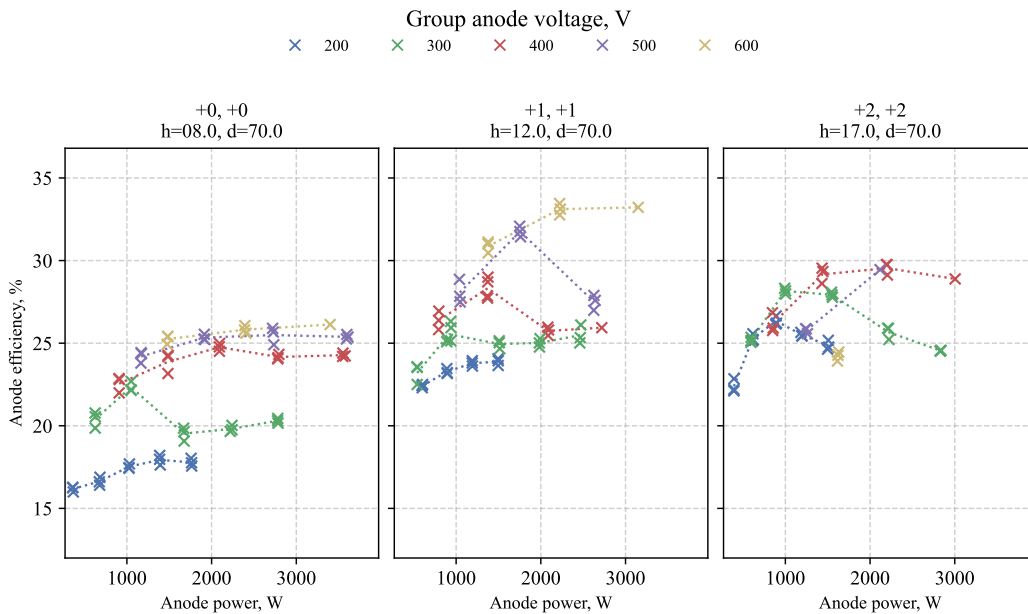


FIGURE 5.31: Comparison of anode efficiency against anode power for xenon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

impulse across the tested power range for these configurations is shown in Fig. 5.30. Across the full power range, it is clear that increasing the channel width enhances performance. This trend is consistent with Eq. 3.32, where the plasma volume-to-surface-area ratio is proportional to channel width. Larger channel widths are expected to reduce plasma-wall interactions, thereby decreasing power losses to the channel walls.

However, as seen in Fig. 5.27, the “+2, +2” configuration exhibits a reduced operational envelope at higher voltages. Consequently, in Fig. 5.31, the “+1, +1” configuration outperforms the others at high voltages but at lower voltages the “+2, +2” configuration enhances efficiency. An additional observation from Fig. 5.31 is that for the “+0, +0” configuration, the anode efficiency plateaus once the discharge voltage exceeds 300 V, whereas the wider channel configurations continue to show enhanced performance at higher voltages. However, at the one flow rate (20 sccm) that was able to sustain operation at 600 V in the “+2, +2” configuration, the anode efficiency was significantly lower than that of “+1, +1” and even slightly lower than the “+0, +0” configuration.

### 5.9.3.2 Xenon: Effect of Channel Mean Diameter

It is particularly useful to compare the diagonal pairs of configurations (lower left to upper right in the grid plot in Fig. 5.20), as in the cases of “+0, +1” versus “+1, +0”, and “+1, +2” versus “+2, +1”. These pairs share the same channel width but differ in channel mean diameter.

Firstly, “+0, +1” and “+1, +0” both have a channel width of 10 mm, but mean diameters of 72 mm and 68 mm, respectively. Across the power range, there is little difference in the measured specific impulse, with “+0, +1” achieving marginally higher values. This becomes clearer in Fig. 5.32, where the two configurations are directly compared. The specific impulse is nearly identical across the full operating regime.

However, Fig. 5.33 shows that the anode efficiency is consistently higher for “+0, +1”, the configuration with the greater mean diameter, with two exceptions. At 300 V, when the anode power exceeded 2 kW, the “+1, +0” configuration outperformed the larger mean diameter case. Similarly, at 600 V the “+1, +0” configuration, although not operated to higher powers, exhibited greater efficiency than “+0, +1” at approximately 1.5 kW.

Now considering the second case of equal channel width, the “+1, +2” and “+2, +1” configurations, which both have a channel width of 14.5 mm but mean diameters of 72.5 mm and 67.5 mm, respectively. Fig. 5.34 shows the specific impulse, where only slight differences are observed between the two mean diameters. In contrast to the 10 mm channel width case, here with a channel width of 14.5 mm, the smaller mean diameter case (“+2, +1”) demonstrates enhanced performance, although the improvement in specific impulse is modest. By contrast, the anode efficiency comparison in Fig. 5.35 reveals a much clearer disparity: the smaller mean diameter configuration significantly outperforms the larger across the entire power range. This performance gap, however, diminishes at higher voltages.




FIGURE 5.32: Comparison of anode specific impulse against anode power for xenon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.




FIGURE 5.33: Comparison of anode efficiency against anode power for xenon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.

### 5.9.3.3 Xenon: Max Performance Grid

Here we introduce a new plot type: the max performance grid. In this format, each major subplot corresponds to a single operating point defined by anode voltage (x-axis) and anode mass flow rate (y-axis). Within each subplot, the grid mirrors the channel configuration layout as shown in Tab. 5.2, with the individual cell values representing the measured performance for that condition. Where multiple measurements were taken, the cell value corresponds to their mean. To highlight




FIGURE 5.34: Comparison of anode specific impulse against anode power for xenon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.




FIGURE 5.35: Comparison of anode efficiency against anode power for xenon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.

trends, a heat map is applied to each subplot, visually emphasising the higher performing configurations.

For xenon, the thrust-to-power ratio performance grid can be seen illustrated in Fig. 5.36. In this plotting configuration, some general trends can be observed. Firstly, at low voltages, the channel configuration does not significantly impact performance, with most configurations for xenon at 200 V and 300 V resulting in similar thrust-to-power ratios. In the 300 V case, some clearer distinctions begin to emerge between good and poor performing configurations, with the “+2, +1” configuration

resulting in the maximum thrust-to-power ratio (with the sole exception of the 30 sccm case, where it is very close to the maximum). At higher voltages, a greater disparity between the best and worst configurations starts to appear, with the “+2, +0” configuration achieving the highest thrust-to-power ratio. This increasing voltage and widening difference likely result from higher plasma temperatures and energies, which cause increased losses in low efficiency configurations.

The same trend is visible in Fig. 5.37, where the same plot structure is used but with the anode efficiency shown. Here the low voltages show a larger performance gap between configurations, with a clear enhancement of efficiency with wider channel widths as well as a slight enhancement with reduced mean diameter. Another way to conceptualise these plots is considering the inner and outer diameters of the channel. Here x-axis “+0”, “+1” and “+2” represent decreasing the inner diameter and the y-axis “+0”, “+1” and “+2” represent increasing the outer diameter. For the 200 V condition regardless of the flow rate increasing both the inner and outer diameter separately or together enhance performance, whilst this trend is absent at the other voltage conditions.

The effect of higher voltages resulting in reduced operational envelope for some configurations is also more apparent where fewer configurations were able to sustain operation. For the 500 V condition, the maximum thrust-to-power was measured consistently in the “+2, +0” configuration, whilst the maximum anode efficiency was recorded in the “+2, +1” configuration, with the exception of the 20 sccm and 30 sccm cases. Interestingly, the 20 sccm case saw the “+2, +0” configuration have a middle-range anode efficiency compared to the other channel geometries. This is also seen at the 600 V 20 sccm case, where the “+2, +0” configuration is on the lower end of measured efficiencies.

For high voltage operation, performance is greatly enhanced with decreasing mean channel diameter or slightly increasing the mean channel diameter without increasing the channel width. This could suggest that moving the plasma closer to either the inner or outer plot enhances performance. However, increasing the channel width decreases the neutral number density to a point where it negatively impacts performance. Hence, performance is enhanced for similar channel widths with either greater or smaller channel mean diameters.

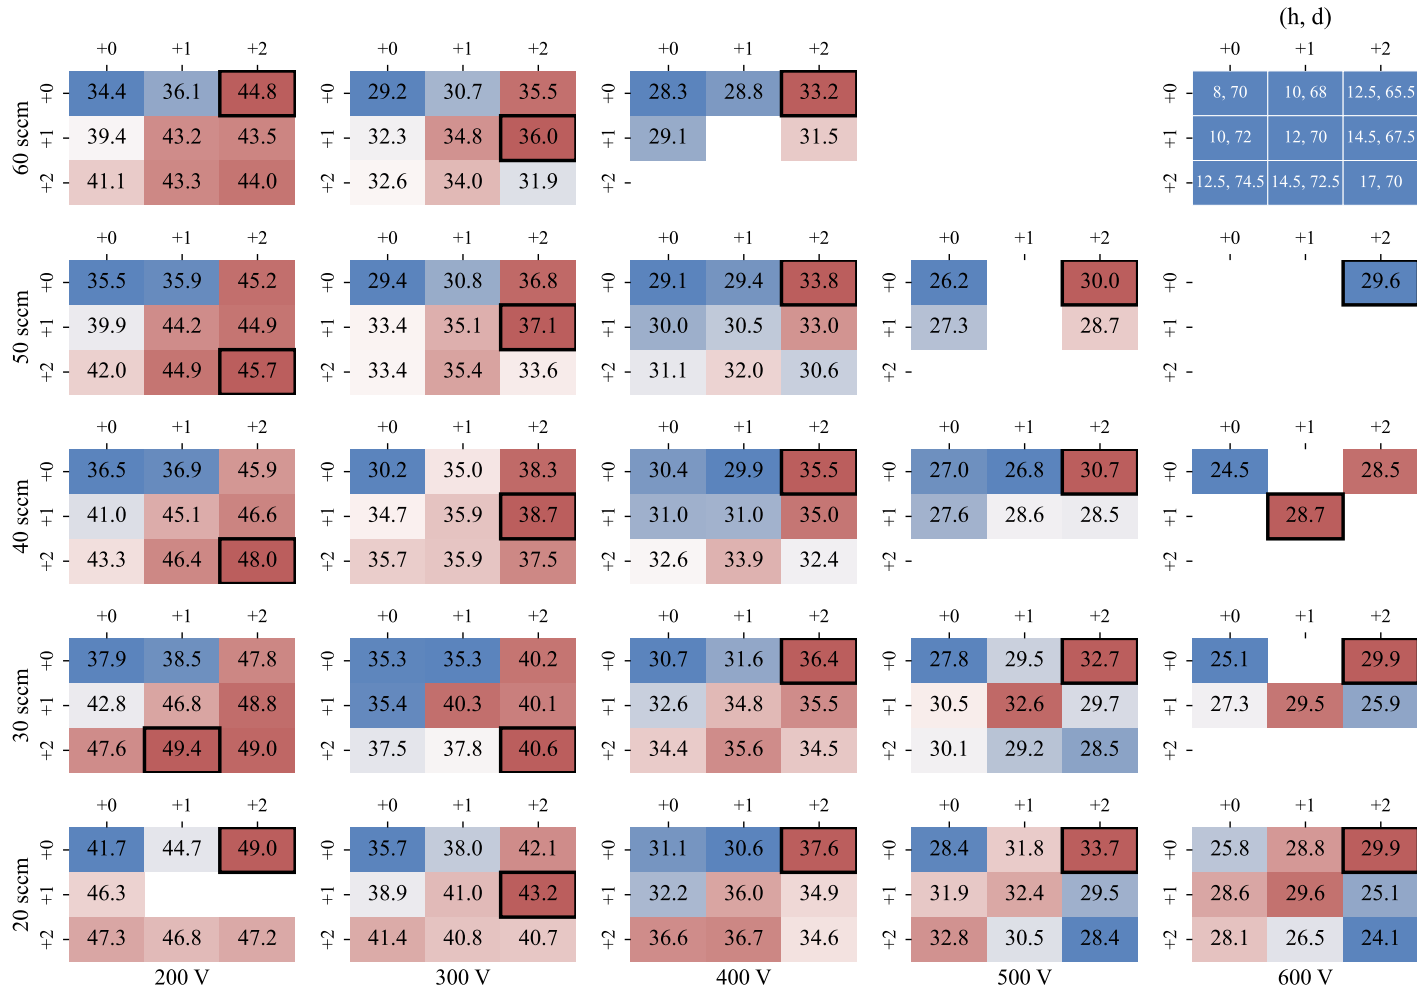


FIGURE 5.36: Performance grid of thrust-to-power ratio (mN/kW) for xenon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.

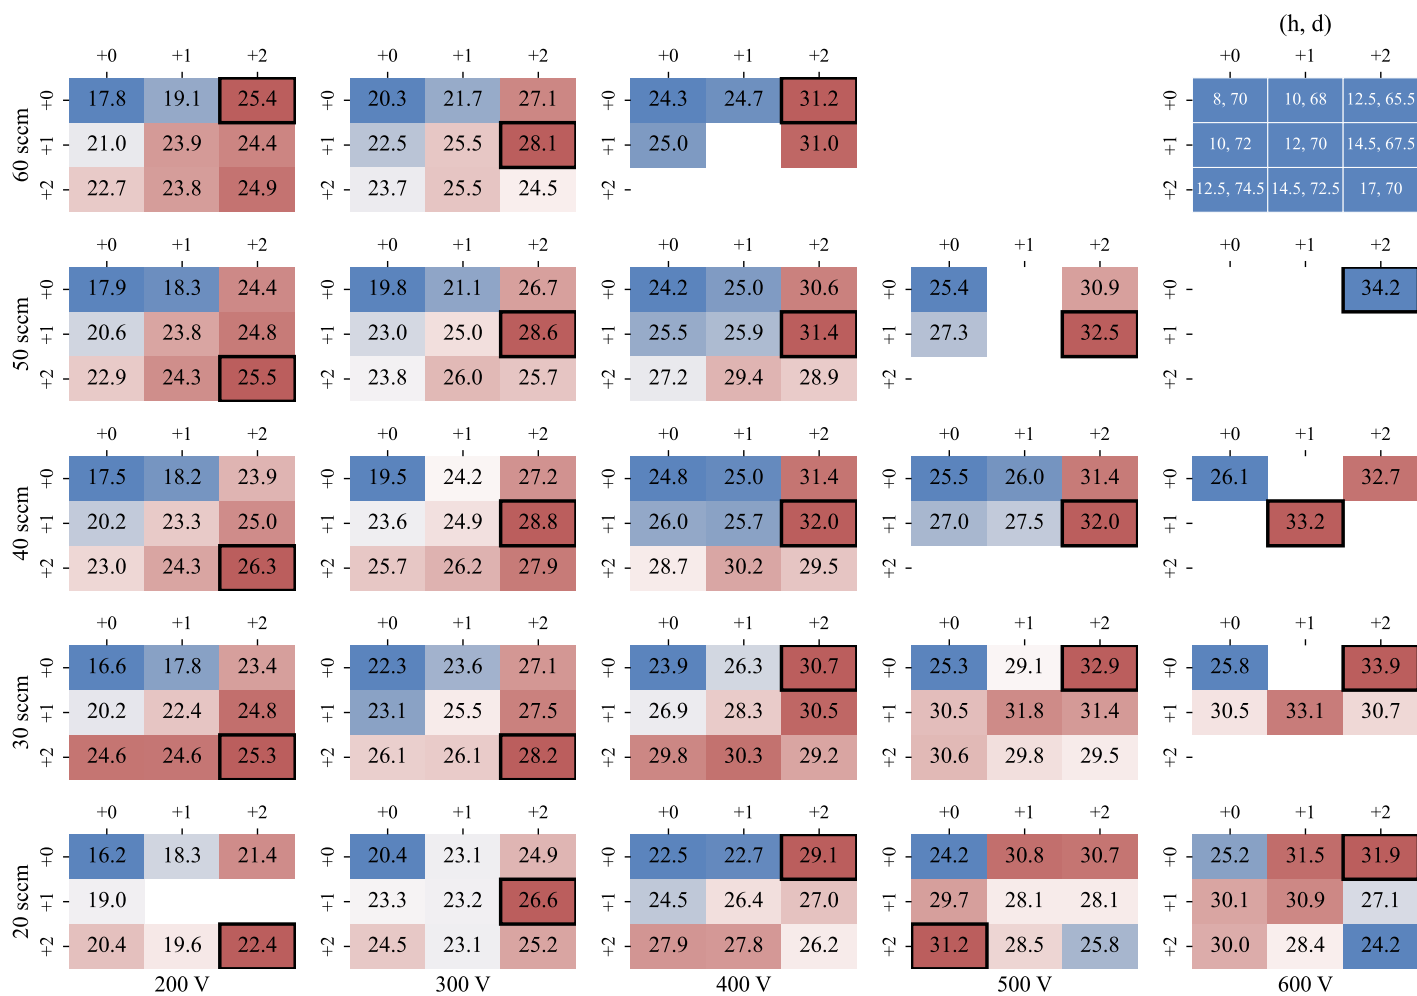


FIGURE 5.37: Performance grid of anode efficiency (%) for xenon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.

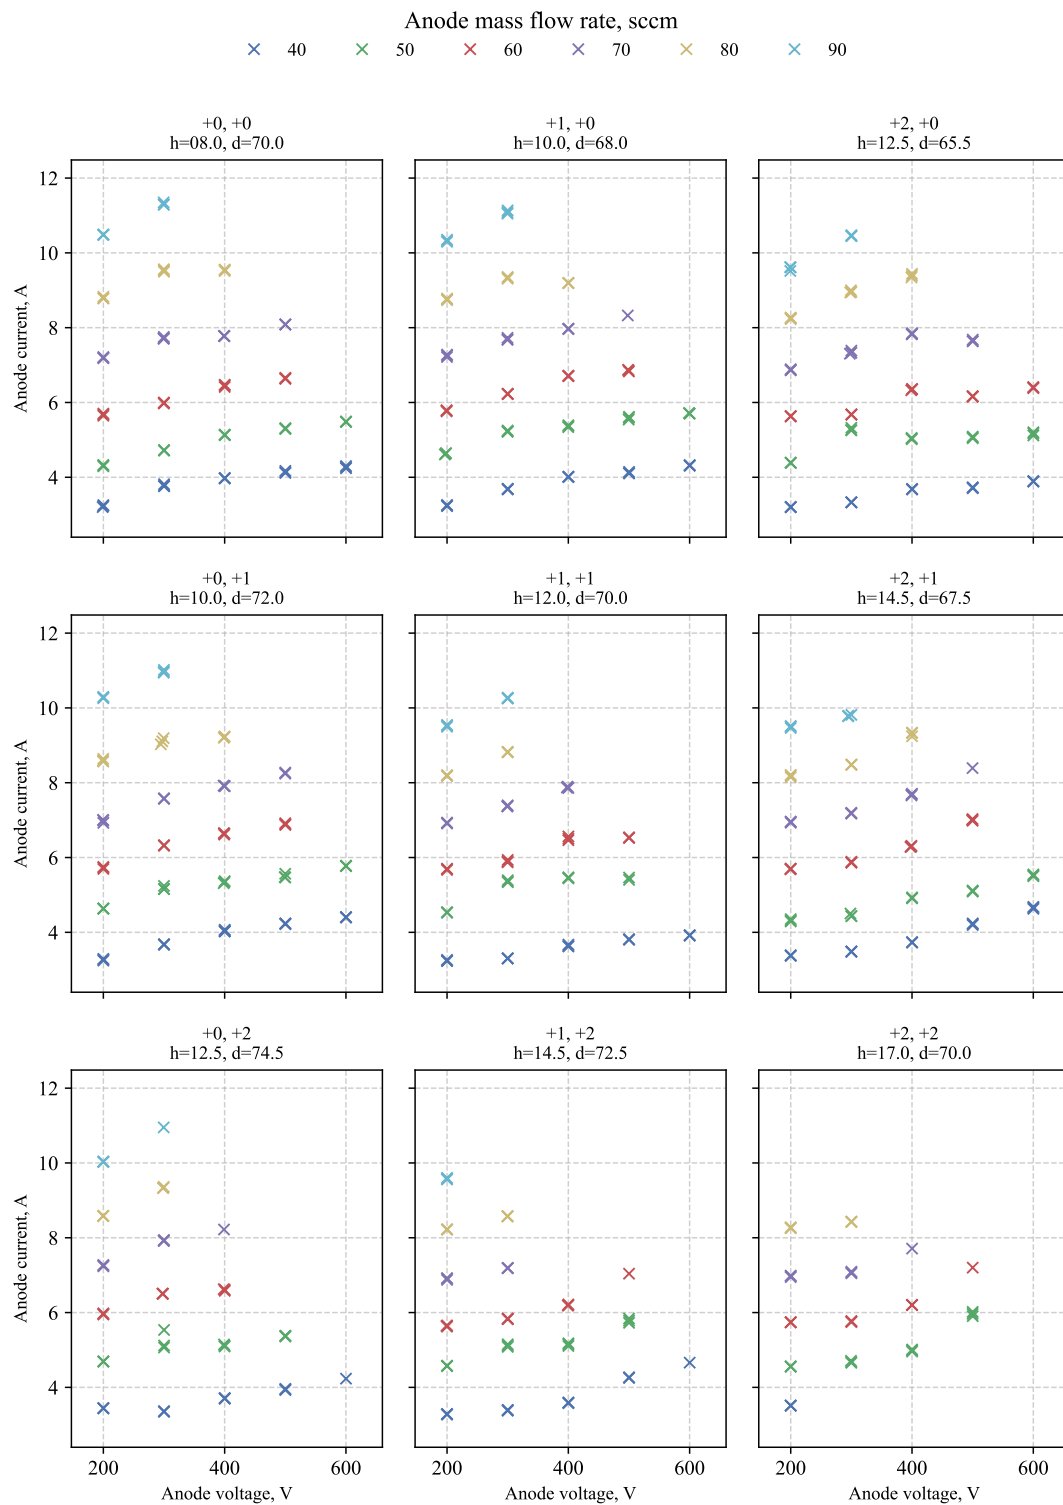


FIGURE 5.38: Anode current draw against anode discharge voltage for krypton, grouped by anode mass flow rate.

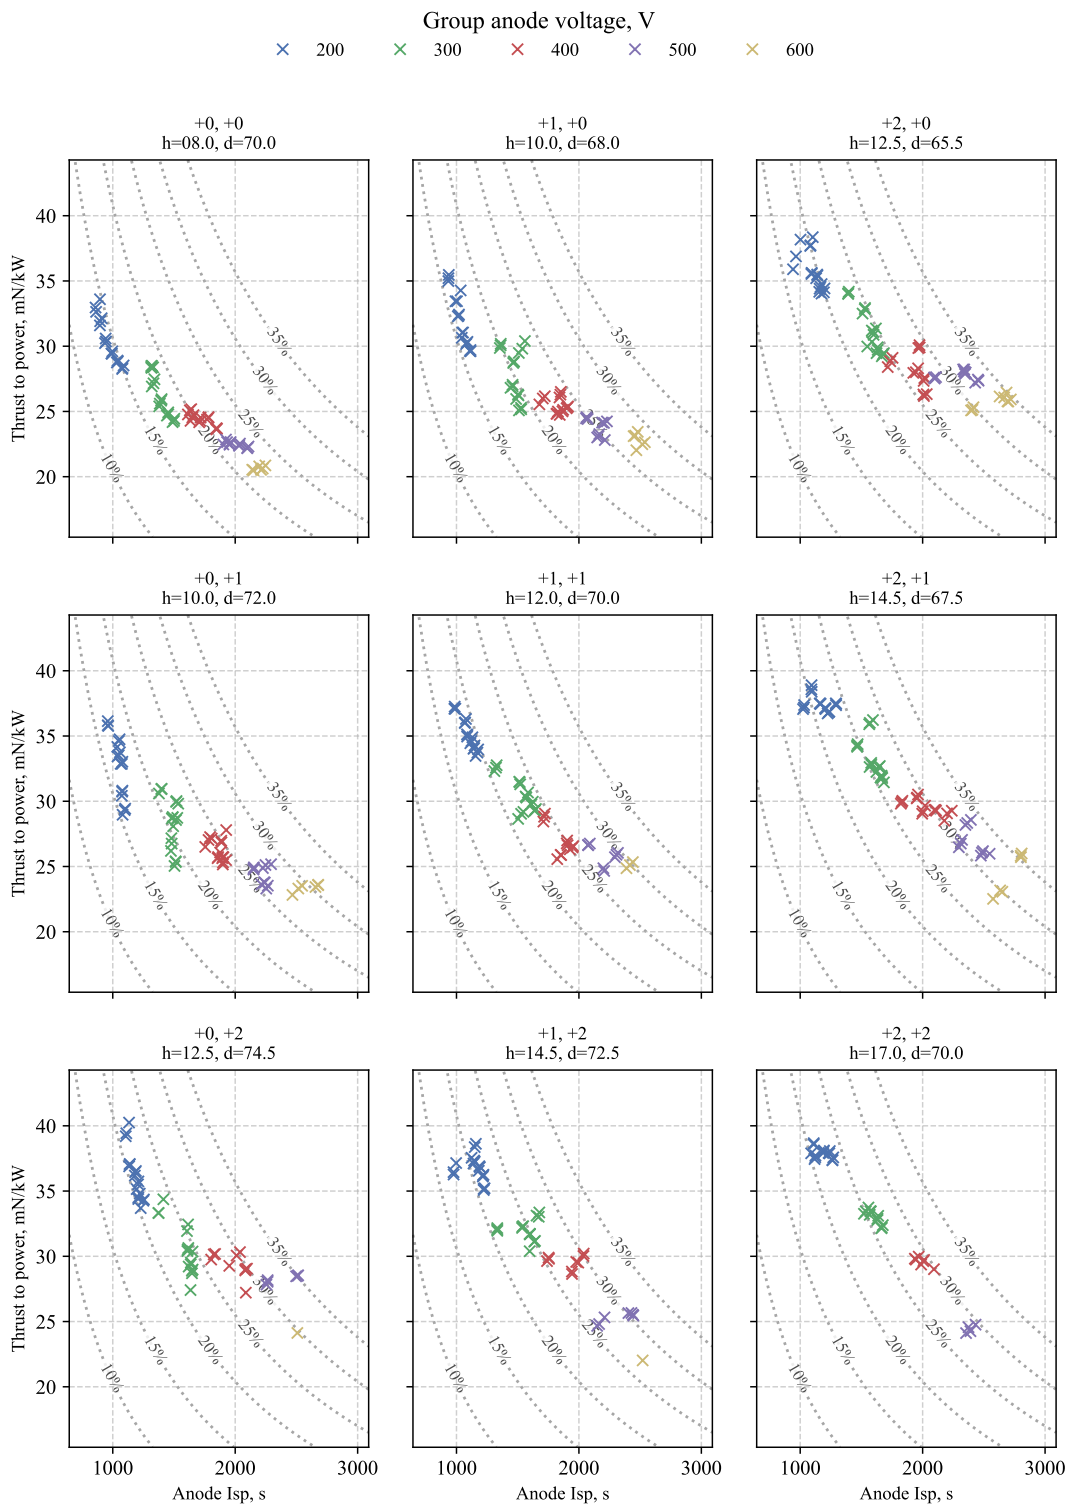


FIGURE 5.39: thrust-to-power ratio against specific impulse anode power for krypton, grouped by discharge voltage.

### 5.9.4 Krypton Performance Trends




FIGURE 5.40: The SHARK-600V operating on krypton in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory.

Krypton operation within the SHARK-600V showed good performance with the highest specific impulses recorded across all other tested propellants. The operational envelope of the thruster showed sensitivity to channel configuration as can be seen in Fig. 5.38 where the bottom row of configurations (representing increase outer channel diameter) showed reduced envelope at high voltages for most of the flow rates tested. This trend is absent in the top row where the largest operational I-V envelope was recorded for the “+2, +0” channel configuration. However, whilst the operation envelope in terms of I-V is reduced, as can be seen in Fig. 5.39 high specific impulse, and anode efficiency, was still achieved in these configurations. Although higher performance was still recorded in the other configurations.

Interestingly, whilst configuration “+0, +0” and “+1, +0” have near identical I-V characteristics in Fig. 5.38 the specific impulse for “+1, +0” exceeds that of “+0, +0” across the power range.

#### 5.9.4.1 Krypton: Effect of Channel Width

Once again, isolating the primary diagonal (top left to bottom right) of the configuration grid allows assessment of the effect of channel width at constant mean diameter. The configurations “+0, +0”, “+1, +1”, and “+2, +2” all share a mean diameter of 70 mm and have channel widths of 8 mm, 12 mm, and 17 mm, respectively.

The krypton specific impulse versus anode power for increasing channel width is shown in Fig. 5.41. The trends broadly mirror those observed for xenon. Notably, the “+0, +0” configuration sustained operation to higher voltages, but this did not translate into improved performance: it delivered consistently lower specific impulse across the power range. A standout feature is the high-voltage behaviour of “+1, +1” at 600 V, which achieved specific impulses in excess of 2400 s at an anode power of approximately 2.3 kW.




FIGURE 5.41: Comparison of anode specific impulse against anode power for krypton across varying channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

Anode efficiency trends in Fig. 5.42 are similar: increasing channel width generally enhances efficiency. The “+0, +0” configuration remains comparatively flat with power, and unlike the wider channels shows no meaningful efficiency increase between 400 V and 600 V. By contrast, the “+1, +1” configuration exhibits an efficiency rise of roughly 7.5 % between 500 V and 600 V.

Overall, while increasing channel width tends to improve performance as can be seen in Figs. 5.41–5.42, the incremental gain in anode efficiency between “+1, +1” and “+2, +2” diminishes at high voltages, most clearly at 500 V. However, the “+1, +1” configuration at 500 V was able to sustain lower power operation than the “+2, +2” configuration.

#### 5.9.4.2 Krypton: Effect of Channel Mean Diameter

Similarly for krypton it is also of interest to compare the diagonal pairs of configurations (lower left to upper right in the grid plot in Fig. 5.20), as in the cases of “+0, +1” versus “+1, +0”, and “+1, +2” versus “+2, +1”. These pairs share the same




FIGURE 5.42: Comparison of anode efficiency against anode power for krypton across varying channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

channel width but differ in channel mean diameter. Firstly, “+0, +1” and “+1, +0” both have a channel width of 10 mm, but mean diameters of 72 mm and 68 mm, respectively. Secondly, the “+1, +2” versus “+2, +1” both have a channel width of 14.5 mm and a channel mean diameter of 72.5 mm and 67.5 mm, respectively.

Firstly, comparing the 10 mm channel width and, 72 mm and 68 mm mean channel diameters (“+0, +1” versus “+1, +0”). Comparing the specific impulse against the tested power range can be seen in Fig. 5.43. The specific impulse difference between the two channel configurations is subtle with the larger mean diameter offering slightly higher specific impulse in all conditions. This is with the exception of the 300 V case where the two channel configurations offer identical specific impulse. In opposition to the xenon case, the difference between the two configurations is greater at higher voltages, suggesting greater sensitivity to channel geometry.

Continuing the 10 mm channel width and, 72 mm and 68 mm mean channel diameters (“+0, +1” versus “+1, +0”) comparison, Fig. 5.44 compares the anode efficiency against anode power for the two mean channel diameters. Here a greater disparity between the performance of these two configurations can be seen. Across the whole power and voltage range it is clear that the 72 mm channel mean diameter case provides increased performance. However, this efficiency gain is small at approximately 1 to 2 % at anode voltages 200 V to 500 V, with the largest difference being at the 600 V case with an efficiency increase of approximately 4 to 5 %. Whilst a small increase, the difference between the two channel configurations is 4 mm of the mean diameter, again suggesting significant sensitivity to channel geometry.




FIGURE 5.43: Comparison of anode specific impulse against anode power for krypton across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.




FIGURE 5.44: Comparison of anode efficiency against anode power for krypton across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.

Comparing the “+1, +2” versus “+2, +1” case, which corresponds to a 14.5 mm channel width and mean diameters of 72.5 mm and 67.5 mm respectively, shows the opposite trend. The specific impulse for both channel configurations across the tested voltages and power range can be seen in Fig. 5.45. Here the reduced operational envelope of the “+1, +2” case can be seen. The smaller mean diameter case was able to be operated to higher powers at all discharge voltages above 200 V. Furthermore, the difference in specific impulse between the two channel configurations is small across the full range. However, the “+2, +1” configuration generally provided higher specific

impulse, with the exception of the 50 sccm case at 200 V, 300 V, and 400 V, where the 72.5 mm mean diameter configuration outperformed the 67.5 mm configuration.




FIGURE 5.45: Comparison of anode specific impulse against anode power for krypton across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.

Continuing the comparison of the “+1, +2” and “+2, +1” configurations, the anode efficiency over the tested power range is shown in Fig. 5.46. A similar trend is observed, with the “+2, +1” case generally outperforming the larger mean diameter configuration. However, apart from the 300 V condition, the increase in anode efficiency is very small yet consistent. The largest discrepancy between the two configurations occurs at the higher voltages, with both the 500 V and 600 V cases showing efficiency gains of approximately 4–5 % for the 67.5 mm mean diameter configuration compared with the 72.5 mm case.

#### 5.9.4.3 Krypton: Max Performance Grid

Here the performance is shown in a different form, with the explanation for the plot structure provided in Sec. 5.9.3.3. Figs. 5.47–5.48 add context to the comparison between the “+0, +1” and “+1, +0” channel configurations discussed in Sec. 5.9.4.2. While in that direct comparison the larger channel mean diameter appeared to offer superior performance, when all channel configurations are considered together it becomes clear that both “+0, +1” and “+1, +0” consistently rank among the lowest performing cases, generally only slightly exceeding the “+0, +0” configuration, which represents the poorest performance across the board.

The general performance trends offered by Figs. 5.47–5.48 are more consistent than the same analysis for xenon, where the highest performing configurations are normally




FIGURE 5.46: Comparison of anode efficiency against anode power for krypton across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.

the larger channel widths with a slight preference for smaller mean channel diameters at the high voltage conditions. For 40 sccm and 50 sccm conditions specifically, the larger mean diameter cases are still high performing. This trend does not hold for higher flow rates and higher voltages where the smaller mean channel diameter configurations are consistently higher performing, if by only a few percent of anode efficiency.

At low voltages (200 V and 300 V), the channel configuration still has strong influence on krypton performance, with most cases having a large range of thrust-to-power ratios and anode efficiency. As the discharge voltage increased, there is less distinctions between several of the configurations, with wider channel configurations consistently offering superior performance. In particular, at 500 V the “+2, +1” and “+2, +0” configurations provided the highest efficiencies, while narrower channels such as “+0, +0”, “+0, +1” and “+1, +0” consistently ranked lowest.

The operational envelope at higher voltages was also reduced for krypton, with some wider channel configurations unable to sustain stable operation above 400 V. Moreover, larger mean channel diameter configurations above 60 sccm and above 300 V failing to sustain operation. Nevertheless, the remaining configurations maintained consistent trends, with wider channels generally outperforming. The distinction is especially clear at 600 V, where the “+2, +1” configuration achieved the highest efficiency.

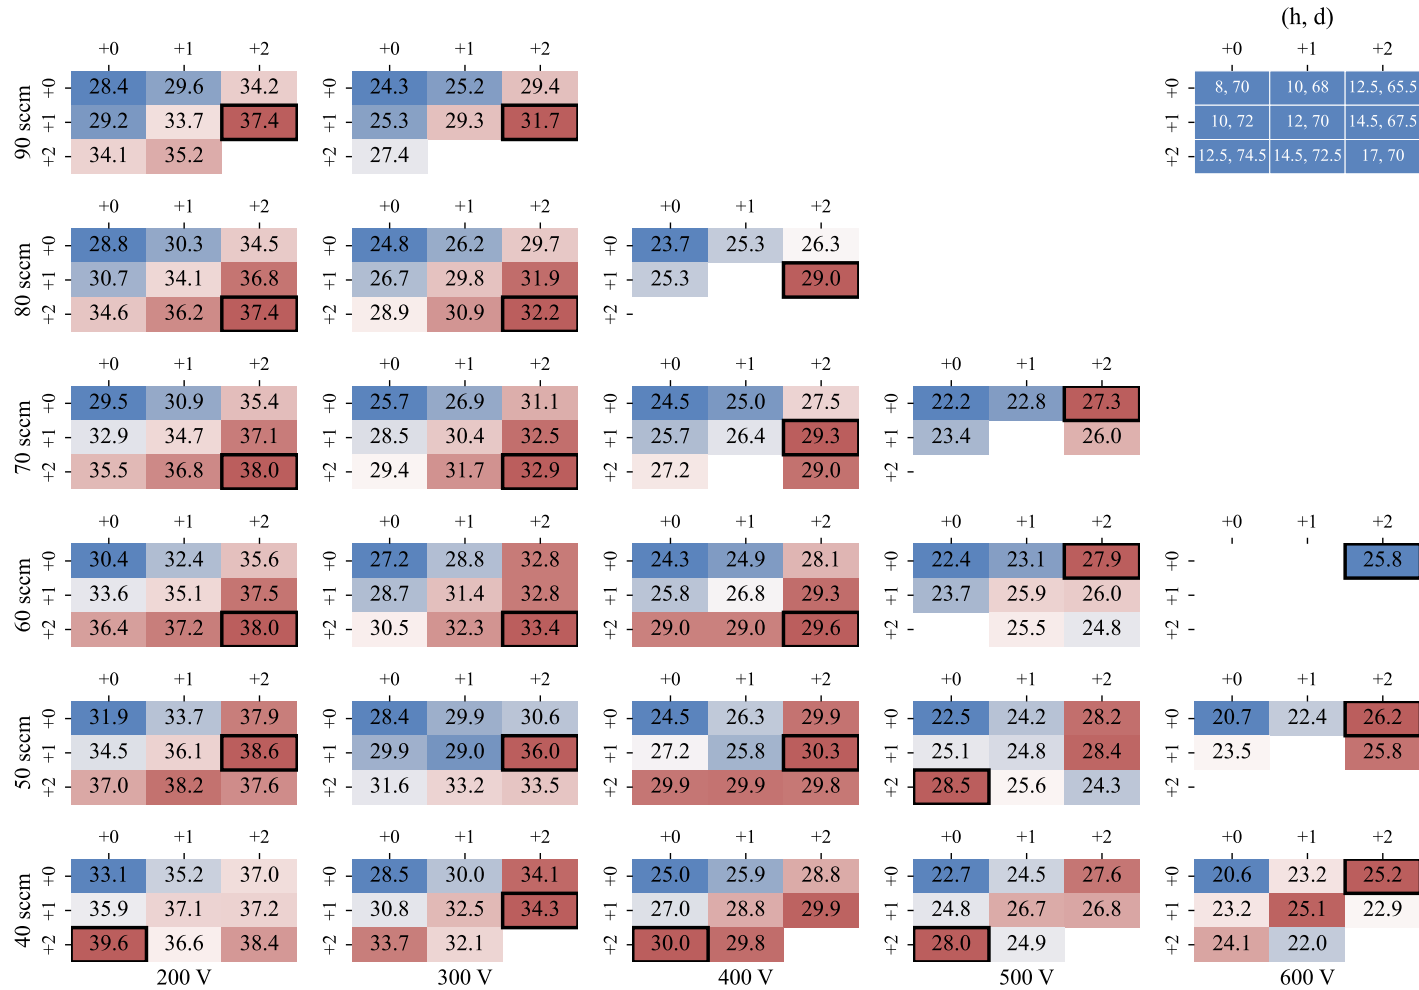


FIGURE 5.47: Performance grid of thrust-to-power (mN/kW) ratio for krypton. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.




FIGURE 5.48: Performance grid of anode efficiency (%) for krypton. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.

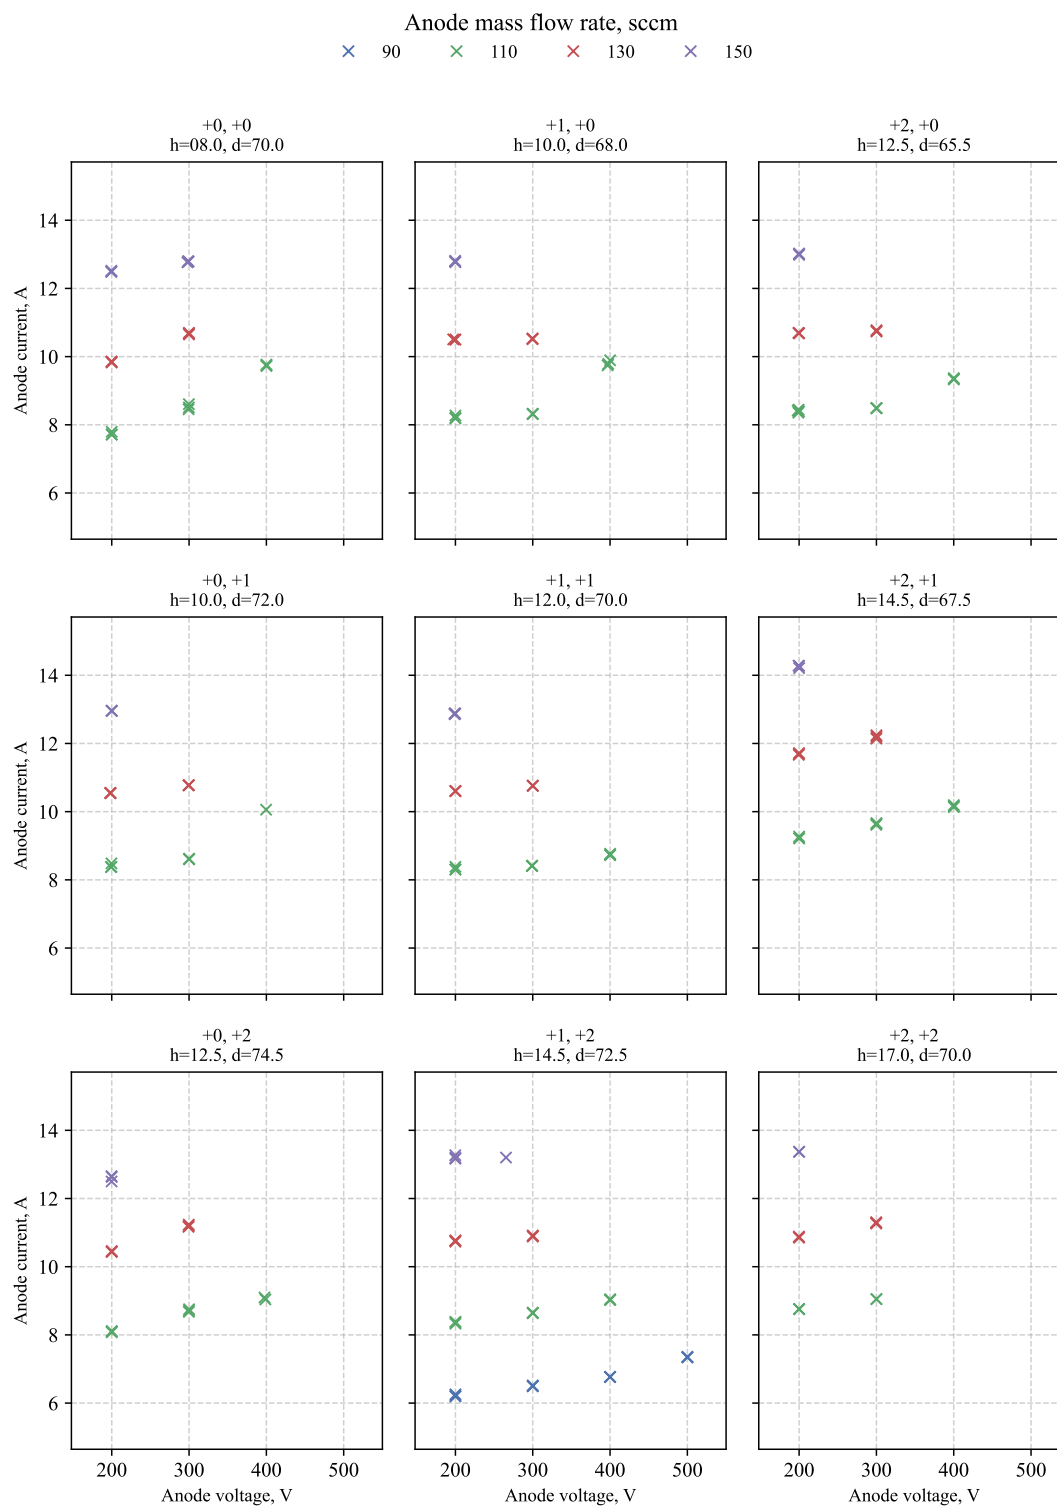


FIGURE 5.49: Anode current draw against anode discharge voltage for argon, grouped by anode mass flow rate.

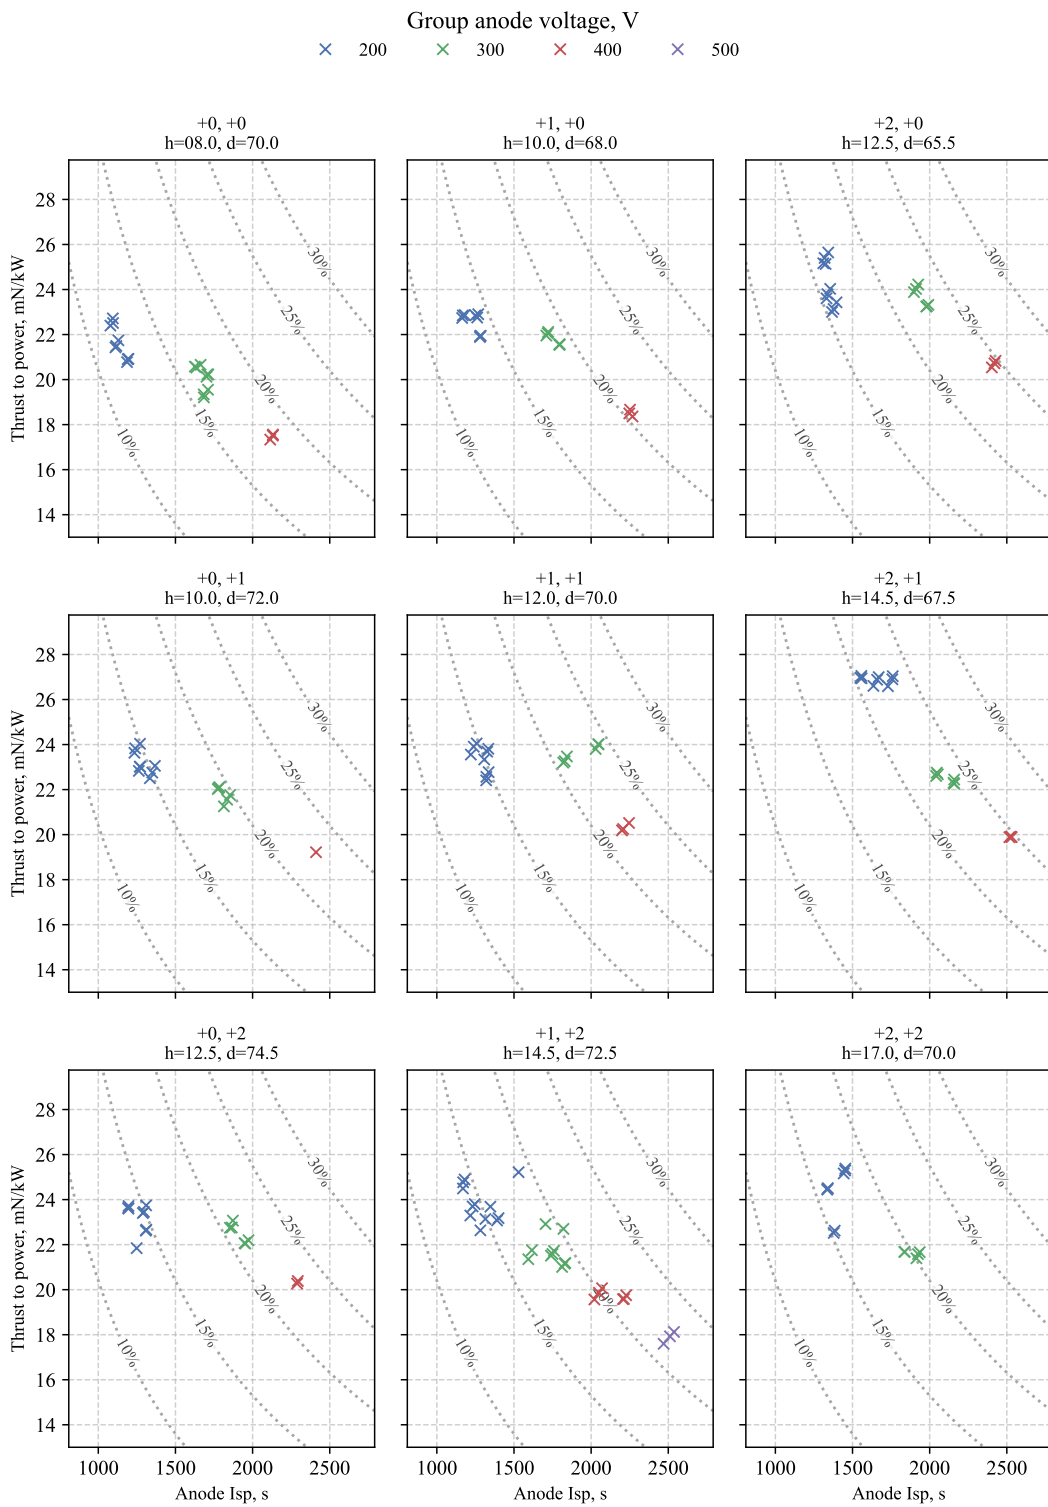


FIGURE 5.50: thrust-to-power ratio verses specific impulse for argon, grouped by discharge voltage.

### 5.9.5 Argon Performance Trends




FIGURE 5.51: The SHARK-600V operating on argon in the large vacuum chamber facility in the David Fearn Electric Propulsion Laboratory.

Argon exhibited the smallest operational envelope of the three propellants tested, as can be seen with the I-V characteristics in Fig. 5.49. This outcome is not unexpected, as relatively high volumetric flow rates are required to achieve sufficient mean free paths for ionisation. For example, largest argon flow test of 150 sccm corresponds to 4.46 mg/s, which is actually less than the 4.88 mg/s corresponding to a krypton flow of 50 sccm. However, due to the lower atomic mass of argon, this mass flow translates into a significantly larger number of neutral atoms than for krypton. These higher neutral number densities in turn drive larger discharge currents. Consequently, operation with argon demanded significantly higher discharge powers than xenon or krypton, frequently approaching the 4.2 kW testing power limit imposed.

Only the “+1, +2” channel configuration was characterised with a mass flow rate of 90 sccm, this was due to the other configurations not sustaining a stable plasma at 200 V for this flow rate such it was not tested to other voltages. Interestingly, this is the only condition that was able to achieve voltages in excess of 400 V, with the 90 sccm case being tested up-to 500 V. Whilst this improved envelope as seen in Fig. 5.49 for the “+1, +2” channel configuration attain higher specific impulse as seen in Fig. 5.50, due to the higher voltages achieved, other channel configurations such as “+2, +1” was able to attain higher specific impulse at 400 V anode voltage suggesting the whilst the envelope of “+1, +2” is enhanced the performance is not.

From Fig. 5.50 a strong relation between channel dimensions and performance can be seen. Here performance can be seen to generally improve moving left to right along the top two rows, for increasing channel width and decreasing channel mean diameter.

### 5.9.5.1 Argon: Effect of Channel Width

For argon, isolating the primary diagonal (top left to bottom right) of the configuration grid allows assessment of the effect of channel width at constant mean diameter. The configurations “+0, +0”, “+1, +1”, and “+2, +2” all share a mean diameter of 70 mm, with channel widths of 8 mm, 12 mm, and 17 mm, respectively.

The argon specific impulse versus anode power for increasing channel width is shown in Fig. 5.52. The general trends broadly mirror those observed for xenon and krypton. However, at 200 V the wider configurations exhibit a reduction in specific impulse with increasing power, while the “+0, +0” configuration shows improvement over the same range. This behaviour does not persist at 300 V and 400 V. At 400 V, the “+1, +1” configuration provides the highest specific impulse among this subset, reaching approximately 2200 s at a discharge power approximately 400 W lower than the corresponding “+0, +0” case.




FIGURE 5.52: Comparison of anode specific impulse against anode power for argon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

Considering the anode efficiency for increasing channel width with argon, some unique behaviour appears. For xenon, the “+2, +2” channel configuration was generally the most efficient across all operating conditions, with the exception of the 500 V and 600 V cases. Krypton followed a similar trend, with the “+2, +2” configuration again providing the highest efficiency at constant mean diameter, except for one 500 V condition and the inability of this configuration to sustain operation at 600 V.

For argon, however, the behaviour diverges. The highest anode efficiency is achieved at 300 V in the “+1, +1” configuration, with a significant reduction in performance

when moving to the wider “+2, +2” channel. Nevertheless, at 200 V the “+2, +2” configuration still outperforms the narrower cases, although its efficiency drops steeply at higher powers (i.e., higher flow rates). Once again, the “+2, +2” channel configuration fails to operate at higher voltages, mirroring krypton’s behaviour, but this limitation occurs at 400 V for argon instead of 600 V as in the krypton case.




FIGURE 5.53: Comparison of anode efficiency against anode power for argon across channel width configurations (“+0, +0” vs. “+1, +1” vs. “+2, +2”), grouped by discharge voltage.

### 5.9.5.2 Argon: Effect of Channel Mean Diameter

For argon, it is again of interest to compare diagonal pairs of configurations, namely “+0, +1” versus “+1, +0” and “+1, +2” versus “+2, +1” (lower left to upper right in the grid plot in Fig. 5.20). These pairs maintain the same channel width but differ in mean channel diameter. The first pair, “+0, +1” and “+1, +0”, both have a channel width of 10 mm with mean diameters of 72 mm and 68 mm, respectively. The second pair, “+1, +2” and “+2, +1”, both have a channel width of 14.5 mm with mean diameters of 72.5 mm and 67.5 mm, respectively.

Beginning with the 10 mm channel width case, the comparison of anode specific impulse against anode power between “+0, +1” and “+1, +0” is shown in Fig. 5.54. Here, the specific impulse difference between the two channel configurations is minor across the operating range. At lower voltages, the smaller mean diameter case (“+1, +0”) appears to plateau, whereas the larger mean diameter case (“+0, +1”) continues to increase. Interestingly, the larger mean diameter case shows slightly

enhanced performance, which is the opposite of the trends observed for krypton and xenon.

These differences become more apparent when considering anode efficiency, as shown in Fig. 5.55. The 72 mm mean diameter case offers an efficiency increase of 0.5–3 % at 200 V and 300 V, and a more significant increase of approximately 5 % at 400 V. This again supports the observation that argon performance is increasingly sensitive to channel geometry at higher voltages. Overall, the larger mean diameter case offers a measurable performance benefit as the discharge voltage is increased.




FIGURE 5.54: Comparison of anode specific impulse against anode power for argon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.




FIGURE 5.55: Comparison of anode efficiency against anode power for argon across channel mean diameter configurations (“+1, +0” vs. “+0, +1”), grouped by discharge voltage.

Now considering the 14.5 mm channel width case for the “+1, +2” and “+2, +1” channel configurations, which have mean diameters of 72.5 mm and 67.5 mm, respectively. The comparison of their specific impulse against anode discharge power is shown in Fig. 5.56. Here, the difference between the two channel configurations is more pronounced than in the “+0, +1” versus “+1, +0” case. This arises because the wider channel width enhances overall performance, thereby making the effect of mean channel diameter a stronger driver of performance. In this case, the smaller mean diameter (“+2, +1”) provides a clear advantage, achieving the same specific impulse as the larger mean diameter case (“+1, +2”) at approximately 100 V lower anode voltage.




FIGURE 5.56: Comparison of anode specific impulse against anode power for argon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.

The impact of this enhanced specific impulse is mirrored in Fig. 5.57, where the “+2, +1” configuration achieves significantly higher anode efficiencies across the full operating range. The performance gap in Fig. 5.57 is more pronounced than in the corresponding xenon and krypton cases. Another noteworthy observation is that the “+2, +1” configuration operated at higher discharge currents than the “+1, +2” case, suggesting that the 5 mm reduction in mean channel diameter substantially increased ionisation efficiency. This resulted in both more thrust-producing ions and greater discharge current. However, the efficiency gap between the two configurations narrows at higher voltages.

### 5.9.5.3 Argon: Max Performance Grid

Here the performance is shown in a different form, with the explanation for the plot structure provided in Sec. 5.9.3.3. Similarly to the krypton analysis, Figs. 5.58–5.59 add context to the comparison between the “+0, +1” and “+1, +0” channel configurations




FIGURE 5.57: Comparison of anode efficiency against anode power for argon across channel mean diameter configurations (“+1, +2” vs. “+2, +1”), grouped by discharge voltage.

discussed in Sec. 5.9.5.2. While in that direct comparison the larger channel mean diameter appeared to offer slightly superior performance, when all channel configurations are considered together it becomes clear that both “+0, +1” and “+1, +0” consistently rank among the lower-performing cases. Their performance is generally comparable to, or only marginally better than, the “+0, +0” configuration, which represents the poorest performance across the set.

It is also apparent from Figs. 5.58–5.59 that argon shows increased sensitivity to channel geometry, with fewer configurations producing high performance. In contrast to xenon and krypton, which showed broader ranges of effective geometries and only slight preferences for wider channels and smaller mean diameters, argon demonstrates a much narrower set of favourable conditions. High performance is only observed in configurations that combine a smaller mean channel diameter with a wider channel width. Moreover, the “+2, +2” configuration, which was often among the highest performance configurations for krypton, consistently performed poorly with argon, further highlighting the propellant’s stronger dependence on channel geometry.




FIGURE 5.58: Performance grid of thrust-to-power (mN/kW) ratio for argon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.

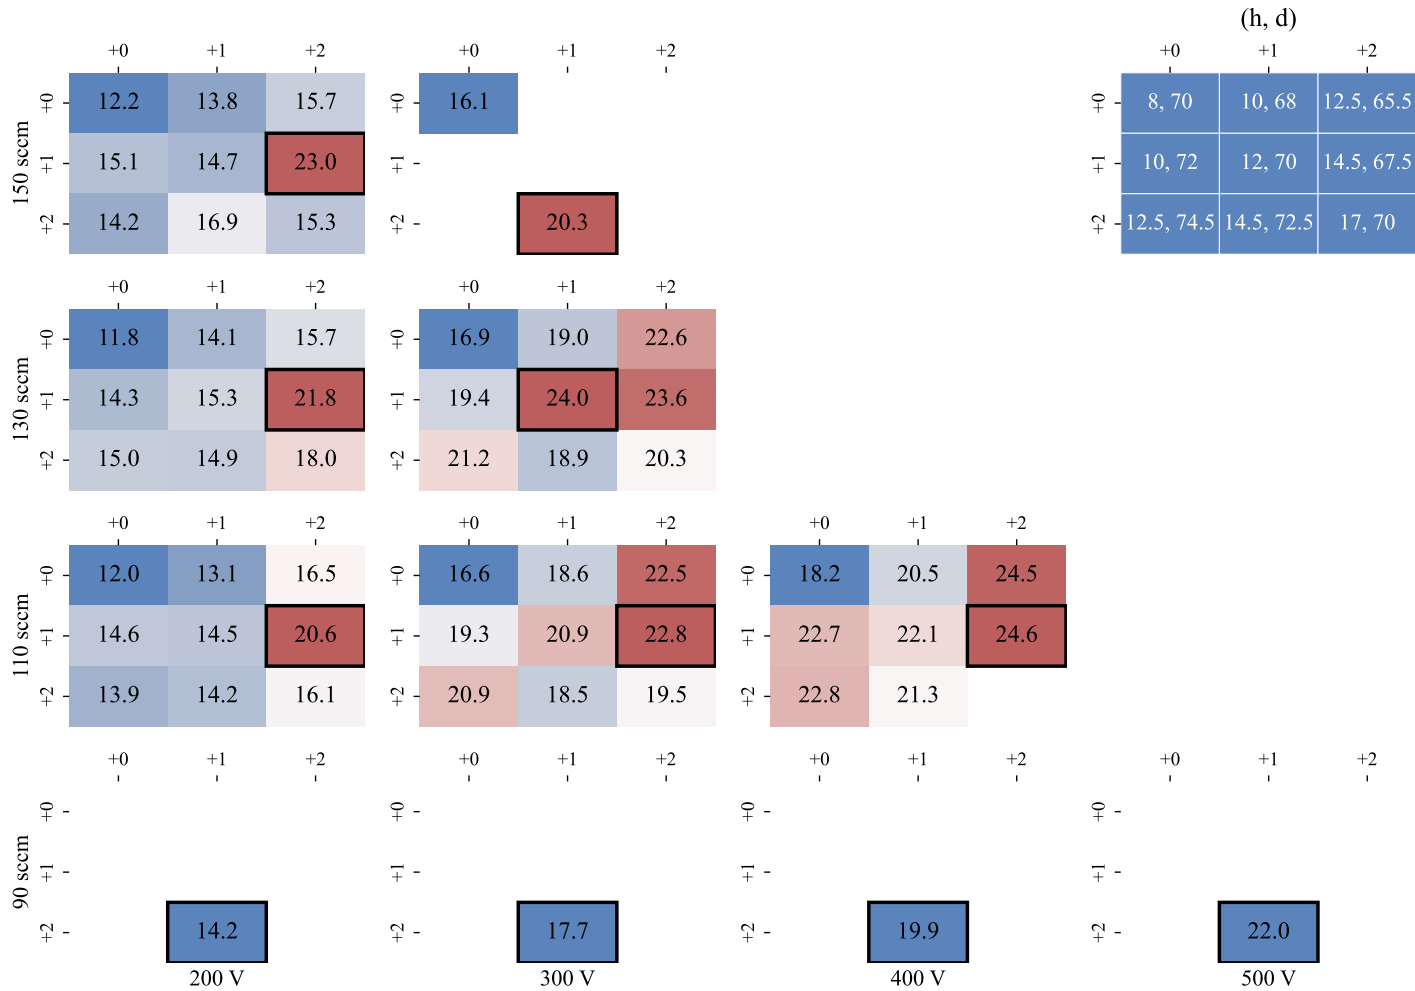


FIGURE 5.59: Performance grid of anode efficiency (%) for argon. The values in each cell represent the average measurement across all data points recorded for that configuration at the given operating condition. The colour map uses a hot-cold scale, with higher values represented by warmer colours. Each voltage-flow rate combination has its own scale.




FIGURE 5.60: A photo of the thruster captured after the plasma was extinguished manually due to overheating concerns where the boron nitride channel is emitting in the visible wavelength due to high temperatures.

### 5.9.6 Anomalous Behaviour

Throughout testing at high voltages several operating conditions resulted in significant asymmetric heating of the channel. This occurred on all propellants and configurations, however the operational point at which this would occur for each configuration and propellant was different showing geometry dependency. The channel is made from boron nitride (BN). Specifically, BN99 a grade of boron nitride that has  $\geq 98.5\%$  BN with  $B_2O_3$  binder which per supplier information is rated for operation up to  $2400^\circ\text{C}$ .

### 5.9.7 Thruster with Anode-Layer Configuration

The initial testing of the SHARK-600V in the anode-layer configuration was conducted early in the campaign. As a result, the full DAQ system described in Sec. 5.7 was not employed. This reduced the ease of operation and introduced greater uncertainty into the measurements. Furthermore, the data were collected at relatively low powers compared to the main test campaign, since these trials were intended primarily to verify that the thruster could operate in the TAL mode at all. Consequently, the results from these tests are not directly comparable to the bulk of the SPT performance data presented in Secs. 5.9.2–5.9.5. This is also the reason why all test data presented for TAL operation were obtained exclusively with krypton propellant.

As with the SPT configuration of the SHARK-600V, which allowed for multiple channel arrangements, the anode-layer configuration also permitted variation. As discussed in Sec. 4.2, four different TAL configurations were designed, defined by



FIGURE 5.61: The TAL-1 configuration operating at approximately 1.3 kW on krypton within the large vacuum chamber facility at the University of Southampton.

varying the length of the anode sheaths. However, due to the initial operational difficulties and generally poor performance observed in TAL mode, only two of these configurations were tested, as shown in Fig. 4.2. These were the “TAL+1” and “TAL-1” configurations, corresponding to anode tip distances from the thruster exit plane of  $-0.5$  mm and 5.5 mm, respectively. Here, negative values indicate that the anode tip extends past the magnetic front of the thruster.

#### 5.9.7.1 TAL: Experimental Results

It is well established that thrusters with anode-layers tend to exhibit lower operational stability than their SPT counterparts [194]. This behaviour was seen during testing, where the SHARK-600V proved more difficult to operate in the TAL configuration with only a few operational points being capable of achieving stable operation at meaningful voltages ( $\geq 200$  V).

A key issue encountered during TAL operation was the increased thermal loading of the anode. This outcome was expected to some extent, due to the higher plasma temperatures in TAL operation, combined with the anode’s extension closer/into the region of peak magnetic field, provided a more favourable path for electrons to reach the anode, and thereby increased the electron current. Despite these challenges, the SHARK-600V was successfully operated in the TAL configuration within the large vacuum chamber facility at the University of Southampton, albeit within a limited operating envelope.




FIGURE 5.62: Anode efficiency for the two TAL configurations tested against anode power grouped by krypton flow rate in sccm and anode configuration.

When compared with the SPT operation shown in Fig. 5.40 from Sec. 5.9.4, the TAL discharge in Fig. 5.61 exhibited a stronger purple glow for krypton. This emission is qualitatively associated with doubly charged ions, indicating higher plasma temperatures as expected in TAL operation. However, without plume diagnostics, this observation remains qualitative only.

### 5.9.7.2 TAL: Performance

Performance data were collected with the “TAL+1” anode prior to its failure, after which the “TAL-1” configuration was tested. In this latter configuration the anode was positioned further from the region of peak magnetic field in order to reduce thermal loading. The best performance of the TAL configurations was obtained with the “TAL-1” case, as shown in Fig. 5.62. In particular, Fig. 5.63 shows that the “TAL-1” configuration achieved modest specific impulse, which is relatively unimpressive due to the significantly lower flow rates achieved with TAL configurations.

The TAL+1 configuration, although tested less extensively, produced comparable performance at lower powers. A direct comparison of the “TAL+1” and “TAL-1” configurations is shown in Fig. 5.62. The 6 mm reduction in anode length significantly improved stability, enabling “TAL-1” to sustain operation at higher powers. However, while “TAL-1” was able to reach higher powers than the “TAL+1” configuration, these conditions were attained only at very low flow rates compared to the performance shown in Sec. 5.9.4. This is likely due to the significantly increased electron currents to the anode in the TAL configurations as a result of reduced electron confinement. This effect appears to be a major driver of the observed poor performance.




FIGURE 5.63: Anode specific impulse against anode power for the “TAL-1” configurations grouped by krypton flow rates.




FIGURE 5.64: Thrust against anode power for the “TAL-1” configurations grouped by krypton flow rates.

The TAL configurations demonstrated the ability to operate at lower flow rates than the magnetic-layer configurations. However, this was not accompanied by comparable specific impulse or anode efficiency. As shown in Fig. 5.64, the apparent efficiency advantage of the “TAL-1” configuration over the “TAL+1” configuration does not correspond to a practically useful operating mode. In particular, the thrust performance of the 30 sccm case in Fig. 5.64 was completely flat. Here, the discharge current increased with voltage without a corresponding increase in performance, suggesting that the higher applied voltage did not translate into increased plasma potential but instead extracted more electrons to the anode without producing higher-velocity ions.




FIGURE 5.65: Thermal radiation from the anode in the “TAL-1” configuration following 1.5 kW operation.




FIGURE 5.66: “TAL+1” firing at 1.3 kW where excessive thermal loading of the anode caused localised melting of the stainless-steel sheaths.

Even in the “TAL-1” configuration, the anode reached very high temperatures, as illustrated in Fig. 5.65, where the latent heat following a 1.5 kW discharge caused visible thermal radiation.

### 5.9.7.3 TAL: Notable Events

Initial testing employed the “TAL+1” anode. This configuration pushes the anode through the region of peak magnetic field. This resulted in several instances of localised melting of the stainless-steel anode sheath, as shown in Fig. 5.66.

Fig. 5.66a shows operation at 1.3 kW and 300 V just before the onset of melting. In Fig. 5.66b the anode surface has liquefied, forming a bead that extended further into the plasma, producing an anode rich exhaust. Finally, Fig. 5.66c captures the aftermath after the discharge was shut down, with residual latent heat visible. The aftermath of the melting event can be seen in Fig. 5.67 once the chamber was brought back to atmosphere.



FIGURE 5.67: The “TAL+1” metallic anode after high thermal load conditions. The image shows the post-operation state and the resultant structural failure after melting, observed once the chamber was returned to atmosphere.

#### 5.9.7.4 TAL: Conclusion

In summary, the TAL tests demonstrated low performance operation in this configuration and suffered from severe thermal loading and reduced efficiency compared to the SPT configuration. The stability of the “TAL-1” configuration indicates a path forward, particularly with the use of refractory metal anodes, but further plume diagnostics are required to fully characterise the performance gap.

## 5.10 Experimental Campaign Discussion

The SHARK-600V campaign represents the first systematic characterisation of a modular Hall effect thruster across nine channel geometries and three noble gas propellants without alteration to the experimental set-up. Operating up to 600 V for xenon and krypton, and 500 V for argon, this study provides a dataset that enables investigation of the effects of discharge channel geometry, operating voltage, and propellant type on thruster performance. For these tests, thrust was measured directly using a thrust balance. From the thrust data, the specific impulse, anode efficiency, and thrust-to-power ratios were calculated using the recorded operational parameters.

Across all propellants, the specific impulse increased with increasing discharge voltage. This behaviour is expected, as a higher anode voltage increases the plasma potential and accelerates the ions to greater velocities. Lighter propellants produced higher specific impulses when operated in efficient configurations. The highest specific impulse recorded was for krypton, reaching 2809 s at 600 V, with an anode flow rate of 50 sccm, and anode power of 3.3 kW. Argon also achieved comparably high values, with a peak of 2537 s at 500 V, 90 sccm, and 3.7 kW. Xenon achieved a maximum specific impulse of 2423 s at 600 V, 30 sccm, and 2.7 kW.

While all three propellants demonstrated the expected scaling of specific impulse with discharge voltage, efficiency and stability were found to be strongly dependent on the discharge channel geometry and propellant type. Certain configurations, notably “+2, +1” and “+2, +0”, consistently provided the highest anode efficiency and thrust-to-power ratios across multiple propellants, particularly at elevated voltages. Conversely, the “+2, +2” configuration, often suffered from reduced stability at high voltages, limiting their operational envelope despite offering good performance at lower discharge voltages. Whilst, the “+2, +1” and “+2, +0” configuration had channel widths of 12.5 mm and 14.5 mm respectively, the “+2, +2” configuration with a channel width 2.5 mm greater than “+2, +1” saw a drastic reduction in stability at high voltages and generally lower performance.

These results highlight both the potential and the limitations of using alternative propellants in Hall effect thrusters. Krypton, when operated in favourable geometries at high voltages, can not only match but in some cases surpass xenon in both efficiency and specific impulse. Argon, although less efficient overall, demonstrated the ability to reach comparable specific impulse to xenon and krypton at lower discharge voltages, Fig. 5.23. However, this required significantly higher volumetric flow rates and discharge powers, narrowing its practical operating envelope.

Channel width emerged as the dominant geometric driver for performance and stability. Increasing width enhanced performance in terms of both specific impulse and anode efficiency, attributed to reduced wall losses via an increased plasma volume-to-surface area ratio. Nevertheless, the widest configuration (“+2, +2”) frequently exhibited reduced stability at higher voltages, demonstrating the existence of an optimal range of channel width. Furthermore, this configuration showed significantly reduced performance for operation on argon. In contrast, the narrow channel width of the “+0, +0” configuration produces a much larger surface-area-to-volume ratio, which increases electron-wall and ion-wall collisions and reduces beam current efficiency. Additionally, the higher neutral number densities present in this configuration raise the frequency of neutral-neutral collisions, further diminishing electron confinement.

Increasing the channel width also has the effect of moving the plasma closer to the magnetic poles. This likely enhances electron confinement at the extremes, as the field gradient towards the walls is greater within the channel, thereby increasing the magnetic mirroring effect that reflects electrons back into the plasma. This observation is consistent with earlier findings by Dannenmayer [167] for a 200 W class Hall effect thruster, the "PPI" (French acronym for "Petit Propulseur Innovant").

The mean channel diameter acted as a second-order control, with its influence most apparent at high discharge voltages. Smaller mean diameters, such as the "+2, +1" configuration, consistently yielded improved performance for krypton and argon at 500 V and 600 V, while xenon exhibited weaker sensitivity. This behaviour suggests that the optimisation of ionisation and acceleration regions relative to the channel boundaries is propellant dependent. One possible explanation is that reducing the mean diameter moves the plasma closer to the central magnetic pole, where the magnetic field gradient between the channel centreline is stronger than between the centreline and the outer pole. This gradient may enhance electron confinement through magnetic mirroring, producing the observed gains. However, this trend is not replicated in the "+2, +2" configuration, where both poles are closer to the plasma, but the channel is at its widest. In this case, the performance falls once the channel width exceeds an optimal value. The weaker magnetic gradient at the outer pole may also limit further confinement improvements, preventing compensation for the performance losses associated with excessive channel width. This could potentially also explain the generally lower performance of the greater mean channel diameter configurations. However, without plume diagnostics that could identify the dominant efficiency loss mechanism in each configuration it is hard to make conclusive arguments.

The SHARK-600V operated in the TAL configuration only within a narrow envelope, with frequent thermal limitations and unstable behaviour compared to the SPT mode. The increased electron current to the anode, driven by reduced electron confinement, led to excessive heating and in some cases structural failure of the metallic sheaths. While the "TAL-1" variant offered improved stability over "TAL+1", its performance was constrained to low flow rates and poor efficiency. In contrast, the SPT mode demonstrated broader operational stability and more consistent scaling with input power, underscoring the relative maturity of SPT design. Previously literature has shown that TALs can be operated to reasonable efficiency and without plasma and thermal instabilities [195, 196]. These results suggest that the poor performance observed here stems less from inherent drawbacks of the TAL concept and more from the configuration choices applied to the SHARK-600V.

## 5.11 Gaussian Process Regression Model of SHARK-600V

Highly dimensional systems such as Hall effect thrusters are hard to describe solely with standard plots as there is high level of interdependency that can be obfuscated when presenting data. This is exacerbated with the SHARK-600V as we add additional dimensions to the data with channel width and channel mean diameter. One such method to extract meaningful results from these datasets is the development of models from experimental data. These methods have been previously applied to electric propulsion in the form of machine learning or computation surrogate models [161, 197–200].

One such model that has been recently applied to Hall effect thrusters is Gaussian Process Regression (GPR). This is a highly adaptable method, capable of producing a surrogate model of any quantifiable aspect of thruster operation from a set of input parameters and a moderate dataset size. The application of this method to the operation of Hall effect thrusters was first introduced to the author in a conversation with Dr. Peter Thoreau who has successfully used this method for real-time modelling and thruster-in-loop optimisation [201]. The application to the dataset from the SHARK-600V thruster offers a deep insight into how channel dimension affect thruster performance as well as allow for investigation of untested geometry that are within or close to the tested values. Together with this model the performance data for each propellant can then be used to infer optimal channel geometry and can be compared to the scaling results.

### 5.11.1 Gaussian Process Regression

Originally developed for sea-floor prospecting of precious metal, GPR is a Bayesian approach to regression that places a probability distribution over functions rather than assuming a fixed parametric form, like linear regression [202, 203]. The key benefits that GPR offers over other surrogate models are that it can handle high-dimensional datasets and provide predictions with an estimated uncertainty. The model replaces physics-based models with a black-box method; this is especially useful here due to the lack of a comprehensive physics model for Hall thruster operation.

The model requires a set of inputs (i.e., anode voltage, anode mass flow rate, channel width, etc.) for a single output value (i.e., thrust, anode current, etc.). To obtain other outputs, a separate model is trained.

Gaussian Process Regression models a system as a distribution over functions. Observations of the system at different input parameters are treated as samples from a joint multivariate Gaussian, with a covariance structure defined by a kernel function. The kernel encodes assumptions about smoothness, continuity, and behaviour. Using

the observed data, the GPR predicts outputs at new input values by computing the posterior distribution, which combines the prior assumptions with the data to give both a mean prediction and an uncertainty estimate [204]. In a Gaussian process regression, we assume our output  $y$  to be defined as a function of a vector of inputs  $\mathbf{x}$  such that

$$y(\mathbf{x}) = f(\mathbf{x}) + \epsilon, \quad (5.8)$$

where  $\mathbf{x}$  is the input vector of our system (specifically:  $\mathbf{x} = [U_d, \dot{m}_a, h, d]$ ), and  $\epsilon$  is the noise which is assumed normal, such that  $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ . This is similar to other regression models that assume our measurement consists of an independent “signal” function and a noise term. The uncertainty regarding  $f(\mathbf{x})$  can be reduced by observing the output of the function at different input points. The noise term  $\epsilon$  reflects the inherent randomness in the observations, which is always present no matter how many observations we make.

Gaussian process regression is a Bayesian method, meaning that it treats the underlying function  $f(\mathbf{x})$  as a random variable with a probability distribution. Predictions are made by updating prior beliefs about the function with observed data according to Bayes’ theorem, the namesake of the Bayesian approach. Formally, for a set of training inputs  $\mathbf{X}_t$  with corresponding outputs  $\mathbf{y}_t$ , the posterior distribution over functions is given by

$$\mathbb{P}(f \mid \mathbf{X}_t, \mathbf{y}_t) = \frac{\mathbb{P}(\mathbf{y}_t \mid f, \mathbf{X}_t) \mathbb{P}(f \mid \mathbf{X}_t)}{\mathbb{P}(\mathbf{y}_t \mid \mathbf{X}_t)}, \quad (5.9)$$

where  $\mathbb{P}(f \mid \mathbf{X}_t)$  represents the prior distribution over functions, defined by the training inputs and the kernel assumptions;  $\mathbb{P}(\mathbf{y}_t \mid f, \mathbf{X}_t)$  is the likelihood, describing the probability of observing the training outputs given a specific function  $f$ ;  $\mathbb{P}(f \mid \mathbf{X}_t, \mathbf{y}_t)$  is the posterior distribution, reflecting the updated belief about the function after observing the training data; and  $\mathbb{P}(\mathbf{y}_t \mid \mathbf{X}_t)$  is the marginal likelihood, which serves to normalize the posterior.

In the context of the SHARK-600V dataset, the input vector  $\mathbf{x}_i = [U_{d,i}, \dot{m}_{a,i}, h_i, d_i]$  defines the operating condition of the thruster  $i$ , and the output  $y_i$  corresponds to a measurable performance metric such as thrust, anode efficiency, or anode current. Here  $\mathbf{X}_t$  is the matrix of all training input vectors, with each row corresponding to one experimental condition of the SHARK-600V thruster, such that

$$\mathbf{X}_t = \begin{bmatrix} U_{d,1} & \dot{m}_{a,1} & h_1 & d_1 \\ U_{d,2} & \dot{m}_{a,2} & h_2 & d_2 \\ \vdots & \vdots & \vdots & \vdots \\ U_{d,N} & \dot{m}_{a,N} & h_N & d_N \end{bmatrix}. \quad (5.10)$$

Bayes' theorem ensures that the GPR combines prior knowledge encoded by the training data with the experimental observations to produce both a predicted mean and an uncertainty estimate for each output at any new input condition.

The Gaussian process regression gets its name from assuming this  $f(\mathbf{x})$  can be approximated with a Gaussian process such that

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (5.11)$$

A Gaussian process  $\mathcal{GP}$  is a distribution over a function and is defined by a mean  $m(\mathbf{x})$ , which represents the expected value,  $\mathbb{E}[f(\mathbf{x})]$ , of the function at  $\mathbf{x}$ , such that  $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ , and a covariance function  $k(\mathbf{x}, \mathbf{x}')$ . The covariance function models the dependency between the function values at different input points  $\mathbf{x}$  and  $\mathbf{x}'$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E} \left[ (f(\mathbf{x}) - m(\mathbf{x})) (f(\mathbf{x}') - m(\mathbf{x}')) \right]. \quad (5.12)$$

The covariance function in GPR is commonly referred to as the kernel. The mathematical description of the covariance function can be chosen for the specific application, such as smoothness or expected trends or patterns. However, here we used the radial basis function that implies the correlation between two points decays with distance between them. The radial basis function provides an expressive kernel to model smooth and stationary functions, such that similar inputs would result in similar outputs. This gives our kernel definition of

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left( - \frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\lambda^2} \right). \quad (5.13)$$

This function introduces two hyper-parameters:  $\lambda$ , the "length-scale", and  $\sigma_f^2$ , the signal variance. Each element of the input vector  $\mathbf{x}$  has its own length scale. As can be seen from Eq. 5.13, the length scale encodes sensitivity to each element within the input vector. Large length scales imply that the function  $f(\mathbf{x})$  is relatively insensitive, and conversely, small length scales imply high sensitivity.

The Gaussian Process model contains hyper-parameters in its kernel function, these control the smoothness, sensitivity, and overall amplitude of the function and must be chosen to reflect the behaviour of the training data  $\mathbf{X}_t, \mathbf{y}_t$ . Rather than selecting them manually, GPR optimises the hyper-parameters by maximizing the log marginal likelihood of the observed outputs given the inputs.

With this kernel description (Eq. 5.13), and measured performance  $f(x_i)$  at the operating point  $\mathbf{x}_i$ , the mean and the standard deviation can be estimated from a dataset of thruster operating conditions, channel geometry and performance,  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ , and their performance,  $\{f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)\}$ . Here, the mean and standard deviation given by the surrogate model are the predicted value and uncertainty at that

operating point. Here we use the data collected earlier within this chapter to form a surrogate model of thruster performance per propellant tested as a function of operating conditions (anode mass flow rate and anode voltage) and channel dimensions (channel mean diameter, and channel width).

This is not a comprehensive description of GPR modelling as that is outside the scope of this thesis. For a deeper and generalisable description of GPR the reader is directed to Schulz, Speekenbrink, and Krause [204].

### 5.11.2 Gaussian Process Regression Model Training

Separate Gaussian process regression models were trained for each propellant (xenon, krypton, and argon) and each output parameter (thrust, specific impulse, anode current, and anode efficiency) using the experimental data from the SHARK-600V campaign, the size of the training and test datasets for each propellant can be seen in Tab. 5.3. The training process maximises the log marginal likelihood of the observed outputs given the inputs. However, this optimisation can converge to local maxima and fail to identify the global optimum if insufficient training data are used or too few random restarts are performed.

To validate each model, a subset of the experimental data was withheld from the training set to serve as a test sample of unseen data. This enables quantification of the model error on data not used during training, providing a measure of validation. As shown in Fig. 5.68, the relative uncertainty of the model decreases and stabilises once a sufficient number of training points are provided.

TABLE 5.3: Dataset sizes used for training and testing the GPR surrogate models for each propellant.

Propellant	Total dataset size	Training dataset size	Test dataset size
Xenon	498	398	100
Krypton	544	435	109
Argon	171	136	35

#### 5.11.2.1 Model Training Results

Once trained, each model was validated against a dataset of unseen experimental points to evaluate its ability to predict performance for configurations outside the training set. The size of the test dataset represents a trade-off: too large a test set leaves insufficient data for effective training, while too small a test set reduces the statistical significance of the validation results. One advantage of GPR models is their ability to maintain reasonable accuracy even with relatively small training datasets. This is illustrated in Fig. 5.68, where the mean relative error for all model outputs falls below




FIGURE 5.68: Mean relative error of the Gaussian process regression model for xenon as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (100 for xenon). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models.

5% once approximately 100 training points are used. The corresponding results for krypton and argon are presented in Appendix B.

An additional consideration in model development is the number of optimisation restarts required to ensure convergence to a global optimum. To determine a sufficient number of restarts for stable output, a similar analysis was undertaken, the results of which are shown in Fig. 5.69.

The Root Mean Squared Error (RMSE) is a standard measure of the average deviation of predictions  $\hat{y}_i$  from the true outputs  $y_i$  over  $n$  samples

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}. \quad (5.14)$$

To provide a dimensionless measure of error relative to the variability of the dataset, we define the Coefficient of Variation of the RMSE, denoted as  $\text{CV}(\text{RMSE})$ , by normalizing the RMSE with the standard deviation of the true outputs  $\sigma_y$

$$\text{CV}(\text{RMSE}) = \frac{\text{RMSE}}{\sigma_y} \times 100\%, \quad (5.15)$$

where  $\sigma_y = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2}$  and  $\bar{y}$  is the mean of the true outputs.  $\text{CV}(\text{RMSE})$  provides an error metric that accounts for the spread of the data, making it more




FIGURE 5.69: Mean relative error of the Gaussian process regression model for xenon as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models.

comparable across different outputs or datasets.

TABLE 5.4: Statistical results of the Gaussian Process Regression test set for xenon dataset against the trained model. Here  $R^2$  is the coefficient of determination, RMSE, is the root mean square error between the tested values and the model predictions,  $\text{RMSE \%}_{\text{mean}}$  is the RMSE relative to the mean value of test value, and the  $\text{CV}(\text{RMSE})$  is Coefficient of Variation of the RMSE and is mathematically described in Eq. 5.15.

Target	$R^2$	RMSE	$\text{RMSE \%}_{\text{mean}}$	$\text{CV}(\text{RMSE})$
Thrust, mN	0.9990	0.7434	1.36%	3.18%
Anode current, A	0.9993	0.0592	1.22%	2.58%
Anode $I_{sp}$ , s	0.9974	19.6957	1.29%	5.12%
Anode efficiency, %	0.9751	0.6233	2.45%	15.78%

TABLE 5.5: Statistical results of the Gaussian Process Regression test set for krypton dataset against the trained model. Here  $R^2$  is the coefficient of determination, RMSE, is the root mean square error between the tested values and the model predictions,  $\text{RMSE \%}_{\text{mean}}$  is the RMSE relative to the mean value of test value, and the  $\text{CV}(\text{RMSE})$  is Coefficient of Variation of the RMSE and is mathamatically described in Eq. 5.15.

Target	$R^2$	RMSE	$\text{RMSE \%}_{\text{mean}}$	$\text{CV}(\text{RMSE})$
Thrust, mN	0.9990	0.6326	1.06%	3.19%
Anode current, A	0.9997	0.0418	0.63%	1.85%
Anode $I_{sp}$ , s	0.9974	21.5951	1.37%	5.10%
Anode efficiency, %	0.9877	0.4798	2.13%	11.08%

TABLE 5.6: Statistical results of the Gaussian Process Regression test set for argon dataset against the trained model. Here  $R^2$  is the coefficient of determination, RMSE, is the root mean square error between the tested values and the model predictions,  $\text{RMSE \%}_{\text{mean}}$  is the RMSE relative to the mean value of test value, and the  $\text{CV}(\text{RMSE})$  is Coefficient of Variation of the RMSE and is mathamatically described in Eq. 5.15.

Target	$R^2$	RMSE	$\text{RMSE \%}_{\text{mean}}$	$\text{CV}(\text{RMSE})$
Thrust, mN	0.9971	1.0213	1.69%	8.59%
Anode current, A	0.9987	0.0636	0.62%	3.59%
Anode $I_{sp}$ , s	0.9930	32.0227	1.88%	8.37%
Anode efficiency, %	0.9735	0.5323	2.89%	16.28%




FIGURE 5.70: Geometry map of the channel configurations from Tab. 5.2 shown in the same format as subplots in Figs. 5.71–5.73.

### 5.11.3 Gaussian Process Regression Model of SHARK-600V Performance

The Gaussian process regression model developed for the SHARK-600V enables continuous interpolation across the geometric dimensions of the thruster, allowing performance predictions for both tested and untested channel configurations. This model facilitates visualisation of thruster performance in a manner similar to the discrete grids shown in Figs. 5.36–5.37, Figs. 5.47–5.48, and Figs. 5.58–5.59. However, unlike the discrete experimental grid, the GPR model enables continuous interpolation between tested channel dimensions and extrapolation beyond them, providing insight into potentially optimal geometries not directly examined during the experimental campaign.

The GPR model predictions for anode efficiency for xenon, krypton, and argon are illustrated in Figs. 5.71–5.73. Each subplot within the array corresponds to a specific combination of anode voltage and anode mass flow rate, with the major axes representing the thruster geometry parameters: channel width (horizontal) and mean channel diameter (vertical). The tested thruster geometries are overlaid on the figures; these points correspond to those shown in Fig. 5.70, which has the same axis dimensions as the subplots shown in Figs. 5.37, 5.48, and 5.59. The colour maps show the continuous GPR-predicted anode efficiency across the geometry space, while the red crosses mark the discrete channel configurations that were experimentally tested. The associated prediction uncertainty for the results shown in Figs. 5.71–5.73 can be found in Appendix B.

An additional point of clarification when comparing these figures is that the channel configuration array (Tab. 5.2) in the GPR plots is effectively rotated and mirrored relative to the configuration tables presented earlier. For example, the “+0, +0” configuration, which appears in the top-left cell of the configuration array in Tab. 4.1, is located at the leftmost position along the middle horizontal axis in the GPR figures. This is further illustrated in Fig. 5.70.

One caveat of the GPR model becomes apparent when comparing Figs. 5.37, 5.48, and 5.59 to Figs. 5.71–5.73: the model allows for evaluation at operating points that were not obtainable during experiments due to plasma instabilities, failure to ignite, or thermal limitations. As this information cannot be encoded into the model, the predictions cannot be used without caution. The GPR model only learns the curvature of the tested parameter space and does not account for the existence of “unobtainable” operating points. This limitation is implicitly reflected in the large uncertainties produced when evaluating the model at these points. The prediction uncertainties corresponding to Figs. 5.71–5.73 are presented in Appendix B, in Figs. B.5–B.7. In these figures, the uncertainty can be seen to increase in the top-right subplots, conditions that exceeded the power limitations of the system and for which no experimental data exist. The model attempts to interpolate to these regions, but with increased uncertainty, as it does not implicitly know that these operating points were unobtainable.

The results of these plots further illustrate the conclusions from the direct analysis of the thruster’s performance, namely that channel width is a first-order driver of performance, whereas mean channel diameter acts as a second-order control with increased importance at higher anode voltages. This is unsurprising, as the GPR representation merely provides a continuous interpolation of the same data used to form the prior analysis.

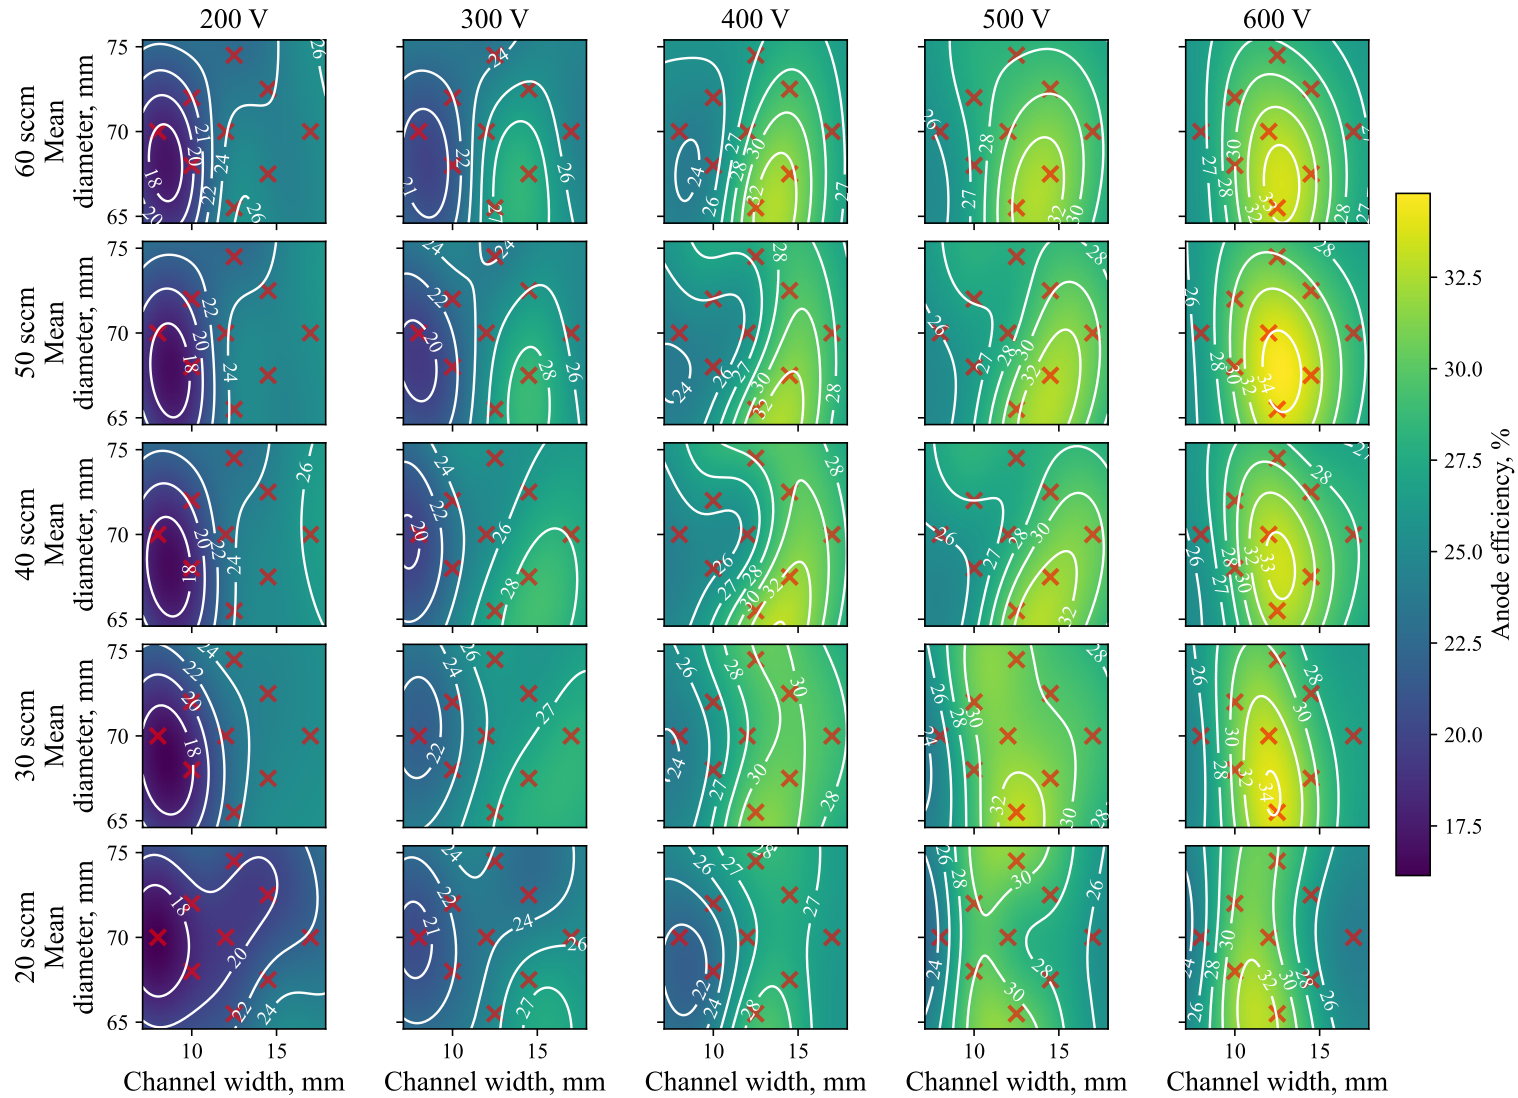


FIGURE 5.71: Array of Gaussian process regression model for xenon, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested.

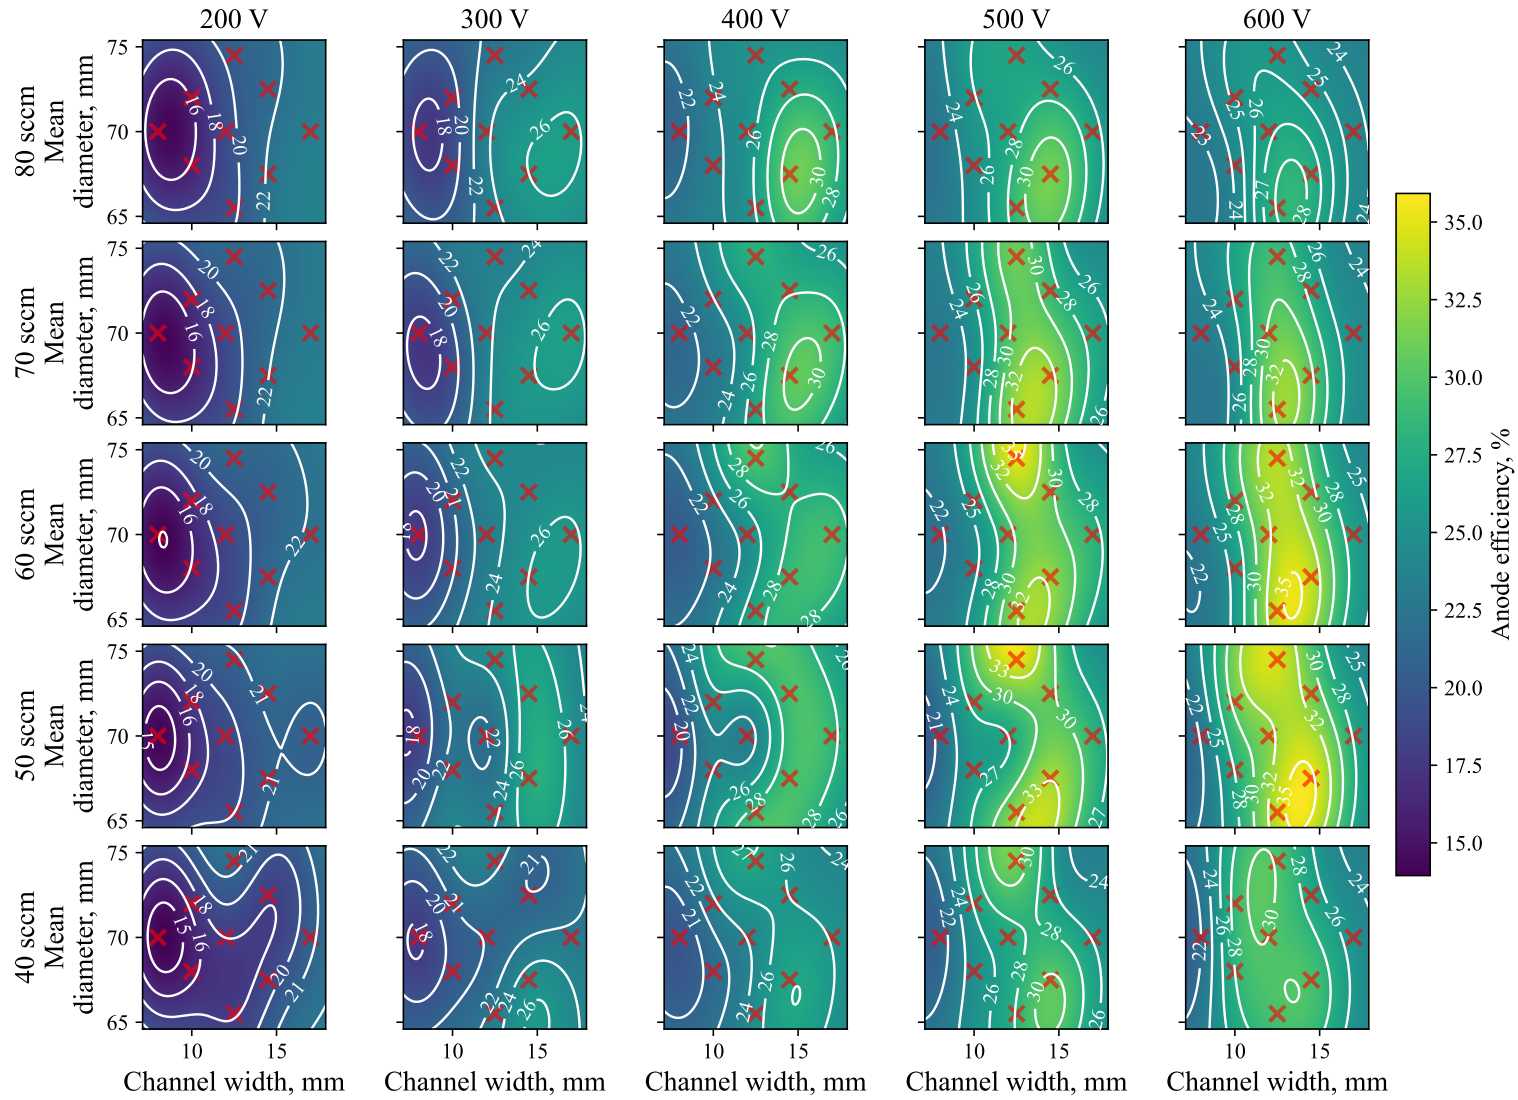


FIGURE 5.72: Array of Gaussian process regression model for krypton, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested.

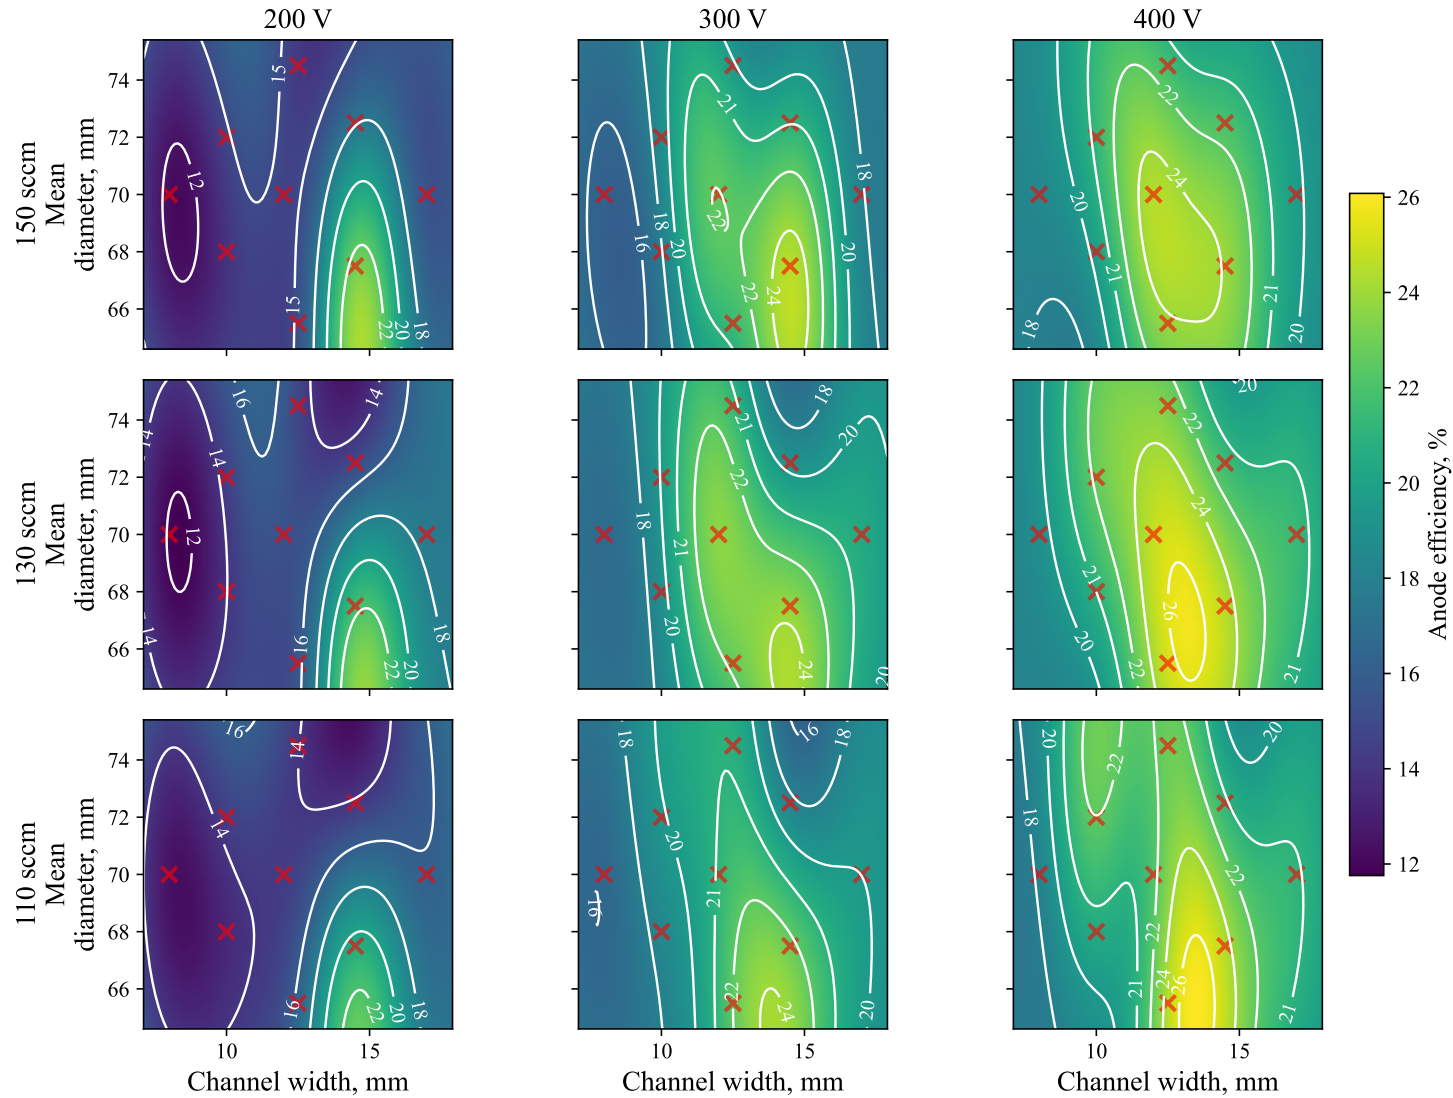


FIGURE 5.73: Array of Gaussian process regression model for argon, showing anode efficiency as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency. Red crosses indicate the discrete channel configurations experimentally tested.

#### 5.11.4 Scaling Implications

Using the Gaussian process regression model, a continuous smooth function that approximately describes the performance of the SHARK-600V thruster was produced. With the addition of a gradient-based optimiser, it becomes possible to search the model parameter space, enabling the evaluation of “optimal” geometries that were not directly tested for xenon, krypton, and argon. However, the ability of the model to extrapolate far beyond the tested ranges is limited, as the GPR uncertainty grows significantly once the predicted point lies far from training data of the experimentally measured values. Furthermore, this method does not account for experimental biases in the data, such as the relatively small test dataset for argon. This limitation is reflected in the high uncertainty at untested values; nevertheless, the optimiser will tend to favour results that lie within the bounds of the dataset.

To compare with the scaling results from Chapter 3 in Sec. 3.6, several operating points are of particular interest. First, the original thesis target: 600 V at 1.35 kW for krypton. Second, the adjusted values based on the updated scaling results with the improved database: 600 V at 2.5 kW on krypton. The same comparison can be made for xenon, and with this dataset, estimates for argon can also be produced—something the scaling methodology could not achieve due to a lack of experimental data. However, the 1.35 kW case for argon carries significant uncertainty, as argon operation was limited to higher powers, with only one tested operating point at 1.2 kW in a single channel configuration. Furthermore, other conditions could be explored, such as 1.35 kW at 300 V, the operating condition of the SPT-100, as this is a well-documented thruster used in the scaling analysis.

##### 1.35 kW and 600 V

To begin this analysis, the GPR model predictions for xenon at 1.35 kW and 600 V are compared with the scaling results in Tab. 5.7. This table presents the optimal channel dimensions and operating conditions identified by the GPR model alongside the scaling results from Chapter 3, with the relative error calculated with respect to the scaling values.

As shown in Tab. 5.7, the optimiser was only able to find a configuration for xenon that satisfied the 1.35 kW and 600 V operating point within the performance dataset of the SHARK-600V. Although xenon was not originally considered in the scaling comparison, the scaling method can still be applied to this configuration with greater confidence due to the larger number of xenon Hall effect thrusters present in the database. Comparing the GPR results with the semi-empirical scaling method for xenon at 1.35 kW and 600 V reveals close agreement. The mean channel diameters are very similar, with an approximate difference of 5 mm between the predicted channel

TABLE 5.7: Comparison of GPR-predicted optimal geometry and performance against the semi-empirical scaling methodology as described in Chapter 3 for the updated database for xenon at 600 V and 1.35 kW. Percentage differences are calculated as relative error of the GPR model with respect to the scaling result. No feasible GPR solutions were found for krypton or argon at this operating point.

Parameter	GPR Optimal	Scaling Results	Rel. Error,%
Geometry and Flow Rates			
Channel width, mm	12.5	7.4	+68.9
Mean channel diameter, mm	65.4	64.3	+1.7
Channel length, mm		6.5	
Anode mass flow rate, mg/s	1.92	2.5	-23.2
Anode mass flow rate, sccm	19.7	25.6	-23.0
Predicted Performance			
Thrust, mN	$40.9 \pm 0.6$	60.4	-32.3
Anode specific impulse, s	$2165.9 \pm 17.8$	2314.4	-6.4
Anode efficiency, %	$31.9 \pm 0.5$	54.5	-41.5
Anode power, W	$1366.1 \pm 11.7$	1350.0	+1.2

widths. This is a notable result, given that the semi-empirical method represents an aggregate of several high-performing xenon thrusters tested across multiple facilities, whereas the GPR model is derived from a single thruster operating in one facility, which yielded below-average anode efficiency. Furthermore, the specific impulse predicted by the two methods is very similar; however, given that specific impulse is strongly driven by discharge voltage, as shown in Eq. 2.18, this close agreement is not unexpected.

### 1.35 kW and 300 V

For the 1.35 kW, 300 V condition, the GPR model was able to identify optimal results for xenon, krypton, and argon within the trained dataset. The GPR predictions are shown in Tab. 5.8, with each column corresponding to a different propellant. As the scaling model cannot be applied to argon due to the lack of experimental data in the literature, Tab. 5.9 presents and compares the results for xenon and krypton only.

Comparison of the GPR-optimised configurations for the 1.35 kW, 300 V condition in Tab. 5.8 with the semi-empirical scaling results in Tab. 5.9 reveals both the capability and the context-dependent nature of the GPR model. For xenon, the GPR model predicts an optimal mean channel diameter of 69.7 mm and channel width of 12.0 mm, whereas the scaling method suggests a larger mean diameter of 90.9 mm and a narrower width of 10.5 mm. For krypton, the GPR model yields a mean diameter of 67.5 mm and width of 14.5 mm, compared to the scaling method's 73.0 mm and 16.9 mm, respectively.

TABLE 5.8: GPR-predicted optimal geometries and performance for xenon, krypton, and argon at 300 V and 1.35 kW.

Parameter	Xenon	Krypton	Argon
Optimal Geometry and Flow Rates			
Channel width, mm	12.0	14.5	14.4
Mean channel diameter, mm	69.7	67.5	72.7
Anode mass flow rate, mg/s	3.6	3.1	2.0
Anode mass flow rate, sccm	37.2	50.0	68.0
GP-Predicted Performance			
Thrust, mN	$49.4 \pm 0.5$	$48.2 \pm 0.6$	$32.9 \pm 3.3$
Anode specific impulse, s	$1383.0 \pm 16.4$	$1578.9 \pm 18.0$	$1570.4 \pm 77.7$
Anode efficiency, %	$25.5 \pm 0.5$	$27.6 \pm 0.5$	$16.9 \pm 1.4$
Anode power, W	$1350.0 \pm 13.7$	$1350.0 \pm 13.0$	$1350.0 \pm 47.9$

TABLE 5.9: Results of the semi-empirical scaling methodology for xenon and krypton at 300 V and 1.35 kW, with relative errors calculated with respect to the GPR-predicted optima shown in Tab. 5.8. Percentage differences are expressed as  $(\text{GPR} - \text{scaling}) / \text{scaling} \times 100\%$ .

Parameter	Xenon		Krypton	
	Scaling	Rel. Error, %	Scaling	Rel. Error, %
Geometry and Flow Rates				
Channel width, mm	10.5	+14.3	16.9	-14.2
Mean channel diameter, mm	90.9	-23.3	73.0	-7.5
Channel length, mm	9.6		12.2	
Anode mass flow rate, mg/s	5.0	-28.0	3.7	-16.2
Anode mass flow rate, sccm	51.2	-27.3	59.4	-15.8
Predicted Performance				
Thrust, mN	85.4	-42.2	73.2	-34.2
Anode specific impulse, s	1636.6	-15.5	1731.4	-8.8
Anode efficiency, %	54.5	-53.2	53.4	-48.3
Anode power, W	1350.0		1350.0	

The differences in predicted performance are more pronounced: the scaling method anticipates anode efficiencies of approximately 54% for both propellants, while the GPR model, trained exclusively on data from the SHARK-600V operating at the University of Southampton, predicts efficiencies of 25.5% for xenon and 27.6% for krypton. This discrepancy underscores a key caveat: the GPR model is inherently limited to the performance envelope demonstrated by a single thruster in a specific facility, which in this case yielded below-average efficiencies. In contrast, the semi-empirical scaling method aggregates data from numerous high-performing xenon and krypton thrusters tested across multiple facilities, providing a more generalised and widely applicable design target.

Nevertheless, the GPR model's strength lies in its ability to identify optimal

geometries within the empirical dataset. While the discrepancy between the dimensions from the GPR model and scaling method is greater for this operational condition, there is still significant agreement, with the channel width within 2 mm of the scaling prediction for both xenon and krypton, implying that the underlying scaling methodology is preserved within the SHARK-600V operation at Southampton.

## 2.5 kW and 600 V

Lastly, the GPR model predictions for xenon, krypton, and argon at 2.5 kW and 600 V are shown in Tab. 5.10. The GPR model was able to identify optima for each propellant, with the caveat that the argon result lies far outside the tested envelope, as the highest tested voltage for argon was 500 V. This results in large model uncertainties. Similar as the 1.35 kW 300 V case the comparison between the GPR-model and the scaling result for xenon and krypton is shown in Tab. 5.11 as the relative error column.

TABLE 5.10: GPR-predicted optimal geometries and performance for xenon, krypton, and argon at 600 V and 2.5 kW. †: Argon exhibits large uncertainty in the GPR model for this operational point due to the flow rate required being much less than the tested range.

Parameter	Xenon	Krypton	Argon <sup>†</sup>
Optimal Geometry and Flow Rates			
Channel width, mm	12.1	10.8	13.7
Mean channel diameter, mm	69.8	70.6	72.8
Anode mass flow rate, mg/s	3.2	2.5	0.7
Anode mass flow rate, sccm	32.8	40.0	23.3
GP-Predicted Performance			
Thrust, mN	$72.4 \pm 0.6$	$60.5 \pm 0.6$	$41.2 \pm 21.7$
Anode specific impulse, s	$2311.5 \pm 17.8$	$2380.4 \pm 19.8$	$2245.1 \pm 455.7$
Anode efficiency, %	$33.5 \pm 0.5$	$29.7 \pm 0.7$	$18.3 \pm 4.0$
Anode power, W	$2500.0 \pm 19.3$	$2500.0 \pm 17.3$	$2500.0 \pm 1001.0$

Comparison of the GPR model “optimised” configurations at 2.5 kW and 600 V in Tab. 5.10 with the semi-empirical scaling results in Tab. 5.11 reveals several notable trends. For xenon, the scaling method predicts an optimal mean channel diameter of 87.5 mm, a value substantially above the tested range of mean diameters for the SHARK-600V (65.5–74.5 mm). This geometry therefore lies outside the domain in which the GPR model, trained exclusively on experimental data from this thruster, is capable of making reliable predictions. The GPR-optimised xenon configuration instead converges on a mean diameter of 69.8 mm, which falls comfortably within the tested range, potentially representing a local maximum accessible within the experimental campaign. For xenon, the channel width predictions show larger

percentage difference: the GPR model predicts a width of 12.1 mm compared to the scaling method's 10.1 mm (+19.8%).

TABLE 5.11: Comparison of semi-empirical scaling methodology results with GPR-predicted optima for xenon and krypton at 600 V and 2.5 kW. Percentage differences are calculated as the relative error of the GPR model with respect to the scaling result. Argon is excluded from this comparison due to the inability of the scaling method to be applied to argon. Percentage differences are expressed as  $(\text{GPR} - \text{scaling}) / \text{scaling} \times 100\%$ .

Parameter	Xenon		Krypton	
	Scaling	Rel. Error, %	Scaling	Rel. Error, %
Geometry and Flow Rates				
Channel width, mm	10.1	+19.8	16.0	-32.5
Mean channel diameter, mm	87.5	-20.2	70.0	+0.9
Channel length, mm	6.5		7.6	
Anode mass flow rate, mg/s	4.6	-30.4	3.4	-26.5
Anode mass flow rate, sccm	47.1	-30.4	54.8	-27.0
Predicted Performance				
Thrust, mN	111.8	-35.2	94.1	-35.7
Anode specific impulse, s	2314.4	-0.1	2398.5	-0.8
Anode efficiency, %	54.5	-38.5	51.9	-42.8
Anode power, W	2500.0		2500.0	

For krypton, the scaling method predicts a mean diameter of 70.0 mm, which is well within the SHARK-600V's tested envelope. Here, the GPR model yields an optimal mean diameter of 70.6 mm, a difference of just 0.6 mm, corresponding to a relative error of +0.9%. This close agreement is striking, particularly given that the thruster was initially scaled for this condition, even if inadvertently. The strong agreement between the scaling result and the GPR-identified optimum suggests that the scaling method, despite its reduced number of input thrusters, produced accurate channel dimensions that have now received experimental validation. However, the channel width prediction for the GPR width of 10.8 mm which differs from the scaling method's 16.0 mm (-32.5%). However, in absolute terms, these differences are modest, only 5 mm, representing small adjustments in the context of physical hardware dimensions.

The predicted performance metrics exhibit larger discrepancies, with the GPR model indicating anode efficiencies of 33.5% for xenon and 29.7% for krypton, compared to the scaling method's 54.5% and 51.9%, respectively. This is expected, as the GPR model reflects the actual performance achieved by the SHARK-600V at Southampton, which was below the state of the art represented by the scaling database. Notably, specific impulses still show strong agreement: the GPR prediction for xenon (2311.5 s) is within 0.1% of the scaling result (2314.4 s), while for krypton the difference is just 0.8%.

These comparisons highlight the complementary nature of the two methods. The semi-empirical scaling methodology aggregates data from numerous high-performing thrusters tested across multiple facilities, providing a generalised design target applicable to a wide range of Hall effect thrusters. However, it may propose geometries that lie beyond the practical design space explored in this experimental campaign. The GPR model, trained exclusively on data from the SHARK-600V operating at the University of Southampton, is constrained to the tested parameter space but offers higher fidelity within that domain, capturing the unique characteristics and performance limitations of this specific thruster and facility.

The closeness of the geometric predictions where comparison is possible, particularly the near-exact agreement for krypton mean diameter, is therefore noteworthy and serves as a form of validation. Despite the fundamental differences in their construction, both methods converge on similar optimal channel dimensions within the overlapping design space. This suggests that the underlying physical relationships captured by the scaling method are reflected in the performance landscape of the SHARK-600V, and conversely, that the GPR model has successfully identified geometries consistent with broader scaling trends. The GPR model thus offers a powerful complement to the scaling method: it provides system-specific optimisation within a constrained experimental envelope, while the scaling method situates these results within the wider context of the field and points toward potentially optimal regions beyond the current design space.

### 5.11.5 Gaussian Process Regression Model Summary

This section has presented a Gaussian process regression model of the SHARK-600V Hall effect thruster, developed to better illustrate and interrogate the highly dimensional nature of Hall effect thruster systems. The model constructs a continuous, smooth function describing thruster performance as a function of four inputs: anode voltage, anode mass flow rate, channel width, and mean channel diameter. Separate models were trained for each propellant (xenon, krypton, and argon) and each output parameter (anode current, thrust, specific impulse, and anode efficiency), together providing a powerful analysis tool capable of predicting thruster behaviour with quantified uncertainty. Validation against held-out test data demonstrated low prediction errors across all outputs, with the anode efficiency exhibiting the largest uncertainty: CV(RMSE) values of 14.8%, 11.9%, and 15.6% for xenon, krypton, and argon, respectively.

The model was subsequently used to compare the scaling results from Chapter 3 with the experimental data obtained from the modular discharge channel of the SHARK-600V. Despite being derived from a single thruster tested in one facility, and therefore more susceptible to facility-specific effects, the GPR model identified optimal

channel dimensions for xenon, krypton, and argon that showed meaningful agreement with the semi-empirical scaling predictions. This agreement is particularly significant for krypton at the 2.5 kW, 600 V condition, where the predicted mean diameters agreed to within 0.6 mm. Such convergence between two fundamentally different methodologies; one a broad empirical aggregate, the other a narrow but high-fidelity surrogate of a single system, serves as a form of experimental validation for the scaling approach, and conversely demonstrates that the GPR model has successfully captured the underlying physical trends encoded in the data.

For argon, the GPR model provides, to the author's knowledge, the first directly "scaled" channel geometry derived from experimental data. The GPR-predicted optima for argon are shown in Tab. 5.8 and Tab. 5.10, with the 2.5 kW, 600 V case suffering from large uncertainty. Traditional semi-empirical scaling methods cannot be applied to argon due to the absence of comprehensive performance data in the literature. While this result is inherently empirical, coming from purely experimental data, and cannot be independently verified, being limited by the small argon dataset, single-facility operation, and the fact that the thruster was not optimised for argon, it nevertheless represents a novel contribution and offers a preliminary design target for future argon Hall thruster development.

## 5.12 Conclusions

This chapter has presented the facilities, auxiliary systems and diagnostics that enabled a systematic experimental characterisation of a modular Hall effect thruster across nine discharge channel geometries and three noble propellants, with direct thrust measurements enabling consistent derivation of specific impulse, anode efficiency, and thrust-to-power. Across all cases, increasing discharge voltage increased specific impulse, while performance and stability depended strongly on channel geometry. Channel width acted as a first order control: wider channels generally reduced wall losses and improved specific impulse and anode efficiency, but excessively wide channels reduced the high voltage operability window. Mean channel diameter provided a second order control that became more influential at high voltage, with smaller mean diameters frequently improving performance, particularly for krypton and argon. Among the tested options, the "+2, +1" and "+2, +0" configurations repeatedly delivered the best high voltage performance, which correspond to channel width of 14.5 mm, and 12.5 mm and mean channel diameters of 67.5 mm and 65.5 mm, respectively. Krypton was capable of exceeding xenon efficiency at higher specific impulse in several high voltage cases, indicating a credible route for alternative propellants when geometry is tuned. Argon achieved equivalent specific impulse to xenon at lower voltage but required higher power and exhibited a narrower stability envelope, with strong sensitivity to geometry. Although

absolute efficiencies were modest, likely due to facility dependent effects, the relative trends were robust. These results provide clear guidance for geometry selection in future designs and motivate targeted diagnostics and lower pressure testing to isolate the dominant loss mechanisms and further improve efficiency with alternative propellants at high voltage.

To identify specific plasma interactions that drive the increase and decrease of efficiency between different channel configurations a comprehensive diagnostic suite is required. Due to the lack of this capability at the University of Southampton this motivated the cross-facility testing of the SHARK-600V Hall effect thruster at the University of Michigan. Due to limited time only the highest performing configuration “+2, +1” operating on krypton was considered for these tests. Here, the thruster was operated purely on krypton across the same 200–600 V anode voltage range, and 40–90 sccm anode mass flow rate range with the same 2 A of current applied to the electromagnets.

Furthermore, the TAL experiments highlight both the promise and the challenges of anode-layer operation. The “TAL-1” configuration did achieve modest specific impulse, but at the cost of severe thermal loading, narrow operating windows, and efficiency well below that of the SPT configuration. Taken together, the findings indicate that TAL should not be regarded as inherently inferior to SPT. Rather, the two approaches are not directly compatible in the same modular architecture of the SHARK-600V and likely require distinct, tailored design solutions. Future iterations must therefore treat TAL as a separate optimisation problem, with refractory materials and configuration-specific tuning.

The Gaussian process regression model developed from the experimental dataset provided a powerful complement to the discrete-geometry performance analysis. By constructing a continuous, smooth function describing thruster performance across the input parameter space, the model enabled interpolation between tested geometries and extrapolation to untested conditions, subject to quantified uncertainty. Validation against withheld test data confirmed low model errors across all propellants and outputs, with the largest uncertainties associated with anode efficiency. Crucially, the optimal geometries identified by the GPR model showed meaningful agreement with the semi-empirical scaling predictions from Chapter 3. This convergence was particularly important for krypton at the 2.5 kW, 600 V condition, the original scaled thruster condition. Where the predicted mean diameters agreed to within 0.6 mm. Such alignment between two fundamentally different methodologies, one a broad empirical aggregate of high-performance thrusters tested at several facilities and the other a narrow but high-fidelity surrogate of a single system, serves as a form of experimental validation for the scaling approach, while simultaneously demonstrating that the GPR model successfully captured the underlying physical trends encoded in the data. For argon, where conventional scaling methods currently

---

cannot be applied due to insufficient literature data, the GPR model provided the first experimentally derived optimal channel geometry. Although this result remains inherently empirical and unverified, constrained by the small argon dataset and single-facility operation, it nonetheless represents a novel contribution and offers a preliminary design target for future argon Hall thruster development.



## Chapter 6

# Michigan Test Campaign

### 6.1 Introduction

Cross-facility testing of Hall thrusters offers a unique opportunity to investigate both the intrinsic behaviour of the thruster and the influence of facility-specific effects. By comparing data collected at multiple sites, it becomes possible to better isolate true thruster physics.

The University of Michigan's testing infrastructure is significantly better in terms of both vacuum chamber pumping and diagnostics capability than the facility available at the University of Southampton. The Michigan facility offers a state-of-the-art suite of plasma diagnostics, including tools that were not accessible during testing at Southampton. Additionally, the vacuum chamber at Michigan the Large Vacuum Test Facility (LVTF) has a significantly lower base pressure and greater pumping capacity. These attributes help reduce background gas interactions, enabling more representative measurements of thruster performance in conditions closer to those expected in space.

This work would not have been possible without the generous support and collaboration of the team at the University of Michigan. In particular, I would like to thank Professor Benjamin Jorns for his support and for hosting the testing campaign, as well as William Hurley, Tate Gill, and Madison Allen for their help during the experiments and to the whole team at Michigan who went out of their way on several occasions to make me feel at home in Ann Arbor.

#### 6.1.1 Funding

Funding for this cross-facility testing was provided by the United Kingdom Space Agency (UKSA) under the International Bi-lateral Fund (IBF) Phase-2 project "US-UK



FIGURE 6.1: Photo of Alex D. Gallimore Large Vacuum Facility taken at the University of Michigan, Ann Arbor (from [81] with permission).

collaboration on the development of Hall thruster spacecraft propulsion.”, specifically under the deliverable “D4.2: Testing of Southampton Thruster at University of Michigan”. This project aimed to build understanding about the legal and financial frameworks enabling future collaboration between international partners.

## 6.2 Experimental Set-up

In this section we will describe the facilities, the thruster configuration tested and the cathode that was used during the testing.

### 6.2.1 Large Vacuum Test Facility

The Alex D. Gallimore Large Vacuum Test Facility (LVTF) is a 6 m long, 9 m diameter vacuum chamber located at the University of Michigan in Ann Arbor, Michigan, USA illustrated in Fig. 6.1. The chamber was originally constructed in 1961 for the Bendix Corporation to support lunar rover testing. It was donated to the University of Michigan in 1982, where it remained unused until Professor Gallimore revitalised the facility, going on to found Plasmadynamics and Electric Propulsion Laboratory (PEPL) in 1992, transforming it into a state-of-the-art electric propulsion laboratory [205]. Today, the LVTF is the largest vacuum chamber testing facility housed within an academic institution in the United States.

The LVTF is equipped with 18 cryogenic panels, comprising both commercial and in-house solutions, and is capable of pumping approximately 500 kL/s of xenon and




FIGURE 6.2: The SHARK-600V thruster in “+2, +1” channel configurations mounted in the LVTF at the University of Michigan prior to testing.

600 kL/s of krypton. A detailed description of the chamber and its capabilities is available in [81, 206].

### 6.2.2 Thruster Configuration

The thruster configuration chosen for the cross-facility test was the “+2, +1” channel configurations that corresponded to a mean diameter of 67.5 mm and a channel width of 14.5 mm, as can be seen in Tab. 6.1. This configuration was chosen due to the largest operation envelope for krypton seen during the testing undertaken at the University of Southampton. The thruster mounted on the thrust balance at Michigan prior to testing can be seen in Fig. 6.2.

TABLE 6.1: Channel width, mm and mean channel diameter, mm for the possible SPT configuration. Here the chosen configurations for the University of Michigan campaign can be seen in **bold** and red.

Outer	Inner		
	SPT +0	SPT +1	<b>SPT +2</b>
SPT +0	8, 70	10, 68	12.5, 65.5
<b>SPT +1</b>	10, 72	12, 70	<b>14.5, 67.5</b>
SPT +2	12.5, 74.5	14.5, 72.5	17, 70

### 6.2.3 Cathode

To maintain consistency between the two test facilities, the same cathode described in Sec. 5.3 was initially used. However, early attempts to ignite this cathode were unsuccessful and ultimately resulted in a short between the insert and the keeper.

Fortunately, the team at Michigan provided the use of their Lanthanum Hexaboride ( $\text{LaB}_6$ ) hollow cathode, as described in [168, 207] and illustrated in Fig. 6.3. While the cathode shown is not the exact unit used during testing, it follows the same design. However, cathode location at Michigan emulates the Southampton configuration to minimise the changes to the testing setup.




FIGURE 6.3: Cathode design used during the Michigan test campaign (image from [168], used with permission).

During testing, the cathode was operated at a constant flow rate of 10.94 sccm, equivalent to 10 sccm<sub>USA</sub>, where sccm<sub>USA</sub> is defined at 25 °C rather than the 0 °C standard used internationally. Throughout this thesis, unless otherwise stated, sccm values are defined using the 0 °C standard. The cathode keeper was operated at a constant 18 V and 1.5 A for all tests. In-comparison, at Southampton the cathode keeper was operated at 25–50 V and 2 A.

This change in flow rate definition will affect the reported performance of the thruster; however, it is expected to result in a slight enhancement due to lower losses associated with the cathode-to-plasma coupling voltage. This would likely yield a 1–2% increase in voltage efficiency. Furthermore, the higher mass flow rate through the cathode might suggest enhanced performance from additional neutral ingestion. However, the cathode was positioned at an angle behind the thruster such that any non-ionised cathode flow would be exhausted away from the channel. Consequently, any neutral ingestion would occur from the background pressure, which was an order of magnitude lower at the Michigan test site than at Southampton. As a result, this would not contribute to improved performance.

## 6.3 Diagnostics

This section outlines the various diagnostics employed during the Michigan testing campaign. For each plasma diagnostic tool, the underlying theory and corresponding analysis methods are also discussed in detail.

### 6.3.1 Thrust Balance

The thrust balance used during testing at the LVTF shares basic similarities with the double inverted pendulum balance described in Sec. 5.6.2, but there are several key differences. The inverted pendulum thrust balance employed at Michigan is a non-displacing design that uses electromagnets to maintain zero displacement while under load from the thruster. As a result, thrust is not measured using a calibration coefficient in units of  $\text{mN}/\mu\text{m}$ , but rather through the calibrated coil current required to maintain zero displacement, with a calibration coefficient in  $\text{mN}/\text{A}$ .

Additionally, the Michigan thrust balance incorporates a water-cooled jacket maintained at  $20^\circ\text{C}$  to minimise thermal drift, thereby improving measurement stability.

A further refinement in this system is the inclusion of an active inclination control system. The thruster shifting on the balance can cause thruster mass to be projected into the thrust measurement axis, introducing erroneous thrust readings. To correct for this, the inclination of the thrust stand's mounting plate is continuously monitored using inclinometers. Any deviations from level are corrected in real time using stepper motors, which adjust the plate orientation to maintain horizontal alignment throughout the test.

For a comprehensive description of the thrust stand design, measurement procedures, error sources, and correction techniques, the reader is referred to Su [81].

### 6.3.2 $E \times B$ Probe

The analysis of charge-to-mass ratios in the far-field plume of a Hall effect thruster is typically conducted using an " $E \times B$ " probe, also known as a Wien filter. This diagnostic operates as a velocity bandpass filter, employing a region of orthogonal electric and magnetic fields to selectively transmit ions of a specific velocity; thus, for a simple mono-atomic propellant, enabling discrimination based on ion charge state. For specific details of the  $E \times B$  probe used within these tests the reader is referred to Su [81].

The use of  $E \times B$  probes for electric propulsion diagnostics dates back to the 1970s, where they were employed to characterise ion charge state distributions within thruster plumes [127]. Ion species with higher charge states are accelerated to higher velocities through the same discharge potential. This behaviour is described by a charge-state-specific form of the ion velocity equation

$$u_{i_n} = \sqrt{\frac{2enU_{b_n}}{m_i}}, \quad (6.1)$$

where  $n$  denotes the charge state,  $u_{i_n}$  is the velocity of an ion at the charge state  $n$ ,  $U_{b_n}$  is the species-specific accelerating voltage, and  $m_i$  is the ion mass. In practice, ions experience a distribution of accelerating voltages rather than a singular, uniform value. It is from Eq. 6.1 that it can be seen that the  $E \times B$  can be used as a mass spectrometer under the assumption of constant acceleration potential that will allow for the determination of the charge efficiency, Eq. 2.26, of the thruster.

A schematic of an  $E \times B$  probe structure is shown in Fig. 6.4. A typical Wien filter consists of four principal components: a collimator, the  $E \times B$  region, a drift tube, and a collector plate. The collimator ensures that only ions travelling approximately parallel to the probe axis enter the filtering region, thereby excluding off-axis charge-exchange ions or plume contributions from other thruster components. The drift tube further narrows the angular acceptance by filtering ions not aligned axially after the  $E \times B$  region.




FIGURE 6.4: Schematic of a Wien filter-style  $E \times B$  probe, illustrating the orthogonal electric and magnetic fields and ion trajectories.

The filtering action is produced by a region where a uniform magnetic field is crossed with a variable electric field. The magnetic field, oriented perpendicular to the probe axis, is typically generated by permanent magnets for simplicity and stability. In some cases, electromagnets are used to allow dynamic adjustment of the field, though this adds design complexity. However, the use of electromagnets avoids the thermal management concerns due to the risk of demagnetisation in permanent magnets at

elevated temperatures. The electric field is generated between two parallel bias plates separated by a known distance  $d$ , across which a voltage  $V_{\text{plate}}$  is applied.

The condition for ion transmission through the Wien filter is given by the so-called Wien velocity,  $v_w$ , which defines the ion velocity for which the electric and magnetic forces balance

$$v_w = \frac{|E|}{|B|} = \frac{\Delta\phi_{\text{plate}}}{d|B|} , \quad (6.2)$$

where  $\Delta\phi_{\text{plate}} = \phi_1 - \phi_2$  is the voltage difference applied across the bias plates. During probe operation,  $\Delta\phi_{\text{plate}}$  is swept over a voltage range, thereby scanning through a spectrum of  $v_w$  values. Ions with velocity matching  $v_w$  traverse the crossed-field region undeflected and strike the collector plate, where they are neutralised by electrons.

The current resulting from this neutralisation process is measured using a high sensitivity picoammeter, which forms the probe's output signal. For a typical  $E \times B$  probe, this signal peaks are on the order of 0.1 nA. As a result, careful alignment of the probe with respect to the thruster is critical if a strong enough signal is to be collected for meaningful measurements. Peaks in the collected current correspond to specific charge-to-mass ratios, enabling the identification of ion species and the determination of their respective charge state distributions.

As the plate bias voltage,  $\Delta\phi_{\text{plate}}$ , is swept, the current collected by the Wien filter changes as the Wien velocity is swept across the ion velocity distribution function's (IVDF) range.

Due to overlapping peaks in the current measurements from the probe, each corresponding to a different species' ion energy distribution function, it is necessary to make assumptions about the form of the ion velocity distribution function. Since the ion mass remains constant within each species, the ion energy distribution function's (IEDF) are directly analogous to the IVDFs. In the literature, several fitting models have been employed, including triangular, Gaussian, twin-Gaussian, and skew-normal profiles [208]. For the analysis presented in this thesis, a twin-Gaussian fitting approach was used, as previous work has shown it best captures the high-velocity tails typically observed in Hall thruster IVDFs [208].

To perform this analysis, the peak corresponding to the singly charged ion was first fitted using the following expression

$$f_i(u) = a_1 \exp[-b_1(u - c_1)] + a_2 \exp[-b_2(u - c_2)] , \quad (6.3)$$

where  $a_1$ ,  $b_1$ ,  $c_1$ ,  $a_2$ ,  $b_2$ , and  $c_2$  are fitting parameters.

Once the first peak, typically the largest under nominal operating conditions, was fitted, the remaining residual signal was successively fitted for the second and third

peaks. This fitting procedure was applied to identify contributions from  $\text{Kr}^+$ ,  $\text{Kr}^{2+}$ , and  $\text{Kr}^{3+}$  ion species for each measurement. An illustration of this process is provided in Fig. 6.5, corresponding to an operating condition of 300 V anode voltage and 90 sccm anode flow rate.

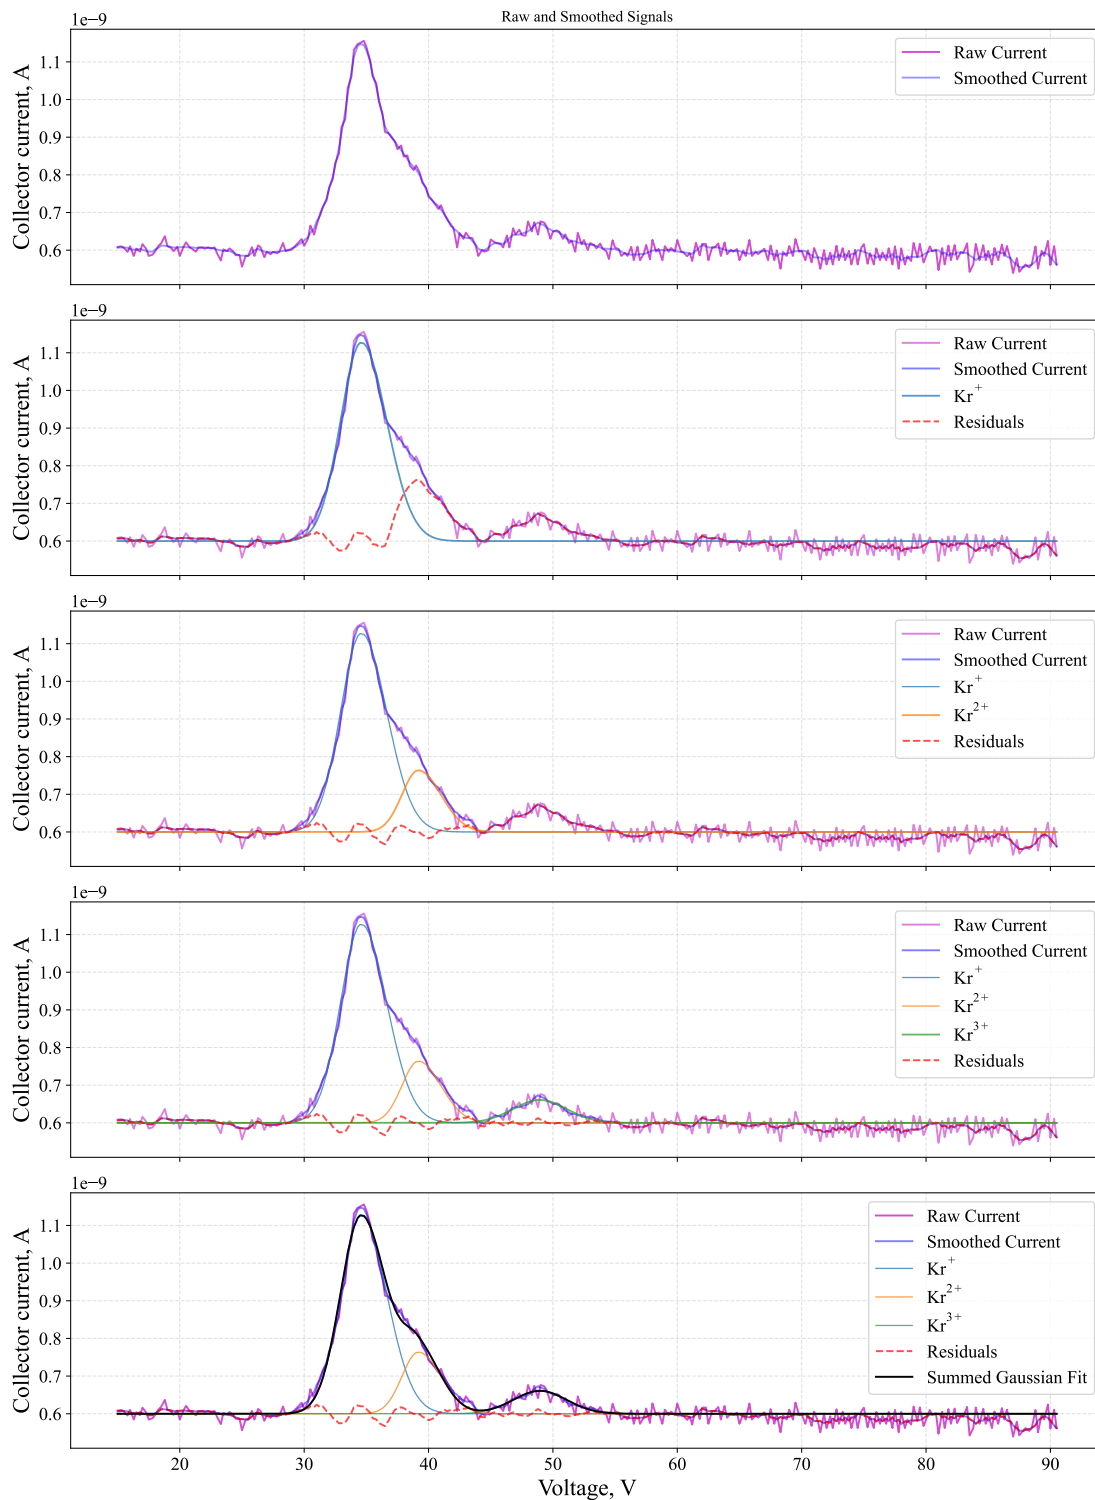


FIGURE 6.5: Step-by-step illustration of the procedure for fitting multiple twin-Gaussian distributions to a representative probe signal. The example corresponds to an operating point of 300 V anode voltage and 90 sccm krypton anode mass flow rate.



The current fractions,  $\Omega_{n,raw}$  (Eq. 2.27), representing the fraction of total current attributed to each ion species  $n$ , were estimated by integrating the fitted twin-Gaussian distribution for each charge state and normalising by the integral of the total fitted signal. These values correspond to uncorrected charge state fractions, as facility effects may influence the measured distribution. During testing in the LVTF, the background pressure was approximately  $5.33 \times 10^{-6}$  mbar with 60 sccm of krypton flow into the chamber.

### 6.3.2.1 E×B Probe: Charge-Exchange Correction

Several collisional effects can alter the measured charge state ratios in the spectra obtained from a Wien filter. Reducing the distance between the probe and the thruster exit plane can help mitigate these effects. However, proximity to the thruster introduces other complications, such as probe heating due to the high-energy plume, which may introduce additional sources of error or even damage the probe.

Prior studies have established that charge-exchange (CEX) collisions within the plume have cross-sections comparable to Coulomb collisions yet exert a significantly greater influence on ion attenuation [208]. CEX interactions predominantly occur between fast ions and background thermal neutrals, and their prevalence is strongly dependent on the vacuum chamber pressure. Consequently, they are often treated as facility effects.

The dominant CEX processes involve fast ions exchanging electrons with slow-moving neutrals and can be described by the following interactions



The background charge-exchange correction method described by Shastry et al. [209] is applied here. This approach corrects the measured current based on the attenuation of high-velocity ion fluxes propagating through a thermal background gas

$$(J/J_0)_n = \exp(-n_0\sigma_n z) , \quad (6.7)$$

where  $J$  is the measured current of the ion species with charge  $n$ ,  $J_0$  is the unattenuated current at the thruster exit plane,  $n_0$  is the background neutral number density

(estimated from chamber pressure),  $\sigma_n$  is the effective CEX cross-sectional area for the  $n^{\text{th}}$  charge state, and  $z$  is the distance between the thruster exit plane and the probe.




FIGURE 6.6: Charge exchange ion correction factor versus acceleration voltage for each ion species considered. A background pressure of  $5.3 \times 10^{-6}$  mbar is assumed, with the shaded region showing  $\pm 25\%$  pressure variation. As can be seen in Eq. 6.7, a higher pressure will result in a lower  $(j/j_0)_n$  value.

For this analysis, the CEX cross sections for  $\text{Kr}^+$  and  $\text{Kr}^{2+}$  are taken from semi-empirical measurements by Hause et al. [210]

$$\sigma_1 = (80.7 - 14.7 \log(V_1)) \times 10^{-20} , \quad (6.8)$$

$$\sigma_2 = (44.6 - 9.8 \log(2V_2)) \times 10^{-20} , \quad (6.9)$$

where  $V_1$  and  $V_2$  are the species-specific acceleration voltages for  $\text{Kr}^+$  and  $\text{Kr}^{2+}$ , respectively. For the analysis here it is assumed that each ion species experiences the same acceleration voltage which is obtained from the retarding potential analyser and will be covered in more detail in Sec. 6.3.4.

No empirical data for  $\text{Kr}^{3+}$  CEX cross sections are available. As recommended by Su [81], the corresponding value for  $\text{Xe}^{3+}$  is used as a surrogate

$$\sigma_3|_{\text{Kr}} \approx \sigma_3|_{\text{Xe}} = (16.9 - 3.0 \log(3V_3)) \times 10^{-20} , \quad (6.10)$$

Given the relatively small contribution of  $\text{Kr}^{3+}$  to the overall ion population, the associated error introduced by this approximation is expected to be minimal.

### 6.3.2.2 $E \times B$ Probe: Current Fractions

To obtain the corrected current fractions,  $\Omega_n$ , of each species in the plume, the raw current fraction,  $\Omega_{n,raw}$ , is corrected using Eq. 6.7 as follows

$$\Omega_n = \frac{\Omega_{n,raw} (J/J_0)_n^{-1}}{\sum_n (\Omega_{n,raw} (J/J_0)_n^{-1})}, \quad (6.11)$$

where  $\Omega_n$  is the corrected species-specific current fraction. The expected value of each correction factor  $(J/J_0)_n$  is illustrated in Fig. 6.6, for an assumed background pressure of  $5.3 \times 10^{-6}$  mbar. The shaded region shows a variation of  $\pm 25\%$  in the assumed pressure.

### 6.3.3 Faraday Probe

Faraday probe measurements of the plasma plume provide the spatial distribution of ion current density. These measurements are crucial for evaluating two key aspects of thruster performance: the total plume current, used in calculating current efficiency (Eq. 2.25), and the spatial distribution of the beam, relevant to beam divergence efficiency (Eq. 2.31).

A Faraday probe typically consists of two main components: an ion-collecting electrode, biased sufficiently to draw ion saturation current, and a guard ring held at the same potential. The guard ring suppresses edge effects by minimising potential gradients at the boundary, ensuring a well-defined and constant effective collecting area. An example of a typical Faraday probe is illustrated in Fig. 6.7.

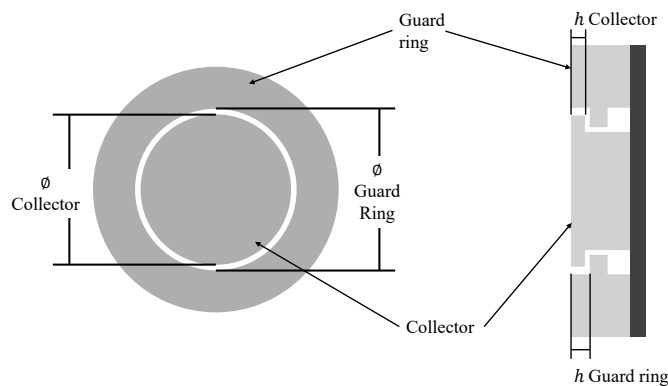


FIGURE 6.7: An illustration of a Faraday probe with important dimensions required for analysis labelled.

The probe is mounted on a rotating sweep arm that performs an azimuthal scan of the plume at a fixed radial distance from the thruster exit plane. During each sweep, the

probe collects time-averaged current readings at different angular positions relative to the thruster centreline. In the Michigan LVTF setup, the Faraday probe is biased to  $-30$  V and swept from  $-90^\circ$  to  $+90^\circ$  in approximately  $1^\circ$  increments, where an angle of  $0^\circ$  represents the thruster centreline. This allows the spatial current density profile to be reconstructed.

From the resulting distribution, the total beam current and plume divergence angle can be derived, offering insight into both propulsion efficiency and beam collimation. The current collected by the probe as a function of angle is converted into a current density using the relation defined by Brown et al. [211]

$$j(\theta_p) = \frac{I_c(\theta_p)}{A_c + \kappa_G} \kappa_{\text{SEE}} , \quad (6.12)$$

where  $j(\theta_p)$  is the current density at sweep angle  $\theta_p$ ,  $I_c$  is the measured collector current,  $A_c$  is the area of the collector plate,  $\kappa_G$  is a geometric correction factor accounting for ions collected in the gap between the collector and the guard ring, and  $\kappa_{\text{SEE}}$  is a correction factor accounting for secondary electron emission from ion impacts.

The correction factor  $\kappa_G$  is geometry-specific and will not be redefined here for the sake of brevity and can be found in Brown et al. [211]. The SEE correction factor,  $\kappa_{\text{SEE}}$ , accounts for the electrons emitted from the collector plate due to impinging high-energy ions. It has separate contributions for each ion species and is defined as

$$\kappa_{\text{SEE}} = \frac{1}{1 + \sum_n \frac{\Omega_n \gamma_n}{Z_n}} , \quad (6.13)$$

where  $\gamma_n$  is the SEE yield of the collector plate material for each ion species. The Faraday probe used in these tests employs a molybdenum collector plate, and the SEE yields used in the analysis are shown in Tab. 6.2. As previously stated, only ion charge states up to  $+3$  are considered. The ion charge fractions,  $\Omega_n$ , obtained from the  $\mathbf{E} \times \mathbf{B}$  probe measurements are assumed invariant with respect to the probe angle relative to the thruster axis.

TABLE 6.2: Secondary electron emission (SEE) yields for molybdenum for impinging krypton ions. Values for  $\text{Kr}^+$  and  $\text{Kr}^{2+}$  from [212]. ‡: No data for  $\text{Kr}^{3+}$ ; 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 krypton impinging on molybdenum. [213].

Bombarding Ion	$Z_n$	SEE Yield of Molybdenum, $\gamma_n$
$\text{Kr}^+$	1	0.069
$\text{Kr}^{2+}$	2	0.30
$\text{Kr}^{3+}$	3	0.87‡

To calculate the total beam current and beam divergence, the current measurements from the Faraday probe must be integrated with respect to the probe position and the thruster's central axis. This yields a total beam current, as described by Brown et al. [211], in the form

$$I_b = 2\pi R^2 \int_0^{\pi/2} j(\theta_p) \frac{\kappa_D}{\kappa_A} \sin(\theta_p) d\theta_p , \quad (6.14)$$

where  $I_b$  is the total beam current calculated over the hemisphere from the thruster centreline (0 rad) to the orthogonal axis ( $\pi/2$  rad),  $R$  is the radial distance of the probe from the thruster's exit plane centre,  $\kappa_D$  accounts for differences in path length from near and far channel exit points to the probe, and  $\kappa_A$  is a geometric correction factor that accounts for the probe's angular misalignment with respect to the thruster axis. A full derivation of these correction factors is omitted here for brevity but can be found in Brown et al. [211] and Tisaev [214].

It is worth noting that Eq. 6.14 corresponds to a  $90^\circ$  segment of the measurement arc. Therefore, a full beam current estimate over the forward hemisphere can be obtained by evaluating the expression for both halves of the  $180^\circ$  sweep. However, in the analysis presented in this thesis, only one half of the sweep was used due to cathode placement obstructing part of the measurement on the opposite side.

To compute the beam divergence angle, as defined in Eq. 2.32, an estimate of the axial component of the beam current is required. This is obtained by modifying Eq. 6.14 to isolate only the axial component

$$I_{\text{axial}} = 2\pi R^2 \int_0^{\pi/2} j(\theta_p) \frac{\kappa_D}{\kappa_A} \cos(\alpha_A(\theta_p)) \sin(\theta_p) d\theta_p , \quad (6.15)$$

where  $\alpha_A(\theta_p)$  is the angle between the direction of ion flux and the thruster's axial direction, which depends on the probe radius  $R$ , sweep angle  $\theta_p$ , and the outer diameter of the thruster channel [211].

### 6.3.3.1 Faraday Probe: Charge-Exchange Correction

As previously discussed in Sec. 6.3.2.1, attenuation and interactions between plume ions and background neutrals affect the measurements of the far-field plume and must be corrected for. The influence of charge-exchange ions on Faraday probe measurements is well known and has an impact several orders of magnitude greater than that of other ion-electron or ion-ion interactions [81, 211, 215]. Furthermore, the correction to remove slow-moving ions at the fringes is required for accurate estimations of beam current, as failure to do so would produce non-physical beam currents.

There are other interactions between ions that can occur in the plume of a Hall thruster, such as ion-ion Coulomb collisions. However, Azziz estimated the mean free




FIGURE 6.8: Faraday probe measurements of the SHARK-600V plume at the University of Michigan for 90 sccm of krypton and an anode voltage of 300 V. Both flat-subtraction and exponential fitting corrections are shown. The plume angle shown is calculated using the exponential fitting correction and Eq. 2.32.

paths of ion-ion and ion-electron Coulomb collisions at one metre from the thruster's exit plane to be on the order of kilometres and hundreds of millions of metres, respectively [140]. These interactions, therefore, do not contribute significantly to plume divergence and can be reasonably neglected in this context.

There are several methods for accounting for the effects of CEX on Faraday probe readings. The most rudimentary of these is the "flat subtraction method", which assumes that the measured current at  $90^\circ$  from the thruster centreline corresponds to zero directed ion current. Any signal measured at this angle is attributed to CEX ions and is subtracted from the entire profile. This correction assumes that the CEX contribution is uniform with respect to angle and does not vary across the plume. While this method has been used in the past, it does not capture the true angular distribution or dynamics of the ion population.

A more refined method of accounting for CEX is derived from experimental observations and plume simulations, which showed an exponential decay of the current beginning between  $5\text{--}10^\circ$  off the centreline. Manzella and Sankovic [216] also observed this exponential trend in the current from between  $5\text{--}10^\circ$  to  $20\text{--}30^\circ$ , with the current measured beyond  $20\text{--}30^\circ$  being primarily attributable to background pressure. Higher background pressures result in elevated current measurements at larger angles. Based on this behaviour, it has been proposed that an exponential fitting correction be applied to the "exponential region" (typically  $5\text{--}10^\circ$  to  $20\text{--}30^\circ$ ) to extrapolate and subtract the CEX contribution in the outer angular regions of the plume.

Fig. 6.8 presents the raw Faraday probe data, alongside results corrected using both the flat-subtraction and exponential fitting methods, for an operating point of 90 sccm krypton and 300 V anode voltage. The impact of cathode placement is clearly visible: the positive angular side of the trace shows suppressed current due to blockage by the cathode, leading the exponential fit to predict a higher current than is actually measured. As a result, all subsequent discussion of Faraday probe results will consider only the left-hand (LH) side of the angular sweep with the exponential fit correction applied, as this region remains unaffected by the cathode obstruction.

### 6.3.4 Retarding Potential Analyser

A retarding potential analyser (RPA) is a gridded energy-to-charge ratio filter used as a far-field plasma diagnostic. In its simplest form, an RPA consists of a gridded electrode and a collector plate. The grid is biased to a specified voltage that retards and rejects ions below a certain energy threshold, while ions with sufficient energy pass through and are measured via the neutralising current they induce upon striking the collector plate.

In practical implementations, RPAs typically consist of four grids, though some designs may include more. The first (plasma-facing) grid is electrically floating and serves to shield the internal bias voltages from perturbing the local plasma. The second grid is negatively biased to reject incoming electrons. The third grid is positively biased and acts as the energy-charge-ratio discriminator. The fourth grid is negatively biased relative to the collector plate to repel secondary electrons and prevent them from escaping the collector, thereby improving measurement fidelity.

An RPA will only transmit ions whose energy-to-charge ratio exceeds the retarding potential set by the discriminator grid. When analysing ions of multiple charge states, an RPA cannot explicitly distinguish between them. This limitation arises because, as indicated by Eq. 6.1, ions of different charge states but the same energy per charge will be retarded identically and therefore appear indistinguishable in the energy spectrum. Nevertheless, experimental studies have shown that higher charge state ions tend to be generated in different regions of the discharge and thus may experience differing acceleration potentials. This results in small variations in the mean energy of different ion species and broadening of the measured RPA signal [217].

Due to the Hall thruster plume containing ions of multiple charge states and the associated variations in acceleration potentials, the RPA spectra are broadened. Each charge state is associated with its own ion energy distribution and mean acceleration energy, which together contribute to the overall spectral shape at the collector plate. This effect becomes more significant when a large fraction of the ions are multiply charged.




FIGURE 6.9: An illustration of the retarding potential analyser measurement for 300 V and 90 sccm anode voltage and flow respectively and the resulting derivative,  $-dC/dV$ , of the signal, with the most probable voltage indicated.

The signal collected by the RPA is the current measured at the collector as a function of the retarding grid voltage. As the retarding bias is increased, ions with progressively higher energy-to-charge ratios are rejected, resulting in a corresponding decrease in the measured current. The derivative of this signal yields the ion energy distribution function. Fig. 6.9 presents an example of this measurement for an operating point of 300 V anode voltage and 90 sccm anode flow, along with its corresponding derivative,  $-dC/dV$ , which approximates the IEDF.

In the IEDF shown in Fig. 6.9, multiple peaks and valleys are observed below a discriminator grid voltage of approximately 250 V. These features are artifacts caused by the auto-ranging behaviour of the picoammeter and are non-physical. As such, they limit the resolution and interpretability of the IEDF in this region. Nevertheless, the most probable energy of the ion population remains discernible and representative. This inferred acceleration voltage is used to estimate the voltage utilisation efficiency, as defined in Eq. 2.30.

## 6.4 Experimental Results

This section presents the results of far-field probe and direct thrust measurements of the SHARK-600V thruster operating on krypton in the Large Vacuum Test Facility (LVTF) at the University of Michigan. The anode efficiency computed from probe data is compared with that derived from direct thrust measurements. Subsequently, the thrust and anode efficiency results are compared to those obtained from the same

thruster configuration operating at the University of Southampton (Soton), followed by a discussion of the observed discrepancies and potential causes.

The thruster was operated using the power supplies and flow controllers provided at the University of Michigan. Care was taken to ensure the same connection points and wiring configuration for the electromagnets as were used during the Southampton campaign.




FIGURE 6.10: The SHARK-600V operating at 300 V anode voltage and 60 sccm anode mass flow rate of krypton in the LVTF at the University of Michigan.

#### 6.4.1 Probe Analysis

Each probe provides measurements that are converted into individual process efficiencies as described in Sec. 2.3.3. These are then combined to produce a probe-based estimate of the anode efficiency, as outlined in Eq. 2.24. The resulting individual efficiencies and the total anode efficiencies from both probe measurements and thrust measurements for 50 sccm are illustrated in Fig. 6.11.

It can be seen that the largest loss in efficiency arises from the mass efficiency (Eq. 2.28), which itself contains contributions from current efficiency and average ion charge state. This result is not unexpected: when operating on krypton, its smaller ionisation cross-section, compared to xenon, means a greater portion of the neutral “mass current” is not converted into ion current.




FIGURE 6.11: Comparison of individual probe anode efficiency, multiplicative probe anode efficiency, and thrust anode efficiency for 50 sccm of krypton.

The probe readings for 600 V are incomplete due to visible heating of the channel. However, based on the trends and the measured thrust, it is reasonable to suspect that the thruster would exhibit high voltage efficiency and high mass efficiency. Moreover, it would be reasonable to expect a marginally lower charge efficiency due to the elevated plasma temperature, which may result in a greater proportion of multiply charged ion species. Nevertheless, the charge efficiency would still be expected to remain similar in magnitude to that of 500 V, at  $\eta_q = 97.3\%$ .

Furthermore, the Faraday probe results for 200 V at 50 sccm suggest an abnormally high beam efficiency, which results in a high anode efficiency. Whilst the anode efficiency for 200 V obtained from the probes is higher than that derived from direct thrust measurements, the discrepancy is relatively small. As a result, the high anode efficiency from the probes likely arises from an error in the beam current correction for that operating point; at low voltages, the charge exchange correction is more sensitive to background pressure measurements. Additionally, the  $E \times B$  probe measurement is taken at a single axial location due to low signal at off-angles, meaning that the secondary electron emission (SEE) correction on which the Faraday probe relies assumes a uniform charge species distribution. As a result, the corrections could be overestimating the measured Faraday current. Further discussion of estimated probe errors can be found in Appendix C.

#### 6.4.2 Charge Efficiency

Charge efficiency quantifies the loss associated with the presence of higher charged ion species in the plume, since their contribution to plume kinetic energy scales with




FIGURE 6.12: Comparison of each “ $E \times B$ ” probe trace at 50 sccm for anode voltages ranging from 200 V to 500 V.

the square root of the charge, while the energy required to produce them increases linearly. This efficiency is determined using the “ $E \times B$ ” probe, as described in Sec. 6.3.2. An example comparison of these traces is shown in Fig. 6.12. Analysis of the probe’s current trace allows reconstruction of ion species populations, as illustrated in Fig. 6.13.




FIGURE 6.13: Resultant ion current fractions expressed as a percentage versus anode mass flow rate of krypton from the  $E \times B$  probe.

The ion species fractions shown in Fig. 6.13 exhibit several notable trends. Firstly, the Kr<sup>3+</sup> population remains nearly constant at approximately 10 % across all operating conditions, suggesting that its formation mechanism is largely independent of the bulk plasma. This result has been seen in literature across a range of input powers the proportion of Kr<sup>3+</sup> is relatively constant [218]. This is likely the result of CEX

interactions or Coulomb collisions occurring in the near-field plume as higher charge species are formed closer the channel exit plane [219, 220].




FIGURE 6.14: Average charge of axial krypton ions against anode mass flow rate in the plume of the SHARK-600V, as measured by the  $E \times B$  probe and calculated using Eq. 6.16.

The  $\text{Kr}^{2+}$  population shows a positive correlation with both increasing anode mass flow rate and discharge voltage. This is expected, as higher discharge voltages lead to higher bulk plasma temperatures, enhancing ionisation. Furthermore, at elevated flow rates, the neutral number density within the channel increases, further facilitating ionisation, as described by Eq. 2.7.4.

A useful way to summarise the effect of ion populations is through the average charge state of the ions in the plume. This parameter,  $\bar{q}$ , is a key quantity in many of the equations presented in Chapter 2 and is defined as

$$\bar{q} = \frac{\sum_n e Z_n \Omega_n}{\sum_n \Omega_n} = e \sum_n Z_n \Omega_n . \quad (6.16)$$

As a result of  $\sum_n \Omega_n \rightarrow 1$  the resulting average charge values, computed using Eq. 6.16, are shown in Fig. 6.14. A clear trend is observed: average charge increases with both anode mass flow rate and discharge voltage. These increases are primarily driven by the increased  $\text{Kr}^{2+}$  current fraction.

The final charge efficiency, as calculated from Eq. 2.26 using Eq. 6.11, is shown in Fig. 6.15 as a function of anode power. A weak negative correlation is observed between charge efficiency and increasing power. However, it is important to note that for 300 V, higher power values correspond to higher flow rates, whereas for the higher-voltage cases, higher powers are achieved at lower flow rates.




FIGURE 6.15: Charge efficiency against anode power grouped by discharge voltage of the SHARK-600V operating on krypton in the LVTF.

This indicates that high charge efficiency can be preserved at high voltages, provided that the mass flow rate remains suitably low. Therefore, low-current, high-voltage operation is capable of maintaining high charge efficiency, as long as an appropriate mass flow rate is maintained.

### 6.4.3 Beam Efficiency

Beam efficiency, or current efficiency, is a measure of the proportion of electrical current delivered to the anode that is converted into useful ion current within the plume. This is described in Eq. 2.25 as the ratio of the beam current to the anode current. The beam current can only be accurately obtained through plume measurements; here, it is measured using Faraday probe data as described in Sec. 6.3.3.

The constructed beam current density profiles for the SHARK-600V thruster at 300 V, across flow rates from 40 sccm to 90 sccm, are shown in Fig. 6.16. As illustrated, increasing flow rate results in a rising peak current density along the centreline from 40 sccm to 60 sccm. This is followed by a reduction in peak current density and a broadening of the profile at 70 sccm, with an increasing current density across the full angular profile from 70 sccm to 90 sccm. This trend is likely due to increased flow rate raising the pressure within the channel, resulting in the plasma moving closer to the exit plane of the thruster at 70 sccm. Beyond this point, further increases in flow rate do not significantly alter plasma confinement but instead enhance the ionisation rate of the less well-confined plasma.




FIGURE 6.16: Faraday sweeps for all anode flow rates tested at 300 V. Due to the effect of the cathode, the trace shown is from  $-90^\circ$  to  $0^\circ$ , mirrored about the centreline to cover  $0^\circ$  to  $90^\circ$ .




FIGURE 6.17: Current efficiency against anode power grouped by discharge voltage as measured from the Faraday probe.

The beam efficiency is calculated from the measured beam current obtained via Faraday probe traces using Eq. 6.14 and the corresponding power supply current delivered to the anode. The results are shown in Fig. 6.17. It is evident that beam efficiency is high during low-power operation at 200 V and 400 V anode voltages but decreases as anode power increases. In contrast, for 300 V and 500 V, beam efficiency exhibits a positive correlation with increasing anode power.

Above approximately 1600 W, the beam efficiency becomes relatively constant for 300 V and 400 V. This behaviour is likely linked to the trend observed in Fig. 6.16, where increasing mass flow rate, and thus increasing anode power, leads to a broader beam. This suggests reduced plasma confinement but also diminished plasma–wall interactions. As a result, beam efficiency remains stable at higher power levels, since further increases in flow rate and voltage primarily raise the overall current density without significantly altering its spatial distribution. The same trend is likely to be true for 500 V and 600 V but would occur at higher powers, as only 40 sccm and 50 sccm were tested up to the high voltages during these tests.

#### 6.4.4 Divergence Efficiency

Beam divergence efficiency is a measure of the efficiency loss due to the non-axial nature of ion acceleration as described by Eq. 2.31. This efficiency value is obtained purely from the Faraday probe measurements as described in Sec. 6.3.3. To evaluate the beam divergence efficiency, the beam divergence half-angle must first be found from Eq. 2.32 as the inverse cosine of the ratio of the axial current to the total beam current.




FIGURE 6.18: Beam divergence angles against anode power grouped by discharge voltage.

The measured beam divergence from Eq. 2.32 is illustrated in Fig. 6.18. As seen, there is a positive correlation between increasing power and beam divergence angle. This trend is likely a result of the same phenomenon described in Sec. 6.4.3, where increasing anode power for the same anode voltage, increasing anode flow rate, results in broader current density distributions and a less-axial current. There is also a local minimum for 400 V and 500 V anode voltages, both occurring at 50 sccm, which corresponds to the flow rate with the most generally efficient operation.




FIGURE 6.19: Divergence efficiency against anode power grouped by discharge voltage.

From the beam divergence angle, the beam divergence efficiency can be obtained and is illustrated in Fig. 6.19. This effectively shows the same trends but inverted, as increased beam angle results in decreased efficiency. Again, more efficient operation is seen at high anode voltages corresponding to lower mass flow rates; this follows as the lower flow rates result in improved plasma confinement within the channel and decreased plume divergence.

#### 6.4.5 Voltage Efficiency

Voltage efficiency is a measure of how effectively the applied anode voltage is being converted into an accelerating potential that the ions experience. As Hall effect thrusters, unlike gridded ion thrusters, do not directly apply the acceleration voltage but rather create a virtual grid with confined electrons the potential drop experienced by the ions during acceleration is not exactly equal to the applied voltage. Furthermore, the process of ionisation and acceleration are overlapping within a Hall thruster discharge such that each ion depending on where it is created can experience a variation of the applied accelerating potential difference.

The voltage efficiency as defined in Eq. 2.30 describes the ratio of the average acceleration voltage experienced by the ions and the applied anode voltage. To measure this a retarding potential analyser as described in Sec. 6.3.4 is used to estimate the most probable voltage that accelerated the ion population from an ion energy distribution function.

The results of the voltage efficiency for the testing are illustrated in Fig. 6.20. As seen, there is a clear positive trend with increasing voltage efficiency with increasing anode




FIGURE 6.20: Voltage efficiency against anode power grouped by discharge voltage.

voltages, increasing power, and increasing anode flow rate. However, this trend is not a result of increasing anode voltage’s reducing the impact that a constant voltage loss,  $\Delta$ , has on the voltage efficiency. Rather, as can be seen in Fig. 6.21, the voltage loss,  $\Delta = U_d - U_b$ , for increasing flow rates and increasing anode voltages reduces by up to approximately 30 V for 40 sccm to 70 sccm at 400 V.




FIGURE 6.21: Measured voltage loss against discharge voltage grouped by anode flow rate.

In Hall thrusters it is known that a portion of the voltage is lost to pull electrons from the cathode into the channel, on the order of tens of volts [81]. This phenomenon should account for a portion of the voltage loss, but it is unknown if this is constant for a range of discharge voltages with the cathode for these tests operating at constant 18.4 V keeper voltage.

### 6.4.6 Mass Efficiency

Mass efficiency, also known as mass utilisation, is a measure of how much of the neutral mass delivered to the thruster is converted into ion current. This metric is particularly important when considering alternative propellants to xenon, such as krypton, due to their smaller ionisation cross-sectional areas and higher ionisation energies. These properties result in reduced ionisation rates for equivalent discharge conditions. Mass efficiency, as described in Eq. 2.28, relies on several far-field probe measurements: the beam current is obtained from the Faraday probe, and the average charge is obtained from the “ $E \times B$ ” probe, as described in Sec. 6.3.3 and Sec. 6.3.2, respectively.




FIGURE 6.22: Mass efficiency against anode mass flow rate grouped by discharge voltage.

Mass efficiency is traditionally the lowest among the individual efficiencies of a Hall thruster. This is due to several factors, one of which is that beam efficiency is inherently included in the definition of mass efficiency fundamentally limiting the maximum mass efficiency to the current efficiency. Furthermore, the ionisation dynamics within a Hall effect thruster play a crucial role in determining mass efficiency, as the mean free path for ionisation must be controlled in efficient designs to ensure it is smaller than the characteristic length scales of the thruster.

The measured mass efficiency of the SHARK-600V thruster operating on krypton can be seen in Fig. 6.22. Here it can be seen that there is a positive trend of increasing mass efficiency with increasing mass flow rate, which plateaus after 70 sccm. There is also a trend of higher mass efficiency at higher voltages for the same mass flow rate. This follows conventional understanding, as increased anode voltage increases the bulk plasma temperature, enhancing the ionisation rate.

### 6.4.7 Anode Efficiency Comparison

Comparing the anode efficiency obtained from both the multiplicative result of the individual efficiencies and the direct thrust measurements provides strong empirical justification for the assumptions made in the probe analysis. While direct thrust measurements cannot resolve specific efficiency losses, they offer a reliable and straightforward method to evaluate thruster performance, being subject to fewer assumptions and sources of error than probe-based methods.

The probe-derived anode efficiency and the anode efficiency obtained from direct thrust measurements are shown in Fig. 6.23. As can be seen, there is relatively good agreement, with the overall trends preserved across the two methods. The largest discrepancies occur at low powers. For the 200 V case, the probes overestimate performance, as the Faraday probe measurements of the plume suggest high current and mass efficiencies. For the 300 V case, the mass efficiency, illustrated in Fig. 6.22, is relatively low and potentially underestimated. These results may be subject to larger-than-normal errors due to weaker signals in low-power plasmas, which would have the secondary effect of increasing sensitivity to corrections applied to small signals, particularly for Faraday probe readings.




FIGURE 6.23: A comparison of anode efficiency as measured with direct thrust measurements (blue) and inferred from probe measurements (red) against anode power.

The performance of the SHARK-600V thruster operating at the University of Michigan in the LVTF demonstrated high efficiency for a krypton discharge, with a peak anode efficiency, as measured by direct thrust measurements, of  $58.1 \pm 3.5\%$ .




FIGURE 6.24: Comparison of thrust as measured by direct thrust measurements from Michigan testing and Southampton testing.

### 6.4.8 Comparison to Southampton Results

The thruster performance recorded at the University of Southampton, as presented in Sec. 5, is here compared with the results obtained at the University of Michigan for the same configuration. This comparison enables the Michigan results to be interpreted within the broader context of the Southampton campaign.

Two datasets from the Michigan campaign are available for comparison with the Southampton results. The first is the direct thrust measurement obtained from the inverted pendulum thrust balance, which can be directly compared to the thrust measurements from Southampton. The second is the anode efficiency derived from probe diagnostics, which will be compared to the anode efficiency calculated from the Southampton direct thrust data.

While good agreement was observed between the probe-based and direct thrust measurements at Michigan, additional insight can be gained through a comparison of these results with those obtained at Southampton.

The direct thrust measurements from both test sites, compared against anode discharge power, are shown in Fig. 6.24. For the same anode flow rate, the Michigan test results indicate higher thrust at slightly lower discharge power. The anode voltages operated at during both test campaigns were the same. The lower power observed at Michigan is attributed to the tests drawing lower anode currents under the same operating conditions. This discrepancy is more clearly illustrated in Fig. 6.25, where thrust is plotted against discharge current.

Here, the trend of slightly enhanced performance at lower discharge current between the two test sites is evident. It is difficult to attribute this difference to a specific cause due to the absence of plume diagnostics at the Southampton test site. However, by




FIGURE 6.25: Comparison of direct thrust measurements against anode discharge current from the Michigan and Southampton testing.




FIGURE 6.26: Comparison of measured anode current difference between Michigan and Southampton testing as a function of anode voltage. Note: negative values indicate a lower anode current measured at Michigan than at Southampton.

comparing the current difference between test sites in Fig. 6.26, a trend with flow rate and the current discrepancy becomes apparent. This suggests that the observed difference is likely due to background pressure variations between the two facilities.

Traditionally, higher background pressures are understood to artificially enhance performance through two primary mechanisms. Firstly, elevated background pressure increases the number of neutrals available for ingestion into the plasma, effectively providing “free” additional anode mass flow, which results in higher thrust and mass efficiency. Secondly, higher background pressure can compress the plasma, shifting it deeper into the discharge channel. In the literature, this confinement effect is




FIGURE 6.27: Comparison of anode efficiency as measured by direct thrust measurements from University of Michigan test campaign and University of Southampton test campaign.

associated with increased performance by promoting a more compact plasma and higher beam divergence efficiency [221–224].

Conversely, the Southampton test site exhibited higher background pressures, on the order of  $3 \times 10^{-5}$  to  $9 \times 10^{-5}$  mbar, compared to the Michigan test site where pressures were on the order of  $5 \times 10^{-6}$  mbar. Based on conventional understanding, it would therefore be expected that the Southampton performance should appear more artificially enhanced than the Michigan results. However, this is not what is observed. The results in Fig. 6.26, which show that the discrepancy between the two test sites decreases with increasing flow rate, suggest a link between background pressure and plasma dynamics. However, even at higher flow rates where the current discrepancy narrows, the overall performance at the University of Michigan remained significantly higher. The increase in thrust observed at Michigan cannot be attributed to pressure thrust effects resulting from the lower background pressure, as the additional thrust generated by this mechanism is estimated to be only on the order of tens of micro-Newtons, well below the measured differences.

Further evidence is provided by the plume evolution illustrated in Fig. 6.16, where Faraday probe measurements indicate a broadening of the plume and a reduction in peak current density between 60 sccm and 70 sccm. As discussed in Sec. 6.4.3, this behaviour was attributed to higher flow rates increasing the neutral density within the thruster, thereby shifting the plasma formation region further downstream, in some cases extending into the near-field plume.

This observation provides insight into the pressure sensitivity of the SHARK-600V. At the University of Southampton test site, the background pressure was approximately an order of magnitude greater than at Michigan, which would almost certainly drive the plasma to form deeper within the channel. Conventionally, such a shift would be




FIGURE 6.28: Comparison of anode specific impulse as measured by direct thrust measurements from Michigan and Southampton testing.

expected to improve performance by increasing plasma confinement and beam divergence efficiency. In this case, however, it is believed that the elevated pressure, combined with the relatively shallow magnetic field gradient of the SHARK-600V, reduced electron confinement near the anode. This in turn promoted enhanced electron current to the anode, leading to degraded overall thruster performance.

The anode efficiency of the SHARK-600V at both test sites is shown in Fig. 6.27, where a clear separation between the datasets can be observed. However, the difference in anode current draw alone cannot fully account for the increased performance measured at the University of Michigan. Another possible source of the disparity arises from differences in the cathode used between the two facilities. At the University of Southampton, as described in Sec. 5.3, the HCES 5000 heater-less hollow cathode by Intlvac Inc. (Georgetown, Canada) was employed. Although this cathode was initially intended for use at the University of Michigan as well, ignition difficulties led to the substitution of the LaB<sub>6</sub> hollow cathode described in Sec. 6.2.3.

The use of the LaB<sub>6</sub> cathode may have reduced the voltage drop required to extract electrons into the bulk plasma, thereby increasing the thruster's voltage efficiency. This would, in turn, raise the plasma potential available for ionisation and acceleration, enhancing overall performance. Without plume diagnostics such as an RPA, however, this effect cannot be quantified directly. A useful proxy is the specific impulse, which is strongly correlated with the acceleration voltage. As shown in Fig. 6.28, the University of Michigan results demonstrate consistently higher specific impulse than those from Southampton, reaching a peak of 3511.2 s ( $\pm 105.4$  s) at 50 sccm and 600 V, as measured by direct thrust.

Whilst specific impulse could also be enhanced by an increase in mass efficiency it is not possible to draw a solid contribution between the use of different cathodes and the

performance of the thruster.

## 6.5 Conclusion

Access to two large vacuum chambers at leading institutions enabled a unique opportunity to expand the experimental scope of this thesis. The additional University of Michigan campaign provided a critical dataset that complements the results obtained at the University of Southampton, particularly by enabling direct cross-facility comparison of thrust measurements and plume characteristics. At the two facilities the thruster was operated on krypton across the same range of anode voltages 200–600 V, and anode flow rates 40–90 sccm with the same electromagnetic currents of 2 A applied. This allows for a like-for-like comparison for voltage and flow rate performance between these two facilities.

The use of a comprehensive suite of far-field diagnostic probes, including the retarding potential analyser, Faraday probe, and  $E \times B$  probe, delivered insights into plume behaviour that were not accessible at Southampton. These measurements not only strengthened confidence in the SHARK-600V results but also highlighted facility-dependent effects.

At the University of Michigan whilst operating on krypton, the SHARK-600V demonstrated consistently higher general- and peak-performance across operating points. Direct thrust measurements showed greater thrust, 107.3 mN, specific impulse, 3511 s, and anode efficiency of 58.1 % at 600 V. Anode efficiency remained higher than at Southampton, with probe diagnostics and direct thrust results in close agreement, lending confidence to the accuracy of the measurements. Furthermore, a clear trend was observed of increased thrust at lower discharge current draw, indicating improved current utilisation and voltage efficiency in the Michigan facility. These results via both direct thrust measurements and indirect efficiency from far-field probes show that the SHARK-600V Hall effect thruster is capable of operating at high efficiency on krypton and at high voltages, providing validation for the key goals of this thesis.

Whilst there are differences in the magnitude of performance measured at the two facilities, the general trends are preserved. Increasing discharge voltage at both test sites leads to increased specific impulse. Furthermore, the highest anode efficiency measured by direct thrust measurements at both facilities occurred at an anode voltage of 600 V and an anode mass flow of 50 sccm of krypton.

However, the disparity in performance measurements between the two test sites asks interesting questions regarding facility effects and cathode choice on performance. Cathode optimisation can significantly improve performance by increasing the

voltage efficiency of the thruster. While this alone cannot account for the magnitude and nature of the performance change, it certainly played a role. One source of difference between the two test sites was the current draw of the anode during operation. The anode current draw at the University of Michigan was lower for all anode flow rates and anode voltages than at Southampton, with the sole exception of the 90 sccm case (Fig. 6.26). This in turn decreased the anode power for the same condition, and this drop in power coincided with an increase in thrust produced. This thrust increase cannot be explained by an increase in pressure thrust due to the low chamber background pressures seen at the University of Michigan. However, the use of a higher quality hollow cathode would produce a greater plasma potential for accelerating and ionising the krypton propellant, slightly enhancing thrust.

Additional insight was provided by the plume diagnostics at Michigan. Faraday probe sweeps revealed a broadening of the plume and reduced peak current density at higher flow rates, suggesting that increased neutral density shifted the plasma formation region downstream. These results emphasise the sensitivity of the SHARK-600V to background pressure and neutral dynamics and highlight the value of comprehensive plume diagnostics in interpreting performance trends.

At this stage, the answers to these questions are still partially speculative due to the lack of a comprehensive diagnostic suite at the University of Southampton that could diagnose the key efficiency loss mechanisms. It is speculated that the lower chamber background pressure seen at the University of Michigan test site resulted in the plasma forming closer to the exit and enhanced electron confinement, thereby enhancing the overall performance of the thruster.

## Chapter 7

# Summary of Key Contributions, Conclusion and, Future Work

### 7.1 Summary of Major Contributions

This thesis makes the following contributions to advancing the understanding and development of Hall effect thrusters operating on alternative propellants:

#### Scaling and Design

An extended semi-empirical scaling methodology was derived that preserves propellant-specific behaviour and enables direct geometric scaling for non-xenon propellants, notably krypton. The method builds on existing semi-empirical approaches, utilising a curated thruster database and yielding practical design targets for mean diameter, channel width and operating points to guide Hall effect thruster design.

Using the scaling outputs as a guideline, the SHARK-600V was designed as a modular Hall effect thruster with interchangeable channel inserts to vary channel width and mean diameter while keeping magnetic, electrical, fluidic and diagnostic systems constant. The design process combined low- and high-fidelity magnetic modelling, laboratory magnetic-field topology mapping, an additively manufactured hollow anode featuring “swirl” injection with print quality validated via X-ray CT, and a zero-dimensional thermal model that accounts for secondary electron emission from the discharge channel for estimating thermal loads. This integrated scaling-and-design outcome provided a platform for the thesis goals of systematic geometry–propellant characterisation.

The scaling methodology received subsequent experimental validation through comparison with the Gaussian process regression model trained on the experimentally obtained dataset of the SHARK-600V operating on xenon, krypton, and argon. Despite the fundamental differences between the two approaches, the scaling predictions and GPR-identified optimal geometries showed meaningful agreement, most notably for krypton at the 2.5 kW, 600 V condition where the predicted mean diameters agreed to within 0.6 mm. This convergence serves as a form of experimental validation for the scaling approach, demonstrating that the semi-empirical method, despite its reliance on a limited krypton database, produced channel dimensions consistent with those identified through systematic experimental characterisation of a modular thruster.

### **Development of Testing Infrastructure and Data Acquisition**

A complete, multi-propellant delivery system was designed and assembled to support rapid and repeatable switching between xenon, krypton and argon, with the capability to test mixed propellants over a range of mass flow rates. A novel inverted double-pendulum thrust balance was designed, manufactured and calibrated. The thrust stand was validated in an inter-laboratory campaign utilising a micro-class thruster, demonstrating the balance's ability to measure a large range of thruster masses and thrust levels with high precision and repeatability.

Complementing the fluidic hardware and the thrust balance, a data-acquisition and control system with an operator GUI was developed to enable continuous measurement and remote operation. This system substantially improved measurement repeatability and enabled robust error analysis of the low-frequency dynamics of thruster operation. Together, the fluidic, thrust measurement and DAQ developments established a laboratory capability suitable for comprehensive Hall effect thruster testing and characterisation.

### **Full Characterisation of SHARK-600V**

The SHARK-600V was experimentally characterised across nine channel geometries and three noble gases, producing a systematic dataset of 1220 operating points from which specific impulse, anode efficiency and thrust-to-power were derived using direct thrust measurements. The results identify channel width as a first-order control on performance and stability, mean channel diameter as a second-order control that becomes important at high voltage, and the "+2, +1" and "+2, +0" geometries as high-performing configurations. These findings provide propellant-dependent guidance for optimising alternative-propellant Hall effect thruster performance.

The experimental characterisation also explored a thruster with anode-layer variant implemented within the same modular platform. TAL operation did not demonstrate clear gains in plasma temperature or specific impulse in this implementation; instead, it revealed severe thermal loading, narrow operating windows and lower anode efficiencies. This suggests TAL operation requires a bespoke thruster architecture and is not well suited to a dual SPT-TAL modular platform.

### **Gaussian Process Regression Modelling**

A Gaussian process regression model was trained from the experimental dataset, creating a continuous, smooth surrogate of SHARK-600V performance across the input parameter space of anode voltage, anode mass flow rate, channel width and mean channel diameter. Independent models were trained for each propellant and each output parameter, providing a powerful analysis tool capable of predicting thruster behaviour with quantified uncertainty. Validation against held-out test data confirmed low prediction errors across all outputs, with coefficient of variation of root mean square error values of 14.8%, 11.9% and 15.6% for anode efficiency for xenon, krypton and argon, respectively.

The GPR model enabled two significant advances. First, it allowed interpolation between tested geometries and extrapolation to untested conditions, identifying optimal channel dimensions that were not directly examined during the experimental campaign. Second, comparison with the semi-empirical scaling methodology revealed meaningful agreement, most notably for krypton at the 2.5 kW, 600 V condition where the predicted mean diameters agreed to within 0.6 mm. This convergence serves as a form of experimental validation for the scaling approach, while demonstrating that the GPR model successfully captured the underlying physical trends. For argon, where conventional scaling methods cannot be applied due to insufficient literature data, the GPR model provided the first experimentally derived optimal channel geometry, offering a preliminary design target for future argon Hall thruster development despite the inherent limitations of the small argon dataset and single-facility operation.

### **International Cross-facility Testing of SHARK-600V**

The cross-facility testing of the SHARK-600V at the University of Michigan enabled validation of thruster performance as well as additional insights from the comprehensive suite of far-field diagnostics, including an  $E \times B$  probe, a Faraday probe and a retarding-potential analyser, and access to the LVTF with substantially higher pumping capacity and lower base pressure to examine facility effects on performance. The Michigan campaign provided direct thrust measurements complemented by

far-field probe data that corroborated the Southampton trends while revealing substantially higher absolute performance and distinct current–voltage characteristics for like-for-like krypton mass flow rates and anode discharge voltages, exposing significant important facility effects.

These cross-facility results both validate the core geometry and highlight the high-performance capability of high-voltage Hall thrusters operating on krypton. Importantly, although absolute performance at Michigan exceeded that at Southampton for the same configuration, both facilities reached peak anode efficiency at the same operating point: 50 sccm, 600 V. At Michigan the peak was 107.3 mN thrust, 3511 s specific impulse and 58.1 % anode efficiency at an anode power of 3180 W; at Southampton, the peak was 85.8 mN thrust, 2809.1 s specific impulse and 35.8 % anode efficiency at an anode power of 3300.8 W.

## 7.2 Conclusion

This thesis has addressed the critical need to reduce reliance on xenon for electric propulsion by providing a comprehensive framework for the development of high-voltage Hall effect thrusters operating on alternative propellants. Through an integrated approach combining extended scaling methodologies, modular thruster design, systematic experimental characterisation, and advanced surrogate modelling, this work has advanced the understanding of how krypton and argon can be optimized for high-performance space propulsion.

The central finding is that the inherent ionisation limitations of lighter propellants can be overcome through a combination of high-voltage operation and precise geometry tuning. Channel width was identified as the primary geometric control, with an optimal range that balances enhanced performance against high-voltage stability. Mean channel diameter emerged as a secondary, propellant-dependent control that becomes influential at elevated voltages. These insights culminated in the identification of the “+2, +1” and “+2, +0” configurations, of the SHARK-600V thruster, as consistently high-performing, demonstrating that krypton, when optimally configured, can match or even exceed xenon in specific impulse and anode efficiency within the same platform. The TAL experiments, while limited, underscored that anode-layer operation requires a fundamentally distinct design approach, reinforcing the maturity and versatility of the SPT architecture explored in this work.

A major contribution of this work is the experimental validation of the semi-empirical scaling methodology through comparison with a Gaussian process regression surrogate model trained on the extensive experimental dataset. The convergence of these two fundamentally different approaches, particularly the strong agreement for krypton at the 2.5 kW, 600 V design point, confirms the utility of the scaling method as

a design tool. Furthermore, the GPR model itself represents a significant advance, providing a validated framework for predicting performance across untested geometries. For argon, where conventional scaling is impossible due to a lack of literature data, the model offers the first experimentally derived design guidance, filling a crucial gap.

Cross-facility testing at the University of Michigan provided both validation and insight. The SHARK-600V achieved performance among the highest reported for krypton Hall effect thrusters, with 58.1% anode efficiency at 600 V, confirming the potential of the design approach. Crucially, while absolute performance differed due to facility effects, the peak efficiency occurred at the same operating point in both facilities, validating the robustness of the underlying performance trends. The accompanying plume diagnostics identified key efficiency loss mechanisms, such as plume divergence and current efficiency, providing a pathway for targeted future improvements. These results, notably at the 50 sccm, 600 V condition, represent some of the highest efficiencies reported in the literature for a krypton Hall effect thruster, achieved with a channel width of 14.5 mm and a mean diameter of 67.5 mm.

Collectively, these contributions bridge the gap between conventional xenon-based scaling and the needs of next-generation, alternative-propellant thrusters. By decoupling geometric and propellant effects and providing validated design tools, this research lays the groundwork for the adoption of krypton and argon in missions where cost and specific impulse are critical. The substantial experimental dataset, the validated scaling methodology, and the GPR surrogate model together form a coherent foundation for the continued development of high-voltage Hall thrusters.

### 7.3 Future Work

The large disparity in performance between test sites has highlighted significant facility effects on thruster operation. Although full plume diagnostics were obtained at the University of Michigan, none were collected during the University of Southampton campaign because the necessary diagnostics were not available at the time. This has since changed, and we are developing a Hall thruster plume diagnostic setup. These probe measurements can non-invasively diagnose efficiency-loss mechanisms at the University of Southampton facilities and provide important guidance for future Hall thruster experimentation.

Furthermore, although the SHARK-600V campaign examined the impact of altering the discharge chamber, the magnetic circuit was held constant at the macroscopic thruster level, so each channel insert experienced a different local portion of the field. A future thruster in which the magnetic circuit can be adjusted to maintain a constant

topology and magnitude across different channel widths and mean diameters will better isolate the effect of channel geometry on performance.

It is also of interest to study lower-power, sub-kilowatt thrusters, since channel geometry has a large impact in this regime. Moreover, small satellites represent a promising market for lower-cost, alternative-propellant Hall thrusters because they are particularly cost sensitive.

## Chapter 8

# Publication History

This section summarises the journal articles and conference contributions produced during the course of this research.

### Journal Contributions

1. T. F. Munro-O'Brien, C. Ryan, *Performance of a low power Hall effect thruster with several gaseous propellants*, *Acta Astronautica*, 2023.
2. T. F. Munro-O'Brien, M. Ahmed, A. Lucca Fabris, C. Ryan, *Inter-Laboratory Characterisation of a Low-Power Channel-Less Hall-Effect Thruster: Performance Comparisons and Lessons Learnt*, *Aerospace*, Vol. 12, No. 7, p.601, 2025.

### Conference Contributions


1. T. Munro-O'Brien, C. Ryan, *Design, manufacture and testing of a magnetically shielded krypton Hall Effect thruster*, *Space Propulsion 2020*, 2021.
2. T. F. Munro-O'Brien, C. Ryan, R. Dinan, J. Lambert, *Experimental Testing of a High Current Lanthanum Hexaboride Hollow Cathode*, *37th International Electric Propulsion Conference*, p.119, 2022.
3. T. Munro-O'Brien, C. Ryan, *Review of SPT and TAL type Hall thrusters' performance and operation through database analysis*, *APS Annual Gaseous Electronics Meeting Abstracts*, DT3.001, 2023.
4. T. Munro-O'Brien, M. Ahmed, C. Ryan, A. Fabris, *Validation of an adjustable sensitivity inverted pendulum thrust stand with a novel external plasma thruster*, *APS Annual Gaseous Electronics Meeting Abstracts*, DT3.006, 2023.

5. V. G. Tirila, T. F. Munro-O'Brien, B. Negre, A. Wittig, C. Ryan, J. Lambert, *Development of a high power Nuclear Electric Propulsion System for interplanetary missions*, International Astronautical Congress, Baku, Azerbaijan, 2023.
6. T. F. Munro-O'Brien, C. Ryan, *Performance Comparison Between Modular SPT and TAL Type Hall Thrusters Operating on Krypton*, 38th International Electric Propulsion Conference, Toulouse, France, 2024.
7. T. F. Munro-O'Brien, C. Ryan, *High-Power Hall effect Thruster Activity at The University of Southampton*, APS Annual Gaseous Electronics Meeting Abstracts, FT3.001, 2024.
8. T. F. Munro-O'Brien, C. Ryan, *Effect of channel width on the performance of a modular Hall effect thruster operating on Xenon, Krypton, and Argon*, 39th International Electric Propulsion Conference, London, United Kingdom, 2025.


## Appendix A

# Magnetic Field Simulation results


Here we have included more example of the intermediate results of the magnetic field simulations used to optimise the magnetic screen.




A. Parametric configuration 1.



B. Parametric configuration 9.




C. Parametric configuration 10.




D. Parametric configuration 22.


FIGURE A.1: Magnetic field simulations for selected geometry sweep 1 showing magnetic field streamlines for 1A inner coil and 1A outer coil currents.




A. Parametric configuration 1.



B. Parametric configuration 10.



C. Parametric configuration 13.



D. Parametric configuration 14.

FIGURE A.2: Magnetic field simulations for selected geometry sweep 2 showing magnetic field streamlines for 2A inner coil and 2A outer coil currents.



## Appendix B

# Gaussian Process Regression Modelling: Additional Figures

### B.1 GPR Training plots

here we have included additional plots with regards to the relative error of the GPR model to training inputs.




FIGURE B.1: Mean relative error of the Gaussian process regression model for krypton as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (109 for krypton). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models.




FIGURE B.2: Mean relative error of the Gaussian process regression model for argon as a function of the number of training samples used. For each training data size, ten independent models were trained using different randomly selected training sets and evaluated on separate randomly selected test data (35 for argon). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95% confidence interval across the ten models.




FIGURE B.3: Mean relative error of the Gaussian process regression model for krypton as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models.




FIGURE B.4: Mean relative error of the Gaussian process regression model for xenon as a function of the number of optimisation restarts. For each restart value, ten independent models were trained using different randomly selected training sets (comprising 80% of the samples) and evaluated on separate randomly selected test data. Each model was evaluated using the same number of test points (the remaining 20% of the samples). The error bars represent the standard deviation of the relative error, while the shaded region indicates the 95 % confidence interval of the relative error across the ten models.

## B.2 GPR Model Error

The error of the inferred performance from the GPR model results for xenon, krypton, and argon shown in Figs. 5.71–5.73 are shown here.

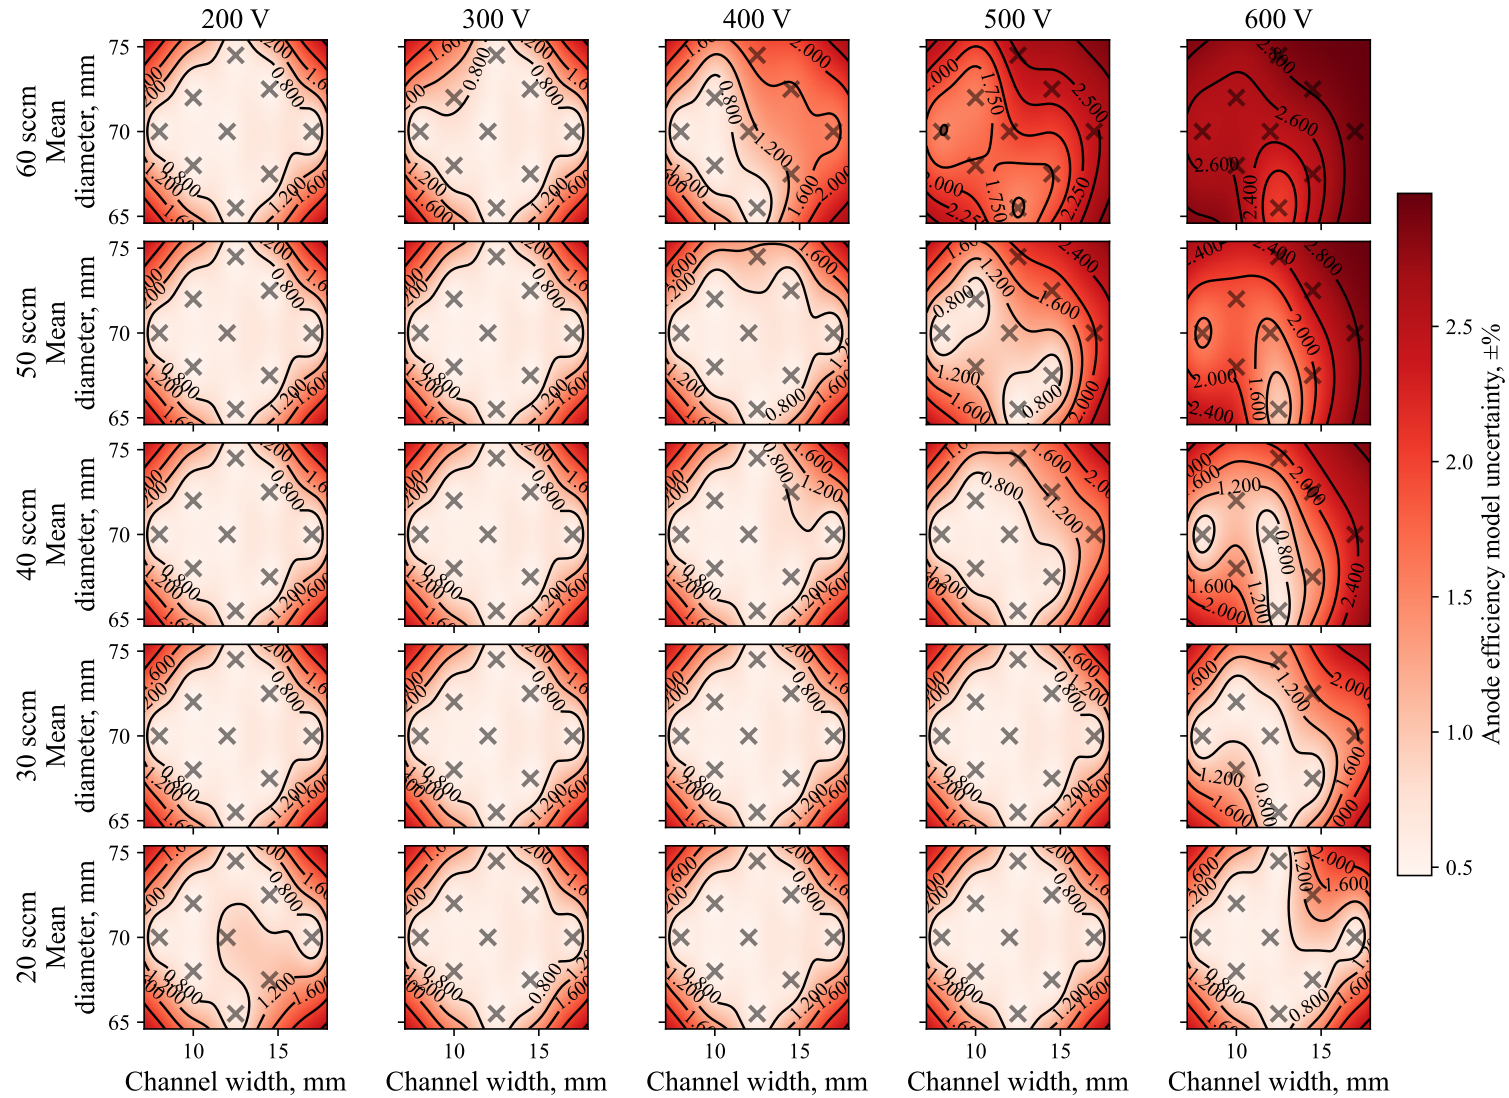


FIGURE B.5: Array of Gaussian process regression model error for xenon from Fig. 5.71, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested.

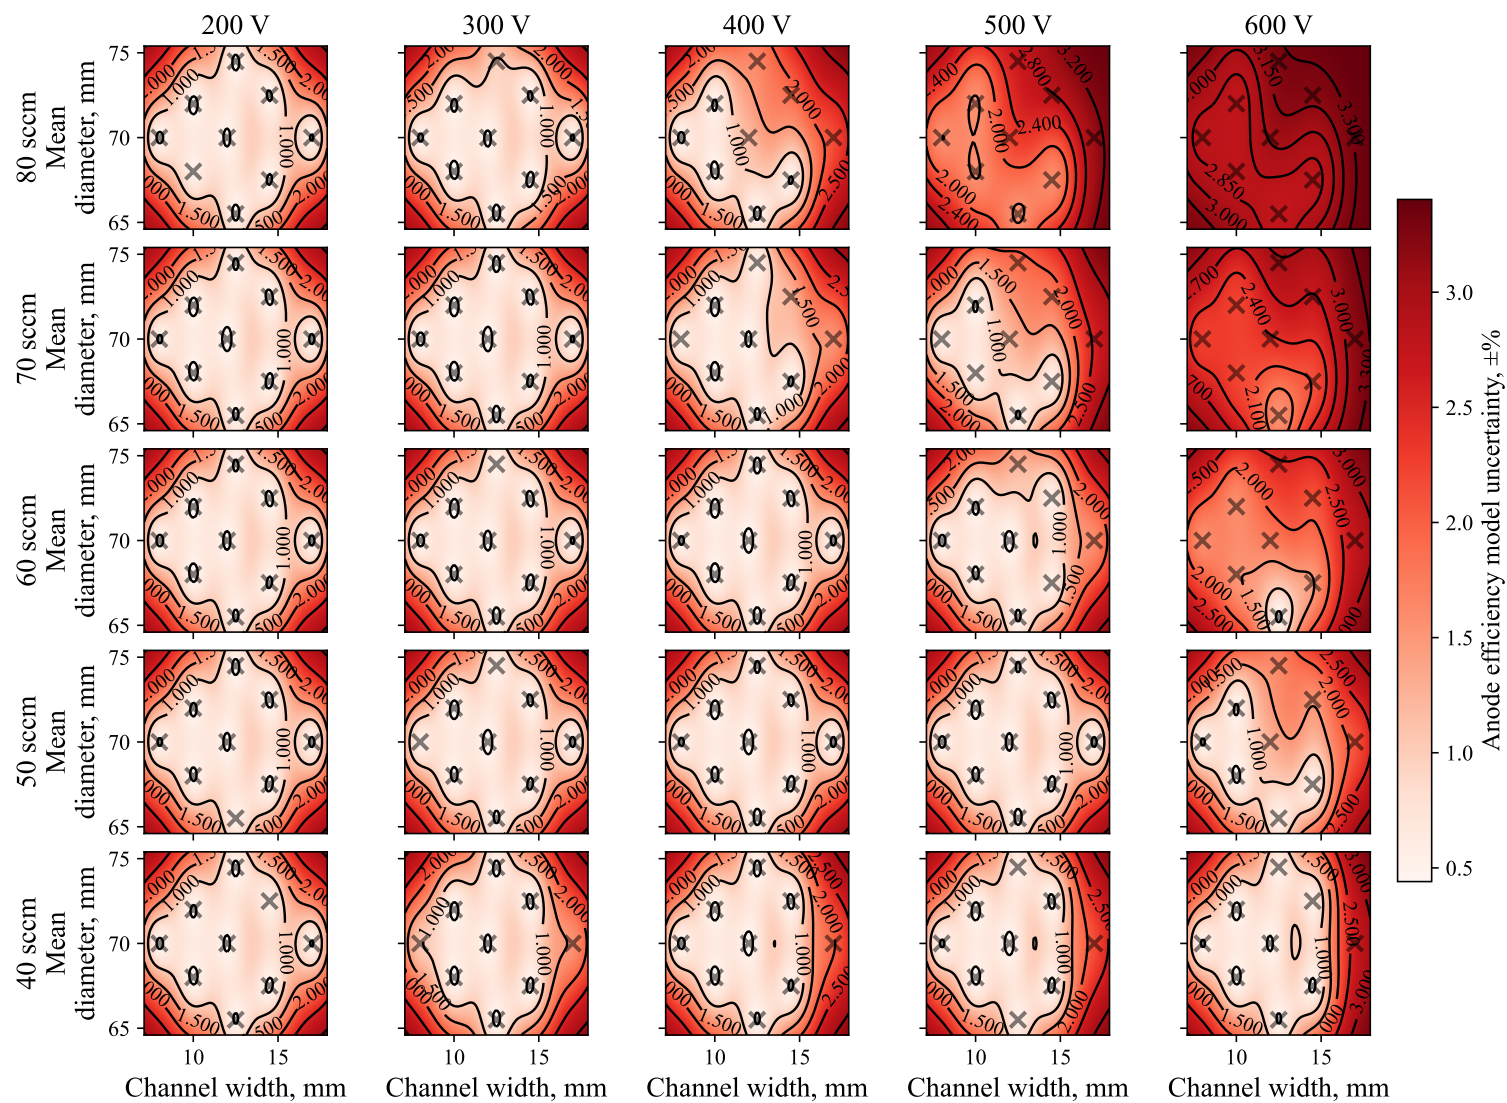


FIGURE B.6: Array of Gaussian process regression model error for krypton from Fig. 5.72, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested.




FIGURE B.7: Array of Gaussian process regression model error for argon from Fig. 5.73, showing the error in the anode efficiency prediction as a function of channel width (horizontal axis) and mean channel diameter (vertical axis) for each voltage and flow rate combination. Each subplot presents a continuous heat map with overlaid iso-lines of constant anode efficiency uncertainty. Red crosses indicate the discrete channel configurations experimentally tested.

## Appendix C

# Michigan Campaign: Uncertainty Analysis & Additional Result

### C.1 Uncertainty Analysis of Far-Field Probes

Error quantification is an integral part of any experimental testing campaign and is especially critical for probe-based diagnostics. Most probes do not directly measure the parameter of interest but instead infer it from secondary quantities. For example, Retarding Potential Analysers,  $E \times B$  probes, and Faraday probes all fundamentally measure current collected on a surface. This current reflects a set of underlying physical processes, such as ion collection, charge separation, or energy filtering. Consequently, final results, such as ion species fractions or ion acceleration voltages, require multiple intermediate steps between raw data and the inferred quantity.

This layered processing introduces challenges in quantifying uncertainty. Specifically, it is difficult to propagate error rigorously through all steps without a full model of each stage. Additionally, since the diagnostics used in this testing campaign were not assembled or calibrated by the author but instead operated as part of the host facility's existing infrastructure, some sources of uncertainty (e.g., power supply resolution, internal probe alignment, or data acquisition fidelity) remain unknown.

As a result, a comprehensive, all-encompassing evaluation of uncertainty is not possible. However, some sources of error are expected to dominate the overall uncertainty budget. This section provides a first-order estimate of those dominant contributions, based on known measurement characteristics and prior literature.

### C.1.1 E×B Probe Error

The E×B probe measurements are inherently subject to uncertainty due to the nature of ion species overlap and the multi-stage data processing required. Only first-order uncertainties are considered here as second order uncertainties are typically two orders of magnitude smaller and negligible in the context of this analysis. The probe signal is calculated from fitted distributions twin-Gaussian models which introduces fitting related uncertainty, particularly in regions where species peaks overlap.

To estimate the contribution of this overlap, the triangular fitting method proposed by Huang et al. [208] was used. While the exact division of current in these regions is not directly measurable, this method provides a geometric upper bound on the uncertainty in the extracted species fractions.

From the twin-Gaussian fitted curves the peak and width are extracted and used to form a geometrically simple triangle. From these triangles the area of the overlap between the peaks is simply described. Following from this the overlap areas for all three species can be found to be:

$$\Delta\Omega_1/\Omega_1 \approx (0.5 \times \frac{s_{12}}{h_1})^2 \quad (C.1)$$

$$\Delta\Omega_2/\Omega_2 \approx (0.5 \times \frac{s_{12}}{h_2})^2 + (0.5 \times \frac{s_{23}}{h_2})^2 \quad (C.2)$$

$$\Delta\Omega_3/\Omega_3 \approx (0.5 \times \frac{s_{23}}{h_3})^2 + (0.5 \times \frac{s_{34}}{h_3})^2 \quad (C.3)$$

where  $\Delta\Omega_n/\Omega_n$  is the relative uncertainty for the  $n^{\text{th}}$  ion species,  $s_{jk}$  is the saddle point of the overlap of the  $j^{\text{th}}$  and the  $k^{\text{th}}$  ion species twin-Gaussians, and  $h_{jk}$  is the height of the saddle point of the overlap of the  $j^{\text{th}}$  and the  $k^{\text{th}}$  ion species twin-Gaussians.

The results from Eqs. C.1–C.3 provide a rough estimate of the upper bound of the uncertainty, as the overlap region will not be composed purely of a single ion species. Using this result as the uncertainty ensures that the actual current fraction lies within the bounds defined by this uncertainty.

Additional uncertainties are introduced into the E×B probe measurement through the applied corrections. As illustrated in Fig. 6.6, the correction factor applied to account for attenuation of ions reaching the probe varies significantly with pressure. This is further emphasised in Fig. C.1, where the calculated attenuation is shown to vary over one order of magnitude of pressure in millibar. This sensitivity is particularly pronounced for singly and doubly charged species, whose impact is increased due to their dominance in the ion population collected by the probe.

These factors present several hurdles to accurately estimating the error of the E×B probe measurement.




FIGURE C.1: Charge exchange ion correction factor versus chamber background pressure for each ion species considered, for a range of discharge voltages. The background neutral temperature is assumed to be 300 K. The dotted vertical line represents  $5.33 \times 10^{-6}$  mbar, the measured pressure during testing.

Firstly, the required charge-exchange correction is strongly dependent on the background pressure. Vacuum pressure measurements are notoriously difficult to obtain with high accuracy, as probe placement, thruster placement, pump configuration, and other factors all play a role. Furthermore, the pressure gauge itself, in the case of the Michigan LVTF, has an error of approximately  $\pm 20\%$  [80, 225]. The charge-exchange correction is also a function of the acceleration voltage, which is determined from the retarding potential analyser and therefore carries the uncertainty associated with that measurement.

Secondly, due to the lower mass of krypton compared with xenon, the peaks in the  $E \times B$  trace are often overlapping, though an upper bound estimate of the error introduced can still be obtained.

Lastly, the signal in an  $E \times B$  probe trace is on the order of nanoamperes; therefore, the current measurement device must be extremely sensitive.

Due to damage to the pressure gauge controller and our unfamiliarity with the setup and requirements for analysing  $E \times B$  probes, the pressure was measured only once, at 60 sccm<sub>USA</sub> total krypton flow into the chamber (50 sccm<sub>USA</sub> anode and 10 sccm<sub>USA</sub> cathode). This single measurement was used for all subsequent analysis and is likely to introduce additional uncertainty into the  $E \times B$  probe measurements, which is difficult to quantify. Based on similar experimental setups, this uncertainty is

estimated to be on the order of  $\pm(10\text{--}15)\%$  as a high estimate, since other facilities have reported errors of approximately  $\pm 8\%$  in charge utilisation measurements [226].

As a result of using a single pressure correction, the attenuation of ions at flow rates greater than 50 sccm<sub>USA</sub> anode is likely under-corrected, lowering the charge efficiency measurement. Conversely, at flow rates less than 50 sccm<sub>USA</sub> anode, the attenuation is likely over-corrected, enhancing the perceived performance. However, given the large pumping capacity of the LVTF, it is unlikely that the pressure varied significantly across the range of flow rates tested.

### C.1.2 Faraday Probe Error

Faraday probe measurements are subject to several sources of error and require specific corrections, as previously discussed in Sec. 6.3.3. This subsection outlines the main corrections applied in this work and the associated sources of uncertainty.

The first consideration is the probe bias, which must be sufficient to reject electron current and ensure that the measured signal represents only collected ions. In the literature, a bias of  $-30$  V is commonly used, and this was the bias applied in these tests.

The second consideration is related to the physical aspects of the probe, such as the ion-collecting region in the gap between the collector and the guard ring. These corrections account for geometry, but uncertainties remain due to manufacturing tolerances or possible misalignments of probe components. In this work, these errors are considered to be at least one order of magnitude smaller than other sources of uncertainty and are therefore treated as negligible.

Two further corrections are applied to the measured signal. The first is the secondary electron emission (SEE) correction, which accounts for electrons ejected from the surface when impinging ions strike. The SEE correction is dependent on material, propellant, and the charge state of the impinging ion. As a result, this value is informed by the  $E \times B$  probe reading for ion species composition. The SEE correction factor,  $\kappa_{\text{SEE}}$ , for each Faraday measurement was approximately 0.90, although it depends on the operating condition. Operating points with a greater proportion of higher charge state species will have a lower  $\kappa_{\text{SEE}}$  factor.

The second applied correction is the charge-exchange correction, which accounts for background neutrals interacting with fast-moving ions and transferring their energy into random low-velocity motion. This interaction contributes to current captured by the Faraday probe at large angles with respect to the thruster centreline. The method employed to account for this is to apply an exponential fit to the near-centreline beam profile and extrapolate outward. This method is susceptible to introducing errors if

the bounds for the fitting region are poorly defined. However, there is no quantitative method for evaluating the error in this corrected current.

Finally, the Faraday probe sweep is performed twice per operational point: once as the arm moves from  $-90^\circ$  to  $+90^\circ$ , and again as the arm moves from  $+90^\circ$  to  $-90^\circ$ . To improve measurement accuracy, these two sweeps are averaged. If the cathode placement did not obscure the portion of the plume in the  $0^\circ$  to  $+90^\circ$  region, as it did in our case, each half of the plume could also be averaged together, assuming a symmetric plume, or compared to provide further confidence in the measurement.

### C.1.3 Retarding Potential Analyser Error

Similar to the other probes used, the RPA also has several sources of error or uncertainty. These include voltage variance on the grids, electrical noise on the collector plate, and possible misalignment of the probe with respect to the thruster plume.

One method used here to quantify the confidence in the acceleration voltage from the RPA trace is bootstrap subsampling. This is a resampling approach used to estimate a confidence range for the extracted most probable voltage (MPV) by repeatedly drawing random subsets of the measurement data and re-performing the MPV extraction procedure.

The process begins by filtering the raw RPA current–voltage data to remove non-finite values and duplicate voltage entries. For each bootstrap iteration, a random subsample containing 80% of the available voltage–current pairs are selected without replacement. This subsample is then smoothed using a moving mean filter of width  $N_{\text{mean}}$ , after which the first derivative,  $-dC/dV$ , is computed to obtain the ion energy distribution function.

Transient artefacts in the derivative trace, often due to electronic noise or signal spikes, are mitigated using a combined spike-detection mask based on the z-score of the derivative and an absolute threshold criterion. Surviving points are median filtered and subsequently passed through a Gaussian filter to further suppress high-frequency noise. A symmetric trimming of the filtered array removes convolution edge effects.

The MPV for each iteration is then identified as the voltage corresponding to the global maximum of the smoothed derivative curve. This process is repeated for  $N_{\text{boot}}$  iterations (here  $N_{\text{boot}} = 1000$ ), producing a set of MPV values. The central tendency and spread of this distribution are reported as the median MPV and its standard deviation, respectively:

$$V_{\text{MPV}} = \text{median}(V_{\text{MPV},i}), \quad \sigma_{\text{MPV}} = \text{std}(V_{\text{MPV},i}) \quad (\text{C.4})$$

This bootstrap subsampling method has the advantage of not requiring any assumptions about the underlying noise distribution and is robust to occasional outliers in the raw dataset. It also naturally incorporates the combined effects of smoothing, filtering, and spike rejection into the reported uncertainty.

The standard deviation calculated from bootstrapping is reported as the uncertainty on the RPA measurement. Whilst this can provide an estimate of the sampling bias error, there are still other sources of error that remain unaccounted for. However, anecdotally, when repeat RPA measurements were taken at the same operational point, the resulting MPV values differed by several volts.

## C.2 Additional Results from Michigan Campaign

The figures in this appendix present supplementary diagnostic results acquired during the University of Michigan campaign.

### C.2.1 $E \times B$ Additional Results

Additional  $E \times B$  probe sweeps recorded during the Michigan campaign for each anode flow rate for all tested voltages.




FIGURE C.2:  $E \times B$  probe data for krypton at 40 sccm for each tested anode voltage.

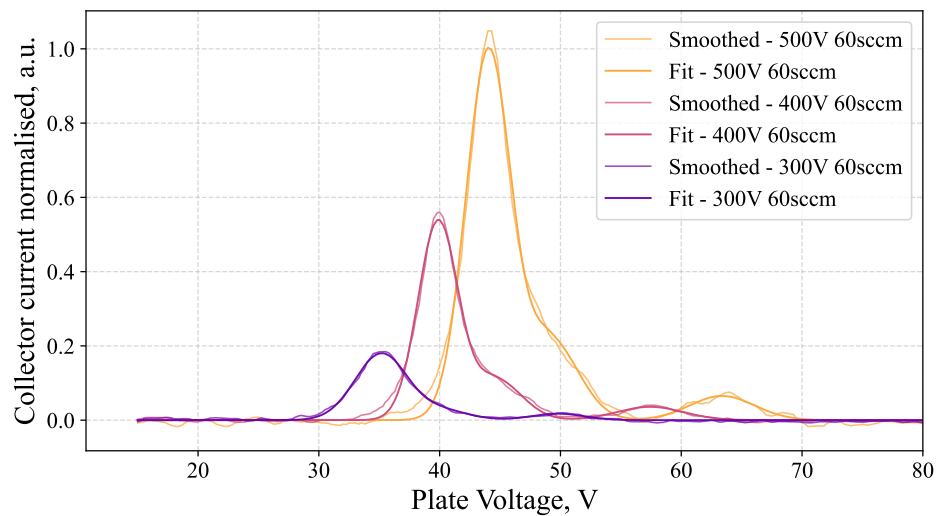


FIGURE C.3:  $E \times B$  probe data for krypton at 60 sccm for each tested anode voltage.




FIGURE C.4: E×B probe data for krypton at 70 sccm for each tested anode voltage.

### C.2.2 Faraday Probe Additional Results

Additional Faraday probe radial sweeps recorded during the Michigan campaign for each anode voltages across each flow rate tested.




FIGURE C.5: Faraday probe radial sweep for krypton at 400 V for all tested flow rates.




FIGURE C.6: Faraday probe radial sweep for krypton at 500 V for all tested flow rates.

### C.2.3 Retarding Potential Analyser Additional Results

Additional RPA sweeps recorded during the Michigan campaign for each anode flow rate across each anode voltage tested.

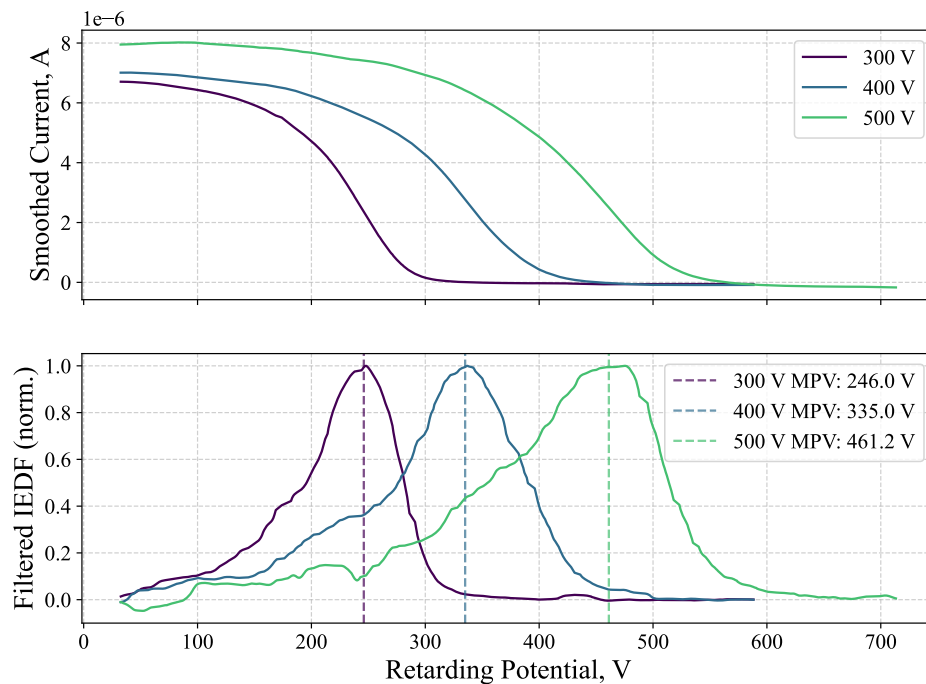


FIGURE C.7: RPA sweep for krypton at 40 sccm across all tested voltages.




FIGURE C.8: RPA sweep for krypton at 50 sccm across all tested voltages.




FIGURE C.9: RPA sweep for krypton at 60 sccm across all tested voltages.




FIGURE C.10: RPA sweep for krypton at 70 sccm across all tested voltages.




FIGURE C.11: RPA sweep for krypton at 300 V across all tested flow rates.



# Bibliography

- [1] A Kosmodemyansky. *Konstantin Tsiolkovsky: His Life and Work*. The Minerva Group, Inc., 2000.
- [2] Edgar Y Choueiri. "A critical history of electric propulsion: The first 50 years (1906-1956)". In: *Journal of propulsion and power* 20.2 (2004), pp. 193–203.
- [3] Igor Levchenko et al. "Space micropropulsion systems for Cubesats and small satellites: From proximate targets to furthestmost frontiers". In: *Applied Physics Reviews* 5.1 (2018).
- [4] Martin JL Turner. *Rocket and spacecraft propulsion: principles, practice and new developments*. Springer Science & Business Media, 2008.
- [5] Jonathan C McDowell. "The Low Earth Orbit Satellite Population and Impacts of the SpaceX Starlink Constellation". In: *The Astrophysical Journal Letters* 892 (2020).
- [6] Dan Lev et al. "The technological and commercial expansion of electric propulsion". In: *Acta Astronautica* 159 (2019), pp. 213–227.
- [7] Lihuan Xie et al. "Effect of plasma initialization on 3d pic simulation of hall thruster azimuthal instability". In: *Physica Scripta* 99.9 (2024).
- [8] Ethan Dale, Benjamin Jorns, and Alec Gallimore. "Future directions for electric propulsion research". In: *Aerospace* 7.9 (2020), p. 120.
- [9] Taisei Nakajima and Akira Kakami. "Performance Evaluation of a Double-Channel TAL-Type Hall Thruster". In: *37th International Electric Propulsion Conference, Boston, USA*. June 2022, p. 354.
- [10] V Zhurin, A Porotnikov, and V Shadov. "Electric propulsion research and development in the USSR". In: *12th International Electric Propulsion Conference*. 1976, p. 1073.
- [11] Shuen-Cheng Hwang, Robert D. Lein, and Daniel A. Morgan. "Noble Gases". In: *Kirk-Othmer Encyclopedia of Chemical Technology*. Vol. 17. American Cancer Society, 2005, pp. 343–383. DOI: <https://doi.org/10.1002/0471238961.0701190508230114.a01.pub2>.

- [12] Thomas Munro-O'Brien and Charles N Ryan. "Performance of a low power Hall effect thruster with several gaseous propellants". In: *Acta Astronautica* 206 (2023), pp. 257–273.
- [13] Daniel A Herman and Kenneth G Unfried. "Xenon Acquisition Strategies for High-Power Electric Propulsion NASA Missions". In: *JANNAF SPS Subcommittee Meeting of the 62nd JANNAF Propulsion Meeting (JPM), Nashville TN*. 2015.
- [14] Melissa L McGuire et al. "NASA GRC Compass Team Conceptual Point Design and Trades of a Hybrid Solar Electric Propulsion (SEP)/Chemical Propulsion Human Mars Deep Space Transport (DST) Vehicle". In: *2018 AIAA SPACE and Astronautics Forum and Exposition*. 2018, p. 5141.
- [15] Clark P Newman et al. "Stationkeeping, orbit determination, and attitude control for spacecraft in near rectilinear halo orbits". In: *American Astronomical Society* (2018).
- [16] Daniel A Herman et al. "Development and Qualification Status of the Electric Propulsion Systems for the NASA PPE Mission and Gateway Program". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022, p. 465.
- [17] KG Unfried et al. "Electric propulsion propellant considerations: can we still use xenon and krypton?" In: *38th International Electric Propulsion Conference, Toulouse, France*. 2024.
- [18] Jason K Hansen et al. *Cost-Benefit Assessment of Krypton and Xenon Recovery from Aqueous Reprocessing*. Tech. rep. Idaho National Laboratory (INL), Idaho Falls, ID (United States), 2023.
- [19] 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*. 2020, p. 3617.
- [20] Jesse Allen Linnell. "An evaluation of krypton propellant in Hall thrusters". PhD thesis. University of Michigan, Jan. 2007.
- [21] Michael R Nakles et al. "A Performance Comparison of Xenon and Krypton Propellant on an SPT-100 Hall Thruster". In: *32nd International Electric Propulsion Conference, Kurhaus, Germany*. 2011.
- [22] AI Bugrova et al. "Experimental investigations of a krypton stationary plasma thruster". In: *International Journal of Aerospace Engineering* 2013 (2013).
- [23] Jacek Kurzyrna et al. "Performance tests of IPPLM's krypton Hall thruster". In: *Laser and Particle Beams* 36.1 (2018), p. 105.
- [24] Jesse A Linnell and Alec D Gallimore. "Efficiency analysis of a Hall thruster operating with krypton and xenon". In: *Journal of Propulsion and Power* 22.6 (2006), pp. 1402–1418.

- [25] Colleen Marrese et al. "An investigation of stationary plasma thruster performance with krypton propellant". In: *31st Joint Propulsion Conference and Exhibit*. 1995, p. 2932.
- [26] Benjamin Vincent, Sedina Tsikata, and Stéphane Mazouffre. "Incoherent Thomson scattering measurements of electron properties in a conventional and magnetically-shielded Hall thruster". In: *Plasma Sources Science and Technology* 29.3 (2020), p. 035015.
- [27] Hani Kamhawi, Wensheng Huang, and Ioannis G Mikellides. "Optimization of the Magnetic Field Topology in the Hall Effect Rocket with Magnetic Shielding". In: *2018 Joint Propulsion Conference*. 2018, p. 4720.
- [28] Junhwi Bak et al. "Plasma formation and cross-field electron transport induced by azimuthal neutral inhomogeneity in an anode layer Hall thruster". In: *Physics of Plasmas* 26.7 (2019).
- [29] EY Choueiri. "Fundamental difference between the two Hall thruster variants". In: *Physics of Plasmas* 8.11 (2001), pp. 5025–5033.
- [30] Richard R Hofer et al. "Effects of internally mounted cathodes on Hall thruster plume properties". In: *IEEE Transactions on Plasma Science* 36.5 (2008), pp. 2004–2014.
- [31] Michael Patterson et al. "NEXT: NASA's evolutionary xenon thruster". In: *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2002, p. 3832.
- [32] J Polk et al. "An overview of the results from an 8200 hour wear test of the NSTAR ion thruster". In: *35th Joint Propulsion Conference and Exhibit*. 1999, p. 2446.
- [33] 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, Boston, USA*. 2022.
- [34] John Snyder et al. "Performance evaluation of the T6 ion engine". In: *Journal of Propulsion and Power* 28.2 (2012), pp. 371–379.
- [35] Kristof Holste et al. "Ion thrusters for electric propulsion: Scientific issues developing a niche technology into a game changer". In: *Review of Scientific Instruments* 91.6 (2020), p. 061101.
- [36] Markus Peukert and Birk Wollenhaupt. "OHB-System's view on electric propulsion needs". In: *EPIC Workshop, Brussels, Belgium*. 2014, pp. 1–48.
- [37] Hiroki Watanabe, Shinatora Cho, and Kenichi Kubota. "Performance and plume characteristics of an 85 W class Hall thruster". In: *Acta Astronautica* 166 (2020), pp. 227–237.
- [38] J Fisher et al. "The Development and Qualification of a 4.5-kW Hall Thruster Propulsion System for GEO Satellite Applications—Status Update". In: *28th International Electric Propulsion Conference, Toulouse, France*. 2003, p. 295.

- [39] J Szabo and Yassir Azziz. "Characterization of a high specific impulse xenon Hall effect thruster". In: *29th International Electric Propulsion Conference, Princeton, USA*. 2005.
- [40] James Szabo et al. "Light metal propellant Hall thrusters". In: *31st International Electric Propulsion Conference, Ann Arbor, USA*. 2009.
- [41] V Hruby et al. "Development of low power Hall thrusters". In: *30th Plasmadynamics and Lasers Conference*. 1999, p. 3534.
- [42] James Szabo et al. "Performance evaluation of an iodine-vapor Hall thruster". In: *Journal of Propulsion and Power* 28.4 (2012), pp. 848–857.
- [43] Michael R Nakles et al. "A plume comparison of xenon and krypton propellant on a 600 W Hall thruster". In: *31st International Electric Propulsion Conference, Ann Arbor, USA*. 2009.
- [44] William A Hargus Jr, Landon J Tango, and Michael R Nakles. "Background pressure effects on krypton Hall effect thruster internal acceleration". In: *33rd International Electric Propulsion Conference, Washington DC, USA*. 2013.
- [45] Gabriel F Benavides et al. "Iodine Hall-effect electric propulsion system research, development, and system durability demonstration". In: *2018 Joint Propulsion Conference*. 2018, p. 4422.
- [46] James Szabo et al. "High density Hall thruster propellant investigations". In: *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2012, p. 3853.
- [47] Lou Grimaud and Stéphane Mazouffre. "Performance comparison between standard and magnetically shielded 200 W Hall thrusters with BN-SiO<sub>2</sub> and graphite channel walls". In: *Vacuum* 155 (2018), pp. 514–523.
- [48] David Jacobson and Robert Jankovsky. "Test results of a 200W class Hall thruster". In: *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. 1998.
- [49] Guojun Xia et al. "Performance optimization of a krypton Hall thruster with a rotating propellant supply". In: *Acta Astronautica* 171 (2020), pp. 290–299.
- [50] Mikhail B Belikov et al. "Development of low-power Hall thruster with lifetime up to 3000 hours". In: *30th International Electric Propulsion Conference, Florence, Italy*. 2007, pp. 2007–129.
- [51] Ryan W Conversano et al. "Performance analysis of a low-power magnetically shielded Hall thruster: experiments". In: *Journal of Propulsion and Power* 33.4 (2017), pp. 975–983.
- [52] James Szabo et al. "Iodine plasma propulsion test results at 1–10 kW". In: *IEEE Transactions on Plasma Science* 43.1 (2014), pp. 141–148.

- [53] AI Bugrova et al. "Integral and spectral characteristics of ATON stationary plasma thruster operating on krypton and xenon". In: *28th International Electric Propulsion Conference, Toulouse, France*. 2003.
- [54] David Manzella et al. "Evaluation of low power Hall thruster propulsion". In: *32nd Joint Propulsion Conference and Exhibit*. 1996.
- [55] Zhong-xi Ning et al. "Research on beam-focusing characteristics of krypton Hall thruster". In: *Plasma Physics Reports* 45.6 (2019), pp. 537–550.
- [56] Stéphane Mazouffre and Lou Grimaud. "Characteristics and performances of a 100-W Hall thruster for microspacecraft". In: *IEEE Transactions on Plasma Science* 46.2 (2018), pp. 330–337.
- [57] B Arkhipov et al. "Small SPT unit development and tests". In: *Spacecraft Propulsion*. Vol. 465. 2000, p. 399.
- [58] Thomas Munro-O'Brien and Charles Ryan. "Design, manufacture and testing of a magnetically shielded krypton Hall Effect thruster". In: *Space propulsion conference 2020*. 7th Edition of the Space Propulsion Conference, 3AF, 2021, p. 402.
- [59] Thomas Malachi Andreano. "Performance and plume characterization of a laboratory krypton Hall thruster". PhD thesis. Colorado State University, 2020.
- [60] James Szabo, Mike Robin, and Vlad Hruby. "Bismuth Vapor Hall Effect Thruster Performance and Plume Experiments". In: *35th International Electric Propulsion Conference, Atlanta, USA*. 2017.
- [61] Jesse A Linnell and Alec D Gallimore. "Internal plasma potential measurements of a Hall thruster using plasma lens focusing". In: *Physics of Plasmas* 13.10 (2006), p. 103504.
- [62] J Kurzyna et al. "Preliminary Tests of HIKHET Laboratory Model at IFPiLM". In: *36th International Electric Propulsion Conference, Vienna, Austria*. 2019, p. 591.
- [63] Richard R Hofer and Robert S Jankovsky. "The influence of current density and magnetic field topography in optimizing the performance, divergence, and plasma oscillations of high specific impulse Hall thrusters". In: *28th International Electric Propulsion Conference, Toulouse, France*. 2003.
- [64] L Biagioni et al. "Design and preliminary characterization of a 5 kW Hall thruster prototype". In: *28th International Electric Propulsion Conference, Toulouse, France*. 2003.
- [65] David Jacobson and David Manzella. "50 kW class krypton Hall thruster performance". In: *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. 2003, p. 4550.
- [66] David H Manzella et al. "Performance Evaluation of the SPT-140". In: *25th International Electric Propulsion Conference, Cleveland, USA*. 1997.

- [67] Richard R Hofer. "Development and characterization of high-efficiency, high-specific impulse xenon Hall thrusters". PhD thesis. University of Michigan, 2004.
- [68] Dan M Goebel and Ira Katz. *Fundamentals of electric propulsion: ion and Hall thrusters*. Vol. 1. John Wiley & Sons, 2008.
- [69] JP Sheehan. "The effects of space-charge limited electron emission on the plasma sheath". PhD thesis. The University of Wisconsin-Madison, 2012.
- [70] Vivien Croes. "Plasma discharge modeling of a Hall-effect thruster". PhD thesis. Ecole Polytechnique, France, 2017.
- [71] Peter Fortescue, Graham Swinerd, and John Stark. *Spacecraft systems engineering*. John Wiley & Sons, 2011.
- [72] RR Lovell and WC Nieberding. *Thrust measurement of SERT I ion thrusters*. Tech. rep. National Aeronautics and Space Administration, 1966.
- [73] Alexander Schwertheim. "The Water Electrolysis Hall Effect Thruster". PhD thesis. Imperial college London, 2022.
- [74] Kristi de Grys et al. "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*. 2010, p. 6698.
- [75] Omri Hamo and Igal Kronhaus. "Experimental and numerical characterization of the narrow channel Hall thruster discharge". In: *Journal of Applied Physics* 130.22 (2021).
- [76] Omri Hamo et al. "Experimental investigation of an external discharge very low anode power (<20 W) Hall thruster". In: *Journal of Electric Propulsion* 1 (2022), p. 112820.
- [77] Benjamin Jorns et al. "Update on the nested Hall thruster subsystem for the NextSTEP XR-100 program". In: *2018 Joint Propulsion Conference*. 2018, p. 4418.
- [78] Scott J Hall et al. "Performance and high-speed characterization of a 100-kW nested Hall thruster". In: *Journal of Propulsion and Power* 38.1 (2022), pp. 40–50.
- [79] Scott J. Hall. "Characterization of a 100-kW Class Nested-Channel Hall Thruster". PhD thesis. University of Michigan, 2018.
- [80] Leanne L Su et al. "Operation and Performance of a Magnetically Shielded Hall Thruster at Ultrahigh Current Densities on Xenon and Krypton". In: *AIAA SCITECH 2023 Forum*. 2023, p. 0842.
- [81] Leanne Su. "Performance of a magnetically shielded Hall thruster operating on krypton at high powers". PhD thesis. University of Michigan, 2023.
- [82] Guido Giammarinaro et al. "A scaling methodology for high-power magnetically shielded Hall thrusters". In: *Journal of Electric Propulsion* 2.1 (2023), p. 17.

- [83] Erica Lopedote, Mario Panelli, and Francesco Battista. "Scaling of Magnetic Circuit for Magnetically Shielded Hall Effect Thrusters: E. Lopedote et al." In: *Aerotecnica Missili & Spazio* 102.2 (2023), pp. 109–125.
- [84] John S Snyder and Richard R Hofer. "Throttled performance of the SPT-140 Hall Thruster". In: *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. 2014, p. 3816.
- [85] S. Mazouffre, K. Dannenmayer, and C. Blank. "Impact of discharge voltage on wall-losses in a Hall thruster". In: *Physics of Plasmas* 18.6 (June 2011), p. 064501. DOI: 10.1063/1.3592251.
- [86] Richard R Hofer et al. "The H10 high power density Hall thruster". In: *Journal of Electric Propulsion* 4.1 (2025), pp. 1–26.
- [87] Maciej Jakubczak et al. "Experimental Optimization of Small Krypton Hall Thruster for Operation at High Voltage". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022.
- [88] AE Solodukhin et al. "Parameters of D-80 anode layer thruster in one-and two-stage operation modes". In: *27th International Electric Propulsion Conference, Pasadena, California, USA*. 2001.
- [89] SO Tverdokhlebov and VI Garkusha. "High-Voltage Mode of a TAL Thruster Operation". In: vol. 23. 1997.
- [90] A Semenkin. "External anode layer thruster parameters at a high specific impulse regimes". In: *25th International Electric Propulsion Conference, Cleveland, USA*. 1997.
- [91] Naoji Yamamoto, Kimiya Komurasaki, and Yoshihiro Arakawa. "Discharge current oscillation in Hall thrusters". In: *Journal of propulsion and power* 21.5 (2005), pp. 870–876.
- [92] James Szabo et al. "Measurements of a krypton fed 1.5 kW Hall effect thruster with a centrally located cathode". In: *35th International Electric Propulsion Conference, Atlanta, USA*. 2017, p. 26.
- [93] James J Szabo et al. "High throughput 600 watt Hall effect thruster for space exploration". In: *52nd AIAA/SAE/ASEE Joint Propulsion Conference*. 2016, p. 4830.
- [94] James Szabo et al. "Eight Kilowatt Hall thruster system characterization". In: *33rd International Electric Propulsion Conference, Washington DC, USA*. 2013.
- [95] Jack Fisher et al. "The development and qualification of a 4.5 kW Hall thruster propulsion system for GEO satellite applications". In: *27th International Electric Propulsion Conference, Pasadena, California, USA*. 2001.
- [96] A Kapulkin et al. "Camila Hall thruster: New results". In: *32nd International Electric Propulsion Conference, Kurhaus, Germany*. 2011.

- [97] David Jacobson et al. "High voltage TAL performance". In: *37th Joint Propulsion Conference and Exhibit*. 2001, p. 3777.
- [98] Eunkwang Lee et al. "Scaling approach for sub-kilowatt Hall-effect thrusters". In: *Journal of Propulsion and Power* 35.6 (2019), pp. 1073–1079.
- [99] Li Jie et al. "Effects of Hydrogen Partial Pressure within the Vacuum Facility on Plume Focusing of a Hall Thruster". In: *32nd International Electric Propulsion Conference, Kurhaus, Germany*. 2011.
- [100] Maciej Jakubczak, Jacek Kurzyzna, and Arsenii Riazantsev. "Experimental verification of the magnetic field topography inside a small Hall thruster". In: *Measurement Science Review* 21.5 (2021), pp. 150–157.
- [101] Stéphane Mazouffre et al. "Rotating spoke instabilities in a wall-less Hall thruster: Experiments". In: *Plasma Sources Science and Technology* 28.5 (2019), p. 054002.
- [102] Richard Hofer, James Haas, and Alec Gallimore. "Ion voltage diagnostics in the far-field plume of a high-specific impulse Hall thruster". In: *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. 2003, p. 4556.
- [103] MW Crofton et al. "Neutral xenon density in the SPT-140 near-field plume". In: *33rd International Electric Propulsion Conference, Washington DC, USA*. 2013.
- [104] AV Loyan and TA Maksymenko. "Performance investigation of SPT-20M low power Hall effect thruster". In: *30th International Electric Propulsion Conference, Florence, Italy*. 2007.
- [105] Konstantin Korolev. "Hall effect thruster design via deep neural network for additive manufacturing". In: *arXiv e-prints* (2023).
- [106] B Arhipov et al. "Investigation of SPT-200 Operating Characteristics at Power Levels up to 12 kw". In: *25th International Electric Propulsion Conference, Cleveland, USA*. 1997, pp. 1997–132.
- [107] A Olano et al. "Modular Design of A Radial Scaled Hall Thruster for Different Magnetic Configurations". In: *2019 IEEE Pulsed Power & Plasma Science (PPPS)*. IEEE. 2019, pp. 1–4.
- [108] L Zakharenkov et al. "Study of low power TAL characteristics". In: *27th International Electric Propulsion Conference, Pasadena, California, USA*. 2001.
- [109] A Olano et al. "Performance of a 100 Watt Radial Scaled Thruster with Anode Layer". In: *36th International Electric Propulsion Conference, Vienna, Austria*. 2019.
- [110] Dibyesh Satpathy et al. "Investigation of correlation between thrust and anode temperature during transient operation of RAIJIN-66". In: *Journal of Electric Propulsion* 2.1 (2023), p. 5.

- [111] Naoji Yamamoto et al. "Thrust Performance in a 5 kW class anode layer type Hall thruster". In: *Transactions of the Japan Society for Aeronautical and Space Sciences, Aerospace Technology Japan* 14 (2016).
- [112] Andrea Kodys and Edgar Choueiri. "A critical review of the state-of-the-art in the performance of applied-field magnetoplasmadynamic thrusters". In: *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2005, p. 4247.
- [113] Jay Polk et al. "Development of high power lithium magnetoplasmadynamic thrusters for human Mars missions". In: *48th Annual AAS Guidance, Navigation, and Control Conference*. 2025.
- [114] Mariano Andrenucci. "Magnetoplasmadynamic Thrusters". In: *Encyclopedia of Aerospace Engineering*. John Wiley & Sons, Ltd, 2010, pp. 1–22. ISBN: 9780470686652. DOI: <https://doi.org/10.1002/9780470686652.eae118>.
- [115] Zhao Yuanzheng et al. "The ablation characteristics of laser-assisted pulsed plasma thruster with metal propellant". In: *Plasma Science and Technology* 23.10 (2021), p. 104007.
- [116] Mahadevan Krishnan, Katherine Velas, and Simon Leemans. "Metal plasma thruster for small satellites". In: *Journal of Propulsion and Power* 36.4 (2020), pp. 535–539.
- [117] James E Polk et al. "A theoretical analysis of vacuum arc thruster and vacuum arc ion thruster performance". In: *IEEE Transactions on Plasma Science* 36.5 (2008), pp. 2167–2179.
- [118] Jonathan Lun and Craig Law. "Influence of cathode shape on vacuum arc thruster performance and operation". In: *IEEE Transactions on Plasma Science* 43.1 (2014), pp. 198–208.
- [119] Vlad-George Tirila, Alain Demairé, and Charles N Ryan. "Review of alternative propellants in Hall thrusters". In: *Acta Astronautica* 212 (2023), pp. 284–306.
- [120] Thomas L. Hardenburger, Matthew Ennis, and Updated by Staff. "Nitrogen". In: *Kirk-Othmer Encyclopedia of Chemical Technology*. American Cancer Society, 2005, pp. 1–23. DOI: <https://doi.org/10.1002/0471238961.1409201808011804.a01.pub2>.
- [121] William P Brabston et al. "Hall Thruster Performance and Efficiency Analysis of a Molecular Propellant". In: *Journal of Propulsion and Power* (2025), pp. 1–14.
- [122] E Rosati Azevedo et al. "Sizing and preliminary design of a 2-kw water propelled hall effect thruster". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022.
- [123] Kento Shirasu et al. "Demonstration and experimental characteristics of a water-vapor Hall thruster". In: *Journal of Electric Propulsion* 2.1 (2023), p. 11.

- [124] JM Tejada and A Knoll. "A water vapour fuelled Hall Effect Thruster: Characterization and comparison with oxygen". In: *Acta Astronautica* 211 (2023), pp. 702–715.
- [125] Seiya Tsuchikawa et al. "Performance evaluation of a low power Hall thruster with carbon dioxide propellant". In: *Acta Astronautica* 224 (2024), pp. 415–426.
- [126] Shunsuke Nosaka et al. "Hall Thruster System Using Dry Ice Propellant at the Triple Point". In: *Journal of Propulsion and Power* (2025), pp. 1–4.
- [127] Antonio Gurciullo, Andrea Lucca Fabris, and Mark A Cappelli. "Ion plume investigation of a Hall effect thruster operating with Xe/N<sub>2</sub> and Xe/air mixtures". In: *Journal of Physics D: Applied Physics* 52.46 (2019), p. 464003.
- [128] V Hrubby et al. "Fullerene fueled electrostatic thrusters-Feasibility and initial experiments". In: *30th Joint Propulsion Conference and Exhibit*. 1994, p. 3240.
- [129] J. Szabo et al. "Fullerene Propellant Hall Thruster Experiment". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022.
- [130] Thimthana Lee et al. "Experimental validation of a naphthalene-fuelled electrothermal plasma thruster". In: *IEEE Access* (2025).
- [131] Michael A Bretti. "Progress and developments of ultra-compact 10 watt class adamantane fueled hall thrusters for picosatellites". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022.
- [132] Alfio E Vinci, Francesco M Bianchi, and Dmytro Rafalskyi. "Modeling and experimental results of low-power iodine-fed Hall thruster propulsion system". In: *Elastic* 1 (2024), 2P.
- [133] Alfio E Vinci et al. "Demonstrating low-power iodine-fed Hall thruster propulsion system". In: *SmallSat Conference Europe*. 2025.
- [134] Mark J Chagnon. "Bismuth and bismuth alloys". In: *Kirk-Othmer Encyclopedia of Chemical Technology* (2000).
- [135] Vlad-George Tirila. "The investigation of alternative solid propellants in Hall effect thrusters". PhD thesis. University of Southampton, 2024.
- [136] Käthe Dannenmayer and Stéphane Mazouffre. "Elementary scaling relations for Hall effect thrusters". In: *Journal of Propulsion and Power* 27.1 (2011), pp. 236–245.
- [137] Kybeom Kwon. "A novel numerical analysis of Hall effect thruster and its application in simultaneous design of thruster and optimal low-thrust trajectory". PhD thesis. Georgia Institute of Technology, 2010.
- [138] Noah Zachary Warner. "Theoretical and experimental investigation of Hall thruster miniaturization". PhD thesis. Massachusetts Institute of Technology, 2007.

- [139] WA Hargus Jr and MA Cappelli. "Laser-induced fluorescence measurements of velocity within a Hall discharge". In: *Applied Physics B* 72 (2001), pp. 961–969.
- [140] Yassir Azziz. "Experimental and theoretical characterization of a Hall thruster plume". PhD thesis. Massachusetts Institute of Technology, 2007.
- [141] Andrey A Shagayda. "On scaling of Hall effect thrusters". In: *IEEE Transactions on Plasma Science* 43.1 (2013), pp. 12–28.
- [142] R Rejoub, BG Lindsay, and RF Stebbings. "Determination of the absolute partial and total cross sections for electron-impact ionization of the rare gases". In: *Physical Review A* 65.4 (2002), p. 042713.
- [143] Robert C Wetzel et al. "Absolute cross sections for electron-impact ionization of the rare-gas atoms by the fast-neutral-beam method". In: *Physical Review A* 35.2 (1987), p. 559.
- [144] Donald Rapp and Paula Englander-Golden. "Total cross sections for ionization and attachment in gases by electron impact. I. Positive ionization". In: *The Journal of Chemical Physics* 43.5 (1965), pp. 1464–1479.
- [145] A Crowe and JW McConkey. "Dissociative ionization by electron impact. II.  $N^+$  and  $N^{++}$  from  $N_2^+$ ". In: *Journal of Physics B: Atomic and Molecular Physics* 6.10 (1973), p. 2108.
- [146] E Krishnakumar and SK Srivastava. "Cross sections for the production of  $N^+$ ,  $N^{++}$ ,  $N_2^+$  and  $N_2^{++}$  by electron impact on  $N_2^+$ ". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 23.11 (1990), p. 1893.
- [147] Robert J LeRoy. "Spectroscopic Reassignment and Ground-State Dissociation Energy of Molecular Iodine". In: *The Journal of Chemical Physics* 52.5 (1970), pp. 2678–2682.
- [148] E Ahedo and V De Pablo. "Combined effects of electron partial thermalization and secondary emission in Hall thruster discharges". In: *Physics of Plasmas* 14.8 (2007), p. 083501.
- [149] V Pigeon et al. "Plasma sheath material induced dependence due to secondary electron emission". In: *Physics of Plasmas* 27.4 (2020), p. 043505.
- [150] GD Hobbs and JA Wesson. *Heat Transmission through a Langmuir Sheath in the presence of Electron Emission*. Tech. rep. Culham Lab., United Kingdom Atomic Energy Authority, Abingdon (England), 1966.
- [151] Manuel Gamero-Castano and Ira Katz. "Estimation of Hall thruster erosion using HPHall". In: (2005).
- [152] AS Bober, V Kim, AS Koroteyev, et al. "State of Work on Electrical Thrusters in the USSR". In: 1991.
- [153] NA Maslenikov. "Russian Electric Propulsion Seminar". In: *Massachusetts Institute of Technology* (1991).

- [154] Frank Stanley Gulczinski III. "Examination of the structure and evolution of ion energy properties of a 5 kw class laboratory Hall effect thruster at various operational conditions". PhD thesis. University of Michigan, 1999.
- [155] Vadim Khayms and M Martinez-Sanchez. "Design of a miniaturized Hall thruster for microsattellites". In: *32nd Joint Propulsion Conference and Exhibit*. 1996, p. 3291.
- [156] E Ahedo and JM Gallardo. "Scaling down Hall thrusters". In: *28th International Electric Propulsion Conference, Toulouse, France*. 2003, pp. 2003–104.
- [157] J Ashkenazy, S Shitrit, and G Appelbaum. "Hall thruster modifications for reduced power operation". In: *29th International Electric Propulsion Conference, Princeton, USA*. 2005.
- [158] Kimiya Komurasaki, Kenji Mikami, and Daisuke Kusamoto. "Channel length and thruster performance of Hall thrusters". In: *32nd Joint Propulsion Conference and Exhibit*, p. 3194.
- [159] Alexander Schwertheim and Aaron Knoll. "Experimental investigation of a water electrolysis Hall effect thruster". In: *Acta Astronautica* 193 (2022), pp. 607–618.
- [160] JM Tejada and A Knoll. "An oxygen-fuelled Hall Effect Thruster: Channel length, ceramic walls and anode material experimental analyses". In: *Acta Astronautica* 203 (2023), pp. 268–279.
- [161] Yegor V Plyashkov et al. "On Scaling of Hall-Effect Thrusters Using Neural Nets". In: *Journal of Propulsion and Power* 38.6 (2022), pp. 935–944.
- [162] Trevor Lafleur and Pascal Chabert. "Analytical model of a Hall thruster". In: *Physics of Plasmas* 31.9 (2024).
- [163] John E Rotter. "An analysis of multiple configurations of next-generation cathodes in a low power Hall thruster". MA thesis. Airforce institute of Technology, 2009.
- [164] Ryan W Conversano et al. "Magnetically shielded miniature Hall thruster: development and initial testing". In: *33rd International Electric Propulsion Conference, Washington DC, USA*. 2013.
- [165] Vladimir Kim. "Main physical features and processes determining the performance of stationary plasma thrusters". In: *Journal of propulsion and power* 14.5 (1998), pp. 736–743.
- [166] Alberto Olano Garcia, Haibin Tang, and Junxue Ren. "Scaling model for SPT and TAL thrusters". In: *IEEE Transactions on Plasma Science* 48.1 (2019), pp. 86–98.
- [167] Käthe Dannenmayer. "Scaling laws and electron properties in Hall effect thrusters". PhD thesis. Université d'Orléans, 2012.

- [168] Richard R Hofer et al. "The H9 magnetically shielded Hall thruster". In: *35th International Electric Propulsion Conference, Atlanta, USA*. Electric Rocket Propulsion Society. 2017, pp. 2017–232.
- [169] Colleen Marre et al. "Analysis of anode layer thruster guard ring erosion". In: *24th International Electric Propulsion Conference, Moscow, Russia*. 1995.
- [170] Bryan Michael Reid. "The influence of neutral flow rate in the operation of Hall thrusters". PhD thesis. University of Michigan, 2009.
- [171] Haotian Fan et al. "Effects of Gas Supply Direction on the Discharge Characteristics of a Low-power Hall Thruster". In: *Vacuum* (2020), p. 109193.
- [172] Minghao Ding et al. "Performance and plume characteristics of a Hall-effect thruster with asymmetrical gas supply". In: *Vacuum* 190 (2021), p. 110285.
- [173] Yongjie Ding et al. "Effect of vortex inlet mode on low-power cylindrical Hall thruster". In: *Physics of Plasmas* 24.8 (2017).
- [174] Philip J Withers et al. "X-ray computed tomography". In: *Nature Reviews Methods Primers* 1.1 (2021), p. 18.
- [175] S Mazouffre, P Echegut, and M Dudeck. "A calibrated infrared imaging study on the steady state thermal behaviour of Hall effect thrusters". In: *Plasma Sources Science and Technology* 16.1 (2006), p. 13.
- [176] Carl F Book and Mitchell LR Walker. "Effect of anode temperature on hall thruster performance". In: *Journal of Propulsion and Power* 26.5 (2010), pp. 1036–1044.
- [177] Rafael A Martinez, Hoang Dao, and Mitchell LR Walker. "Power deposition into the discharge channel of a Hall effect thruster". In: *Journal of Propulsion and Power* 30.1 (2014), pp. 209–220.
- [178] John Michael Fife. "Hybrid-PIC modeling and electrostatic probe survey of Hall thrusters". PhD thesis. Massachusetts Institute of Technology, 1998.
- [179] John Fife et al. "A numerical study of low-frequency discharge oscillations in Hall thrusters". In: *33rd Joint Propulsion Conference and Exhibit*. 1997, p. 3052.
- [180] Sergey V Irishkov, Oleg A Gorshkov, and Andrey A Shagayda. "Fully kinetic modeling of low-power Hall thrusters". In: *29th International Electric Propulsion Conference, Princeton, USA*. 2005.
- [181] Horatiu C Dragnea et al. "Simulation of a Hall Effect Thruster Using Krypton Propellant". In: *Journal of Propulsion and Power* 36.3 (2020), pp. 1–11.
- [182] Francesco Taccogna et al. "Plasma flow in a Hall thruster". In: *Physics of plasmas* 12.4 (2005), p. 043502.
- [183] Thomas Munro-O'Brien et al. "Inter-Laboratory Characterisation of a Low-Power Channel-Less Hall-Effect Thruster: Performance Comparisons and Lessons Learnt". In: *Aerospace* 12.7 (2025), p. 601.

- [184] Thomas Munro-O'Brien et al. "Validation of an adjustable sensitivity inverted pendulum thrust stand with a novel external plasma thruster". In: *APS Annual Gaseous Electronics Meeting Abstracts*. 2023, DT3-006.
- [185] James E Polk et al. "Recommended practices in thrust measurements". In: *33rd International Electric Propulsion Conference*. 2013.
- [186] Stefano Cesare et al. "Nanobalance: the European balance for micro-propulsion". In: *31st International Electric Propulsion Conference, Ann Arbor, USA*. Vol. 182. 2009.
- [187] Kunning G Xu and Mitchell LR Walker. "High-power, null-type, inverted pendulum thrust stand". In: *Review of Scientific Instruments* 80.5 (2009), p. 055103.
- [188] Kean How Cheah and Kay Soon Low. "Torsional thrust stand for characterization of microthrusters". In: *IOP Conference Series: Materials Science and Engineering*. Vol. 152. 1. IOP Publishing. 2016, p. 012019.
- [189] J Soni and S Roy. "Design and characterization of a nano-Newton resolution thrust stand". In: *Review of Scientific Instruments* 84.9 (2013).
- [190] Thomas W Haag. "Thrust stand for pulsed plasma thrusters". In: *Review of Scientific Instruments* 68.5 (1997), pp. 2060-2067.
- [191] Vlad-George Tirila et al. "Performance investigation of zinc propellant in sub kw class Hall thrusters". In: *37th International Electric Propulsion Conference, Boston, USA*. 2022.
- [192] Ugur Kokal, Emin Saridede, and Murat Celik. "Development and tests of a thrust stand with an in-situ null position adjustment and calibration method for low power plasma thrusters". In: *Results in Engineering* 18 (2023), p. 101219.
- [193] Brett Robert Tartler. "Construction and performance of an inverted pendulum thrust balance". MA thesis. Massachusetts Institute of Technology, 2010.
- [194] Vladimir Kim et al. "History of the Hall thrusters development in USSR". In: *30th International Electric Propulsion Conference, Florence, Italy*. 2007.
- [195] Yushi Hamada et al. "Characterization of acceleration zone shifting in an anode-layer-type Hall thruster RAIJIN66". In: *Vacuum* 186 (2021), p. 110040.
- [196] DV Dukhopelnikov, SO Shilov, and DK Alekseev. "Comparison of the characteristics of the laboratory model of a anode layer thruster operating with krypton and xenon". In: *AIP Conference Proceedings*. Vol. 2318. 1. AIP Publishing LLC. 2021, p. 040005.
- [197] Sridevi Bhat and Ashish Mishra. "Prediction of liner erosion and life estimation of Stationary Plasma Thrusters using Machine Learning". In: *36th International Electric Propulsion Conference, Vienna, Austria*. 2019, pp. 15-20.

- [198] Benjamin Jorns. "Predictive, data-driven model for the anomalous electron collision frequency in a Hall effect thruster". In: *Plasma Sources Science and Technology* 27.10 (2018), p. 104007.
- [199] Masato Kawazu et al. "Prediction of thruster performance in hall thrusters using neural network with auto encoder". In: *TRANSACTIONS OF THE JAPAN SOCIETY FOR AERONAUTICAL AND SPACE SCIENCES, AEROSPACE TECHNOLOGY JAPAN* 19.5 (2021), pp. 760–765.
- [200] Peter Thoreau et al. "Rapid thruster-in-the-loop optimization for Hall thrusters". In: *Journal of Electric Propulsion* 4.1 (2025), p. 56.
- [201] Peter Thoreau. "Advanced Techniques for Hall Thruster Research and Development". PhD thesis. University of Washington, 2025.
- [202] Peter Thoreau et al. "Surrogate modeling and real-time optimization of propellant mixtures for hall thrusters". In: *Journal of Electric Propulsion* 5.1 (2026), p. 14.
- [203] JP Chiles and P Chauvet. "Kriging: a method for cartography of the sea floor". In: *The International Hydrographic Review* (1975).
- [204] Eric Schulz, Maarten Speekenbrink, and Andreas Krause. "A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions". In: *Journal of mathematical psychology* 85 (2018), pp. 1–16.
- [205] Accessed: 2025-05-26. URL: <https://pepl.engin.umich.edu/about/>.
- [206] Eric A Vigés et al. "University of Michigan's upgraded large vacuum test facility". In: *36th International Electric Propulsion Conference, Vienna, Austria*. 2019, pp. 1–18.
- [207] Sarah E Cusson et al. "Performance of the H9 magnetically shielded Hall thrusters". In: *35th International Electric Propulsion Conference, Atlanta, USA*. 2017.
- [208] Wensheng Huang and Rohit Shastry. "Analysis of Wien filter spectra from Hall thruster plumes". In: *Review of Scientific Instruments* 86.7 (2015).
- [209] Rohit Shastry et al. "Method for analyzing  $E \times B$  probe spectra from Hall thruster plumes". In: *Review of Scientific Instruments* 80.6 (2009), p. 063502.
- [210] Michael L Hause, Benjamin D Prince, and Raymond J Bemish. "Krypton charge exchange cross sections for Hall effect thruster models". In: *Journal of Applied Physics* 113.16 (2013).
- [211] Daniel L Brown et al. "Recommended practice for use of Faraday probes in electric propulsion testing". In: *Journal of Propulsion and Power* 33.3 (2017), pp. 582–613.
- [212] Homer D Hagstrum. "Auger ejection of electrons from molybdenum by noble gas ions". In: *Physical Review* 104.3 (1956), p. 672.

- [213] Homer D Hagstrum. "Auger ejection of electrons from tungsten by noble gas ions". In: *Physical Review* 96.2 (1954), p. 325.
- [214] Mansur Tisaev. "Enabling air-breathing electric propulsion for VLEO: Computational analysis of orbital envelope and experimental development of microwave plasma cathode". PhD thesis. University of Surrey, 2024.
- [215] Mitchell Walker, Richard Hofer, and Alec Gallimore. "The effects of nude Faraday probe design and vacuum facility backpressure on the measured ion current density profile of Hall thruster plumes". In: *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2002, p. 4253.
- [216] David Manzella and John Sankovic. "Hall thruster ion beam characterization". In: *31st Joint Propulsion Conference and Exhibit*, p. 2927.
- [217] Ira Katz, Richard R Hofer, and Dan M Goebel. "Ion current in Hall thrusters". In: *IEEE Transactions on plasma science* 36.5 (2008), pp. 2015–2024.
- [218] Seth J Thompson et al. "Combined electrostatic analyzer—Wien filter probe for characterization of species distributions in Hall thrusters". In: *Journal of Applied Physics* 130.23 (2021).
- [219] Sang-Wook Kim and Alec D Gallimore. "Plume study of a 1.35-kW SPT-100 using an ExB probe". In: *Journal of Spacecraft and Rockets* 39.6 (2002), pp. 904–909.
- [220] Bryan Reid et al. "Angularly-Resolved ExB Probe Spectra in the Plume of a 6-kW Hall Thruster". In: *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2008, p. 5287.
- [221] Matthew P Byrne, Benjamin A Jorns, and A Arbor. "Data-driven models for the effects of background pressure on the operation of Hall thrusters". In: *36th International Electric Propulsion Conference, Vienna, Austria*. 2019.
- [222] Kevin D Diamant, Raymond Liang, and Ron L Corey. "The effect of background pressure on SPT-100 Hall thruster performance". In: *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. 2014, p. 3710.
- [223] John Steven Snyder et al. "Effects of background pressure on SPT-140 Hall thruster performance". In: *Journal of Propulsion and Power* 36.5 (2020), pp. 668–676.
- [224] Kevin Diamant et al. "The effects of background pressure on Hall thruster operation". In: *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. 2012, p. 3735.
- [225] John W Dankanich et al. "Recommended practice for pressure measurement and calculation of effective pumping speed in electric propulsion testing". In: *Journal of Propulsion and Power* 33.3 (2017), pp. 668–680.
- [226] Ryan Conversano. "Low-power magnetically shielded Hall thrusters". PhD thesis. UCLA, 2015.